

Oracle® Fusion Middleware

High Availability Guide

11g Release 1 (11.1.1)

E10106-17

December 2011

Documentation for administrators, developers, and others that describes high availability concepts as well as administration and configuration procedures to deploy and manage Oracle Fusion Middleware with high availability requirements.

Primary Author: Christine Ford

Contributing Author: Fermin Castro, Shailesh Dwivedi, Michael Blevins, Pradeep Bhat, Michael Rhys, Samrat Ray, Jeni Ferns, Richard Delval, Bharath K. Reddy, Susan Kornberg, Joe Paul, Ajay Keni, Amit Sharma, Vasuki Ashok, Olaf Stullich, Jingjing Wei, Olfat Aly, Eileen He, Ramaprakash Sathyanarayan, Apurv Chandra, Atika Jain, Buddhika Kottahachchi, Arun Theebaparakasam, Rajiv Jaisankar, Vinay Kalra, Daniel Shih, Pratima Gogineni, Russ Hodgson, Christopher Johnson, Pramodini Gattu, Gururaj BS, Ben Gelernter, Peter Jacobsen.

Contributor: Van Sioung Ng Yan Tun, Eric Cloney, Fiona Zheng, Ruchica Behl, David J. Jones, Ellen Desmond, Don Biasotti, Vinaye Misra, Dhaval Shah, Herbert Stiel, Albert Tam, Pushkar Kapasi, Shail Goel, Rahul Menezes, Satheesh Amilineni, Ratheesh Pai, Jan Carlin.

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

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

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

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

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

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

Contents

Preface	xliii
Intended Audience.....	xliii
Documentation Accessibility	xliii
Related Documentation	xliii
Conventions	xliii
 1 Introduction to High Availability	
1.1 What is High Availability	1-1
1.1.1 High Availability Problems.....	1-1
1.1.2 High Availability Solutions.....	1-2
1.2 High Availability Information in Other Documentation.....	1-5
 2 Oracle Fusion Middleware High Availability Framework	
2.1 Understanding Key Oracle Fusion Middleware Concepts	2-1
2.1.1 What is a WebLogic Server Domain?	2-2
2.1.1.1 What Is the Administration Server?	2-3
2.1.1.2 Understanding Managed Servers and Managed Server Clusters	2-4
2.1.1.3 What Is Node Manager?	2-4
2.1.2 What Is a System Component Domain?.....	2-5
2.1.3 What Is a Middleware Home?	2-5
2.1.4 What Is an Oracle Home?	2-5
2.1.4.1 What Is an Oracle Common Home?	2-5
2.1.5 What Is a WebLogic Server Home?.....	2-5
2.2 Oracle Fusion Middleware High Availability Terminology	2-6
2.3 Oracle Fusion Middleware High Availability Solutions.....	2-8
2.3.1 Local High Availability.....	2-8
2.3.2 Oracle Fusion Middleware High Availability Technologies.....	2-9
2.3.2.1 Server Load Balancing	2-11
2.3.3 Active-Passive Deployment	2-13
2.3.4 About Active-Active and Active-Passive Solutions	2-14
2.3.5 Disaster Recovery	2-16
2.4 Protection from Planned and Unplanned Down Time	2-17
 3 High Availability for WebLogic Server	
3.1 What Is a WebLogic Server Cluster?.....	3-1

3.2	WebLogic Server Clusters and WebLogic Server Domains.....	3-2
3.3	Benefits of Clustering	3-2
3.4	Key Capabilities of a Cluster	3-2
3.4.1	Application Failover.....	3-3
3.4.2	Migration.....	3-3
3.4.3	Load Balancing.....	3-3
3.5	Types of Objects That Can Be Clustered	3-4
3.6	Communications in a Cluster.....	3-4
3.7	Cluster-Wide JNDI Naming Service	3-5
3.8	Failover and Replication in a Cluster.....	3-5
3.8.1	Session Replication	3-6
3.9	Whole Server Migration.....	3-6
3.9.1	Node Manager's Role in Whole Server Migration	3-7
3.9.2	Server Migration Processes and Communications	3-7
3.9.2.1	Startup Process in a Cluster with Migratable Servers.....	3-7
3.9.2.2	Automatic Whole Server Migration Process	3-9
3.9.2.3	Manual Whole Server Migration Process	3-10
3.9.2.4	Administration Server's Role in Whole Server Migration.....	3-11
3.9.2.5	Migratable Server Behavior in a Cluster	3-12
3.9.2.6	Cluster Master's Role in Whole Server Migration	3-12
3.10	JMS and JTA High Availability.....	3-13
3.10.1	User-Preferred Servers and Candidate Servers.....	3-14
3.10.2	Considerations for Using File Stores on NFS	3-14
3.11	Administration Server and Node Manager High Availability	3-17
3.11.1	Administration Server Failure	3-18
3.11.2	Node Manager Failure	3-18
3.12	Load Balancing.....	3-18
3.13	GridLink Data Sources	3-19
3.14	Multi Data Sources.....	3-19
3.15	Cluster Configuration and config.xml.....	3-19
3.16	About Singleton Services	3-20
3.17	WebLogic Server and LDAP High Availability	3-20

4 Considerations for High Availability Oracle Database Access

4.1	Oracle Real Application Clusters and Fusion Middleware	4-1
4.1.1	Java-Based Oracle Fusion Middleware Components Deployed to Oracle WebLogic Server 4-2	
4.1.2	Using GridLink Data Sources with Oracle RAC	4-2
4.1.3	Using Multi Data Sources with Oracle RAC	4-3
4.1.3.1	Configuring Multi Data Sources for MDS Repositories	4-4
4.1.3.2	Oracle RAC Configuration Requirements	4-5
4.1.3.3	Configuring Schemas for Transactional Recovery Privileges.....	4-6
4.1.4	Configuring GridLink Data Sources with Oracle RAC.....	4-6
4.1.5	Configuring Multi Data Sources with Oracle RAC	4-7
4.1.6	JDBC Clients	4-9
4.1.7	System Clients	4-10
4.1.7.1	Oracle Internet Directory.....	4-10

4.1.7.2	Oracle Forms	4-11
4.1.7.3	Oracle Portal.....	4-11
4.1.7.4	Oracle Reports and Oracle Discoverer	4-13
4.2	Protecting Idle Connections from Firewall Timeouts.....	4-13
4.3	Troubleshooting	4-13
4.4	Oracle Fusion Middleware Products are Certified to be Used with 11.2 RDBMS Oracle RAC 4-14	
4.4.1	SCAN Run Time Implications and Limitations	4-15

5 Configuring High Availability for Oracle SOA Suite

5.1	Introduction to Oracle SOA Suite.....	5-2
5.2	Oracle SOA Service Infrastructure High Availability	5-3
5.2.1	Oracle SOA Service Infrastructure Single-Instance Characteristics	5-4
5.2.1.1	Oracle SOA Service Infrastructure Application Characteristics.....	5-5
5.2.1.2	Oracle SOA Service Infrastructure Startup and Shutdown Lifecycle.....	5-5
5.2.1.3	Oracle SOA Service Infrastructure External Dependencies	5-6
5.2.1.4	Oracle SOA Service Infrastructure Startup and Shut Down of Processes	5-6
5.2.1.5	Oracle SOA Service Infrastructure Configuration Artifacts.....	5-7
5.2.1.6	Oracle SOA Service Infrastructure Log File Locations	5-8
5.2.2	Oracle SOA Service Infrastructure High Availability Architecture and Failover Considerations 5-9	
5.2.2.1	Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior 5-12	
5.2.2.1.1	WebLogic Server Crash	5-12
5.2.2.1.2	Node Failure.....	5-12
5.2.2.1.3	Database Failure	5-13
5.2.2.2	Oracle SOA Service Infrastructure Cluster-Wide Deployment	5-13
5.2.2.3	Online Redeployment of Oracle SOA Service Infrastructure Composites in a Cluster 5-13	
5.2.2.4	Oracle SOA Service Infrastructure Cluster-Wide Configuration Changes.....	5-14
5.3	Oracle BPEL Process Manager and High Availability Concepts.....	5-14
5.3.1	Oracle BPEL Process Manager Single-Instance Characteristics	5-14
5.3.1.1	BPEL Process Manager Component Characteristics	5-15
5.3.1.2	Oracle BPEL Process Manager Startup and Shutdown Lifecycle.....	5-16
5.3.1.3	Oracle BPEL Process Manager Request Flow and Recovery	5-17
5.3.1.4	Oracle BPEL Process Manager Configuration Artifacts	5-18
5.3.2	Oracle BPEL Process Manager High Availability Architecture and Failover Considerations 5-19	
5.3.2.1	Oracle BPEL Process Manager Protection from Failures and Expected Behavior..... 5-19	
5.3.2.1.1	Recovering Failed BPEL and Mediator Instances.....	5-21
5.3.2.2	Oracle BPEL Process Manager Cluster-Wide Configuration Changes.....	5-21
5.4	Oracle BPM Suite and High Availability Concepts	5-21
5.4.1	Oracle BPM Suite Single Instance Concepts	5-22
5.4.1.1	Oracle BPM Suite Component Characteristics.....	5-23
5.4.1.2	Oracle BPM Suite Component Interaction.....	5-25
5.4.1.3	Oracle BPM Suite Startup and Shutdown Lifecycle	5-27

5.4.1.4	Oracle BPM Suite Configuration Artifacts.....	5-27
5.4.2	Oracle BPMN Service Engine High Availability	5-28
5.4.2.1	Oracle BPMN Service Engine Single Instance Characteristics.....	5-29
5.4.2.1.1	Oracle BPMN Service Engine Single Instance Architecture	5-29
5.4.2.1.2	Oracle BPMN Service Engine External Dependencies	5-30
5.4.2.1.3	Oracle BPMN Service Engine Startup and Shutdown Lifecycle	5-30
5.4.2.1.4	Oracle BPMN Service Engine Log Files	5-31
5.4.2.2	Oracle BPMN Service Engine High Availability Considerations.....	5-31
5.4.2.2.1	Oracle BPMN Service Engine High Availability Architecture and Failover Considerations	5-31
5.4.2.2.2	Configuring Oracle BPMN Service Engine for High Availability	5-32
5.4.2.2.3	Cluster-Wide Configuration Changes for Oracle BPMN Service Engine ..	5-32
5.4.3	Oracle Business Process Web Applications High Availability	5-33
5.4.3.1	Oracle Business Process Web Applications Single Instance Characteristics	5-33
5.4.3.1.1	Oracle Business Process Web Applications Single Instance Architecture .	5-33
5.4.3.1.2	Oracle Business Process Web Applications External Dependencies	5-33
5.4.3.1.3	Oracle Business Process Web Applications Startup and Shutdown Lifecycle.....	5-33
5.4.3.1.4	Oracle Business Process Web Applications Log Files.....	5-34
5.4.3.2	Oracle Business Process Web Applications High Availability Considerations	5-34
5.4.3.2.1	Oracle Business Process Web Applications High Availability Architecture and Failover Considerations	5-34
5.4.3.2.2	Configuring Oracle Business Process Web Applications for High Availability ..	5-34
5.4.3.2.3	Cluster-Wide Configuration Changes for Oracle Business Process Web Applications	5-34
5.4.4	Oracle Business Process Analytics High Availability	5-34
5.4.4.1	Oracle Business Process Analytics Single Instance Characteristics	5-34
5.4.4.1.1	Oracle Business Process Analytics Single Instance Architecture	5-34
5.4.4.1.2	Oracle Business Process Analytics External Dependencies	5-35
5.4.4.1.3	Oracle Business Process Analytics Startup and Shutdown Lifecycle	5-36
5.4.4.1.4	Oracle Business Process Analytics Log Files	5-36
5.4.4.2	Oracle Business Process Analytics High Availability Considerations	5-36
5.4.4.2.1	Oracle Business Process Analytics High Availability Architecture and Failover Considerations	5-36
5.4.4.2.2	Configuring Oracle Business Process Analytics for High Availability	5-36
5.4.4.2.3	Cluster-Wide Configuration Changes for Oracle Business Process Analytics.....	5-36
5.5	Oracle Mediator and High Availability Concepts	5-36
5.5.1	Oracle Mediator Single-Instance Characteristics	5-36
5.5.1.1	Oracle Mediator Component Characteristics.....	5-37
5.5.1.2	Oracle Mediator Startup and Shutdown Lifecycle	5-37
5.5.1.3	Oracle Mediator Request Flow	5-38
5.5.1.4	Oracle Mediator Configuration Artifacts.....	5-38
5.5.2	Oracle Mediator High Availability Architecture and Failover Considerations	5-39
5.5.2.1	Oracle Mediator Protection from Failures and Expected Behavior	5-39
5.5.2.1.1	Recovering Failed Mediator Instances	5-40
5.5.2.1.2	Oracle Mediator's Resequencing in a Cluster	5-41

5.5.2.2	Troubleshooting Oracle Mediator High Availability	5-42
5.6	Oracle Human Workflow and High Availability Concepts	5-42
5.6.1	Oracle Human Workflow Single-Instance Characteristics	5-42
5.6.1.1	Oracle Human Workflow Startup and Shutdown Lifecycle	5-42
5.6.1.2	Oracle Human Workflow Request Processing	5-43
5.6.1.3	Oracle Human Workflow Configuration Artifacts.....	5-43
5.6.1.3.1	Managing the URI of the Human Task Service Component Task Details Application 5-43	
5.6.2	Oracle Human Workflow High Availability Architecture and Failover Considerations. 5-44	
5.6.2.1	Oracle Human Workflow Protection from Failures and Expected Behavior ...	5-44
5.6.2.2	Manual Recovery Required for Human Workflow Task in Rejected MSG Table	5-44
5.6.3	Troubleshooting Oracle Human Workflow High Availability	5-44
5.7	Oracle B2B and High Availability Concepts	5-45
5.7.1	Oracle B2B Single-Instance Characteristics.....	5-45
5.7.1.1	Oracle B2B Component Characteristics	5-46
5.7.1.2	Oracle B2B Startup and Shutdown Lifecycle.....	5-46
5.7.1.3	Oracle B2B Request Flow.....	5-46
5.7.1.4	Oracle B2B Configuration Artifacts	5-47
5.7.2	Oracle B2B High Availability Architecture and Failover Considerations.....	5-47
5.7.2.1	Oracle B2B Protection from Failures and Expected Behavior.....	5-47
5.7.2.2	Oracle B2B Cluster-Wide Configuration Changes	5-48
5.7.2.3	Oracle B2B Deployments in a Cluster	5-48
5.7.2.4	Troubleshooting Oracle B2B Active-Active Configuration.....	5-48
5.7.2.4.1	Purge, Import, or Deployment of B2B Metadata	5-48
5.7.2.4.2	Error While Retrieving Oracle B2B Document Definitions.....	5-49
5.8	Oracle Web Services Manager and High Availability Concepts	5-49
5.8.1	Oracle WSM Single-Instance Characteristics.....	5-49
5.8.1.1	Oracle WSM Component Characteristics	5-51
5.8.1.2	Oracle WSM Startup and Shutdown Lifecycle.....	5-52
5.8.1.3	Oracle WSM Request Flow.....	5-52
5.8.1.4	Oracle WSM Configuration Artifacts	5-52
5.8.2	Oracle WSM High Availability Architecture and Failover Considerations.....	5-53
5.8.2.1	Oracle WSM Protection from Failures and Expected Behavior.....	5-53
5.8.2.2	Oracle WSM Cluster-Wide Configuration Changes	5-54
5.8.2.3	Configuring the Java Object Cache for Oracle WSM	5-54
5.8.2.4	Configuring Distributed Notifications for the MDS Repository	5-54
5.9	Oracle User Messaging Service and High Availability Concepts.....	5-55
5.9.1	Oracle User Messaging Service Single-Instance Characteristics.....	5-55
5.9.1.1	Oracle User Messaging Service Component Characteristics	5-56
5.9.1.2	Oracle User Messaging Service Startup and Shutdown Lifecycle	5-57
5.9.1.3	Oracle User Messaging Service Request Flow	5-58
5.9.1.4	Oracle User Messaging Service Configuration Artifacts	5-58
5.9.2	Oracle User Messaging Service High Availability Architecture and Failover Considerations 5-59	
5.9.2.1	Oracle User Messaging Service Protection from Failures and Expected Behavior.....	5-59

5.9.2.2	Oracle User Messaging Service Cluster-Wide Configuration Changes	5-60
5.10	Oracle JCA Adapters and High Availability Concepts	5-61
5.10.1	Oracle JCA Adapters Single-Instance Characteristics	5-61
5.10.1.1	Oracle JCA Adapters Component Lifecycle	5-62
5.10.1.2	Oracle JCA Adapters Reliability and Transactional Behavior	5-62
5.10.1.3	Oracle JCA Adapters - Rejected Message Handling	5-63
5.10.2	Oracle JCA Adapters High Availability Architecture and Failover Considerations.....	5-64
5.10.2.1	Oracle JCA Adapters High Availability Error Handling	5-64
5.10.2.2	Oracle File and FTP Adapters High Availability	5-65
5.10.2.3	Oracle Database Adapters High Availability	5-67
5.10.2.4	Oracle JMS Adapters High Availability	5-68
5.10.2.4.1	Message Redelivery	5-68
5.10.2.4.2	Use of WebLogic Server Distributed Destinations	5-69
5.10.2.5	Oracle JCA Adapters Log File Locations	5-71
5.11	Oracle Business Activity Monitoring and High Availability Concepts	5-72
5.11.1	Oracle Business Activity Monitoring Single-Instance Characteristics	5-72
5.11.1.1	Oracle Business Activity Monitoring Component Characteristics	5-73
5.11.1.2	Oracle Business Activity Monitoring Startup/Shutdown Lifecycle	5-75
5.11.1.3	Oracle Business Activity Monitoring Startup and Shutdown of Processes	5-76
5.11.1.4	Oracle Business Activity Monitoring Configuration Artifacts	5-77
5.11.2	Oracle Business Activity Monitoring High Availability Architecture and Failover Considerations	5-78
5.11.2.1	Oracle Business Activity Monitoring Protection from Failures and Expected Behavior	5-80
5.11.2.2	Oracle Business Activity Monitoring Cluster-Wide Configuration Changes ..	5-81
5.11.2.3	Considerations for BAM Client Retries	5-82
5.12	Oracle Service Bus and High Availability Concepts	5-82
5.12.1	Oracle Service Bus Single Node Characteristics	5-83
5.12.1.1	Oracle Service Bus Session State	5-85
5.12.1.2	Oracle Service Bus External Dependencies	5-85
5.12.1.3	Oracle Service Bus Configuration Artifacts	5-86
5.12.1.4	Oracle Service Bus Deployment Artifacts	5-86
5.12.1.5	Oracle Service Bus Startup and Shutdown	5-87
5.12.1.6	Oracle Service Bus Log File Locations	5-87
5.12.2	Oracle Service Bus High Availability Architecture and Failover Considerations ..	5-87
5.12.2.1	Oracle Service Bus Protection from Failures and Expected Behavior	5-90
5.12.2.1.1	WebLogic Server Failure	5-90
5.12.2.1.2	Node Failure	5-91
5.12.2.1.3	Database Failure	5-91
5.12.2.2	Oracle Service Bus Cluster-Wide Deployment	5-92
5.12.2.3	Online Redeployment of Oracle Service Bus in a Cluster	5-92
5.12.2.4	Oracle Service Bus Cluster-Wide Configuration Changes	5-92
5.13	Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines	5-92
5.13.1	Preparing the Environment: Prerequisite Steps Before Setting up a SOA High Availability Configuration	5-95
5.13.1.1	Database Prerequisites	5-95

5.13.1.2	VIP and IP Prerequisites	5-95
5.13.1.3	Shared Storage Prerequisites	5-96
5.13.1.4	Installing and Configuring an LDAP Provider	5-96
5.13.1.5	Synchronizing System Clocks.....	5-97
5.13.1.6	Terminology for Directories and Directory Environment Variables	5-97
5.13.1.7	Installing and Configuring the Database Repository	5-98
5.13.1.8	Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database 5-99	
5.13.1.8.1	Running RCU.....	5-99
5.13.1.8.2	Configuring SOA Schemas for Transactional Recovery Privileges	5-100
5.13.1.9	Configuring Virtual Server Names and Ports for the Load Balancer	5-100
5.13.1.10	Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2.....	5-102
5.13.1.10.1	Validating Oracle HTTP Server.....	5-103
5.13.2	Installing Oracle Fusion Middleware Home	5-104
5.13.2.1	Installing Oracle WebLogic Server	5-104
5.13.2.2	Installing Oracle Fusion Middleware for Oracle SOA	5-104
5.13.3	Enabling VIP1 in SOAHOST1 and VIP2 in SOAHOST2.....	5-105
5.13.4	Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA Domain 5-106	
5.13.5	Creating boot.properties for the Administration Server on SOAHOST1.....	5-112
5.13.6	Starting and Validating the Administration Server in SOAHOST1.....	5-113
5.13.6.1	Starting the Administration Server on SOAHOST1	5-113
5.13.6.2	Validating the Administration Server	5-113
5.13.7	Disabling Host Name Verification for the Administration Server and the WLS_SOAn Managed Servers 5-113	
5.13.8	Configuring Oracle Coherence for Deploying Composites	5-114
5.13.9	Starting the System in SOAHOST1	5-116
5.13.9.1	Starting Node Manager on SOAHOST1	5-116
5.13.9.2	Starting and Validating the WLS_SOA1 Managed Server	5-117
5.13.10	Propagating the Domain Configuration to SOAHOST2 with pack/unpack Utilities..... 5-117	
5.13.11	Extracting XEngine Files in the Second Node	5-118
5.13.12	Starting the System in SOAHOST2.....	5-118
5.13.12.1	Starting Node Manager on SOAHOST2	5-118
5.13.12.2	Starting and Validating the WLS_SOA2 Managed Server	5-118
5.13.13	Configuring Oracle HTTP Servers for the Administration Server and the WLS_SOAn Managed Servers 5-119	
5.13.14	Validating Access Through Oracle HTTP Server.....	5-121
5.13.15	Configuring JMS Persistence Store as Shared Across the Servers.....	5-122
5.13.16	Configuring a Default Persistent Store for Transaction Recovery	5-122
5.13.17	Setting the Front End HTTP Host and Port	5-123
5.13.18	Setting the WLS Cluster Address for Direct Binding/RMI Invocations to Composites... 5-124	
5.13.19	Deploying Applications.....	5-125
5.13.20	Configuring Server Migration for the WLS_SOA Servers.....	5-127
5.13.20.1	Setting Up a User and Tablespace for the Server Migration Leasing Table ...	5-127
5.13.20.2	Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console 5-128	

5.13.20.3	Editing Node Manager's Properties File	5-130
5.13.20.4	Setting Environment and Superuser Privileges for the wlsifconfig.sh Script	5-132
5.13.20.5	Configuring Server Migration Targets	5-133
5.13.20.6	Testing the Server Migration	5-133
5.13.21	Scaling the Topology	5-136
5.13.21.1	Scaling Up the Topology (Adding Managed Servers to Existing Nodes)	5-136
5.13.21.2	Scaling Out the Topology (Adding Managed Servers to New Nodes)	5-141
5.14	Configuring High Availability for Oracle Service Bus, with SOA Service Infrastructure and Component Service Engines 5-147	
5.14.1	Preparing the Environment: Prerequisite Steps Before Setting up a SOA High Availability Configuration 5-150	
5.14.1.1	Database Prerequisites	5-150
5.14.1.2	VIP and IP Prerequisites	5-151
5.14.1.3	Shared Storage Prerequisites	5-151
5.14.1.4	Installing and Configuring an LDAP Provider	5-152
5.14.1.5	Synchronizing System Clocks	5-153
5.14.1.6	Terminology for Directories and Directory Environment Variables	5-153
5.14.1.7	Installing and Configuring the Database Repository	5-154
5.14.1.8	Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database 5-155	
5.14.1.8.1	Running RCU	5-155
5.14.1.8.2	Configuring SOA Schemas for Transactional Recovery Privileges	5-156
5.14.1.9	Configuring Virtual Server Names and Ports for the Load Balancer	5-156
5.14.1.10	Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2	5-158
5.14.1.10.1	Validating Oracle HTTP Server	5-160
5.14.2	Installing Oracle Fusion Middleware Home	5-160
5.14.2.1	Installing Oracle WebLogic Server	5-160
5.14.2.2	Installing Oracle Fusion Middleware for Oracle SOA	5-161
5.14.2.3	Installing Oracle Service Bus	5-161
5.14.3	Enabling VIP1 and VIP3 in SOAHOST1 and VIP2 and VIP4 in SOAHOST2	5-162
5.14.4	Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA and OSB Domain 5-163	
5.14.5	Creating boot.properties for the Administration Server on SOAHOST1	5-169
5.14.6	Starting and Validating the Administration Server in SOAHOST1	5-170
5.14.6.1	Starting the Administration Server on SOAHOST1	5-170
5.14.6.2	Validating the Administration Server	5-170
5.14.7	Disabling Host Name Verification for the Administration Server and the WLS_SOAn/WLS_OSBn Managed Servers 5-170	
5.14.8	Configuring Oracle Coherence for Deploying Composites	5-171
5.14.9	Configure Oracle Coherence for the Oracle Service Bus Result Cache	5-173
5.14.10	Setting Connection Destination Identifiers for B2B Queues	5-174
5.14.11	Starting the System in SOAHOST1	5-174
5.14.11.1	Starting Node Manager on SOAHOST1	5-174
5.14.11.2	Starting and Validating the WLS_SOA1 Managed Server	5-175
5.14.12	Propagating the Domain Configuration to SOAHOST2, OSBHOST1, and OSBHOST2 with pack/unpack Utilities 5-175	
5.14.13	Extracting XEngine Files in the Second Node	5-176
5.14.14	Starting the System in SOAHOST2, OSBHOST1, and OSBHOST2	5-176

5.14.14.1	Starting Node Manager on SOAHOST2, OSBHOST1, and OSBHOST2	5-176
5.14.14.2	Starting and Validating the WLS_SOA2, WLS_OSB1, and WLS_OSB2 Managed Server 5-176	
5.14.15	Configuring Oracle HTTP Servers for the Administration Server, WLS_SOAn, and WLS_OSBn Managed Servers 5-177	
5.14.16	Validating Access Through Oracle HTTP Server.....	5-179
5.14.17	Setting the Front End HTTP Host and Port	5-181
5.14.18	Configuring a Default Persistent Store for Transaction Recovery	5-182
5.14.19	Deploying Applications.....	5-182
5.14.20	Configuring Server Migration for the WLS_SOA Servers.....	5-184
5.14.21	Configuring Server Migration for the WLS_OSB Servers	5-187
5.14.22	Scaling the Topology	5-193
5.14.22.1	Scaling Up the Topology (Adding Managed Servers to Existing Nodes).....	5-193
5.14.22.2	Scaling Out the Topology (Adding Managed Servers to New Nodes)	5-200
5.15	Configuring High Availability for Oracle BAM	5-208
5.15.1	Preparing the Environment: Prerequisite Steps Before Setting up a High Availability Configuration for Oracle BAM 5-211	
5.15.1.1	Database Prerequisites	5-211
5.15.1.2	VIP and IPs Prerequisites	5-211
5.15.1.3	Installing and Configuring the Database Repository	5-211
5.15.1.4	Using Oracle Fusion Middleware Repository Creation Utility to Load Oracle Fusion Middleware Schemas 5-212	
5.15.1.4.1	Running RCU.....	5-212
5.15.1.5	Configuring Virtual Server Names and Ports for the Load Balancer	5-214
5.15.2	Installing Oracle HTTP Server on WEBHOST1	5-215
5.15.2.1	Validating Oracle HTTP Server	5-217
5.15.3	Installing Oracle Fusion Middleware Home	5-217
5.15.3.1	Installing Oracle WebLogic Server	5-217
5.15.3.2	Installing Oracle BAM Using the Oracle SOA Suite Installer	5-218
5.15.4	Enabling VIP0 and VIP1 on BAMHOST1	5-218
5.15.5	Running Oracle Fusion Middleware Configuration Wizard on BAMHOST1 to Create the WebLogic Server Oracle BAM Domain 5-218	
5.15.6	Creating boot.properties for the Administration Server and for WLS_BAM1 on BAMHOST1 5-224	
5.15.7	Starting the Administration Server on BAMHOST1	5-225
5.15.8	Disabling Host Name Verification for the Servers	5-225
5.15.9	Configuring a JMS Persistence Store for BAM UMS.....	5-226
5.15.10	Configuring a Default Persistence Store for Transaction Recovery	5-226
5.15.11	Untargeting the BAM Server System from BAMHOST2.....	5-227
5.15.12	Propagating the Domain Configuration from BAMHOST1 with pack/unpack Utilities 5-228	
5.15.13	Starting Node Manager on BAMHOST1 and BAMHOST2.....	5-228
5.15.14	Starting the Oracle BAM System	5-229
5.15.15	Configuring Oracle RAC Failover for the WLS_BAM Servers	5-229
5.15.16	Configuring the BAM Web Applications to Use the BAM Server in BAMHOST1	5-230
5.15.17	Configuring Oracle HTTP Servers for the Administration Server and the WLS_BAMn Managed Servers 5-230	

5.15.18	Validating Access through Oracle HTTP Server.....	5-232
5.15.19	Configuring Server Migration for the WLS_BAM Servers.....	5-232
5.15.19.1	Setting Up the User and Tablespace for the Server Migration Leasing Table	5-233
5.15.19.2	Creating a GridLink or Multi Data Source from the WebLogic Server Administration Console	5-233
5.15.19.3	Edit the Node Manager's Properties File	5-236
5.15.19.4	Set Environment and Superuser Privileges for the wlsifconfig.sh Script	5-237
5.15.19.5	Configure Server Migration Targets.....	5-238
5.15.19.6	Test Server Migration.....	5-238
5.15.20	Configuring Clients Connecting to the BAM System.....	5-239

6 Configuring High Availability for Oracle ADF and WebCenter Portal Applications

6.1	Oracle ADF and High Availability Concepts	6-1
6.1.1	Understanding Oracle ADF	6-1
6.1.1.1	Oracle ADF Components	6-2
6.1.1.1.1	ADF Business Components	6-3
6.1.1.1.2	ADF Model Layer.....	6-4
6.1.1.1.3	ADF Controller	6-4
6.1.1.1.4	ADF Faces Rich Client	6-4
6.1.1.2	Oracle ADF Single Node Architecture	6-5
6.1.1.3	Oracle ADF External Dependencies	6-5
6.1.1.4	Oracle ADF Log File.....	6-6
6.1.2	Oracle ADF High Availability Considerations	6-6
6.1.2.1	Oracle ADF Scope and Session State	6-6
6.1.2.2	Oracle ADF Failover and Expected Behavior.....	6-8
6.1.2.3	Oracle ADF Active Data Services.....	6-8
6.1.2.4	Configuring the ADF Application Module for Oracle RAC	6-9
6.1.3	Configuring Oracle ADF for High Availability	6-9
6.1.3.1	Configuring Application Modules.....	6-9
6.1.3.2	Configuring weblogic.xml.....	6-10
6.1.3.3	Configuring adf-config.xml	6-10
6.1.3.4	Configuring org.apache.myfaces.trinidad.CHECK_FILE_MODIFICATION..	6-11
6.1.4	Troubleshooting Oracle ADF High Availability	6-11
6.1.4.1	Troubleshooting Oracle ADF Development Issues	6-11
6.1.4.2	Troubleshooting Oracle ADF Deployment Issues	6-12
6.1.4.3	Troubleshooting Oracle ADF Replication and Failover Issues.....	6-12
6.2	Configuring an Oracle ADF High Availability Deployment	6-16
6.2.1	Terminology for Directories and Directory Environment Variables	6-16
6.2.2	Using RCU to Load Fusion Middleware Schemas in the Database	6-17
6.2.2.1	Running RCU	6-17
6.2.3	Installing Oracle HTTP Server on WEBHOST1.....	6-18
6.2.3.1	Validating Oracle HTTP Server	6-20
6.2.4	Installing the Oracle Fusion Middleware Home.....	6-20
6.2.4.1	Installing Oracle WebLogic Server	6-20
6.2.4.2	Installing Oracle Fusion Middleware for Oracle ADF Applications	6-21
6.2.5	Administration Server High Availability.....	6-21

6.2.6	Running the Configuration Wizard on APPHOST1 to Create the WebLogic Server ADF Domain	6-22
6.2.6.1	Creating boot.properties for the Administration Server and Managed Servers on APPHOST1	6-23
6.2.7	Starting the System in APPHOST1	6-24
6.2.7.1	Starting the Administration Server on APPHOST1	6-24
6.2.7.2	Validating the Administration Server	6-24
6.2.7.3	Disabling Host Name Verification for the Administration Server and Managed Servers for APPHOST1 and APPHOST2	6-24
6.2.7.4	Starting Node Manager on APPHOST1	6-25
6.2.8	Installing Oracle WebLogic Server and Oracle ADF on APPHOST2	6-25
6.2.9	Propagating the Domain Configuration to APPHOST2 with pack/unpack Utilities	6-25
6.2.9.1	Creating boot.properties for the Administration Server and Managed Servers on APPHOST2	6-26
6.2.9.2	Starting Node Manager on APPHOST2	6-26
6.2.9.3	Configuring the ADF Application for Replication	6-26
6.2.9.4	Deploying the ADF Application	6-27
6.2.9.5	Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers	6-27
6.2.9.6	Validating Access through Oracle HTTP Server	6-28
6.2.10	Scaling the Topology	6-28
6.2.10.1	Scaling Up the Topology (Adding Managed Servers to Existing Nodes)	6-29
6.2.10.2	Scaling Out the Topology (Adding Managed Servers to New Nodes)	6-29
6.3	WebCenter Portal and High Availability Concepts	6-31
6.3.1	Understanding WebCenter Portal	6-31
6.3.1.1	WebCenter Portal Components	6-31
6.3.1.2	WebCenter Portal Single-node Architecture	6-32
6.3.1.3	WebCenter Portal: State and Configuration Persistence	6-33
6.3.1.4	WebCenter Portal Log File Locations	6-33
6.3.2	WebCenter High Availability Architecture and Failover Considerations	6-33
6.3.2.1	WebCenter Portal Applications	6-34
6.3.2.2	Oracle WebCenter Startup Order	6-35
6.3.2.3	Deploying WebCenter Portal Application on a Cluster	6-35
6.3.2.4	WebCenter Portal Analytics Collector	6-35
6.3.2.5	WebCenter Portal State Replication	6-36
6.3.2.6	Understanding the Distributed Java Object Cache	6-36
6.3.2.7	WebCenter Portal Protection from Failover and Expected Behavior	6-38
6.3.2.8	Expected Behavior for Application Failover	6-39
6.3.2.9	Monitoring Logging of Application Deployments	6-41
6.3.2.10	WebCenter Portal Cluster-wide Configuration Changes	6-41
6.3.2.11	Maintaining Configuration in a Clustered Environment	6-41
6.4	Configuring High Availability for WebCenter Portal	6-41
6.4.1	Preparing the Environment: Prerequisite Steps Before Setting up a WebCenter Portal High Availability Configuration	6-42
6.4.1.1	Database Prerequisites	6-43
6.4.1.2	VIP and IP Prerequisites	6-43
6.4.1.3	Installing and Configuring the Database Repository	6-43

6.4.1.4	Installing and Configuring an LDAP Provider	6-44
6.4.1.5	Terminology for Directories and Directory Environment Variables	6-44
6.4.1.6	Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database 6-45	
6.4.1.6.1	Running RCU.....	6-45
6.4.2	Installing Oracle HTTP Server on WEBHOST1	6-46
6.4.2.1	Validating Oracle HTTP Server	6-48
6.4.3	Installing Oracle Fusion Middleware Home	6-48
6.4.3.1	Installing Oracle WebLogic Server	6-48
6.4.3.2	Installing Oracle Fusion Middleware for WebCenter Portal	6-49
6.4.4	Enabling the Administration Server VIP.....	6-50
6.4.5	Running Oracle Fusion Middleware Configuration Wizard on APPHOST1 to Create the WebLogic Server WebCenter Domain 6-50	
6.4.6	Creating boot.properties for the Administration Server and for Managed Servers on APPHOST1 6-56	
6.4.7	Starting the System in APPHOST1	6-57
6.4.7.1	Starting the Administration Server on APPHOST1	6-57
6.4.7.2	Validating the Administration Server	6-57
6.4.7.3	Disabling Host Name Verification for the Administration Server and the Managed Servers for APPHOST1 and APPHOST2 6-58	
6.4.7.4	Starting Node Manager on APPHOST1.....	6-58
6.4.8	Install WebLogic Server and WebCenter Portal on APPHOST2.....	6-58
6.4.9	Propagating the Domain Configuration to APPHOST2 with pack/unpack Utilities.....	6-59
6.4.10	Starting Node Manager on APPHOST2.....	6-59
6.4.11	Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers 6-59	
6.4.11.1	Configuring a Virtual Host for Oracle Pagelet Producer and Sharepoint	6-61
6.4.11.1.1	Virtual Hosts Requirement	6-61
6.4.11.1.2	Additional Configuration.....	6-62
6.4.11.2	Validating Access through Oracle HTTP Server	6-62
6.4.12	Configuring Manual Failover of the Administration Server to APPHOST2	6-63
6.4.13	Configuring the Java Object Cache	6-63
6.4.14	Configuring the Distributed Notifications for the MDS Repository.....	6-63
6.4.15	Configuring WebCenter Portal for Replication.....	6-63
6.4.16	Configuring Analytics.....	6-64
6.4.17	Configuring Activity Graph.....	6-64
6.4.18	Configuring Clustering for Discussions Server.....	6-65
6.4.19	Scaling the Topology.....	6-65
6.4.19.1	Scaling Up the Topology (Adding Managed Servers to Existing Nodes).....	6-65
6.4.19.2	Scaling Out the Topology (Adding Managed Servers to New Nodes)	6-66
6.4.20	Troubleshooting WebCenter Portal High Availability	6-68
6.4.20.1	Troubleshooting WebCenter Portal Deployment Issues	6-68
6.4.20.2	Troubleshooting WebCenter Portal Replication and Failover Issues	6-68
6.4.20.3	Troubleshooting Lost Changes to Policies	6-70
6.4.20.4	Troubleshooting JOC Configuration	6-70
6.4.21	Converting Discussions from Multicast to Unicast	6-70
6.5	Configuring High Availability for WebCenter Portal Applications	6-71

6.5.1	Configuring a Cluster for WebCenter Portal Applications	6-71
6.5.2	Adding More WebCenter Portal Applications Servers.....	6-72
6.5.3	Configuring Distributed Notifications for the MDS Repository	6-72

7 High Availability for Oracle Data Integrator

7.1	Introduction to Oracle Data Integrator	7-1
7.2	Oracle Data Integrator Single Instance Characteristics	7-1
7.2.1	Oracle Data Integrator Sessions Lifecycle and Recovery.....	7-3
7.2.1.1	Sessions Interruption	7-3
7.2.1.2	Recovering Sessions	7-3
7.2.2	Agent Startup and Shutdown Cycle	7-4
7.2.3	Oracle Data Integrator External Dependencies.....	7-5
7.2.4	Oracle Data Integrator Startup and Shutdown Process	7-5
7.2.5	Oracle Data Integrator Configuration Artifacts	7-5
7.2.5.1	Java EE Agent Configuration.....	7-5
7.2.5.2	Standalone Agent Configuration	7-6
7.2.5.3	Oracle Data Integrator Console Configuration	7-7
7.2.5.4	Oracle Data Integrator Log Locations and Configuration.....	7-7
7.2.5.4.1	Oracle Data Integrator Session Logs	7-7
7.2.5.4.2	Java EE Agent Log Files	7-7
7.2.5.4.3	Standalone Agent Log Files	7-8
7.2.5.4.4	Oracle Data Integrator Console Log Files.....	7-8
7.3	Oracle Data Integrator High Availability and Failover Considerations.....	7-8
7.3.1	Oracle Data Integrator Clustered Deployment	7-8
7.3.2	Standalone Agent High Availability with OPMN.....	7-10
7.3.3	Oracle Data Integrator Protection from Failure and Expected Behavior	7-10
7.3.3.1	WebLogic Server or Standalone Agent Crash.....	7-10
7.3.3.2	Repository Database Failure	7-11
7.3.3.3	Scheduler Node Failure	7-11
7.4	Configuring High Availability for Oracle Data Integrator	7-12
7.4.1	Running RCU to Create the Master and Work Repositories.....	7-12
7.4.2	Installation and Configuration of the First Oracle Data Integrator Host	7-12
7.4.2.1	Installing the Oracle WebLogic Server on APPHOST1	7-12
7.4.2.2	Install Oracle Data Integrator on APPHOST1	7-13
7.4.2.3	Create the High Availability Domain.....	7-14
7.4.2.4	Start the Administration Server.....	7-17
7.4.2.5	Configure the Credential Store.....	7-17
7.4.2.5.1	Configuring Credentials Using WLST	7-17
7.4.2.5.2	Configuring Credentials using Enterprise Manager.....	7-17
7.4.2.6	Configure the Default Agent	7-18
7.4.2.7	Configure Coherence for the Cluster	7-19
7.4.2.8	Configure Node Manager and Start odi_server1.....	7-19
7.4.2.9	Verify the Oracle Data Integrator Agent is Running	7-19
7.4.3	Installation and Configuration of the Second Oracle Data Integrator Host	7-19
7.4.3.1	Installing the Oracle WebLogic Server on APPHOST2	7-19
7.4.3.2	Pack and Unpack the Domain from APPHOST1 to APPHOST2	7-20
7.4.3.3	Configure Node Manager and Start odi_server2.....	7-20

7.4.3.4	Verify the Oracle Data Integrator Agent is Running	7-20
7.4.4	Installing Oracle HTTP Server	7-20
7.4.4.1	Upgrading the Oracle HTTP Server Oracle Home	7-21
7.4.4.2	Configuring Oracle HTTP Server	7-21
7.4.4.3	Configuring the Load Balancer	7-22
7.4.4.4	Verify the Oracle Data Integrator Agent is Running	7-22
7.4.4.5	Reconfigure Agents	7-22

8 Configuring High Availability for Identity Management Components

8.1	Identity Management Product Components and High Availability Concepts	8-2
8.1.1	About the 11g Oracle Identity Management Products	8-5
8.2	Prerequisites for Oracle Identity Management High Availability Configuration	8-6
8.2.1	Oracle Home Requirement	8-6
8.2.2	Database Prerequisites	8-7
8.2.3	Installing and Configuring the Database Repository	8-7
8.2.4	Obtaining the Repository Creation Utility Software	8-8
8.2.4.1	Executing the Repository Creation Utility	8-8
8.2.5	Configuring the Database for Oracle Fusion Middleware 11g Metadata	8-9
8.2.5.1	Database Examples in This Chapter	8-9
8.2.5.2	Database Services	8-10
8.2.5.3	Verifying Transparent Application Failover (TAF)	8-11
8.2.5.4	Configuring Virtual Server Names and Ports for the Load Balancer	8-11
8.2.5.4.1	Load Balancers	8-11
8.2.5.4.2	Virtual Server Names	8-13
8.3	Oracle Internet Directory High Availability	8-14
8.3.1	Oracle Internet Directory Component Architecture	8-14
8.3.1.1	Oracle Internet Directory Component Characteristics	8-16
8.3.1.1.1	Runtime Processes	8-16
8.3.1.1.2	Process Lifecycle	8-17
8.3.1.1.3	Request Flow	8-18
8.3.1.1.4	Configuration Artifacts	8-18
8.3.1.1.5	External Dependencies	8-18
8.3.1.1.6	Oracle Internet Directory Log File	8-19
8.3.2	Oracle Internet Directory High Availability Concepts	8-19
8.3.2.1	Oracle Internet Directory High Availability Architecture	8-19
8.3.2.1.1	Starting and Stopping the Cluster	8-20
8.3.2.1.2	Cluster-Wide Configuration Changes	8-20
8.3.2.2	Protection from Failures and Expected Behavior	8-21
8.3.2.2.1	Oracle Internet Directory Process Failure	8-21
8.3.2.2.2	Expected Client Application Behavior When Failure Occurs	8-21
8.3.2.2.3	External Dependency Failure	8-22
8.3.2.3	Oracle Internet Directory Prerequisites	8-22
8.3.2.3.1	Synchronizing the Time on Oracle Internet Directory Nodes	8-23
8.3.2.3.2	Using RCU to Create Oracle Internet Directory Schemas in the Repository	8-23
8.3.2.3.3	Load Balancer Virtual Server Names for Oracle Internet Directory	8-24
8.3.3	Oracle Internet Directory High Availability Configuration Steps	8-24

8.3.3.1	Installing Oracle Fusion Middleware Components	8-24
8.3.3.1.1	Install Oracle WebLogic Server	8-24
8.3.3.1.2	Installing Oracle Fusion Middleware for Identity Management	8-25
8.3.3.1.3	Upgrading Oracle Identity Management	8-27
8.3.3.2	Configuring Oracle Internet Directory Without a WebLogic Domain	8-27
8.3.3.2.1	Configuring Oracle Internet Directory on OIDHOST1	8-27
8.3.3.2.2	Oracle Internet Directory Component Names Assigned by Oracle Identity Management Installer 8-30	
8.3.3.2.3	Configuring Oracle Internet Directory on OIDHOST2	8-30
8.3.3.2.4	Registering Oracle Internet Directory with a WebLogic Domain	8-33
8.3.3.3	Configuring Oracle Internet Directory With a WebLogic Domain	8-34
8.3.3.3.1	Configuring Oracle Internet Directory on OIDHOST1	8-34
8.3.3.3.2	Creating boot.properties for the Administration Server on OIDHOST1 ...	8-37
8.3.3.3.3	Configuring Oracle Internet Directory on OIDHOST2	8-38
8.3.4	Validating Oracle Internet Directory High Availability	8-41
8.3.5	Oracle Internet Directory Failover and Expected Behavior	8-42
8.3.5.1	Performing an Oracle Internet Directory Failover	8-42
8.3.5.2	Performing an Oracle RAC Failover	8-43
8.3.6	Troubleshooting Oracle Internet Directory High Availability	8-44
8.3.7	Additional Oracle Internet Directory High Availability Issues	8-45
8.3.7.1	Changing the Password of the ODS Schema Used by Oracle Internet Directory	8-45
8.4	Oracle Virtual Directory High Availability	8-46
8.4.1	Oracle Virtual Directory Component Architecture	8-46
8.4.1.1	Oracle Virtual Directory Runtime Considerations	8-47
8.4.1.2	Oracle Virtual Directory Component Characteristics	8-48
8.4.1.2.1	Oracle Virtual Directory Log File	8-49
8.4.2	Oracle Virtual Directory High Availability Concepts	8-49
8.4.2.1	Oracle Virtual Directory High Availability Architecture	8-49
8.4.2.1.1	Oracle Virtual Directory High Availability Connect Features	8-50
8.4.2.2	Oracle Virtual Directory Prerequisites	8-51
8.4.2.2.1	Load Balancer Virtual Server Names for Oracle Virtual Directory	8-51
8.4.3	Oracle Virtual Directory High Availability Configuration Steps	8-51
8.4.3.1	Configuring Oracle Virtual Directory Without a WebLogic Domain	8-52
8.4.3.1.1	Configuring Oracle Virtual Directory on OVDHOST1	8-52
8.4.3.1.2	Configuring Oracle Virtual Directory on OVDHOST2	8-55
8.4.3.1.3	Registering Oracle Virtual Directory with a WebLogic Domain	8-57
8.4.3.2	Configuring Oracle Virtual Directory With a WebLogic Domain	8-58
8.4.3.2.1	Configuring Oracle Virtual Directory on OVDHOST1	8-58
8.4.3.2.2	Creating boot.properties for the Administration Server on OVDHOST1..	8-61
8.4.3.2.3	Configuring Oracle Virtual Directory on OVDHOST2	8-62
8.4.3.3	Configuring Oracle Virtual Directory with Highly Available Data Sources....	8-65
8.4.3.3.1	Configuring Oracle Virtual Directory with an Oracle RAC Database	8-65
8.4.3.3.2	Configuring Oracle Virtual Directory with LDAP	8-65
8.4.4	Validating Oracle Virtual Directory High Availability	8-65
8.4.4.1	Validating Oracle Virtual Directory High Availability Using SSL	8-67
8.4.5	Oracle Virtual Directory Failover and Expected Behavior	8-68

8.4.5.1	Performing an Oracle Virtual Directory Failover	8-68
8.4.5.2	Performing an Oracle RAC Failover	8-68
8.4.6	Troubleshooting Oracle Virtual Directory High Availability	8-68
8.4.6.1	Troubleshooting LDAP Adapter Creation.....	8-69
8.5	Oracle Directory Integration Platform High Availability	8-69
8.5.1	Oracle Directory Integration Platform Component Architecture	8-69
8.5.1.1	Oracle Directory Integration Platform Component Characteristics	8-70
8.5.1.1.1	Runtime Processes.....	8-71
8.5.1.1.2	Process Lifecycle.....	8-71
8.5.1.1.3	Request Flow	8-72
8.5.1.1.4	Configuration Artifacts.....	8-75
8.5.1.1.5	External Dependencies	8-76
8.5.1.1.6	Oracle Directory Integration Platform Log File	8-76
8.5.2	Oracle Directory Integration Platform High Availability Concepts	8-76
8.5.2.1	Oracle Directory Integration Platform High Availability Architecture	8-76
8.5.2.1.1	Starting and Stopping the Cluster.....	8-78
8.5.2.1.2	Cluster-Wide Configuration Changes.....	8-78
8.5.2.2	Protection from Failures and Expected Behavior	8-79
8.5.2.2.1	Process Failure	8-79
8.5.2.2.2	Expected Client Application Behavior When Failure Occurs.....	8-80
8.5.2.2.3	External Dependency Failure	8-80
8.5.2.3	Oracle Directory Integration Platform Prerequisites	8-80
8.5.3	Oracle Directory Integration Platform and Oracle Directory Services Manager High Availability Configuration Steps 8-80	
8.5.3.1	Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST1 8-81	
8.5.3.2	Creating boot.properties for the Administration Server on IDMHOST1.....	8-84
8.5.3.3	Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST2 8-85	
8.5.3.4	Post-Installation Steps for Oracle Directory Integration Platform and Oracle Directory Services Manager 8-86	
8.5.3.4.1	Copy the Oracle Directory Integration Platform Configuration from IDMHOST1 to IDMHOST2 8-86	
8.5.3.4.2	Restart the Managed Server on IDMHOST2 in a Cluster.....	8-87
8.5.3.5	Installing Oracle Fusion Middleware Components on WEBHOST1 and WEBHOST2 8-87	
8.5.3.5.1	Installing Oracle HTTP Server for the Web Tier.....	8-87
8.5.3.5.2	Upgrading the Oracle HTTP Server Oracle Home to Patch Set 3	8-88
8.5.3.6	Configuring Oracle HTTP Server on WEBHOST1 and WEBHOST2.....	8-88
8.5.3.6.1	Configuring Oracle HTTP Server for Oracle Directory Services Manager High Availability 8-90	
8.5.4	Oracle Directory Integration Platform Failover and Expected Behavior	8-92
8.5.5	Troubleshooting Oracle Directory Integration Platform High Availability	8-92
8.5.5.1	Managed Server Log File Exceptions Received for Oracle Directory Integration Platform During an Oracle RAC Failover 8-93	
8.5.5.2	Dealing with Error Messages Received After Starting WebLogic Node Manager.....	8-93
8.5.5.3	If WebLogic Node Manager Fails to Start.....	8-94

8.5.5.4	Configuration Changes Do Not Automatically Propagate to All Oracle Directory Integration Platform Instances in a Highly Available Topology	8-94
8.5.5.5	Operation Cannot Be Completed for Unknown Errors Message.....	8-94
8.6	Oracle Directory Services Manager High Availability	8-95
8.6.1	Oracle Directory Services Manager Component Architecture	8-95
8.6.1.1	Oracle Directory Services Manager Component Characteristics	8-96
8.6.1.1.1	Lifecycle Management.....	8-96
8.6.1.1.2	Oracle Directory Services Manager Log File	8-97
8.6.2	Oracle Directory Services Manager High Availability Concepts	8-97
8.6.2.1	Oracle Directory Services Manager High Availability Architecture	8-97
8.6.2.1.1	Starting and Stopping the Cluster.....	8-99
8.6.2.2	Protection from Failures and Expected Behaviors.....	8-99
8.6.2.2.1	Process Failure	8-99
8.6.2.2.2	Expected Client Application Behavior When Failure Occurs.....	8-99
8.6.2.2.3	Expected Dependency Failure.....	8-100
8.6.2.3	Oracle Directory Services Manager Prerequisites	8-100
8.6.3	Oracle Directory Services Manager High Availability Configuration Steps	8-100
8.6.4	Validating Oracle Directory Services Manager High Availability	8-101
8.6.4.1	Performing a WebLogic Server Instance Failover	8-101
8.6.4.2	Performing an Oracle RAC Database Failover.....	8-102
8.6.5	Oracle Directory Services Manager Failover and Expected Behavior	8-103
8.6.5.1	Using Oracle Directory Services Manager to Validate a Failover of a Managed Server	8-103
8.6.5.2	Using Oracle Directory Services Manager to Validate a Failover of an Oracle Internet Directory Instance	8-103
8.6.5.3	Using Oracle Directory Services Manager to Validate an Oracle RAC Failover.....	8-104
8.6.6	Troubleshooting Oracle Directory Services Manager	8-104
8.6.6.1	Dealing with Error Messages Received After Starting WebLogic Node Manager.....	8-104
8.6.6.2	If WebLogic Node Manager Fails to Start.....	8-105
8.6.6.3	Oracle Directory Services Manager Failover Using Oracle HTTP Server is Not Transparent	8-105
8.6.6.4	Oracle Directory Services Manager Displays "LDAP Server is down" Message During Oracle Internet Directory Failover	8-106
8.6.6.5	Oracle Directory Services Manager Temporarily Loses Its Connection During Oracle RAC Failover	8-106
8.6.7	Additional Considerations for Oracle Directory Services Manager High Availability	8-106
8.7	Collocated Architecture High Availability	8-106
8.7.1	Collocated Architecture Overview.....	8-106
8.7.2	Collocated Architecture High Availability Deployment.....	8-107
8.7.2.1	Collocated Architecture Prerequisites.....	8-108
8.7.2.2	Configuring Collocated Components for High Availability.....	8-108
8.7.3	Validating the Collocated Components High Availability	8-110
8.7.3.1	Validation Tests	8-110
8.7.3.2	Failures and Expected Behaviors	8-110
8.7.4	Troubleshooting Collocated Components Manager High Availability	8-110

8.7.5	Additional Considerations for Collocated Components High Availability	8-111
8.8	Oracle Access Manager High Availability	8-111
8.8.1	Oracle Access Manager Component Architecture.....	8-112
8.8.1.1	Oracle Access Manager Component Characteristics.....	8-112
8.8.1.1.1	Oracle Access Manager State Information.....	8-114
8.8.1.1.2	Oracle Access Manager Request Flow	8-115
8.8.1.1.3	Oracle Access Manager Process Lifecycle.....	8-116
8.8.1.1.4	Oracle Access Manager Configuration Artifacts	8-116
8.8.1.1.5	Oracle Access Manager External Dependencies.....	8-117
8.8.1.1.6	Oracle Access Manager Log File Location.....	8-117
8.8.2	Oracle Access Manager High Availability Concepts	8-117
8.8.2.1	Oracle Access Manager High Availability Architecture.....	8-117
8.8.2.1.1	Starting and Stopping the Cluster.....	8-120
8.8.2.1.2	Cluster-Wide Configuration Changes.....	8-121
8.8.2.2	Protection from Failures and Expected Behaviors.....	8-122
8.8.2.2.1	WebLogic Server Crash	8-122
8.8.2.2.2	Node Failure.....	8-123
8.8.2.2.3	Database Failure	8-123
8.8.3	Oracle Security Token Service High Availability	8-123
8.8.3.1	Oracle Security Token Service High Availability Architecture	8-124
8.8.3.1.1	Clients and Client Connections	8-125
8.8.3.1.2	Cluster Wide Configuration Changes	8-125
8.8.3.2	Oracle Security Token Service Component Characteristics	8-126
8.8.3.2.1	Oracle Security Token Service Component Lifecycle	8-126
8.8.3.2.2	Runtime Processes.....	8-126
8.8.3.2.3	Starting and Stopping Oracle Security Token Service	8-127
8.8.3.2.4	J2EE Components and Subcomponents.....	8-127
8.8.3.2.5	Session State Information.....	8-127
8.8.3.2.6	Configuration Artifacts.....	8-127
8.8.3.2.7	External Dependencies	8-128
8.8.3.3	Oracle Security Token Service High Availability Configuration Steps.....	8-128
8.8.3.4	Validating Oracle Security Token Service High Availability.....	8-128
8.8.3.5	Oracle Security Token Service Failover and Expected Behavior.....	8-128
8.8.3.5.1	Death Detection and Restart	8-129
8.8.3.5.2	Node Failure.....	8-129
8.8.3.6	Disabling and Enabling Oracle Security Token Service.....	8-129
8.8.3.7	Troubleshooting Oracle Security Token Service	8-129
8.8.3.8	Log File Location	8-130
8.8.3.9	Additional Considerations	8-130
8.8.4	Oracle Access Manager High Availability Configuration Steps	8-130
8.8.4.1	Prerequisites for Oracle Access Manager Configuration.....	8-131
8.8.4.2	Run the Repository Creation Utility to Create the OAM Schemas in a Database	8-131
8.8.4.3	Installing Oracle WebLogic Server	8-131
8.8.4.4	Install and Configure the Oracle Access Manager Application Tier	8-132
8.8.4.4.1	Install Oracle Fusion Middleware for Identity Management	8-132
8.8.4.4.2	Configure Oracle Identity Management on OAMHOST1	8-133
8.8.4.5	Creating boot.properties for the Administration Server on OAMHOST1	8-136

8.8.4.6	Start OAMHOST1.....	8-137
8.8.4.6.1	Create the Node Manager Properties File on OAMHOST1.....	8-137
8.8.4.6.2	Start Node Manager.....	8-137
8.8.4.6.3	Start Oracle Access Manager on OAMHOST1.....	8-137
8.8.4.7	Validating OAMHOST1.....	8-137
8.8.4.8	Configure OAM on OAMHOST2.....	8-137
8.8.4.9	Start OAMHOST2.....	8-138
8.8.4.9.1	Create the Node Manager Properties File on OAMHOST2.....	8-138
8.8.4.9.2	Start Node Manager.....	8-138
8.8.4.9.3	Start Oracle Access Manager on OAMHOST2.....	8-138
8.8.4.10	Validating OAMHOST2.....	8-138
8.8.4.11	Configure Oracle Access Manager to Work with Oracle HTTP Server.....	8-138
8.8.4.11.1	Update Oracle HTTP Server Configuration.....	8-138
8.8.4.11.2	Restart Oracle HTTP Server.....	8-139
8.8.4.11.3	Make Oracle Access Manager Server Aware of the Load Balancer.....	8-139
8.8.4.12	Configure Oracle Access Manager to use an External LDAP Store.....	8-139
8.8.4.12.1	Extending Directory Schema for Oracle Access Manager.....	8-140
8.8.4.12.2	Create Users and Groups in LDAP.....	8-141
8.8.4.12.3	Create a User Identity Store.....	8-143
8.8.4.12.4	Set LDAP to System and Default Store.....	8-144
8.8.4.12.5	Set Authentication to Use External LDAP.....	8-144
8.8.4.13	Validating the Oracle Access Manager Configuration.....	8-144
8.8.4.14	Configuring Oracle Coherence to Keep Configuration Files in Sync.....	8-144
8.8.4.15	Scaling Up and Scaling Out the Oracle Access Manager Topology.....	8-145
8.8.4.15.1	Scaling Up Oracle Access Manager.....	8-145
8.8.4.15.2	Scaling Out Oracle Access Manager.....	8-147
8.9	Oracle Identity Manager High Availability.....	8-149
8.9.1	Oracle Identity Manager Component Architecture.....	8-150
8.9.1.1	Oracle Identity Manager Component Characteristics.....	8-151
8.9.1.2	Runtime Processes.....	8-152
8.9.1.3	Component and Process Lifecycle.....	8-152
8.9.1.4	Starting and Stopping Oracle Identity Manager.....	8-153
8.9.1.5	Configuration Artifacts.....	8-153
8.9.1.6	External Dependencies.....	8-153
8.9.1.7	Oracle Identity Manager Log File Locations.....	8-154
8.9.2	Oracle Identity Manager High Availability Concepts.....	8-154
8.9.2.1	Oracle Identity Manager High Availability Architecture.....	8-155
8.9.2.2	Starting and Stopping the Oracle Identity Manager Cluster.....	8-157
8.9.2.3	Cluster-Wide Configuration Changes.....	8-158
8.9.2.4	Considerations for Synchronizing with LDAP.....	8-158
8.9.3	Oracle Identity Manager High Availability Configuration Steps.....	8-158
8.9.3.1	Prerequisites for Oracle Identity Manager Configuration.....	8-159
8.9.3.1.1	Running RCU to Create the OIM Schemas in a Database.....	8-160
8.9.3.1.2	Installing Oracle WebLogic Server.....	8-161
8.9.3.1.3	Installing the Oracle SOA Suite on OIMHOST1 and OIMHOST2.....	8-162
8.9.3.1.4	Installing the Oracle Identity Manager on OIMHOST1 and OIMHOST2.....	8-162

8.9.3.2	Creating and Configuring the WebLogic Domain for OIM and SOA on OIMHOST1 8-163	
8.9.3.3	Post-Installation Steps on OIMHOST1	8-167
8.9.3.3.1	Creating boot.properties for the Administration Server on OIMHOST1	8-167
8.9.3.3.2	Update Node Manager on OIMHOST1	8-168
8.9.3.3.3	Start Node Manager on OIMHOST1	8-168
8.9.3.3.4	Start the Administration Server on OIMHOST1.....	8-168
8.9.3.4	Configuring Oracle Identity Manager on OIMHOST1	8-168
8.9.3.4.1	Prerequisites for Configuring Oracle Identity Manager.....	8-169
8.9.3.4.2	Running the Oracle Identity Management Configuration Wizard.....	8-176
8.9.3.5	Post-Configuration Steps for the Managed Servers	8-179
8.9.3.5.1	Start the WLS_SOA1 and WLS_OIM1 Managed Servers on OIMHOST1	8-179
8.9.3.5.2	Updating the Coherence Configuration for the Coherence Cluster	8-179
8.9.3.6	Validate the Oracle Identity Manager Instance on OIMHOST1.....	8-179
8.9.3.7	Propagating Oracle Identity Manager to OIMHOST2	8-180
8.9.3.8	Post-Installation Steps on OIMHOST2.....	8-180
8.9.3.8.1	Update Node Manager on OIMHOST2	8-180
8.9.3.8.2	Start Node Manager on OIMHOST2.....	8-180
8.9.3.8.3	Start the WLS_SOA2 and WLS_OIM2 Managed Servers on OIMHOST2	8-181
8.9.3.9	Validate the Oracle Identity Manager Instance on OIMHOST2.....	8-181
8.9.3.10	Configuring Oracle Internet Directory using the LDAP Configuration Post-setup Script 8-181	
8.9.3.11	Configuring Server Migration for the OIM and SOA Managed Servers	8-182
8.9.3.11.1	Setting Up a User and Tablespace for the Server Migration Leasing Table	8-182
8.9.3.11.2	Creating a Multi Data Source Using the Oracle WebLogic Administration Console 8-183	
8.9.3.11.3	Editing Node Manager's Properties File.....	8-184
8.9.3.11.4	Setting Environment and Superuser Privileges for the wlsifconfig.sh Script	8-186
8.9.3.11.5	Configuring Server Migration Targets.....	8-186
8.9.3.11.6	Testing the Server Migration	8-187
8.9.3.12	Configuring a Shared JMS Persistence Store	8-188
8.9.3.13	Configuring a Default Persistence Store for Transaction Recovery	8-189
8.9.3.14	Install Oracle HTTP Server on WEBHOST1 and WEBHOST2	8-190
8.9.3.15	Configuring Oracle Identity Manager to Work with the Web Tier	8-190
8.9.3.15.1	Prerequisites.....	8-190
8.9.3.15.2	Configuring Oracle HTTP Servers to Front End the OIM and SOA Managed Servers 8-190	
8.9.3.16	Validate the Oracle HTTP Server Configuration	8-192
8.9.3.17	Oracle Identity Manager Failover and Expected Behavior	8-193
8.9.3.18	Troubleshooting Oracle Identity Manager High Availability	8-193
8.9.3.19	Scaling Up and Scaling Out the Oracle Identity Manager Topology	8-194
8.9.3.19.1	Scaling Up Oracle Identity Manager	8-194
8.9.3.19.2	Scaling Out Oracle Identity Manager.....	8-198
8.10	Authorization Policy Manager High Availability.....	8-204
8.11	Oracle Identity Navigator High Availability	8-204
8.12	Oracle Adaptive Access Manager High Availability.....	8-205

8.12.1	Oracle Adaptive Access Manager Component Architecture	8-205
8.12.1.1	Oracle Adaptive Access Manager Component Characteristics	8-207
8.12.1.1.1	Oracle Adaptive Access Manager State Information	8-207
8.12.1.1.2	Oracle Adaptive Access Manager Runtime Processes	8-208
8.12.1.1.3	Oracle Adaptive Access Manager Process Lifecycle	8-208
8.12.1.1.4	Oracle Adaptive Access Manager Configuration Artifacts	8-209
8.12.1.1.5	Oracle Adaptive Access Manager Deployment Artifacts	8-210
8.12.1.1.6	Oracle Adaptive Access Manager External Dependencies	8-210
8.12.1.1.7	Oracle Adaptive Access Manager Log File Location	8-210
8.12.2	Oracle Adaptive Access Manager High Availability Concepts	8-210
8.12.2.1	Oracle Adaptive Access Manager High Availability Architecture	8-210
8.12.2.1.1	Starting and Stopping the Cluster	8-212
8.12.2.1.2	Cluster-Wide Configuration Changes	8-213
8.12.2.2	Protection from Failures and Expected Behaviors	8-213
8.12.3	Oracle Adaptive Access Manager High Availability Configuration Steps	8-213
8.12.3.1	Prerequisites for Oracle Adaptive Access Manager Configuration	8-214
8.12.3.2	Run the Repository Creation Utility to Create the OAAM Schemas in a Database ... 8-215	
8.12.3.3	Installing Oracle WebLogic Server	8-215
8.12.3.4	Install and Configure the Oracle Adaptive Access Manager Application Tier	8-215
8.12.3.5	Creating boot.properties for the Administration Server on OAAMHOST1 ...	8-221
8.12.3.6	Create the Oracle Adaptive Access Manager Administration User	8-222
8.12.3.7	Start OAAMHOST1	8-222
8.12.3.7.1	Create the Node Manager Properties File on OAAMHOST1	8-223
8.12.3.7.2	Start Node Manager	8-223
8.12.3.7.3	Start Oracle Adaptive Access Manager on OAAMHOST1	8-223
8.12.3.8	Validating OAAMHOST1	8-223
8.12.3.9	Configure Oracle Adaptive Access Manager on OAAMHOST2	8-223
8.12.3.10	Start OAAMHOST2	8-224
8.12.3.10.1	Create the Node Manager Properties File on OAAMHOST2	8-224
8.12.3.10.2	Start Node Manager	8-224
8.12.3.10.3	Start Oracle Adaptive Access Manager on OAAMHOST2	8-224
8.12.3.11	Validating OAAMHOST2	8-224
8.12.3.12	Configure Oracle Adaptive Access Manager to Work with Oracle HTTP Server	8-225
8.12.3.12.1	Update Oracle HTTP Server Configuration	8-225
8.12.3.12.2	Restart Oracle HTTP Server	8-225
8.12.3.12.3	Change Host Assertion in WebLogic	8-225
8.12.3.13	Validating the Oracle Adaptive Access Manager Configuration	8-227
8.12.3.14	Scaling Up and Scaling Out the Oracle Adaptive Access Manager Topology	8-227
8.12.3.14.1	Scaling Up Oracle Adaptive Access Manager	8-227
8.12.3.14.2	Scaling Out Oracle Adaptive Access Manager	8-228
8.13	Oracle Identity Federation High Availability	8-230
8.13.1	Oracle Identity Federation Component Architecture	8-230
8.13.1.1	Oracle Identity Federation Component Characteristics	8-231
8.13.1.1.1	Runtime Processes	8-232

8.13.1.1.2	Process Lifecycle.....	8-232
8.13.1.1.3	Request Flow	8-232
8.13.1.1.4	Configuration Artifacts.....	8-233
8.13.1.1.5	External Dependencies	8-233
8.13.1.1.6	Oracle Identity Federation Log File Location	8-237
8.13.2	Oracle Identity Federation High Availability Concepts	8-237
8.13.2.1	Oracle Identity Federation High Availability Architecture	8-239
8.13.2.1.1	Starting and Stopping the Cluster.....	8-240
8.13.2.1.2	Cluster-Wide Configuration Changes.....	8-241
8.13.2.2	High Availability Considerations for Integration with Oracle Access Manager	8-242
8.13.2.3	Oracle Identity Federation Prerequisites	8-242
8.13.2.3.1	Using RCU to Create Oracle Identity Federation Schemas in the Repository	8-243
8.13.3	Oracle Identity Federation High Availability Configuration Steps	8-244
8.13.3.1	Configuring Oracle Identity Federation on OIFHOST1	8-244
8.13.3.2	Creating boot.properties for the Administration Server on OIFHOST1	8-252
8.13.3.3	Configuring Oracle Identity Federation on OIFHOST2	8-253
8.13.3.4	Post-Installation Steps for Oracle Identity Federation	8-255
8.13.3.4.1	Copy the Oracle Identity Federation Configuration Directory from OIFHOST1 to OIFHOST2	8-255
8.13.3.4.2	Start the Managed Server on OIFHOST2 in a Cluster	8-256
8.13.3.4.3	Configure Oracle HTTP Server	8-256
8.13.3.5	Configuring the Load Balancer	8-256
8.13.3.5.1	Load Balancer Virtual Server Name Setup	8-257
8.13.3.5.2	Oracle Identity Federation Configuration	8-257
8.13.3.6	Validating Oracle Identity Federation High Availability	8-257
8.13.3.7	Enabling Oracle Identity Federation Integration with Highly Available LDAP Servers	8-257
8.13.4	Oracle Identity Federation Failover and Expected Behavior	8-258
8.13.4.1	Performing an Oracle Identity Federation Failover	8-258
8.13.4.2	Performing an Oracle RAC Failover.....	8-258
8.13.5	Troubleshooting Oracle Identity Federation High Availability	8-259
8.14	Starting and Stopping Oracle Identity Management Components	8-259
8.14.1	Oracle Internet Directory	8-260
8.14.2	Oracle Virtual Directory	8-260
8.14.3	Oracle HTTP Server.....	8-260
8.14.4	Node Manager.....	8-261
8.14.5	WebLogic Administration Server.....	8-261
8.14.6	Oracle Identity Manager.....	8-262
8.14.7	Oracle Access Manager Managed Servers	8-263
8.14.8	Oracle Adaptive Access Manager Managed Servers.....	8-263
8.14.9	Oracle Identity Federation.....	8-264

9 Configuring High Availability for Web Tier Components

9.1	About the Web Tier	9-1
9.2	Oracle HTTP Server and High Availability Concepts.....	9-2
9.2.1	Oracle HTTP Server Single-Instance Characteristics.....	9-2

9.2.1.1	Oracle HTTP Server and Oracle WebLogic Server	9-4
9.2.1.2	Oracle HTTP Server External Dependencies	9-4
9.2.2	Oracle HTTP Server Startup and Shutdown Lifecycle	9-4
9.2.3	Starting and Stopping Oracle HTTP Server	9-5
9.2.3.1	Understanding the PID File	9-5
9.2.3.2	Starting and Stopping Oracle HTTP Server Using Oracle Fusion Middleware Control 9-6	
9.2.3.3	Starting and Stopping Oracle HTTP Server Using opmnctl	9-6
9.2.4	Oracle HTTP Server Configuration Artifacts	9-6
9.2.5	Oracle HTTP Server Log File Locations	9-6
9.2.6	Oracle HTTP Server High Availability Architecture and Failover Considerations ..	9-7
9.2.7	Oracle HTTP Server Protection from Failures and Expected Behaviors	9-8
9.2.8	Oracle HTTP Server Cluster-Wide Configuration Changes	9-8
9.2.9	Configuring Oracle HTTP Server for High Availability	9-8
9.2.9.1	Prerequisites	9-8
9.2.9.1.1	Load Balancer	9-8
9.2.9.1.2	Associating Oracle HTTP Server with a WebLogic Domain	9-9
9.2.9.2	Install Oracle HTTP Server on WEBHOST1	9-9
9.2.9.2.1	Configure Virtual Host(s)	9-11
9.2.9.2.2	Configure mod_wl_ohs	9-11
9.2.9.2.3	Restart Oracle HTTP Server	9-12
9.2.9.2.4	Validate the Oracle HTTP Server Configuration	9-12
9.2.9.3	Install Oracle HTTP Server on WEBHOST2	9-12
9.2.9.3.1	Configure Virtual Host(s)	9-13
9.2.9.3.2	Configure mod_wl_ohs	9-13
9.2.9.3.3	Restart Oracle HTTP Server	9-13
9.2.9.3.4	Validate the Oracle HTTP Server Configuration	9-13
9.3	Oracle Web Cache and High Availability Concepts	9-13
9.3.1	Oracle Web Cache Single-Node Characteristics	9-13
9.3.1.1	Oracle Web Cache Component Characteristics	9-14
9.3.1.2	Oracle Web Cache Process Monitoring	9-15
9.3.1.3	Oracle Web Cache Startup and Shutdown Lifecycle	9-15
9.3.1.4	Oracle Web Cache Request Flow	9-15
9.3.1.5	Oracle Web Cache Configuration Artifacts	9-17
9.3.1.6	Log File Locations	9-17
9.3.2	Oracle Web Cache High Availability Considerations	9-18
9.3.2.1	Oracle Web Cache Stateless Load Balancing	9-18
9.3.2.2	Oracle Web Cache Backend Failover	9-20
9.3.2.3	Oracle Web Cache Session Binding	9-21
9.3.2.4	Oracle Web Cache Cluster-Wide Configuration Changes	9-22
9.3.2.5	Oracle Web Cache as a Software Load Balancer	9-24
9.3.3	Configuring Oracle Web Cache High Availability Solutions	9-25
9.3.3.1	Configure Oracle Web Cache Session Binding	9-25
9.3.3.2	Configuring a Cache Cluster	9-26
9.3.3.2.1	Configuration Prerequisites	9-26
9.3.3.2.2	Understanding Failover Threshold and Capacity Settings	9-27
9.3.3.2.3	Task 1: Add Caches to the Cluster and Configure Properties	9-29

9.3.3.2.4	Task 2: Enable Tracking of Session Binding	9-30
9.3.3.2.5	Task 3: Synchronize Configuration to Cluster Members	9-30
9.3.3.2.6	Removing a Cache Member from a Cluster	9-30
9.3.3.2.7	Configuring Administration and Invalidation-Only Clusters.....	9-31
9.3.3.3	Configure Oracle Web Cache as a Software Load Balancer.....	9-31

10 Configuring Identity Management for Maximum High Availability

10.1	Introduction to the Maximum High Availability Identity Management Deployment .	10-1
10.2	Overview of Replication	10-3
10.3	Setting up Multimaster Replication	10-3
10.3.1	Setting Up LDAP Multimaster Replication	10-4
10.3.1.1	Adding a Node in LDAP Multimaster Replication	10-5
10.3.1.2	Deleting a Node in LDAP Multimaster Replication.....	10-6
10.3.2	Setting Up Oracle Advanced Database Multimaster Replication	10-6
10.3.2.1	Adding a Node in Oracle Advanced Database Multimaster Replication	10-7
10.3.2.2	Deleting a Node in Oracle Advanced Database Multimaster Replication.....	10-8

11 Configuring High Availability for Enterprise Content Management

11.1	Oracle Imaging and Process Management High Availability	11-1
11.1.1	Oracle I/PM Component Architecture.....	11-1
11.1.1.1	Oracle I/PM Component Characteristics	11-3
11.1.1.1.1	Oracle I/PM State Information	11-3
11.1.1.1.2	Oracle I/PM Runtime Processes	11-3
11.1.1.1.3	Oracle I/PM Process Lifecycle	11-4
11.1.1.1.4	Oracle I/PM Configuration Artifacts	11-4
11.1.1.1.5	Oracle I/PM External Dependencies.....	11-4
11.1.1.1.6	Oracle I/PM Log File Location.....	11-5
11.1.2	Oracle I/PM High Availability Concepts	11-5
11.1.2.1	Oracle I/PM High Availability Architecture	11-5
11.1.2.1.1	Starting and Stopping the Cluster.....	11-7
11.1.2.1.2	Cluster-Wide Configuration Changes.....	11-7
11.1.2.2	Protection from Failures and Expected Behaviors.....	11-7
11.1.2.3	Creation of Oracle I/PM Artifacts in a Cluster	11-10
11.1.2.4	Troubleshooting Oracle I/PM	11-10
11.2	Oracle Universal Content Management High Availability	11-11
11.2.1	Oracle Universal Content Management Component Architecture.....	11-11
11.2.1.1	Oracle UCM Component Characteristics.....	11-12
11.2.1.1.1	Oracle UCM State Information.....	11-12
11.2.1.1.2	Oracle UCM Runtime Processes	11-12
11.2.1.1.3	Oracle UCM Process Lifecycle.....	11-12
11.2.1.1.4	Oracle UCM Configuration Artifacts	11-12
11.2.1.1.5	Oracle UCM Deployment Artifacts	11-12
11.2.1.1.6	Oracle UCM External Dependencies.....	11-12
11.2.1.1.7	Oracle UCM Log File Locations	11-13
11.2.2	Oracle UCM High Availability Concepts.....	11-13
11.2.2.1	Oracle UCM High Availability Architecture.....	11-13
11.2.2.1.1	Starting and Stopping the Cluster.....	11-14

11.2.2.1.2	Cluster-Wide Configuration Changes.....	11-15
11.2.2.2	Oracle UCM and Inbound Refinery High Availability Architecture	11-15
11.2.2.2.1	Content Server and Inbound Refinery Communication	11-15
11.2.2.2.2	Content Server Clusters and Inbound Refinery Instances	11-16
11.2.2.2.3	Inbound Refinery Instances and Load Balancers	11-16
11.2.2.2.4	Inbound Refinery Availability.....	11-16
11.2.2.3	Oracle URM High Availability.....	11-16
11.2.2.4	Protection from Failure and Expected Behaviors	11-16
11.2.2.5	Troubleshooting Oracle UCM High Availability	11-17
11.3	Oracle ECM High Availability Configuration Steps	11-17
11.3.1	Shared Storage.....	11-20
11.3.2	Configuring the Oracle Database	11-20
11.3.3	Installing and Configuring ECMHOST1	11-21
11.3.3.1	Installing Oracle WebLogic Server on ECMHOST1	11-22
11.3.3.2	Installing Oracle ECM on ECMHOST1	11-22
11.3.3.3	Create a High Availability Domain	11-23
11.3.3.4	Start the Administration Server and the Managed Servers on ECMHOST1 ..	11-29
11.3.3.5	Disabling Host Name Verification for the Administration Server and the Managed Servers for ECMHOST1 and ECMHOST2	11-29
11.3.3.6	Configure the WLS_UCM1 Managed Server	11-30
11.3.3.7	Configure the WLS_URM1 Managed Server	11-30
11.3.4	Installing and Configuring WEBHOST1	11-30
11.3.4.1	Installing Oracle HTTP Server on WEBHOST1	11-30
11.3.4.2	Configuring Oracle HTTP Server on WEBHOST1	11-33
11.3.5	Configuring the Load Balancer.....	11-34
11.3.6	Installing and Configuring ECMHOST2.....	11-35
11.3.6.1	Installing Oracle WebLogic Server on ECMHOST2.....	11-35
11.3.6.2	Installing ECM on ECMHOST2.....	11-35
11.3.6.3	Using pack and unpack to Join the Domain on ECMHOST1	11-35
11.3.6.4	Start Node Manager and the WLS_UCM2 Server on ECMHOST2.....	11-35
11.3.6.5	Start the Managed Servers on ECMHOST2.....	11-35
11.3.6.6	Configure the WLS_UCM2 Managed Server	11-35
11.3.6.7	Configure the WLS_URM2 Managed Server	11-36
11.3.7	Installing and Configuring WEBHOST2	11-36
11.3.7.1	Installing Oracle HTTP Server on WEBHOST2	11-36
11.3.7.2	Configuring Oracle HTTP Server on WEBHOST2	11-39
11.3.8	Configuring the I/PM Managed Servers	11-39
11.3.8.1	Configuring a JMS Persistence Store for I/PM JMS	11-39
11.3.8.2	Configuring a Default Persistence Store for Transaction Recovery	11-40
11.3.8.3	Configuring Oracle I/PM with Oracle UCM	11-40
11.3.8.3.1	Enabling I/PM in UCM.....	11-40
11.3.8.3.2	Upgrading the Default File Store	11-41
11.3.8.3.3	Adding the I/PM Server Listen Addresses to the List of Allowed Hosts in UCM	11-41
11.3.8.3.4	Creating a Connection to the UCM System.....	11-41
11.3.8.4	Configuring BPEL CSF Credentials	11-42
11.3.8.5	Configuring the BPEL Connection.....	11-42

11.3.8.6	Setting the Front End HTTP Host and Port.....	11-43
11.3.8.7	Configuring Server Migration for Oracle I/PM Instances	11-43
11.3.8.7.1	About Configuring Server Migration.....	11-43
11.3.8.7.2	Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console 11-44	
11.3.8.7.3	Testing the Server Migration	11-47
11.3.9	Configuration of Inbound Refinery Instances	11-48
11.3.9.1	Inbound Refinery and Inbound Refinery Cluster Concepts	11-48
11.3.9.2	Content Server and Inbound Refinery Configuration	11-48
11.3.9.3	Inbound Refinery Instances and Oracle HTTP Server	11-49

12 Active-Passive Topologies for Oracle Fusion Middleware High Availability

12.1	Oracle Fusion Middleware Cold Failover Cluster Topology Concepts.....	12-1
12.2	Configuring Oracle Fusion Middleware for Active-Passive Deployments	12-4
12.2.1	General Requirements for Cold Failover Cluster.....	12-5
12.2.1.1	Terminology for Directories and Directory Environment Variables	12-6
12.2.2	Transforming Oracle Fusion Middleware Infrastructure Components	12-7
12.2.2.1	Administration Server Topology 1	12-7
12.2.2.1.1	Topology 1 Installation Procedure.....	12-8
12.2.2.2	Administration Server Topology 2	12-11
12.2.2.2.1	Topology 2 Installation Procedure.....	12-12
12.2.2.3	Transforming the Administration Server for Cold Failover Cluster	12-14
12.2.2.4	Transforming Oracle WebLogic Managed Servers	12-15
12.2.2.4.1	Transforming an Oracle WebLogic Managed Server using the Fusion Middleware Administration Console 12-16	
12.2.2.4.2	Transforming an Oracle WebLogic Managed Server using the WLST Command Line 12-17	
12.2.2.5	Transforming Node Manager	12-17
12.2.2.6	Transforming Oracle Process Management and Notification Server	12-18
12.2.2.7	Transforming Oracle Enterprise Manager for an Oracle Instance	12-19
12.2.2.8	Transforming Web Tier Components and Clients.....	12-20
12.2.2.8.1	Transforming Oracle HTTP Server.....	12-20
12.2.2.8.2	Transforming Oracle Web Cache.....	12-20
12.2.2.9	Instance-specific considerations	12-22
12.2.2.9.1	UNIX Platforms	12-22
12.2.2.9.2	Windows Platform	12-22
12.2.3	Transforming Oracle Fusion Middleware Components	12-24
12.2.3.1	Transforming Oracle Internet Directory and Its Clients.....	12-24
12.2.3.1.1	Transforming Oracle Internet Directory	12-24
12.2.3.1.2	Transforming Oracle Internet Directory Clients.....	12-25
12.2.3.2	Transforming Oracle Virtual Directory and Its Clients	12-25
12.2.3.2.1	Transforming Oracle Virtual Directory.....	12-25
12.2.3.2.2	Transforming Oracle Virtual Directory Clients	12-26
12.2.3.3	Transforming Oracle Directory Integration Platform and Oracle Directory Services Manager and Their Clients 12-26	
12.2.3.3.1	Transforming Oracle Directory Integration Platform and Oracle Directory Services Manager 12-26	

12.2.3.3.2	Transforming Oracle Directory Integration Platform and Oracle Directory Services Manager Clients	12-26
12.2.3.4	Transforming Oracle Identity Federation and Its Client	12-26
12.2.3.4.1	Transforming Oracle Identity Federation	12-26
12.2.3.4.2	Transforming Oracle Identity Federation Clients	12-27
12.2.3.5	Transforming an Oracle SOA Suite	12-27
12.2.3.6	Transforming Oracle Access Manager and Its Clients	12-29
12.2.3.6.1	Transforming Oracle Access Manager	12-29
12.2.3.6.2	Transforming Oracle Access Manager Clients	12-30
12.2.3.7	Transforming Oracle Adaptive Access Manager and Its Clients	12-30
12.2.3.7.1	Transforming Oracle Adaptive Access Manager	12-30
12.2.3.7.2	Transforming Oracle Adaptive Access Manager Clients	12-31
12.2.3.8	Transforming Oracle Identity Manager and Its Clients	12-31
12.2.3.8.1	Transforming Oracle Identity Manager	12-31
12.2.3.8.2	Transforming Oracle Identity Manager Clients	12-32
12.2.3.9	Transforming an Oracle WebCenter Portal Suite	12-32
12.2.3.10	Transforming Oracle Portal, Forms, Reports, and Discoverer	12-33
12.2.3.10.1	Transforming Oracle Forms for Cold Failover Cluster	12-33
12.2.3.10.2	Transforming Oracle Reports for Cold Failover Cluster	12-33
12.2.3.10.3	Transforming Oracle Discoverer for Cold Failover Cluster	12-35
12.2.3.10.4	Transforming Oracle Portal for Cold Failover Cluster	12-36
12.2.3.10.5	Transforming Oracle Business Activity Management (BAM)	12-38
12.2.3.10.6	Transforming a Custom ADF Deployment	12-39
12.2.3.11	Transforming an Oracle WebCenter Content	12-39
12.2.3.12	Transforming Oracle Business Intelligence	12-41
12.2.3.12.1	Transforming Oracle Business Intelligence Enterprise Edition and its Clients	12-41
12.2.3.12.2	Transforming Oracle Business Intelligence Publisher and its Clients	12-44
12.2.3.12.3	Transforming Oracle Real-Time Decisions and its Clients	12-46
12.2.3.13	Transforming Oracle BI for Microsoft Office and its Clients	12-47
12.2.3.14	Single Sign-On Re-registration (If required)	12-48
12.2.4	Transforming an Oracle Database	12-48
12.2.4.1	Database Instance Platform-Specific Considerations	12-50
12.3	Oracle Fusion Middleware Cold Failover Cluster Example Topologies	12-52
12.3.1	Example Topology 1	12-52
12.3.2	Example Topology 2	12-53
12.3.3	Example Topology 3	12-54
12.4	Transforming the Administration Server in an Existing Domain for Cold Failover Cluster	12-55
12.4.1	Destination Topologies	12-56
12.4.2	Cold Failover Cluster Transformation Procedure	12-57

13 Using Oracle Cluster Ready Services

13.1	Introduction to Oracle Clusterware	13-1
13.2	Cluster Ready Services and Oracle Fusion Middleware	13-1
13.3	Installing and Configuring Oracle Clusterware with CRS	13-2
13.3.1	Upgrading Older Versions of ASCRS to the Current ASCRS Version	13-3

13.3.2	Installing ASCRS.....	13-3
13.3.3	Configuring ASCRS with Oracle Fusion Middleware	13-4
13.4	Using ASCRS to Manage Resources.....	13-7
13.4.1	Creating CRS Managed Resources	13-7
13.4.1.1	Creating a Virtual IP Resource	13-8
13.4.1.2	Creating a Shared Disk Resource	13-8
13.4.1.3	Creating an Oracle Database Listener Resource	13-10
13.4.1.4	Creating an Oracle Database Resource	13-10
13.4.1.5	Creating a Middleware Resource.....	13-11
13.4.1.5.1	Creating a Resource for OPMN Managed Components	13-11
13.4.1.5.2	Creating a Resource for WebLogic Servers	13-12
13.4.2	Updating Resources	13-15
13.4.3	Starting Up Resources.....	13-15
13.4.4	Shutting Down Resources	13-16
13.4.5	Resource Switchover	13-16
13.4.6	Deleting Resources	13-16
13.4.7	Checking Resource Status.....	13-16
13.4.8	Configuring and Using Health Monitors	13-18
13.5	Example Topologies	13-19
13.6	Troubleshooting Oracle CRS	13-23
13.6.1	OPMN Resource Depends on a Virtual IP Resource.....	13-24
13.6.2	ASCRS Logging.....	13-24

14 Configuring High Availability for Oracle Portal, Forms, Reports, and Discoverer

14.1	Overview of Oracle Portal, Forms, Reports, and Discoverer	14-1
14.1.1	Oracle Portal, Forms, Reports, and Discoverer Architecture	14-2
14.1.2	Common Log Files.....	14-4
14.1.3	Common Component Failures and Expected Behaviors.....	14-4
14.1.3.1	Oracle Web Cache and Oracle HTTP Server Process Failures.....	14-4
14.1.3.2	Common Component Node Failures	14-4
14.1.3.3	Common Component WebLogic Managed Server Failures	14-5
14.1.3.4	Common Component Database Failures	14-5
14.1.4	Oracle Portal, Forms, Reports, and Discoverer Cluster-Wide Configuration Changes	14-5
14.1.5	Common Component Log File Information	14-6
14.2	Oracle Portal and High Availability Concepts	14-6
14.2.1	Oracle Portal Single-Instance Characteristics	14-6
14.2.1.1	Oracle Portal Request Flow	14-6
14.2.1.2	Oracle Portal Component Characteristics.....	14-6
14.2.1.3	Oracle Portal Startup and Shutdown of Processes and Lifecycle.....	14-7
14.2.1.4	Oracle Portal Deployment Artifacts.....	14-7
14.2.1.5	Oracle Portal Configuration Information.....	14-7
14.2.1.6	Oracle Portal Logging and Log Configuration	14-8
14.2.1.6.1	Oracle Portal Log Files.....	14-8
14.2.1.7	Oracle Portal External Dependencies	14-8
14.2.2	Oracle Portal Protection from Failures and Expected Behavior	14-8

14.2.2.1	Oracle Portal Process Failures	14-10
14.2.2.2	Oracle Portal Node Failures.....	14-10
14.2.2.3	Oracle Portal WebLogic Managed Server Failures.....	14-10
14.2.2.4	Oracle Portal Protection from Database Failures.....	14-10
14.3	Oracle Reports and High Availability Concepts	14-10
14.3.1	Oracle Reports Single-Instance Characteristics	14-10
14.3.1.1	Oracle Reports State Information.....	14-11
14.3.1.2	Oracle Reports External Dependencies.....	14-11
14.3.1.3	Oracle Reports Specific Configuration Files.....	14-11
14.3.1.4	Oracle Reports Connection Retry.....	14-12
14.3.1.5	Oracle Reports Process Flow.....	14-12
14.3.1.6	Oracle Reports Log Files.....	14-13
14.3.2	Oracle Reports Protection from Failure and Expected Behavior.....	14-13
14.3.2.1	Oracle Reports Process Failures	14-15
14.3.2.2	Oracle Reports Node Failures.....	14-15
14.3.2.3	Oracle Reports WebLogic Managed Server Failures.....	14-15
14.3.2.4	Oracle Reports Database Failures	14-15
14.4	Oracle Forms and High Availability Concepts	14-15
14.4.1	Oracle Forms Single-Instance Component Characteristics	14-15
14.4.1.1	Oracle Forms State Information	14-15
14.4.1.2	Oracle Forms Database Requirements	14-15
14.4.1.3	Oracle Forms Request Flow	14-16
14.4.1.4	Oracle Forms Configuration Persistence.....	14-17
14.4.1.5	Oracle Forms Runtime Considerations	14-18
14.4.1.6	Oracle Forms Process Flow	14-18
14.4.1.7	Oracle Forms Configuration Files	14-18
14.4.1.8	Oracle Forms External Dependencies.....	14-19
14.4.1.9	Oracle Forms Log Files	14-19
14.4.2	Oracle Forms Protection from Failover and Expected Behavior	14-20
14.4.2.1	Oracle Forms N+1 Redundancy	14-21
14.4.2.2	Oracle Forms N+M Redundancy	14-22
14.4.2.3	Oracle Forms Virtual Machines.....	14-22
14.4.2.4	Oracle Forms Configuration Cloning.....	14-22
14.4.2.5	Oracle Forms Process Failures.....	14-22
14.4.2.6	Oracle Forms Node Failures	14-22
14.4.2.7	Oracle Forms WebLogic Managed Server Failures	14-22
14.4.2.8	Oracle Forms Database Failures.....	14-22
14.5	Oracle Discoverer and High Availability Concepts	14-22
14.5.1	Oracle Discoverer Single-Instance Characteristics	14-23
14.5.1.1	Oracle Discoverer Runtime Considerations	14-23
14.5.1.2	Oracle Discoverer Viewer and Web Cache.....	14-24
14.5.1.3	Oracle Discoverer Configuration Considerations	14-24
14.5.1.4	Oracle Discoverer Deployment Considerations	14-25
14.5.1.5	Oracle Discoverer Log File Locations.....	14-25
14.5.1.6	Discoverer Log Files.....	14-25
14.5.2	Oracle Discoverer Protection from Failures and Expected Behavior.....	14-25
14.5.2.1	Preference Server Failover.....	14-26

14.5.2.2	Session State Replication and Failover	14-26
14.5.2.3	Performance Recommendation	14-27
14.5.2.4	Propagation of Configuration Changes Across the Cluster.....	14-27
14.5.2.5	Cluster-Wide Application Deployment	14-27
14.5.2.6	Online Application Deployment	14-27
14.5.2.7	Oracle Discoverer Process Failures.....	14-27
14.5.2.8	Oracle Discoverer Node Failures	14-28
14.5.2.9	Oracle Discoverer WebLogic Managed Server Failures	14-28
14.5.2.10	Oracle Discoverer Database Failures	14-28
14.6	Configuring Oracle Portal, Forms, Reports, and Discoverer for High Availability	14-28
14.6.1	Prerequisites	14-28
14.6.1.1	Dependencies	14-28
14.6.1.2	Network Requirements.....	14-28
14.6.1.2.1	Load Balancer.....	14-29
14.6.1.2.2	Load Balancer Configuration - Virtual Server Names and Ports	14-29
14.6.1.3	Databases	14-30
14.6.1.4	Shared Directories	14-31
14.6.1.5	Managed Port Numbers	14-31
14.6.1.6	Site Names	14-31
14.6.2	Assumptions.....	14-32
14.6.2.1	Ports	14-32
14.6.3	Creating the Metadata Repository	14-33
14.6.3.1	Install the Repository Creation Utility (RCU)	14-33
14.6.3.2	Run Repository Creation Utility.....	14-33
14.6.4	Install and Configure Application Tier on APPHOST1	14-34
14.6.4.1	Install Oracle WebLogic Server	14-34
14.6.4.2	Install Oracle Portal, Forms, Reports, and Discoverer	14-34
14.6.4.3	Configure Oracle Portal, Forms, Reports, and Discoverer Software	14-35
14.6.4.4	Validation	14-37
14.6.4.5	Generic Configuration	14-37
14.6.4.5.1	Set Admin Server Listen Address.....	14-37
14.6.4.5.2	Configure Virtual Hosts	14-37
14.6.4.5.3	Create boot.properties File.....	14-38
14.6.4.5.4	Configure sqlnet.ora	14-38
14.6.4.5.5	Configure Web Cache	14-38
14.6.4.5.6	Change the Web Cache Passwords.....	14-40
14.6.4.5.7	Restart Web Tier (Oracle HTTP Server and Web Cache)	14-40
14.6.4.5.8	Register with Single Sign-On Server	14-40
14.6.4.5.9	Enabling the WebLogic Plug-in	14-41
14.6.4.5.10	Change Host Assertion in WebLogic	14-41
14.6.4.6	Configure Oracle Portal for High Availability	14-42
14.6.4.6.1	Rewire Portal Repository	14-42
14.6.4.6.2	Configure Parallel Page Engine Loop-Back with Load Balancer	14-43
14.6.4.6.3	Database Wallets and Portal	14-44
14.6.4.6.4	Restart All Components	14-45
14.6.4.6.5	Post-installation Step for Portal Installation with Oracle RAC	14-45
14.6.4.6.6	Validate Configuration	14-46

14.6.4.7	Configure Oracle Forms for High Availability	14-46
14.6.4.7.1	Create TNSNAMES Entries for Customer Databases.....	14-46
14.6.4.7.2	Restart WLS_FORMS.....	14-46
14.6.4.7.3	Validate Configuration	14-47
14.6.4.8	Configure Oracle Reports for High Availability.....	14-47
14.6.4.8.1	Create Reports Queue in Database	14-47
14.6.4.8.2	Create a TNSNAMES Entry for Reports Queue	14-47
14.6.4.8.3	Create a Security Key for the Reports Queue.....	14-48
14.6.4.8.4	Configure the Database Job Repository for In-Process Reports Servers..	14-48
14.6.4.8.5	Configure the Reports Server to Access Shared Output Directory.....	14-49
14.6.4.8.6	Restart WLS_REPORTS.....	14-49
14.6.4.8.7	Validate Configuration	14-50
14.6.4.9	Configure Oracle Discoverer for High Availability	14-50
14.6.4.9.1	Create TNSNAMES Entries for Customer Databases.....	14-50
14.6.4.9.2	Update configuration.xml.....	14-51
14.6.4.9.3	Discoverer Viewer and Web Cache	14-51
14.6.4.9.4	Enable Single Sign On.....	14-51
14.6.4.9.5	Restart All Components	14-51
14.6.4.9.6	Validate Configuration	14-52
14.6.5	Install and Configure Application Tier on APPHOST2.....	14-52
14.6.5.1	Install Oracle WebLogic Server	14-52
14.6.5.2	Install Oracle Portal, Forms, Reports, and Discoverer Software	14-53
14.6.5.3	Configure Oracle Portal, Forms, Reports, and Discoverer Software	14-53
14.6.5.4	Generic Configuration	14-54
14.6.5.4.1	Copy Configuration Information from APPHOST1	14-54
14.6.5.4.2	Configure Virtual Hosts	14-55
14.6.5.4.3	Update Oracle HTTP Server Configuration to be Cluster Aware.....	14-55
14.6.5.4.4	Change the Web Cache Passwords.....	14-56
14.6.5.4.5	Configure Web Cache.....	14-57
14.6.5.4.6	Restart Web Processes on APPHOST1 and APPHOST2	14-58
14.6.5.5	Configure Oracle Portal for High Availability.....	14-58
14.6.5.5.1	Copy Configuration Information from APPHOST1	14-58
14.6.5.5.2	Create Portal Directories	14-59
14.6.5.5.3	Update Instance Paths	14-59
14.6.5.5.4	Restart the Web Processes.....	14-59
14.6.5.5.5	Start WLS_PORTAL1.....	14-59
14.6.5.5.6	Validate the Configuration	14-60
14.6.5.5.7	Best Practices.....	14-60
14.6.5.6	Configure Oracle Forms for High Availability	14-60
14.6.5.6.1	Create a TNSNAMES entries for Customer Databases	14-60
14.6.5.6.2	Copy Forms Configuration Files.....	14-60
14.6.5.6.3	Update default.env.....	14-61
14.6.5.6.4	Restart WLS_FORMS1.....	14-61
14.6.5.6.5	Validate the Configuration	14-62
14.6.5.6.6	Best Practices.....	14-62
14.6.5.7	Configure Oracle Reports for High Availability.....	14-62
14.6.5.7.1	Create TNSNAMES Entries for Customer Databases.....	14-62

14.6.5.7.2	Configure the Reports Server to Access Shared Output Directory.....	14-62
14.6.5.7.3	Configure the Database Job Repository for In-process Reports Servers..	14-63
14.6.5.7.4	Creating an Oracle Reports Server Cluster.....	14-63
14.6.5.7.5	Restart WLS_REPORTS and WLS_REPORTS1.....	14-64
14.6.5.7.6	Validate the Configuration	14-64
14.6.5.7.7	Managing Connection Availability for Oracle Reports Services.....	14-65
14.6.5.8	Configure Oracle Discoverer for High Availability	14-65
14.6.5.8.1	Create TNSNAMES Entries for Customer Databases.....	14-65
14.6.5.8.2	Copy Discoverer Configuration Files.....	14-65
14.6.5.8.3	Update configuration.xml.....	14-66
14.6.5.8.4	Changing the Preference Store	14-66
14.6.5.8.5	Restart WLS_DISCO and WLS_DISCO1	14-66
14.6.5.8.6	Validate the Configuration	14-67
14.6.5.8.7	Failover of the Preference Server	14-67
14.6.5.8.8	Setting up Discoverer WSRP Portlet Producer in a Clustered Environment.....	14-68
14.6.5.8.9	Best Practices.....	14-70
14.6.6	Scaling Out the Deployment.....	14-70

15 Configuring High Availability for Oracle Business Intelligence and EPM

15.1	High Availability for Oracle Business Intelligence Enterprise Edition and Enterprise Performance Management	15-1
15.1.1	Oracle BI EE Component Architecture.....	15-2
15.1.1.1	Oracle BI EE Component Characteristics.....	15-3
15.1.1.1.1	Process Lifecycle.....	15-5
15.1.1.1.2	External Dependencies	15-5
15.1.1.1.3	Configuration Artifacts.....	15-5
15.1.1.1.4	Deployment Artifacts.....	15-6
15.1.1.1.5	Log File Locations	15-6
15.1.2	Oracle BI EE High Availability Concepts.....	15-7
15.1.2.1	Oracle BI EE and EPM High Availability Architecture	15-8
15.1.2.1.1	Web Server High Availability Considerations.....	15-11
15.1.2.1.2	Oracle BI Presentation Services Plug-in High Availability Considerations.....	15-11
15.1.2.1.3	Presentation Services High Availability Considerations.....	15-12
15.1.2.1.4	BI Cluster Controller High Availability Considerations	15-12
15.1.2.1.5	BI Server High Availability Considerations.....	15-13
15.1.2.1.6	Administration Tool High Availability Considerations.....	15-13
15.1.2.1.7	Oracle BI Scheduler High Availability Considerations.....	15-14
15.1.2.1.8	BI JavaHost High Availability Considerations	15-14
15.1.2.2	Shared Files and Directories	15-14
15.1.2.2.1	Oracle BI Presentation Catalog Shared Files	15-14
15.1.2.2.2	Repository Publishing Directory Shared Files	15-15
15.1.2.2.3	Global Cache Shared Files.....	15-15
15.1.2.2.4	Oracle BI Scheduler Scripts Shared Files	15-15
15.1.2.3	Starting and Stopping the Cluster.....	15-15
15.1.2.4	Cluster-Wide Configuration Changes.....	15-16

15.1.2.5	Protection From Failures and Expected Behaviors.....	15-16
15.1.2.5.1	Machine Failure	15-16
15.1.2.5.2	WebLogic Administration Server Failure	15-17
15.1.2.5.3	WebLogic Managed Server Failure	15-17
15.1.2.5.4	Oracle BI Scheduler Failure	15-17
15.1.2.5.5	Cluster Controller Failure	15-18
15.1.2.5.6	Presentation Services Failure	15-18
15.1.2.5.7	BI Server Failure	15-19
15.1.2.5.8	Troubleshooting Oracle BI EE	15-19
15.1.3	Oracle Essbase Component Architecture.....	15-20
15.1.3.1	Oracle Essbase Component Characteristics.....	15-20
15.1.3.1.1	State Information	15-21
15.1.3.1.2	Runtime Processes.....	15-21
15.1.3.1.3	Process Lifecycle.....	15-21
15.1.3.1.4	Request Flow	15-22
15.1.3.1.5	External Dependencies	15-22
15.1.3.1.6	Configuration Artifacts.....	15-22
15.1.3.1.7	Deployment Artifacts.....	15-22
15.1.3.1.8	Log Files.....	15-22
15.1.4	Oracle Essbase High Availability Concepts.....	15-23
15.1.4.1	Oracle Essbase High Availability Architecture.....	15-23
15.1.4.1.1	Shared Files and Directories	15-23
15.1.4.1.2	Cluster-Wide Configuration Changes.....	15-23
15.1.4.2	Protection from Failures and Expected Behaviors.....	15-23
15.1.4.3	Troubleshooting.....	15-24
15.1.5	Oracle Hyperion Provider Services Component Architecture	15-24
15.1.5.1	Oracle Hyperion Provider Services Component Characteristics	15-24
15.1.5.1.1	State Information	15-24
15.1.5.1.2	Runtime Processes.....	15-24
15.1.5.1.3	Process Lifecycle.....	15-25
15.1.5.1.4	Request Flow	15-25
15.1.5.1.5	External Dependencies	15-25
15.1.5.1.6	Configuration Artifacts.....	15-25
15.1.5.1.7	Deployment Artifacts.....	15-25
15.1.5.1.8	Log Files.....	15-25
15.1.6	Oracle Hyperion Provider Services High Availability Concepts	15-25
15.1.6.1	Oracle Hyperion Provider Services High Availability Architecture	15-26
15.1.6.1.1	Hyperion Registry Structure for a Provider Services Cluster.....	15-26
15.1.6.1.2	Cluster-Wide Configuration Changes.....	15-26
15.1.6.1.3	OPMN Essbase Cluster Support	15-26
15.1.6.1.4	Essbase Database Clustering by Provider Services	15-26
15.1.6.2	Protection from Failures and Expected Behaviors.....	15-26
15.1.7	Oracle EPM Workspace Component Architecture	15-27
15.1.7.1	Workspace Component Characteristics	15-27
15.1.7.1.1	State Information	15-28
15.1.7.1.2	Runtime Processes.....	15-28
15.1.7.1.3	Process Lifecycle.....	15-28

15.1.7.1.4	Request Flow	15-28
15.1.7.1.5	External Dependencies	15-28
15.1.7.1.6	Configuration Artifacts.....	15-28
15.1.7.1.7	Deployment Artifacts.....	15-28
15.1.7.1.8	Log Files.....	15-28
15.1.8	Oracle EPM Workspace High Availability Concepts.....	15-28
15.1.8.1	Workspace High Availability Architecture	15-28
15.1.8.1.1	Shared Files and Directories	15-29
15.1.8.1.2	Cluster-Wide Configuration Changes.....	15-29
15.1.8.2	Protection from Failures and Expected Behaviors.....	15-29
15.1.8.3	Troubleshooting.....	15-29
15.1.9	Oracle Hyperion Financial Reporting Component Architecture	15-29
15.1.9.1	Oracle Hyperion Financial Reporting Component Characteristics	15-30
15.1.9.1.1	State Information	15-30
15.1.9.1.2	Runtime Processes.....	15-31
15.1.9.1.3	External Dependencies	15-33
15.1.9.1.4	Configuration Artifacts.....	15-33
15.1.9.1.5	Deployment Artifacts.....	15-33
15.1.9.1.6	Log Files.....	15-34
15.1.10	Oracle Hyperion Financial Reporting High Availability Concepts	15-34
15.1.10.1	Oracle Hyperion Financial Reporting High Availability Architecture	15-34
15.1.10.1.1	Shared Files and Directories	15-34
15.1.10.1.2	Cluster-Wide Configuration Changes.....	15-34
15.1.10.2	Protection from Failures and Expected Behaviors.....	15-34
15.1.11	Allocation Manager Component Architecture and Characteristics	15-35
15.1.12	Oracle BI EE High Availability Configuration Steps.....	15-35
15.1.12.1	Prerequisite Steps Before Setting Up a High Availability Configuration for Oracle BI Enterprise Edition and BI Publisher 15-35	
15.1.12.1.1	Database Prerequisites.....	15-36
15.1.12.1.2	VIP and IP Prerequisites.....	15-36
15.1.12.1.3	Shared Storage Prerequisites	15-36
15.1.12.1.4	Clock Synchronization.....	15-37
15.1.12.1.5	Installing and Configuring the Database Repository	15-37
15.1.12.1.6	Using RCU to Load the Business Intelligence Schemas in the Database .	15-38
15.1.12.1.7	Configuring Virtual Server Names and Ports for the Load Balancer	15-39
15.1.12.1.8	Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2	15-39
15.1.12.1.9	Validating Oracle HTTP Server.....	15-42
15.1.12.2	Installing Oracle Business Intelligence Enterprise Edition for High Availability	15-42
15.1.12.2.1	Installing Oracle WebLogic Server	15-42
15.1.12.2.2	Installing Oracle Fusion Middleware for Business Intelligence Enterprise Edition and Oracle BI Publisher 15-43	
15.1.12.3	Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2	15-44
15.1.12.4	Creating a Domain with the Administration Server and the First BI_SERVER1 Managed Server 15-44	
15.1.12.5	Creating boot.properties for the Administration Server on APPHOST1	15-45
15.1.12.6	Starting and Validating the Administration Server on APPHOST1	15-46
15.1.12.6.1	Starting the Administration Server on APPHOST1	15-46

15.1.12.6.2	Validating the Administration Server	15-46
15.1.12.7	Setting the Listen Address for BI_SERVER1 Managed Server	15-46
15.1.12.7.1	Updating the Oracle BI Publisher Scheduler Configuration	15-47
15.1.12.8	Disabling Host Name Verification for the BI_SERVER1 Managed Server	15-47
15.1.12.9	Starting the System in APPHOST1	15-48
15.1.12.9.1	Starting Node Manager on APPHOST1.....	15-48
15.1.12.9.2	Starting and Validating the BI_SERVER1 Managed Server.....	15-48
15.1.12.9.3	Starting and Validating the BI EE System Components on APPHOST1 .	15-48
15.1.12.10	Configuring Oracle BI EE	15-49
15.1.12.10.1	Setting the Location of the Shared Oracle BI Repository	15-49
15.1.12.10.2	Setting the Shared Global Cache for BI Server.....	15-49
15.1.12.10.3	Setting the Scheduler Script Path and Default Script Path.....	15-50
15.1.12.10.4	Setting the Location of the Shared Oracle BI Presentation Catalog	15-50
15.1.12.11	Setting Server Configuration Options	15-51
15.1.12.12	Scaling Out the BI System on APPHOST2.....	15-51
15.1.12.13	Scaling Out BI System Components	15-52
15.1.12.14	Making Singleton Components Active-Passive	15-53
15.1.12.15	Setting the Listen Address for the BI_SERVER2 Managed Server.....	15-53
15.1.12.15.1	Updating the Oracle BI Publisher Scheduler Configuration on APPHOST1 and APPHOST2	15-54
15.1.12.16	Disabling Host Name Verification for the BI_SERVER2 Managed Server	15-54
15.1.12.17	Configuring Oracle BI for Microsoft Office SSO Properties.....	15-54
15.1.12.17.1	Validating Oracle BI for Microsoft Office Configuration	15-56
15.1.12.18	Configuring Oracle BI Publisher	15-58
15.1.12.18.1	Setting Oracle BI Presentation Services Options	15-58
15.1.12.18.2	Setting Scheduler Configuration Options.....	15-58
15.1.12.18.3	Setting the Oracle BI EE Data Source	15-58
15.1.12.18.4	Configuring JMS for Oracle BI Publisher.....	15-59
15.1.12.18.5	Configuring a Default Persistence Store for Transaction Recovery	15-60
15.1.12.19	Starting the System in APPHOST2	15-60
15.1.12.19.1	Starting Node Manager on APPHOST2.....	15-61
15.1.12.19.2	Starting and Validating the BI_SERVER2 Managed Server.....	15-61
15.1.12.19.3	Starting and Validating the Business Intelligence Enterprise Edition System Components	15-61
15.1.12.20	Configuring Oracle HTTP Server for the BI_SERVERn Managed Servers.....	15-61
15.1.12.20.1	Validating Access Through Oracle HTTP Server	15-63
15.1.12.21	Setting the Frontend HTTP Host and Port	15-63
15.1.12.22	Configuring Server Migration for the BI_SERVERn Servers	15-64
15.1.12.22.1	Setting Up a User and Tablespace for the Server Migration Leasing Table	15-64
15.1.12.22.2	Creating a Multi Data Source Using the Oracle WebLogic Server Administration Console	15-65
15.1.12.22.3	Editing Node Manager's Properties File.....	15-66
15.1.12.22.4	Setting Environment and Superuser Privileges for the wlsifconfig.sh Script	15-68
15.1.12.22.5	Configuring Server Migration Targets.....	15-68
15.1.12.22.6	Testing the Server Migration	15-69
15.1.12.23	Scaling Up the Oracle BI EE Topology	15-70

15.1.12.24	Scaling Out the Oracle BI EE Topology to a New Node (APPHOST3)	15-71
15.2	High Availability for Oracle Business Intelligence Publisher	15-73
15.2.1	Oracle BI Publisher Component Architecture	15-73
15.2.1.1	Oracle BI Publisher Component Characteristics	15-74
15.2.1.1.1	State Information	15-74
15.2.1.1.2	Runtime Processes	15-74
15.2.1.1.3	Process Lifecycle	15-75
15.2.1.1.4	Request Flow	15-75
15.2.1.1.5	External Dependencies	15-75
15.2.1.1.6	Configuration Artifacts	15-75
15.2.1.1.7	Deployment Artifacts	15-75
15.2.1.1.8	Log Files	15-75
15.2.2	Oracle BI Publisher High Availability Concepts	15-76
15.2.2.1	Oracle BI Publisher High Availability Architecture	15-76
15.2.2.1.1	Shared Files and Directories	15-77
15.2.2.1.2	Cluster-Wide Configuration Changes	15-78
15.2.2.2	Protection from Failures and Expected Behaviors	15-78
15.2.2.3	Troubleshooting	15-78
15.2.3	Oracle BI Publisher High Availability Configuration Steps	15-78
15.2.3.1	Preparing the Environment: Prerequisite Steps Before Setting Up an Oracle BI Publisher High Availability Configuration 15-78	
15.2.3.1.1	Database Prerequisites	15-78
15.2.3.1.2	VIP and IP Prerequisites	15-79
15.2.3.1.3	Shared Storage Prerequisites	15-79
15.2.3.1.4	Clock Synchronization	15-80
15.2.3.1.5	Installing and Configuring the Database Repository	15-80
15.2.3.1.6	Using RCU to Load the Business Intelligence Schemas in the Database .	15-80
15.2.3.1.7	Configuring Virtual Server Names and Ports for the Load Balancer	15-82
15.2.3.1.8	Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2	15-82
15.2.3.1.9	Validating Oracle HTTP Server	15-84
15.2.3.2	Installing Oracle Fusion Middleware Home	15-84
15.2.3.2.1	Installing Oracle WebLogic Server	15-85
15.2.3.2.2	Installing Oracle Fusion Middleware for BI Publisher	15-85
15.2.3.3	Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2	15-86
15.2.3.4	Creating a Domain with the Administration Server and the First BI_SERVER1 Managed Server 15-86	
15.2.3.5	Creating boot.properties for the Administration Server on APPHOST1	15-88
15.2.3.6	Starting and Validating the Administration Server on APPHOST1	15-88
15.2.3.6.1	Starting the Administration Server on APPHOST1	15-88
15.2.3.6.2	Validating the Administration Server	15-88
15.2.3.7	Setting the Listen Address for BI_SERVER1 Managed Server	15-89
15.2.3.8	Disabling Host Name Verification for the BI_SERVER1 Managed Server	15-89
15.2.3.9	Starting the System in APPHOST1	15-89
15.2.3.9.1	Starting Node Manager on APPHOST1	15-89
15.2.3.9.2	Starting and Validating the BI_SERVER1 Managed Server	15-90
15.2.3.10	Prerequisites for Scaling Out the BI System on APPHOST2	15-90
15.2.3.10.1	Setting Server Configuration Options	15-90
15.2.3.10.2	Setting Scheduler Configuration Options	15-90

15.2.3.10.3	Updating the Oracle BI Publisher Scheduler Configuration	15-91
15.2.3.11	Scaling Out the BI System on APPHOST2.....	15-91
15.2.3.12	Setting the Listen Address for the BI_SERVER2 Managed Server.....	15-92
15.2.3.13	Disabling Host Name Verification for the BI_SERVER2 Managed Server	15-92
15.2.3.14	Configuring Oracle BI Publisher	15-93
15.2.3.14.1	Configuring JMS Persistence Store for BI Publisher	15-93
15.2.3.14.2	Configuring a Default Persistence Store for Transaction Recovery	15-94
15.2.3.15	Starting the System in APPHOST2	15-95
15.2.3.15.1	Starting Node Manager on APPHOST2.....	15-95
15.2.3.15.2	Starting and Validating the BI_SERVER2 Managed Server.....	15-95
15.2.3.16	Configuring Oracle HTTP Server for the BI_SERVERn Managed Servers	15-95
15.2.3.17	Validating Access Through Oracle HTTP Server	15-96
15.2.3.18	Configuring Server Migration for the BI_SERVERn Servers	15-97
15.2.3.19	Scaling Out the Oracle BI Publisher Topology to a New Node (APPHOST3)	15-97
15.3	High Availability for Oracle Real-Time Decisions.....	15-98
15.3.1	Oracle RTD Component Architecture	15-98
15.3.1.1	Oracle RTD Component Characteristics	15-101
15.3.1.1.1	Component Lifecycle	15-102
15.3.1.1.2	Process Flow.....	15-102
15.3.1.1.3	External Dependencies	15-102
15.3.1.1.4	Configuration Artifacts.....	15-102
15.3.1.1.5	Deployment Artifacts.....	15-102
15.3.1.1.6	Log File Locations	15-102
15.3.2	Oracle RTD High Availability Concepts	15-103
15.3.2.1	Oracle RTD High Availability Architecture	15-103
15.3.2.1.1	Starting and Stopping the Cluster.....	15-106
15.3.2.1.2	Cluster-Wide Configuration Changes.....	15-106
15.3.2.2	Protection from Failures and Expected Behaviors.....	15-106
15.3.2.2.1	Decision Server Failure.....	15-106
15.3.2.2.2	Cluster Coordinator Failure.....	15-106
15.3.2.2.3	Learning Service Failure.....	15-106
15.3.2.2.4	Decision Center Failure	15-107
15.3.2.2.5	Batch Manager Failure.....	15-107
15.3.3	Oracle RTD High Availability Configuration Steps	15-107
15.3.3.1	Prerequisite Steps Before Setting up an Oracle RTD High Availability Configuration 15-107	
15.3.3.1.1	Database Prerequisites.....	15-107
15.3.3.1.2	Installing and Configuring the Database Repository	15-108
15.3.3.1.3	Using RCU to Load the Business Intelligence Schemas into the Database.....	15-108
15.3.3.1.4	Configuring Virtual Server Names and Ports for the Load Balancer	15-109
15.3.3.1.5	Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2	15-110
15.3.3.1.6	Validating Oracle HTTP Server.....	15-112
15.3.3.2	Installing Oracle Fusion Middleware Home	15-112
15.3.3.2.1	Installing Oracle WebLogic Server	15-113
15.3.3.2.2	Installing Oracle RTD	15-113
15.3.3.3	Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2	15-114

15.3.3.4	Creating a Domain with the Administration Server and the First BI_SERVER1 Managed Server 15-114	
15.3.3.5	Creating boot.properties for the Administration Server on APPHOST1	15-116
15.3.3.6	Starting and Validating the Administration Server on APPHOST1	15-116
15.3.3.6.1	Starting the Administration Server on APPHOST1	15-116
15.3.3.6.2	Validating the Administration Server	15-116
15.3.3.7	Setting the Listen Address for BI_SERVER1 Managed Server	15-116
15.3.3.8	Disabling Host Name Verification for the BI_SERVER1 Managed Server ...	15-117
15.3.3.9	Starting the System in APPHOST1	15-117
15.3.3.9.1	Starting Node Manager on APPHOST1.....	15-117
15.3.3.9.2	Starting and Validating the BI_SERVER1 Managed Server.....	15-117
15.3.3.10	Scaling Out the BI System on APPHOST2.....	15-118
15.3.3.11	Setting the Listen Address for the BI_SERVER2 Managed Server.....	15-119
15.3.3.12	Disabling Host Name Verification for the BI_SERVER2 Managed Server ...	15-119
15.3.3.13	Configuring Oracle RTD	15-120
15.3.3.13.1	Configuring RTD Cluster-Specific Properties.....	15-120
15.3.3.14	Starting the System in APPHOST2	15-121
15.3.3.14.1	Configuring a Default Persistence Store for Transaction Recovery	15-121
15.3.3.14.2	Starting Node Manager on APPHOST2.....	15-122
15.3.3.14.3	Starting and Validating the BI_SERVER2 Managed Server.....	15-122
15.3.3.15	Configuring Oracle HTTP Server for the BI_SERVERn Managed Servers ...	15-122
15.3.3.16	Validating Access Through Oracle HTTP Server	15-123
15.3.3.17	Configuring Server Migration for the BI_SERVERn Servers	15-124
15.3.3.18	Scaling Out the Oracle RTD Topology to a New Node (APPHOST3).....	15-124

16 Using HA Power Tools

16.1	Enabling HA Power Tools in the Oracle WebLogic Administration Console	16-1
16.2	Configuring Java Object Cache for a Cluster Using HA Power Tools	16-1
16.3	Configuring Java Object Cache for Managed Servers Using HA Power Tools	16-2

A Setting Up Auditing with an Oracle RAC Database Store

A.1	Using WebLogic Server to Configure Audit Data Sources and Multi Data Sources	A-2
A.2	Configuring the JDBC String for the Audit Loader	A-3

B Recommended Multi Data Sources

B.1	JDBC Multi Data Source-0	B-1
B.2	JDBC Data Source-0 (non-XA)	B-1
B.3	JDBC Data Source-0 (XA).....	B-2

C Whole Server Migration for Windows

C.1	Using Windows Control Panel	C-1
C.2	Using the netsh Command Line	C-1

D Component Workbooks

D.1	Oracle SOA Suite Workbook.....	D-1

D.1.1	Workbook Tables for Oracle SOA Suite	D-1
D.2	Oracle Identity Management Workbook.....	D-4
D.2.1	Workbook Tables for Oracle Identity Management	D-4
D.3	Oracle WebCenter Portal Workbook	D-10
D.3.1	Workbook Tables for Oracle WebCenter Portal.....	D-10
D.4	Oracle Portal, Forms, Reports, and Discoverer Workbook.....	D-13
D.4.1	Workbook Tables for Oracle Portal, Forms, Reports, and Discoverer	D-13
D.5	Oracle Enterprise Content Management Workbook	D-15
D.5.1	Workbook Tables for Oracle WebCenter Content	D-16
D.6	Oracle Data Integrator Workbook.....	D-18
D.6.1	Workbook Tables for Oracle Data Integrator	D-18
D.7	Oracle Business Intelligence Platform Workbook.....	D-20
D.7.1	Workbook Tables for Oracle Business Intelligence Platform	D-20

E ascrsctl Online Help

E.1	start.....	E-1
E.2	stop	E-2
E.3	status	E-2
E.4	switch.....	E-3
E.5	delete.....	E-4
E.6	create/disk.....	E-4
E.7	update/disk.....	E-6
E.8	create/vip.....	E-8
E.9	update/vip.....	E-9
E.10	create/dblsnr	E-10
E.11	update/dblsnr	E-11
E.12	create/db.....	E-12
E.13	update/db	E-14
E.14	create/as.....	E-16
E.15	update/as.....	E-17

F Configuring the Distributed Java Object Cache

F.1	Configuring the Java Object Cache	F-1
F.2	Running CacheWatcher	F-3

G Configuring Distributed Notifications for MDS

Preface

This preface contains these sections:

- [Intended Audience](#)
- [Documentation Accessibility](#)
- [Related Documentation](#)
- [Conventions](#)

Intended Audience

The *Oracle Fusion Middleware High Availability Guide* is intended for administrators, developers, and others whose role is to deploy and manage Oracle Fusion Middleware with high availability requirements.

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/support/contact.html> or visit <http://www.oracle.com/accessibility/support.html> if you are hearing impaired.

Related Documentation

For more information, see these Oracle resources:

- *Oracle Fusion Middleware Concepts*
- *Oracle Fusion Middleware Administrator's Guide*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Introduction to High Availability

A high availability architecture is one of the key requirements for any Enterprise Deployment. Oracle Fusion Middleware has an extensive set of high availability features, which protect its components and applications from unplanned down time and minimize planned downtime.

The solutions and procedures described in this book are designed to eliminate single points of failure for Oracle Fusion Middleware components with no or minimal down time. These solutions help ensure that applications that deployed with Oracle Fusion Middleware meet the required availability to achieve your business goals.

This guide discusses the architecture, interaction, and dependencies of Oracle Fusion Middleware components, and explains how they can be deployed in a high availability architecture.

This chapter explains high availability and its importance from the perspective of Oracle Fusion Middleware. This chapter includes the following sections:

- [Section 1.1, "What is High Availability"](#)
- [Section 1.2, "High Availability Information in Other Documentation"](#)

1.1 What is High Availability

High availability refers to the ability of users to access a system without loss of service. Deploying a high availability system minimizes the time when the system is down, or unavailable and maximizes the time when it is running, or available. This section provides an overview of high availability from a problem-solution perspective. This section includes the following topics:

- [Section 1.1.1, "High Availability Problems"](#)
- [Section 1.1.2, "High Availability Solutions"](#)

1.1.1 High Availability Problems

Mission critical computer systems need to be available 24 hours a day, 7 days a week, and 365 days a year. However, part or all of the system may be down during planned or unplanned downtime. A system's availability is measured by the percentage of time that it is providing service in the total time since it is deployed. [Table 1–1](#) provides an example.

Table 1–1 *Availability Percentages and Corresponding Downtime Values*

Availability Percentage	Approximate Downtime Per Year
95%	18 days

Table 1–1 (Cont.) Availability Percentages and Corresponding Downtime Values

Availability Percentage	Approximate Downtime Per Year
99%	4 days
99.9%	9 hours
99.99%	1 hour
99.999%	5 minutes

System downtime may be categorized as planned or unplanned. Unplanned downtime is any sort of unexpected failure. Planned downtime refers to scheduled operations that are known in advance and that render the system unavailable. The effect of planned downtime on end users is typically minimized by scheduling operational windows when system traffic is slow. Unplanned downtime may have a larger effect because it can happen at peak hours, causing a greater impact on system users.

These two types of downtimes (planned and unplanned) are usually considered separately when designing a system's availability requirements. A system's needs may be very restrictive regarding its unplanned downtimes, but very flexible for planned downtimes. This is the typical case for applications with high peak loads during working hours, but that remain practically inactive at night and during weekends. You may choose different high availability features depending on the type of failure is being addressed.

1.1.2 High Availability Solutions

High availability solutions can be categorized into local high availability solutions that provide high availability in a single data center deployment, and disaster recovery solutions, which are usually geographically distributed deployments that protect your applications from disasters such as floods or regional network outages.

Among possible types of failures, process, node, and media failures as well as human errors can be protected by local high availability solutions. Local physical disasters that affect an entire data center can be protected by geographically distributed disaster recovery solutions.


To solve the high availability problem, a number of technologies and best practices are needed. The most important mechanism is redundancy. High availability comes from redundant systems and components. You can categorize local high availability solutions by their level of redundancy, into active-active solutions and active-passive solutions (see [Figure 1–1](#)):

- **Active-active solutions** deploy two or more active system instances and can be used to improve scalability and provide high availability. In active-active deployments, all instances handle requests concurrently.
- **Active-passive solutions** deploy an active instance that handles requests and a passive instance that is on standby. In addition, a heartbeat mechanism is set up between these two instances. This mechanism is provided and managed through operating system vendor-specific clusterware. Generally, vendor-specific cluster agents are also available to automatically monitor and failover between cluster nodes, so that when the active instance fails, an agent shuts down the active instance completely, brings up the passive instance, and application services can successfully resume processing. As a result, the active-passive roles are now switched. The same procedure can be done manually for planned or unplanned

downtime. Active-passive solutions are also generally referred to as cold failover clusters.

You can use Oracle Cluster Ready Services (CRS) to manage the Fusion Middleware Active-Passive (CFC) solutions.

Figure 1–1 Active-Active and Active-Passive High Availability Solutions



In addition to architectural redundancies, the following local high availability technologies are also necessary in a comprehensive high availability system:

- **Process death detection and automatic restart**

Processes may die unexpectedly due to configuration or software problems. A proper process monitoring and restart system should monitor all system processes constantly and restart them should problems appear.

A system process should also maintain the number of restarts within a specified time interval. This is also important since continually restarting within short time periods may lead to additional faults or failures. Therefore a maximum number of restarts or retries within a specified time interval should also be designed as well.

- **Clustering**

Clustering components of a system together allows the components to be viewed functionally as a single entity from the perspective of a client for runtime processing and manageability. A cluster is a set of processes running on single or multiple computers that share the same workload. There is a close correlation between clustering and redundancy. A cluster provides redundancy for a system.

If failover occurs during a transaction in a clustered environment, the session data is retained as long as there is at least one surviving instance available in the cluster.

- **State replication and routing**

For stateful applications, client state can be replicated to enable stateful failover of requests in the event that processes servicing these requests fail.

- **Failover**

With a load-balancing mechanism in place, the instances are redundant. If any of the instances fail, requests to the failed instance can be sent to the surviving instances.

- **Server load balancing**

When multiple instances of identical server components are available, client requests to these components can be load balanced to ensure that the instances have roughly the same workload.

- **Server Migration**

Some services can only have one instance running at any given point of time. If the active instance becomes unavailable, the service is automatically started on a different cluster member. Alternatively, the whole server process can be automatically started on a different system in the cluster.

- **Integrated High Availability**

Components depend on other components to provide services. The component should be able to recover from dependent component failures without any service interruption.

- **Rolling Patching**

Patching product binaries often requires down time. Patching a running cluster in a rolling fashion can avoid downtime. Patches can be uninstalled in a rolling fashion as well.

- **Configuration management**

A clustered group of similar components often need to share common configuration. Proper configuration management ensures that components provide the same reply to the same incoming request, allows these components to synchronize their configurations, and provides high availability configuration management for less administration downtime.


- **Backup and Recovery**

User errors may cause a system to malfunction. In certain circumstances, a component or system failure may not be repairable. A backup and recovery facility should be available to back up the system at certain intervals and restore a backup when an unrepairable failure occurs.

Disaster Recovery

Disaster recovery solutions typically set up two homogeneous sites, one active and one passive. Each site is a self-contained system. The active site is generally called the production site, and the passive site is called the standby site. During normal operation, the production site services requests; if a site failover or switchover occurs, the standby site takes over the production role and all requests are routed to that site. To maintain the standby site for failover, not only must the standby site contain homogeneous installations and applications, data and configurations must also be synchronized constantly from the production site to the standby site.

Figure 1–2 Geographically Distributed Disaster Recovery



Oracle Fusion Middleware Components Protected by High Availability Solutions

The Oracle Fusion Middleware High Availability Guide discusses high availability solutions for the following components:

- Oracle WebLogic Server
- Oracle SOA Suite
- Oracle ADF
- Oracle WebCenter Portal
- Oracle Identity Management Components
- Oracle HTTP Server
- Oracle Web Cache
- Oracle Portal, Forms, Reports, and Discoverer

1.2 High Availability Information in Other Documentation

[Table 1–2](#) lists Oracle Fusion Middleware guides (other than this guide) that contain high availability information. This information pertains to high availability of various Oracle Fusion Middleware components.

Table 1–2 High Availability Information in Oracle Fusion Middleware Documentation

Component	Location of Information
Oracle SOA Suite	<i>Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite</i> <i>Oracle Fusion Middleware Installation Guide for Oracle SOA Suite</i> <i>Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite</i>
Oracle WebCenter Portal	<i>Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal</i> <i>Oracle Fusion Middleware Installation Guide for Oracle WebCenter</i> <i>Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal</i>
Oracle ADF	<i>Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework</i> <i>Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework</i>
Oracle Data Integrator	<i>Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator</i> <i>Oracle Fusion Middleware Connectivity and Knowledge Modules Guide for Oracle Data Integrator</i> <i>Oracle Fusion Middleware Knowledge Module Developer's Guide for Oracle Data Integrator</i>
Oracle WebLogic Server Clusters	<i>Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server</i>
Oracle Fusion Middleware Backup and Recovery	<i>Oracle Fusion Middleware Administrator's Guide</i>
Oracle Web Cache	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache</i>
Oracle Identity Management	<i>Oracle Fusion Middleware Installation Guide for Oracle Identity Management</i> <i>Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management</i>
Oracle Virtual Directory	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory</i>
Oracle HTTP Server	<i>Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server</i>

Table 1–2 (Cont.) High Availability Information in Oracle Fusion Middleware Documentation

Component	Location of Information
Oracle Internet Directory	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory</i>
Oracle Access Manager	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service</i>
Oracle Authorization Policy Manager	<i>Oracle Fusion Middleware Authorization Policy Manager Administrator's Guide</i>
Oracle Identity Manager	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Identity Manager</i>
Oracle Adaptive Access Manager	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Adaptive Access Manager</i>
Oracle Real Application Clusters (Oracle RAC)	<i>Oracle Real Application Clusters Installation Guide</i>
Oracle WebCenter Content	<i>Oracle Fusion Middleware Overview Guide for Oracle Enterprise Content Management</i>
Oracle WebCenter Content: Imaging	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Imaging and Process Management</i>
Oracle Universal Content Management	<i>Oracle WebCenter Content System Administrator's Guide for Content Server</i>
Oracle Universal Records Management	<i>Oracle Fusion Middleware Administrator's Guide for Universal Records Management</i>
Oracle Repository Creation Utility (RCU)	<i>Oracle Fusion Middleware Repository Creation Utility User's Guide</i>
Oracle Portal	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Portal</i>
Oracle Forms	<i>Oracle Fusion Middleware Forms Services Deployment Guide</i>
Oracle Reports	<i>Oracle Fusion Middleware Oracle Reports User's Guide to Building Reports</i>
Oracle Business Intelligence Discoverer	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Discoverer</i>
Oracle Business Intelligence Enterprise Edition	<i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Oracle Real-Time Decisions	<i>Oracle Fusion Middleware Administrator's Guide for Oracle Real-Time Decisions</i>

Oracle Fusion Middleware High Availability Framework

This chapter describes the Oracle Fusion Middleware features that are important in high availability topologies. It contains the following topics:

- [Section 2.1, "Understanding Key Oracle Fusion Middleware Concepts"](#)
- [Section 2.2, "Oracle Fusion Middleware High Availability Terminology"](#)
- [Section 2.3, "Oracle Fusion Middleware High Availability Solutions"](#)
- [Section 2.4, "Protection from Planned and Unplanned Down Time"](#)

2.1 Understanding Key Oracle Fusion Middleware Concepts

Oracle Fusion Middleware provides two types of components:


- **Java components:** A Java component is a peer of a system component but is deployed as one or more Java EE applications and a set of resources. Java components are deployed to an Oracle WebLogic Server domain as part of a domain template. Examples of Java components are Oracle SOA Suite and Oracle WebCenter Portal: Spaces.
- **System components:** A system component is a manageable process that is not deployed as a Java application. Instead, a system component is managed by the Oracle Process Manager and Notification (OPMN). Examples of system components include Oracle HTTP Server, Oracle Internet Directory, Oracle Web Cache, and Oracle Forms Services.

A Java component and a system component are peers.

After you install and configure Oracle Fusion Middleware, your Oracle Fusion Middleware environment contains the following:

- An Oracle WebLogic Server domain, which contains one Administration Server and one or more managed servers.
- If your environment includes system components, one or more system component domains.
- An Oracle Metadata Repository, if the components you installed require one. For example, Oracle SOA Suite requires an Oracle Metadata Repository.

[Figure 2-1](#) shows an Oracle Fusion Middleware environment with an Oracle WebLogic Server domain with an Administration Server and two managed servers, a system component domain, and an Oracle Metadata Repository.

Figure 2–1 Oracle Fusion Middleware Environment

Your environment also includes a Middleware home, which consists of the Oracle WebLogic Server home, and, optionally, one or more Oracle homes.

2.1.1 What is a WebLogic Server Domain?

A WebLogic Server administration **domain** is a logically related group of Java components. A domain includes a special WebLogic Server instance called the **Administration Server**, which is the central point from which you configure and manage all resources in the domain. Usually, you configure a domain to include additional WebLogic Server instances called *managed servers*. You deploy Java components, such as Web applications, EJBs, and Web services, and other resources to the managed servers and use the Administration Server for configuration and management purposes only.


Managed servers in a domain can be grouped together into a cluster.

An Oracle WebLogic Server Domain is a peer of a system component domain. Both contain specific configurations outside of their Oracle homes.

The directory structure of an WebLogic Server domain is separate from the directory structure of the WebLogic Server Home. It can reside anywhere; it need not be within the Middleware home directory.

Figure 2–2 shows a Oracle WebLogic Server domain with an Administration Server, three standalone managed servers, and three managed servers in a cluster.

Figure 2–2 Oracle WebLogic Server Domain



See Also: *Oracle Fusion Middleware Understanding Domain Configuration for Oracle WebLogic Server* for more information about domain configuration

The following topics describe entities in the domain:

- [What Is the Administration Server?](#)
- [Understanding Managed Servers and Managed Server Clusters](#)
- [What Is Node Manager?](#)

2.1.1.1 What Is the Administration Server?

The **Administration Server** operates as the central control entity for the configuration of the entire domain. It maintains the domain's configuration documents and distributes changes in the configuration documents to managed servers. You can use the Administration Server as a central location from which to monitor all resources in a domain.

Each WebLogic Server domain must have one server instance that acts as the Administration Server.

To interact with the Administration Server, you can use the Oracle WebLogic Server Administration Console, Oracle WebLogic Scripting Tool (WLST), or create your own JMX client. In addition, you can use Oracle Enterprise Manager Fusion Middleware Control for some tasks.

Fusion Middleware Control and the WebLogic Administration Console run in the Administration Server. Fusion Middleware Control is a Web-based administration console used to manage Oracle Fusion Middleware, including components such as Oracle HTTP Server, Oracle SOA Suite and Oracle WebCenter Portal, Oracle Portal, Forms, Reports, and Discoverer, and the Oracle Identity Management components. Oracle WebLogic Server Administration Console is the Web-based administration

console used to manage the resources in an Oracle WebLogic Server domain, including the Administration Server and managed servers in the domain.

2.1.1.2 Understanding Managed Servers and Managed Server Clusters

Managed servers host business applications, application components, Web services, and their associated resources. To optimize performance, managed servers maintain a read-only copy of the domain's configuration document. When a managed server starts up, it connects to the domain's Administration Server to synchronize its configuration document with the document that the Administration Server maintains.

When you create a domain, you create it using a particular domain template. That template supports a particular component or group of components, such as the Oracle SOA Suite. The Managed Servers in the domain are created specifically to host those particular Oracle Fusion Middleware system components.

Java-based Oracle Fusion Middleware system components (such as Oracle SOA Suite, Oracle WebCenter Portal, and some Identity Management components) and customer-developed applications are deployed to Managed Servers in the domain.

If you want to add other components, such as Oracle WebCenter Portal, to a domain that was created using a template that supports another component, you can extend the domain by creating additional Managed Servers in the domain, using a domain template for the component which you want to add.

For production environments that require increased application performance, throughput, or high availability, you can configure two or more Managed Servers to operate as a cluster. A **cluster** is a collection of multiple WebLogic Server server instances running simultaneously and working together to provide increased scalability and reliability. In a cluster, most resources and services are deployed identically to each Managed Server (as opposed to a single Managed Server), enabling failover and load balancing. A single domain can contain multiple WebLogic Server clusters and multiple Managed Servers that are not configured as clusters. The key difference between clustered and non-clustered Managed Servers is support for failover and load balancing. These features are available only in a cluster of Managed Servers.

See Also: "Understanding WebLogic Server Clustering" in *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*

2.1.1.3 What Is Node Manager?

Node Manager is a Java utility that runs as separate process from Oracle WebLogic Server and allows you to perform common operations for a Managed Server, regardless of its location with respect to its Administration Server. While use of Node Manager is optional, it provides valuable benefits if your WebLogic Server environment hosts applications with high-availability requirements.

If you run Node Manager on a system that hosts Managed Servers, you can start and stop the Managed Servers remotely using the Administration Console or the command line. Node Manager can also automatically restart a Managed Server after an unexpected failure.

See Also: *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*

2.1.2 What Is a System Component Domain?

A **system component domain** contains one or more system components, such as Oracle Web Cache, Oracle HTTP Server, or Oracle Internet Directory. The system components in a system component domain must reside on the same system. A system component domain directory contains files that can be updated, such as configuration files, log files, and temporary files.

A system component domain is a peer of a Oracle WebLogic Server domain. Both contain specific configurations outside of their Oracle homes.

The directory structure of a system component domain is separate from the directory structure of the Oracle home. It can reside anywhere; it need not be within the Middleware home directory.

2.1.3 What Is a Middleware Home?

A **Middleware home** consists of the Oracle WebLogic Server home, and, optionally, one or more Oracle homes.

A Middleware home can reside on a local file system or on a remote shared disk that is accessible through NFS.

See [Section 2.1.4, "What Is an Oracle Home?"](#) for information about Oracle homes. See [Section 2.1.1, "What is a WebLogic Server Domain?"](#) for information about Oracle WebLogic Server homes.

In a high availability installation where two or more hosts are clustered, the Middleware Home binaries must be installed into directories with the exact same directory paths on all the hosts in the cluster.

2.1.4 What Is an Oracle Home?

An **Oracle home** contains installed files necessary to host a specific product. For example, the SOA Oracle home contains a directory that contains binary and library files for Oracle SOA Suite.

An Oracle home resides within the directory structure of the Middleware home. Each Oracle home can be associated with multiple system component domains or Oracle WebLogic Server domains.

In a high availability installation where two or more hosts are clustered, the Oracle Home binaries must be installed into directories with the exact same directory paths on all the hosts in the cluster.

2.1.4.1 What Is an Oracle Common Home?

The **Oracle Common home** contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF). There can be only one Oracle Common home within each Middleware home.

2.1.5 What Is a WebLogic Server Home?

A WebLogic Server home contains installed files necessary to host a WebLogic Server. The WebLogic Server home directory is a peer of Oracle home directories and resides within the directory structure of the Middleware home.

2.2 Oracle Fusion Middleware High Availability Terminology

The definitions of terms listed in this section are useful in helping to understand the concepts presented in this book:

- **failover:** When a member of a high availability system fails unexpectedly (unplanned downtime), the system undergoes a failover operation to continue offering services to its consumers. For an *active-active* system, the load balancer entity serving requests to the active members performs the failover. If an active member fails, the load balancer detects the failure and automatically redirects requests for the failed member to surviving active members.

For an *active-passive* system, the passive member activates during the failover operation and consumers are directed to it. You can perform the failover process or it automate it by setting up hardware cluster services to detect failures and move cluster resources from the failed node to the standby node.

- **failback:** A planned operation after unplanned downtime. After a system undergoes a successful failover operation, you can repair the original failed member and re-introduce into the system as a standby member. You can initiate a failback process to activate this member and deactivate the other. This process reverts the system back to its pre-failure configuration.
- **shared storage:** Although each node in a cluster is a standalone server that runs its own set of processes, there are some file system based data and configuration elements which need uniform access from all nodes in a cluster. Shared storage refers to the ability of the cluster to be able to access the same storage, usually disks, from any node in the cluster.

For SAN based deployments, a clustered file system, such as OCFS, may also be needed. Some examples of this usage are, JMS file based persistence store, transaction logs persistence store, domain configuration in case of cold failover cluster setup.

- **primary node:** The node that is actively running Oracle Fusion Middleware at any given time in a cluster. If this node fails, Oracle Fusion Middleware fails over to the secondary node. Because the primary node runs the active Oracle Fusion Middleware installation(s), if this node fails, Oracle Fusion Middleware fails over to the secondary node. See the definition for secondary node in this section.
- **secondary node:** When the primary node that runs the active Oracle Fusion Middleware installation(s) fails, Oracle Fusion Middleware fails over to this node. See the definition for primary node in this section.
- **network hostname:** A name assigned to an IP address either through the `/etc/hosts` file (on UNIX), `C:\WINDOWS\system32\drivers\etc\hosts` file (on Windows) or through DNS resolution. This name is visible in the network that the system to which it refers to is connected. Often, the network hostname and physical hostname are identical. However, each system has only one physical hostname but may have multiple network hostnames. Thus, a system's network hostname may not always be its physical hostname.
- **physical hostname:** This guide differentiates between the terms physical hostname and network hostname. This guide uses physical hostname to refer to the internal name of the current system. On UNIX, this is the name returned by the `hostname` command.
- **switchover and switchback:** Planned operations. During normal operation, active members of a system may require maintenance or upgrading. You can start a switchover process to allow a substitute member to take over the workload

performed by the member that requires maintenance, upgrading, or any planned downtime. The switchover operation ensures continued service to system consumers.

When the maintenance or upgrade is complete, you can perform a switchback operation to activate the upgraded member and bring the system back to the pre-switchover configuration.

- **virtual IP:** To present a single system view of a cluster to network clients, a virtual IP serves as an entry point IP address to the group of servers that are cluster members. You can assign a virtual IP to a server load balancer or a hardware cluster.

A load balancer also uses a virtual IP as the entry point to a set of servers. These servers tend to be active at the same time. This virtual IP address is not assigned to any individual server but to the load balancer which acts as a proxy between servers and their clients.

- **virtual hostname:** In a cluster, a network hostname assigned to virtual IP bound to one of the nodes in the cluster at any given time.

Note: When this document uses the term *virtual hostname*, it is assumed to be associated with a virtual IP address. In cases where just the IP address is needed or used, it is explicitly stated.

- **hardware cluster:** A collection of computers that provides a single view of network services (such as an IP address) or application services (such as databases, Web servers) to clients of these services. Each node in a hardware cluster is a standalone server that runs its own processes. These processes can communicate with one another to form what looks like a single system that cooperatively provides applications, system resources, and data to users.

A hardware cluster achieves high availability and scalability through the use of specialized hardware (cluster interconnect, shared storage) and software (health monitors, resource monitors). (The cluster interconnect is a private link that the hardware cluster uses for heartbeat information to detect node death.) Due to the need for specialized hardware and software, hardware clusters are commonly provided by hardware vendors such as Sun, HP, IBM, and Dell. While the number of nodes that can be configured in a hardware cluster is vendor dependent, for the purpose of Oracle Fusion Middleware high availability, only two nodes are required. Hence, this document assumes a two-node hardware cluster for high availability solutions employing a hardware cluster.

- **cluster agent:** The software that runs on a node member of a hardware cluster that coordinates availability and performance operations with other nodes. Clusterware provides resource grouping, monitoring, and the ability to move services. A cluster agent can automate the service failover.
- **clusterware:** A software that manages the operations of cluster members as a system. It enables you to define a set of resources and services to monitor using a heartbeat mechanism between cluster members and to move these resources and services to a different member in the cluster as efficiently and transparently as possible.

2.3 Oracle Fusion Middleware High Availability Solutions


This section describes local high availability concepts, and Oracle Fusion Middleware high availability technologies

2.3.1 Local High Availability

Local high availability solutions can be categorized as either active-active or active-passive solutions. Oracle Fusion Middleware supports both active-active deployments and active-passive deployments.

Figure 2–3 shows an Oracle Fusion Middleware high availability active-active deployment topology.

Figure 2–3 Oracle Fusion Middleware Enterprise Deployment Architecture



As shown in Figure 2–3, this topology represents a multi-tiered architecture. Users access the system from the client tier. Requests go through a hardware load balancer, which then routes them to a Web server cluster that is running Oracle HTTP Server and Oracle Web Cache in the web tier. Web servers use Proxy Plug-in (mod_wl_ohs) to route the requests to the WebLogic cluster in the application tier. Applications running

on the WebLogic cluster in the application tier then interact with the database cluster in the data tier to service the request.

There is no single point of failure in the entire architecture. WebLogic Administration Server is configured in Cold Failover Cluster mode, as described in [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster,"](#) and is protected using external clusterware.

2.3.2 Oracle Fusion Middleware High Availability Technologies

The Oracle Fusion Middleware infrastructure has these high availability features:

- **Process death detection and automatic restart**

For Java EE components running on WebLogic Server, Node Manager monitors the Managed Servers. If a Managed Server goes down, it attempts to restart the Managed Server for a configured number of times.

For system components, OPMN monitors the processes. If a system component process goes down, OPMN attempts to restart the process for a configurable number of times.

- **Clustering**

Oracle Fusion Middleware Java EE components leverage underlying powerful WebLogic Server clustering capabilities to provide clustering. Oracle Fusion Middleware uses WebLogic clustering capabilities, such as redundancy, failover, session state replication, cluster-wide JNDI services, Whole Server Migration, and cluster wide configuration.

These capabilities provide for seamless failover of all Java EE Oracle Fusion Middleware system components transparent to the client preserving session and transaction data as well as ensuring data consistency. For further description of these features, see [Chapter 3, "High Availability for WebLogic Server."](#)

System components can also be deployed in a run time cluster. They are typically front-ended by a load balancer to route traffic.

- **State replication and routing**

Oracle WebLogic Server can be configured for replicating the state of stateful applications. It does so by maintaining a replica of the state information on a different Managed Server, which is a cluster member. Oracle Fusion Middleware components, such as ADF and WebCenter, which are stateful, leverage this feature to ensure seamless failover to other members of the cluster.

System components, such as Oracle Internet Directory, Oracle HTTP Server, Oracle Web Cache are stateless.

Some Oracle Fusion Middleware components, which have part of the functionality implemented in C, such as Oracle Forms and Oracle Reports, are stateful and do not have state replication capabilities. Please refer to following paragraph for information about failover of these components.

- **Failover**

Typically, a Managed Server running Oracle Fusion Middleware Java EE components has a Web server, such as Oracle HTTP Server, clustered in front of it. The Web server proxy plug-in (mod_wl_ohs) is aware of the run time availability of different Managed Servers and the location of the Managed Server on which the state replica is maintained. If the primary Managed Server becomes unavailable, the plug-in routes the request to the server where the application is available. If

stateful, applications such as Oracle ADF and WebCenter Portal, the location of the replica is also taken into account while routing to the new Managed Server.

For stateless system components, their multiple instances deploy as a runtime cluster behind a load balancer. The load balancer is configured to do a periodic health check of the component instances. If an instance is unavailable, the load balancer routes the subsequent requests to another available instance and the failover is seamless.

For stateful components, which have parts based on C, and do not have state replication, sticky routing ensures that the subsequent requests go to the cluster member where the state was initially established. This is ensured by a Web server proxy plug-in and the Java EE parts of the components. If the component instance fails, subsequent requests route to another available member in the cluster. In this situation, the state information is lost and the user must recreate the session.

Some of the internal implementation of components use EJBs. EJB failover is seamlessly handled by replica aware WebLogic Server stubs.

Where needed, components are JTA compliant and data consistency is preserved in case of failover.

Singleton services leverage built-in failover capabilities, such as singleton SOA adapters, or use the underlying WebLogic Server infrastructure, such as Whole Server Migration.

- **Server Migration**

Oracle Fusion Middleware components, such as SOA, which uses pinned services, such as JMS and JTA, leverage WebLogic Server capabilities to provide failover an automatic restart on a different cluster member.

- **Integrated High Availability**

Oracle Fusion Middleware has a comprehensive feature set around load balancing and failover to leverage availability and scalability of Oracle RAC databases. All Oracle Fusion Middleware components have built-in protection against loss of service, data or transactions as a result of Oracle RAC instance unavailability due to planned or unplanned downtime. This is achieved by using Oracle WebLogic Server multi data sources. Additionally, components have proper exception handling and configurable retry logic for seamless failover of in-flight transactions at the time of failure.

For XA compliant applications, such as Oracle SOA components, WebLogic server acts as a Transaction coordinator and ensures that all branches of a transaction are pinned to one of the Oracle RAC instances.

In case of a Managed Server failure the transaction service is automatically migrated over to another node in the cluster and performs the transaction recovery.

For communication between Web servers and application servers, the proxy plug-in has a built-in load balancing and failover capability to seamlessly reroute client requests to an available cluster member.

- **Rolling Patching**

Oracle WebLogic Server allows for rolling patching where a minor maintenance patch can be applied to the product binaries in a rolling fashion without having to shut down the entire cluster.

During the rolling patching of a cluster, each server in the cluster is individually patched and restarted while the other servers in the cluster continue to host your

application. You can also uninstall a patch, maintenance pack, or minor release in a rolling fashion.

- **Configuration Management**

Most of the Oracle Fusion Middleware component configuration can be done at the cluster level. Oracle Fusion Middleware uses WebLogic Server's cluster wide configuration capabilities for server configuration, such as data sources, EJBs, and JMS, as well as component application artifacts, and ADF and WebCenter custom applications.

A central MDS repository available to all members of the cluster stores additional application level components. This includes component level configuration for components, such as Oracle SOA, WebCenter Portal, and application artifacts, such as SOA composites.

- **Backup and Recovery**

Oracle Fusion Middleware backup and recovery is a simple solution based on file system copy for Middle-tier components. RMAN is used for Oracle databases. There is also support for online backups. With Oracle Fusion Middleware, you can integrate with existing backup and recovery tools, or use scheduled backup tasks through Oracle Fusion Middleware Enterprise Manager or cron jobs.

2.3.2.1 Server Load Balancing

Typically, Oracle Fusion Middleware high availability deployments are front ended by a load balancer which can be configured to distribute incoming requests using various algorithms.

Oracle Fusion Middleware also has built-in load balancing capabilities for intra component interaction. For example, Web server to application server, or application server to database server.

Oracle Fusion Middleware 11g does not provide external load balancers. To ensure that your external load balancer is compatible with Oracle Fusion Middleware, check that your external load balancer meets the requirements listed below:

- **Virtual servers and port configuration:** The load balancer should have the ability to configure virtual server names and ports on your external load balancer. The virtual server names and ports must meet the following requirements.
 - The load balancer should allow configuration of multiple virtual servers. For each virtual server, the load balancer should allow configuration of traffic management on more than one port. For example, for Oracle Fusion Middleware Identity Management, the load balancer needs to be configured with a virtual server and port for HTTP / HTTPS traffic, and separate virtual servers and ports for LDAP and LDAPS traffic.
 - The virtual server names must be associated with IP addresses and be part of your DNS. Clients must be able to access the external load balancer through the virtual server names.
- **Persistence/stickiness:** Some Oracle Fusion Middleware components use persistence or stickiness in an external load balancer. If your external load balancer does not allow you to set cookie persistence at the URI level, set the cookie persistence for all HTTP traffic. In either case, set the cookie to expire when the browser session expires. Refer to your external load balancer documentation for details.

The recommended architecture for Oracle Fusion Middleware is a load balancer fronting Oracle HTTP Servers in the web tier, with Oracle WebLogic Server behind the Oracle HTTP Servers in the application tier.

If Oracle WebLogic Server is deployed directly behind a load balancer in the web tier, then review the information in the "Load Balancers and the WebLogic Session Cookie" in the *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*. Note that this is not a recommended deployment architecture for Oracle Fusion Middleware.

- Resource monitoring/port monitoring/process failure detection: Configure the external load balancer to detect service and node failures (through notification or some other means) and to stop directing traffic to the failed node. Your external load balancer may have the ability to automatically detect failures.

For example, for Oracle Fusion Middleware Identity Management, the external load balancer should monitor Oracle Internet Directory, Oracle Fusion Middleware Single Sign-On, and Oracle Delegated Administration Services. To monitor these components, set up monitors for the following protocols:

- LDAP and LDAPS listen ports
- HTTP and HTTPS listen ports (depending on the deployment type)

These monitors use the respective protocols to monitor the services, meaning they use LDAP for the LDAP port, LDAP over SSL for the LDAP SSL port, and HTTP/HTTPS for the Oracle HTTP Server port. If your external load balancer does not offer these monitors, consult your external load balancer documentation for the best method of configuring it to automatically stop routing incoming requests to a service that is unavailable.

- Network Address Translation (NAT): The load balancer should have the capability to perform network address translation (NAT) for traffic being routed from clients to the Oracle Fusion Middleware nodes.
- Port translation configuration: The load balancer should have the ability to perform port translation, where it allows incoming requests received on one port to be routed to a server process running on a different port. For example, a request received on port 80 can be routed to port 7777.
- Protocol translation: The load balancer should support protocol translation between systems running different protocols. It enables users on one network to access hosts on another network, despite differences in the native protocol stacks associated with the originating device and the targeted host. For example, incoming requests can be HTTPS, and outgoing requests can be HTTP.

This feature is recommended but not required.

- SSL acceleration: SSL acceleration is a method of offloading the processor-intensive public key encryption algorithms involved in SSL transactions to a hardware accelerator.

This feature is recommended but not required.

- Fault tolerant mode: Oracle highly recommends configuring the load balancer to be in fault-tolerant mode, otherwise the load balancer becomes a single point of failure for the system. This rules out most software load balancers that are based on a single process/interceptor as reliable solutions.
- Ability to preserve the client IP addresses: The load balancer must have the capability to insert the original client IP address of a request in an

X-Forwarded-For HTTP header or a similar feature to preserve the client IP address.

- Other: Oracle highly recommends configuring the load balancer virtual server to return immediately to the calling client when the back-end services to which it forwards traffic are unavailable. This configuration is preferred over the client disconnecting on its own after a timeout, based on the TCP/IP settings on the client system.


You may not need to meet all of the requirements in the previous listed. The requirements for external load balancers depend on the topology you are considering, and on the Oracle Fusion Middleware components you are load balancing.

2.3.3 Active-Passive Deployment

Oracle Fusion Middleware provides an active-passive model for all its components using Oracle Fusion Middleware Cold Failover Clusters. In an Oracle Fusion Middleware Cold Failover Cluster configuration, two or more application server instances are configured to serve the same application workload but only one is active at any particular time.

Figure 2–4 illustrates an example active-passive deployment.

Figure 2–4 Example Active-Passive Cold Failover Cluster Deployment



In Figure 2–4, the Administration Server runs on a two-node hardware cluster: Node 1 and Node 2. The Administration Server is listening on the Virtual IP or hostname. The Middleware Home and the domain directory is on a shared disk that is mounted on Node 1 or Node 2 at any given point. Both the Middleware home and the domain directory should be on the same shared disk or shared disks that can fail over together. If an enterprise has multiple Fusion Middleware domains for multiple applications or environments, this topology is well suited for Administration Server high availability. A single hardware cluster can be deployed to host these multiple Administration Servers. Each Administration Server can use its own virtual IP and set of shared disks to provide high availability of domain services.

For details about active-passive concepts, and configuration procedures for Oracle Cold Failover Clusters, see [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability."](#)

2.3.4 About Active-Active and Active-Passive Solutions

Oracle recommends using active-active solutions when possible. This is the primary recommendation for maximum availability. Active-active solutions provide faster failover, scalability, and protection against node, instance and component failures. In addition, active-active solutions also offer transparent failover and the easy addition of resources for scaling up vertically and horizontally.

Scalability requirements are an important consideration when designing an Oracle Fusion Middleware high availability solution. Active-active solutions scale up (vertically) by adding more instances or components inside the same node.

Adding multiple redundant services in the same node improves the availability of a system, but only against instance failures and not against node failures. In addition, as described in [Section 2.3.2, "Oracle Fusion Middleware High Availability Technologies,"](#) Oracle Fusion Middleware provides death detection and automatic restart of components. Active-active high availability solutions include multiple active instances installed on different nodes. As a result, when a node is completely lost other instances are available to keep the system going, uninterrupted. Using multiple instances in different nodes provides what is known as horizontal scalability, or scaling out.

Active-active solutions require logic to load balance and failover requests among the active Oracle Fusion Middleware instances. Load balancing is provided by distributing the incoming requests to different service providers. Failover is achieved by detecting any failures in this service providers and re-routing the request to other available service providers. This logic is implemented in different ways:

- **Direct implementation:** The logic is implemented directly by the client making a request to the system. For example, a JDBC client implements load balancing and failover logic to connect to multiple instances of an Oracle database (Real Application Cluster). It can be implemented by an external hardware load balancer.
- **Third party implementation:** The logic is provided by third party components that intercept the client requests and distribute the load to the multiple Oracle Instances. When several Oracle Instances are grouped to work together, they present themselves as a single virtual entry point to the system, which hides the multiple instance configuration. External load balancers can send requests to any application server instance in a cluster, as any instance can service any request.

Unlike the scalability properties of an active-active configuration, in active-passive configurations the passive component is used only when the active component fails. In active-active solutions all instances handle requests concurrently. As a result, active-active systems provide higher transparency and have greater scalability than an active-passive system.

Active-passive solutions are limited to vertical scalability, with just one node remaining active. Active-passive solutions also have an implicit failover time when failure occurs. This failover time is usually determined by the time it takes to restart the components in the node that becomes active post-failure. However, the operational and licensing costs of an active-passive model are lower than that of an active-active deployment.

There are situations where active-passive solutions are appropriate. Oracle recommends using hardware-cluster based active-passive solutions in the following scenarios:

- The licensing, management, and the total cost of ownership of a load balancer excludes an active-active solution, particularly if there is a hardware cluster

available. Hardware clusters require two nodes, a switch for connecting the nodes, and shared storage that can be reached from both hardware nodes.

- You may have concurrency issues with Singleton services. With Singleton services, only one active instance can exist at runtime. Singleton services may be important in relation to other components. They typically provide basic services to multiple components, so if they are not available, then many other services or processes may not be available. Here are some issues to consider when protecting Singleton services:
 - Recovery Time: Singleton services or components can not run in active-active configurations. Client requests are not transparently load balanced to multiple instances of the service. This implies that, in case of a failure, there is an implicit recovery time. This recovery time varies depending on the type of Singleton protection model you adopt.
 - Reliability in failure detection: The system must prevent "false positives," or, at minimum, they the system should analyze their affect on different singleton services. Most singleton services access data repositories that may or may not be designed for concurrent access. If a singleton service is reported as 'dead' and the system decides to start a new instance, a 'split brain' scenario could arise. The dead service must be analyzed for implications of this concurrency and how likely a false positive is to happen based on the failure detection mechanism.
 - Consistency in service across restarts: Singleton components must provide consistent service after a failover. These components must maintain the same behavior after recovering from a crash. The configuration and persistent repositories used by the service must be available during failures. Also, start dependencies must be accounted for upon failover. For example, if the singleton service needs to be restarted it may have start dependencies on other services and these must be preserved.
 - Cost (hardware/software resources required): Different protection mechanisms may require a pure software based solution, or a hardware based solution with the implicit costs.
 - Installation/Configuration/Management: The different protection mechanisms for singleton services should not add complexity to the system
 - Maintenance (patches, upgrades): Protection models for singleton services should enable easily and allow minimum downtime for applying patches and upgrades.

Based on these criteria, different solutions may be used for Oracle Fusion Middleware Singleton Components depending on the pertaining requirements to the specific singleton service:

- Cold Failover Cluster Solution: This solution requires a shared storage and a connection to detect hardware failures. The re-routing of requests by migration of the VHN through the failover procedure does not need intelligence on clients or load balancers.
- Whole Server Migration- This is the process of moving a clustered WebLogic Server instance or a component running on a clustered instance elsewhere if a failure occurs. In the case of whole server migration, the server instance is migrated to a different physical system upon failure. In the case of service-level migration, the services are moved to a different server instance within the cluster.
- Custom active-passive models based on software blocking mechanism: This logic is included in a component to prevent other instances of the same component from

becoming active at the same time. Typical solutions use locks in a database, or custom in-memory active notifications that prevent concurrency.

In many cases, reliability of failure detection is an important factor for adopting one solution over another. This is especially true when concurrency can cause corruption of resources that are used by the singleton service. Typically, files may be written concurrently by different active instances.


You may adopt other solutions for different components for the issues explained in this section.

2.3.5 Disaster Recovery

[Figure 2–5](#) illustrates an Oracle Fusion Middleware architecture configured for Disaster Recovery. For Oracle Fusion Middleware product binaries, configuration files, and metadata files, the disk replication-based solution involves deploying Oracle Fusion Middleware on NAS/SAN devices. Product binaries and configuration data, stored in Oracle Homes, are stored on NAS/SAN devices using mounted locations from host systems. In addition, disk replication technologies are used to replicate product binaries and configuration from a production site shared storage system to a standby site shared storage system on a periodic basis. Standby site servers are also mounted to the disks on the standby site. If a failure or planned outage of the production (active) site occurs, replication to the standby (passive) site is stopped. The services and applications are subsequently started on the standby site. The network traffic is then be routed to the standby site.

For Oracle Database content, because of its superior level of protection and high availability, Oracle Data Guard is the recommended solution for disaster protection of Oracle Databases. This includes the databases Oracle Fusion Middleware Repositories uses and customer data.

For detailed information about disaster recovery for Oracle Fusion Middleware components, refer to *Oracle Fusion Middleware Disaster Recovery Guide*.

Figure 2–5 Production and Standby Site for Oracle Fusion Middleware Disaster Recovery Topology

2.4 Protection from Planned and Unplanned Down Time

The following tables list possible planned and unplanned downtime and suggested solutions for these downtime possibilities. [Table 2–1](#) describes planned downtime:

Table 2–1 Planned Down Time Solutions

Operations	Solutions
Deploying and redeploying applications	Hot Deployment
Patching	Rolling Patching
Configuration Changes	Online configuration Changes Change Notification Batching of changes Deferred Activation
Scalability and Topology Extensions	Cluster Scale-Out

Table 2–2 describes unplanned downtime:

Table 2–2 Unplanned Down Time Solutions

Failures	Solutions
Software Failure	Death Detection and restart using Node Manager for Java EE and OPMN for system components. Server Clusters & Load Balancing Cold Failover Clusters Server Migration Service Migration State Replication and Replica aware Stubs
Hardware Failure	Server Clusters & Load Balancing Server Migration Clusterware Integration
Data Failure	Backup and Recovery
Human Error	
Site Disaster	Oracle Fusion Middleware Disaster Recovery Solution

Note: The architectures and deployment procedures defined in this guide enable simple clustered deployments. The procedures described in these chapters can be used as a building block to enable this and other similar high availability topologies for these Fusion Middleware components. It is also expected that production deployments will use other required procedures, such as associating security policies with a centralized LDAP server. For complete details of secured, multi-tiered architecture, and deployment procedures, please refer to the Enterprise Deployment Guide for the component you are configuring.

High Availability for WebLogic Server

This chapter describes the Oracle WebLogic Server high availability capabilities that provide Oracle Fusion Middleware high availability.

- [Section 3.1, "What Is a WebLogic Server Cluster?"](#)
- [Section 3.2, "WebLogic Server Clusters and WebLogic Server Domains"](#)
- [Section 3.3, "Benefits of Clustering"](#)
- [Section 3.4, "Key Capabilities of a Cluster"](#)
- [Section 3.5, "Types of Objects That Can Be Clustered"](#)
- [Section 3.6, "Communications in a Cluster"](#)
- [Section 3.7, "Cluster-Wide JNDI Naming Service"](#)
- [Section 3.8, "Failover and Replication in a Cluster"](#)
- [Section 3.9, "Whole Server Migration"](#)
- [Section 3.10, "JMS and JTA High Availability"](#)
- [Section 3.11, "Administration Server and Node Manager High Availability"](#)
- [Section 3.12, "Load Balancing"](#)
- [Section 3.13, "GridLink Data Sources"](#)
- [Section 3.14, "Multi Data Sources"](#)
- [Section 3.15, "Cluster Configuration and config.xml"](#)
- [Section 3.16, "About Singleton Services"](#)
- [Section 3.17, "WebLogic Server and LDAP High Availability"](#)

For complete documentation of Oracle WebLogic Server clustering, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

3.1 What Is a WebLogic Server Cluster?

A WebLogic Server cluster consists of multiple WebLogic Server server instances running simultaneously and working together to provide increased scalability and reliability. A cluster appears to clients to be a single WebLogic Server instance. The server instances that constitute a cluster can run on the same system, or be located on different systems. You can increase a cluster's capacity by adding additional server instances to the cluster on an existing system, or you can add systems to the cluster to host the incremental server instances. Each server instance in a cluster must run the same version of WebLogic Server.

3.2 WebLogic Server Clusters and WebLogic Server Domains

A cluster is part of a particular WebLogic Server domain. A domain is an interrelated set of WebLogic Server resources that are managed as a unit. A domain includes one or more WebLogic Server instances, which can be clustered, non-clustered, or a combination of clustered and non-clustered instances. A domain can include multiple clusters. A domain also contains the application components deployed in the domain and the resources and services required by those application components and the server instances in the domain. Examples of the resources and services used by applications and server instances include system definitions, optional network channels, connectors, and startup classes.

In each domain, one WebLogic Server instance acts as the Administration Server—the server instance which configures, manages, and monitors all other server instances and resources in the domain. Each Administration Server manages one domain only. If a domain contains multiple clusters, each cluster in the domain has the same Administration Server.

All server instances in a cluster must reside in the same domain; you cannot "split" a cluster over multiple domains. Similarly, you cannot share a configured resource or subsystem between domains. For example, if you create a JDBC connection pool in one domain, you cannot use it with a server instance or cluster in another domain. (Instead, you must create a similar connection pool in the second domain.)

Clustered WebLogic Server instances behave similarly to non-clustered instances, except that they provide failover and load balancing. The process and tools used to configure clustered WebLogic Server instances are the same as those used to configure non-clustered instances. However, to achieve the load balancing and failover benefits that clustering enables, you must adhere to certain guidelines for cluster configuration.

3.3 Benefits of Clustering

A WebLogic Server cluster provides these benefits:

- Scalability

The capacity of an application deployed on a WebLogic Server cluster can be increased dynamically to meet demand. You can add server instances to a cluster without interruption of service—the application continues to run without impact to clients and end users.

- High Availability

In a WebLogic Server cluster, application processing can continue when a server instance fails. You “cluster” application components by deploying them on multiple server instances in the cluster—so, if a server instance on which a component is running fails, another server instance on which that component is deployed can continue application processing.

The choice to cluster WebLogic Server instances is transparent to application developers and clients. However, understanding the technical infrastructure that enables clustering helps programmers and administrators maximize the scalability and availability of their applications.

3.4 Key Capabilities of a Cluster

The following sections define, in non-technical terms, the key clustering capabilities that enable scalability and high availability.

3.4.1 Application Failover

Failover means that when an application component (typically referred to as an object in the following sections) doing a particular job—some set of processing tasks—becomes unavailable for any reason, a copy of the failed object finishes the job.

For the new object to be able to take over for the failed object:

- There must be a copy of the failed object available to take over the job.
- There must be information, available to other objects and the program that manages failover, defining the location and operational status of all objects—so that it can be determined that the first object failed before finishing its job.
- There must be information, available to other objects and the program that manages failover, about the progress of jobs in process—so that an object taking over an interrupted job knows how much of the job was completed before the first object failed, for example, what data has been changed, and what steps in the process were completed.

WebLogic Server uses standards-based communication techniques and facilities—including IP sockets and the Java Naming and Directory Interface (JNDI)—to share and maintain information about the availability of objects in a cluster. These techniques allow WebLogic Server to determine that an object stopped before finishing its job, and where there is a copy of the object to complete the interrupted job.

Information about what has been done on a job is called state. WebLogic Server maintains state information using techniques called *session replication* and *replica-aware stubs*. When an object unexpectedly stops doing its job, replication techniques enable a copy of the object to pick up where the failed object stopped and finish the job.

3.4.2 Migration

WebLogic Server supports automatic and manual migration of a clustered server instance from one system to another. A Managed Server that can be migrated is referred to as a migratable server. This feature is designed for environments with requirements for high availability. Server migration is useful for:

- Ensuring uninterrupted availability of singleton services—services that must run on only a single server instance at any given time, such as JMS and the JTA transaction recovery system, when the hosting server instance fails. A Managed Server configured for automatic migration automatically migrate to another system if a failure occurs.
- Facilitating the process of relocating a Managed Server and all services it hosts, as part of a planned system administration process. You can initiate the Managed Server migration from the Administration Console or command line.

The server migration process relocates a Managed Server in its entirety—including IP addresses and hosted applications—to one of a predefined set of available host systems.

3.4.3 Load Balancing

Load balancing is the even distribution of jobs and associated communications across the computing and networking resources in your environment. For load balancing to occur:

- There must be multiple copies of an object that can do a particular job.

- Information about the location and operational status of all objects must be available.

WebLogic Server allows objects to be clustered—deployed on multiple server instances—so that there are alternative objects to do the same job. WebLogic Server shares and maintains the availability and location of deployed objects using unicast, IP sockets, and JNDI.

3.5 Types of Objects That Can Be Clustered

A clustered application or application component is one that is available on multiple WebLogic Server instances in a cluster. If an object is clustered, failover and load balancing for that object is available. Deploy objects homogeneously—to every server instance in your cluster—to simplify cluster administration, maintenance, and troubleshooting.

Web applications can consist of different types of objects, including Enterprise Java Beans (EJBs), servlets, and Java Server Pages (JSPs). Each object type has a unique set of behaviors related to control, invocation, and how it functions within an application. For this reason, the methods that WebLogic Server uses to support clustering—and hence to provide load balancing and failover—can vary for different types of objects. The following types of objects can be clustered in a WebLogic Server deployment:

- Servlets
- JSPs
- EJBs
- Remote Method Invocation (RMI) objects
- Java Messaging Service (JMS) destinations

Different object types can have certain behaviors in common. When this is the case, the clustering support and implementation considerations for those similar object types may be same. In the sections that follow, explanations and instructions for the following types of objects are generally combined:

- Servlets and JSPs
- EJBs and RMI objects

3.6 Communications in a Cluster

WebLogic Server instances in a cluster communicate with one another using two basic network technologies:

- IP sockets, which are the conduits for peer-to-peer communication between clustered server instances.
- IP unicast or multicast, which server instances use to broadcast availability of services and heartbeats that indicate continued availability. When creating a new cluster, Oracle recommends that you use unicast for messaging within a cluster. For backward compatibility with previous versions of WebLogic Server, you must use multicast for communications between clusters.

Note: When using the unicast protocol for a WebLogic Server cluster, servers that are part of the cluster must specify a listen address. Therefore, the servers cannot be listening on ANY (which is equivalent to leaving the **Listen Address** field in the Oracle WebLogic Administration Console blank.)

3.7 Cluster-Wide JNDI Naming Service

Clients of a non-clustered WebLogic Server server instance access objects and services by using a JNDI-compliant naming service. The JNDI naming service contains a list of the public services that the server instance offers, organized in a tree structure. A WebLogic Server instance offers a new service by binding into the JNDI tree a name that represents the service. Clients obtain the service by connecting to the server instance and looking up the bound name of the service.

Server instances in a cluster use a cluster-wide JNDI tree. A cluster-wide JNDI tree is similar to a single server instance JNDI tree, insofar as the tree contains a list of available services. In addition to storing the names of local services, however, the cluster-wide JNDI tree stores the services offered by clustered objects (EJBs and RMI classes) from other server instances in the cluster.

Each WebLogic Server instance in a cluster creates and maintains a local copy of the logical cluster-wide JNDI tree. Creation of a cluster-wide JNDI tree begins with the local JNDI tree bindings of each server instance. As a server instance boots (or as new services are dynamically deployed to a running server instance), the server instance first binds the implementations of those services to the local JNDI tree. The implementation is bound into the JNDI tree only if no other service of the same name exists.

Once the server instance successfully binds a service into the local JNDI tree, additional steps are performed for clustered objects that use replica-aware stubs. After binding the clustered object's implementation into the local JNDI tree, the server instance sends the object's stub to other members of the cluster. Other members of the cluster monitor the multicast or unicast address to detect when remote server instances offer new services.

3.8 Failover and Replication in a Cluster

In order for a cluster to provide high availability it must be able to recover from service failures.

WebLogic Server instances in a cluster detect failures of their peer server instances by monitoring:

- Socket connections to a peer server

WebLogic Server instances monitor the use of IP sockets between peer server instances as an immediate method of detecting failures. If a server connects to one of its peers in a cluster and begins transmitting data over a socket, an unexpected closure of that socket causes the peer server to be marked as "failed," and its associated services are removed from the JNDI naming tree.

- Regular server heartbeat messages

If clustered server instances do not have opened sockets for peer-to-peer communication, failed servers may also be detected via the WebLogic Server

heartbeat. All server instances in a cluster use multicast or unicast to broadcast regular server heartbeat messages to other members of the cluster.

Each heartbeat message contains data that uniquely identifies the server that sends the message. Servers broadcast their heartbeat messages at regular intervals of 10 seconds. In turn, each server in a cluster monitors the multicast or unicast address to ensure that all peer servers' heartbeat messages are being sent.

If a server monitoring the multicast or unicast address misses three heartbeats from a peer server (i.e., if it does not receive a heartbeat from the server for 30 seconds or longer), the monitoring server marks the peer server as "failed." It then updates its local JNDI tree, if necessary, to retract the services that were hosted on the failed server.

3.8.1 Session Replication

User session data can be stored in two standard ways in a Java EE application: stateful session EJBs or HTTP sessions. By themselves, they rarely impact cluster scalability. However, when coupled with a session replication mechanism required to provide high-availability, bottlenecks are introduced. If a Java EE application has Web and EJB components, you should store user session data in HTTP sessions:

- HTTP session management provides more options for handling fail-over, such as replication, a shared database or file.
- Superior scalability.
- Replication of the HTTP session state occurs outside of any transactions. Stateful session bean replication occurs in a transaction which is more resource intensive.
- The HTTP session replication mechanism is more sophisticated and provides optimizations a wider variety of situations than stateful session bean replication.

3.9 Whole Server Migration

In a WebLogic Server cluster, most services are deployed homogeneously on all server instances in the cluster, enabling transparent failover from one server to another. In contrast, "pinned services" such as JMS and the JTA transaction recovery system are targeted at individual server instances within a cluster—for these services, WebLogic Server supports failure recovery with migration as opposed to failover.

Migration in WebLogic Server is the process of moving a clustered WebLogic Server instance or a component running on a clustered instance elsewhere if failure occurs. In the case of whole server migration, the server instance is migrated to a different physical system upon failure. In the case of service-level migration, the services are moved to a different server instance within the cluster.

WebLogic Server provides a feature for making JMS and the JTA transaction system highly available: migratable servers. Migratable servers provide for both automatic and manual migration at the server-level, rather than the service level.

When a migratable server becomes unavailable for any reason, for example, if it hangs, loses network connectivity, or its host system fails—migration is automatic. Upon failure, a migratable server automatically restarts on the same system if possible. If the migratable server cannot restart on the system it failed on, it is migrated to another system. In addition, an administrator can manually initiate migration of a server instance.

3.9.1 Node Manager's Role in Whole Server Migration

Server migration requires Node Manager—it must run on each system that hosts, or is intended to host.

Node Manager supports server migration in these ways:

- You must use Node Manager for the initial startup of migratable servers.

When you initiate the startup of a Managed Server from the Administration Console, the Administration Server uses Node Manager to start the server instance. You can also invoke Node Manager to start the server instance using the stand-alone Node Manager client; however, the Administration Server must be available so that the Managed Server can obtain its configuration.

Note: Migration of a server instance that was not initially started with Node Manager will fail.

- You must use Node Manager to suspend, shutdown, or force shutdown of migratable servers.
- Node Manager tries to restart a migratable server with an expired lease on the system where it was running at the time of failure.

Node Manager performs the steps in the server migrate process by running customizable shell scripts, provided with WebLogic Server, that start, restart and stop servers; migrate IP addresses; and mount and unmount disks. The scripts are available for Solaris and Linux.

- In an automatic migration, the cluster master invokes Node Manager to perform the migration.
- In a manual migration, the Administration Server invokes Node Manager to perform the migration.

3.9.2 Server Migration Processes and Communications


The following sections describe key processes in a cluster that contains migratable servers:

- [Section 3.9.2.1, "Startup Process in a Cluster with Migratable Servers"](#)
- [Section 3.9.2.2, "Automatic Whole Server Migration Process"](#)
- [Section 3.9.2.3, "Manual Whole Server Migration Process"](#)
- [Section 3.9.2.4, "Administration Server's Role in Whole Server Migration"](#)
- [Section 3.9.2.5, "Migratable Server Behavior in a Cluster"](#)
- [Section 3.9.2.6, "Cluster Master's Role in Whole Server Migration"](#)

3.9.2.1 Startup Process in a Cluster with Migratable Servers

[Figure 3–1](#) illustrates the processing and communications that occur during startup of a cluster that contains migratable servers.

The example cluster contains two Managed Servers, both of which are migratable. The Administration Server and the two Managed Servers each run on different machines. A fourth machine is available as a backup—in the event that one of the migratable servers fails. Node Manager is running on the backup machine and on each machine with a running migratable server.

Figure 3–1 Startup of Cluster with Migratable Servers

These are the key steps that occur during startup of the cluster illustrated in [Figure 3–1](#):


1. The administrator starts up the cluster.
2. The Administration Server invokes Node Manager on Machines B and C to start Managed Servers 1 and 2, respectively. See [Section 3.9.2.4, "Administration Server's Role in Whole Server Migration."](#)
3. The Node Manager on each machine starts up the Managed Server that runs there. See [Section 3.9.1, "Node Manager's Role in Whole Server Migration."](#)
4. Managed Servers 1 and 2 contact the Administration Server for their configuration. See [Section 3.9.2.5, "Migratable Server Behavior in a Cluster."](#)
5. Managed Servers 1 and 2 cache the configuration they started up.

6. Managed Servers 1 and 2 each obtain a migratable server lease in the lease table. Because Managed Server 1 starts up first, it also obtains a cluster master lease. See [Section 3.9.2.6, "Cluster Master's Role in Whole Server Migration."](#)
7. Managed Server 1 and 2 periodically renew their leases in the lease table, proving their health and liveness.

3.9.2.2 Automatic Whole Server Migration Process

Figure 3–2 illustrates the automatic migration process after the failure of the machine hosting Managed Server 2.

Figure 3–2 Automatic Migration of a Failed Server



1. Machine C, which hosts Managed Server 2, fails.
2. Upon its next periodic review of the lease table, the cluster master detects that Managed Server 2's lease has expired. See [Section 3.9.2.6, "Cluster Master's Role in Whole Server Migration."](#)
3. The cluster master tries to contact Node Manager on Machine C to restart Managed Server 2, but fails, because Machine C is unreachable.

Note: If the Managed Server 2's lease had expired because it was hung, and Machine C was reachable, the cluster master would use Node Manager to restart Managed Server 2 on Machine C.

4. The cluster master contacts Node Manager on Machine D, which is configured as an available host for migratable servers in the cluster.
5. Node Manager on Machine D starts Managed Server 2. See [Section 3.9.1, "Node Manager's Role in Whole Server Migration."](#)
6. Managed Server 2 starts up and contacts the Administration Server to obtain its configuration.
7. Managed Server 2 caches the configuration it started up with.
8. Managed Server 2 obtains a migratable server lease.

During migration, the clients of the migrating Managed Server may have a brief interruption in service; it may be necessary to reconnect. On Solaris and Linux operating systems, you can reconnect using `ifconfig` command. The clients of a migrated server do not need to know the particular machine the server migrates to.

When a machine that previously hosted a server instance that was migrated becomes available again, the reversal of the migration process—migrating the server instance back to its original host machine—is known as *failback*. WebLogic Server does not automate the failback process. You can accomplish failback by manually restoring the server instance to its original host.


The general procedures to restore a server to its original host are:

- Gracefully shutdown the new instance of the server.
- After you restart the failed machine, restart Node Manager and the managed server.

The exact procedures you follow depend on your server and network environment.

3.9.2.3 Manual Whole Server Migration Process

[Figure 3–3](#) illustrates what happens when an administrator manually migrates a migratable server.

Figure 3–3 Manual Whole Server Migration

1. An administrator uses the Administration Console to initiate the migration of Managed Server 2 from Machine C to Machine B.
2. The Administration Server contacts Node Manager on Machine C. See [Section 3.9.2.4, "Administration Server's Role in Whole Server Migration."](#)
3. Node Manager on Machine C stops Managed Server 2.
4. Managed Server 2 removes its row from the lease table.
5. The Administration Server invokes Node Manager on Machine B.
6. Node Manager on Machine B starts Managed Server 2.
7. Managed Server 2 obtains its configuration from the Administration Server.
8. Managed Server 2 caches the configuration it started up with.
9. Managed Server 2 adds a row to the lease table.

3.9.2.4 Administration Server's Role in Whole Server Migration

In a cluster that contains migratable servers, the Administration Server invokes Node Manager on each system:

- That hosts cluster members to start the migratable servers. This is a prerequisite for server migratability—if Node Manager did not initially start a server instance, you cannot migrate the server.
- Involved in a manual migration process to stop and start the migratable server.

- That hosts cluster members to stop server instances during a normal shutdown. This is a prerequisite for server migratability—if a server instance is shut down directly without using Node Manager, when the cluster master detects that the server instance is not running, it will call Node Manager to restart it.

The Administration Server also provides its regular domain management functionality, persisting configuration updates issued by an administrator and providing a run-time view of the domain, including the migratable servers it contains.

3.9.2.5 Migratable Server Behavior in a Cluster

A migratable server is a clustered Managed Server that is configured as migratable. These are the key behaviors of a migratable server:

- If you are using a database to manage leasing information, during startup and restart by Node Manager, a migratable server adds a row to the lease table. The row for a migratable server contains a timestamp, and the machine where it is running.
- When using a database to manage leasing information, a migratable server adds a row to the database as a result of startup, it tries to take on the role of cluster master, and succeeds if it is the first server instance to join the cluster.
- Periodically, the server renews its "lease" by updating the timestamp in the lease table.

By default a migratable server renews its lease every 30,000 milliseconds—the product of two configurable `ServerMBean` properties:

- `HealthCheckIntervalMillis`, which by default is 10,000.
- `HealthCheckPeriodsUntilFencing`, which by default is 3.
- If a migratable server fails to reach the lease table and renew its lease before the lease expires, it terminates as quickly as possible using a Java `System.exit`—in this case, the lease table still contains a row for that server instance. For information about how this relates to automatic migration, see [Section 3.9.2.6, "Cluster Master's Role in Whole Server Migration."](#)
- During operation, a migratable server listens for heartbeats from the cluster master. When it detects that the cluster master is not sending heartbeats, it attempts to take over the role of cluster master, and succeeds if no other server instance has claimed that role.

3.9.2.6 Cluster Master's Role in Whole Server Migration

In a cluster that contains migratable servers, one server instance acts as the cluster master. Its role is to orchestrate the server migration process. Any server instance in the cluster can serve as the cluster master. When you start a cluster that contains migratable servers, the first server to join the cluster becomes the cluster master and starts up the cluster manager service. If a cluster does not include at least one migratable server, it does not require a cluster master, and the cluster master service does not start up. In the absence of a cluster master, migratable servers can continue to operate, but server migration is not possible. These are the key functions of the cluster master:

- Issues periodic heartbeats to the other servers in the cluster.
- Periodically reads the lease table to verify that each migratable server has a current lease. An expired lease indicates to the cluster master that the migratable server should be restarted.

- Upon determining that a migratable server's lease is expired, waits for period specified by the `FencingGracePeriodMillis` on the `ClusterMBean`, and then tries to invoke the Node Manager process on the machine that hosts the migratable server whose lease is expired, to restart the migratable server.
- If unable to restart a migratable server whose lease has expired on its current machine, the cluster master selects a target machine in this fashion:
- If you have configured a list of preferred destination machines for the migratable server, the cluster master chooses a machine on that list, in the order the machines are listed.
- Otherwise, the cluster master chooses a machine on the list of those configured as available for hosting migratable servers in the cluster.

A list of machines that can host migratable servers can be configured at two levels: for the cluster as a whole, and for an individual migratable server. You can define a machine list at both levels. You must define a machine list at least one level.

- To accomplish the migration of a server instance to a new machine, the cluster master invokes the Node Manager process on the target machine to create a process for the server instance.

The time required to perform the migration depends on the server configuration and startup time.

- The maximum time taken for cluster master to restart the migratable server is $(\text{HealthCheckPeriodsUntilFencing} * \text{HealthCheckIntervalMillis}) + \text{FencingGracePeriodMillis}$.
- The total time before the server becomes available for client requests depends on the server startup time and the application deployment time.

3.10 JMS and JTA High Availability

You can configure JMS and JTA services for high availability by using a migratable target, a special target that can migrate from one server in a cluster to another. A migratable target provides a way to group migratable services that should move together. When a migratable target migrates, all services the target hosts also migrate.

To configure a migratable JMS service for migration, it must be deployed to a migratable target. A migratable target specifies a set of servers that can host a target, and can optionally specify a user-preferred host for the services and an ordered list of candidate backup servers should the preferred server fail. Only one of these servers can host the migratable target at any one time.

After you configure a service to use a migratable target, the service is independent from the server member that is currently hosting it. For example, if a JMS server with a deployed JMS queue is configured to use a migratable target then the queue is independent of when a specific server member is available. That is, the queue is always available when the migratable target is hosted by any server in the cluster.

An administrator can manually migrate pinned migratable services from one server instance to another in the cluster in response to a server failure or as part of regularly scheduled maintenance. If you do not configure a migratable target in the cluster, migratable services can be migrated to any WebLogic Server instance in the cluster.

3.10.1 User-Preferred Servers and Candidate Servers

When deploying a JMS service to the migratable target, you can select a user-preferred server (UPS) target to host the service. When configuring a migratable target, you can also specify constrained candidate servers (CCS) that can host the service if the user-preferred server fails. If the migratable target does not specify a constrained candidate server, you can migrate the JMS server to any available server in the cluster.

WebLogic Server enables you to create separate migratable targets for JMS services. This allows you to always keep each service running on a different server in the cluster, if necessary. Conversely, you can configure the same selection of servers as the constrained candidate servers for both JTA and JMS, to ensure that the services remain co-located on the same server in the cluster.

3.10.2 Considerations for Using File Stores on NFS

If JMS messages and transaction logs are stored on an NFS-mounted directory, Oracle strongly recommends that you verify the behavior of a server restart after abrupt machine failures. Depending on the NFS implementation, different issues can arise post failover/restart.

To verify server restart behavior, abruptly shut down the node that hosts WebLogic servers while the servers are running.

- If you configured the server for server migration, it should start automatically in failover node after the failover period.
- If you did not configure the server for server migration, you can manually restart the WebLogic Server on the same host after the node completely reboots.

If Oracle WebLogic Server does not restart after abrupt machine failure, the following error entry may appear in server log files:

```
<MMM dd, yyyy hh:mm:ss a z> <Error> <Store> <BEA-280061> <The persistent
store "_WLS_server_soal" could not be deployed:
weblogic.store.PersistentStoreException: java.io.IOException:
[Store:280021]There was an error while opening the file store file
"_WLS_SERVER_SOA1000000.DAT"
weblogic.store.PersistentStoreException: java.io.IOException:
[Store:280021]There was an error while opening the file store file
"_WLS_SERVER_SOA1000000.DAT"
    at weblogic.store.io.file.Heap.open(Heap.java:168)
    at weblogic.store.io.file.FileStoreIO.open(FileStoreIO.java:88)
...
java.io.IOException: Error from fcntl() for file locking, Resource
temporarily unavailable, errno=11
```

This error occurs when the NFSv3 system does not release locks on the file stores. WebLogic Server maintains locks on files that store JMS data and transaction logs to prevent data corruption that can occur if you accidentally start two instances of the same managed server. Because the NFSv3 storage device doesn't track lock owners, NFS holds the lock indefinitely if a lock owner crashes. As a result, after abrupt machine failure followed by a restart, subsequent attempts by WebLogic Server to acquire locks may fail.

How you resolve this error depends on your NFS environment: (See *Oracle Fusion Middleware Release Notes* for updates on this topic.)

- **For NFSv4 environments**, you can set a tuning parameter on the NAS server to release locks within the approximate time required to complete server migration; you do not need to follow the procedures in this section. See your storage vendor's

documentation for information on locking files stored in NFS-mounted directories on the storage device, and test the results.

- **For NFSv3 environments**, the following sections describe how to disable WebLogic file locking mechanisms for: the default file store, a custom file store, a JMS paging file store, a diagnostics file store.

WARNING: NFSv3 file locking prevents severe file corruptions that occur if more than one managed server writes to the same file store at any point in time.

If you disable NFSv3 file locking, you must implement administrative procedures / policies to ensure that only one managed server writes to a specific file store. Corruption can occur with two managed servers in the same cluster or different clusters, on the same node or different nodes, or on the same domain or different domains.

Your policies could include: never copy a domain, never force a unique naming scheme of WLS-configured objects (servers, stores), each domain must have its own storage directory, no two domains can have a store with the same name that references the same directory.

If you configure a managed server using a file store for server migration, always configure the database- based leasing option. This option enforces additional locking mechanisms using database tables and prevents automated restart of more than one instance of a particular managed server.

Disabling File Locking for the Default File Store

To disable file locking for the default file store using the WebLogic Server Administration Console:

1. If necessary, click **Lock & Edit** in the Change Center (upper left corner) of the Administration Console to get an Edit lock for the domain.
2. In the **Domain Structure** tree, expand the **Environment** node and select **Servers**.
3. In the **Summary of Servers** list, select the server you want to modify.
4. Select the **Configuration > Services** tab.
5. Scroll down to the **Default Store** section and click **Advanced**.
6. Scroll down and deselect the **Enable File Locking** check box.
7. Click **Save**. If necessary, click **Activate Changes** in the Change Center.
8. **Restart** the server you modified for the changes to take effect.

The resulting `config.xml` entry will look like the following:

```
<server>
  <name>examplesServer</name>
  ...
  <default-file-store>
    <synchronous-write-policy>Direct-Write</synchronous-write-policy>
    <io-buffer-size>-1</io-buffer-size>
    <max-file-size>1342177280</max-file-size>
    <block-size>-1</block-size>
    <initial-size>0</initial-size>
```

```
        <file-locking-enabled>false</file-locking-enabled>
    </default-file-store>
</server>
```

Disabling File Locking for a Custom File Store

To disable file locking for a custom file store using the WebLogic Server Administration Console:

1. If necessary, click **Lock & Edit** in the Change Center (upper left corner) of the Administration Console to get an Edit lock for the domain.
2. In the **Domain Structure** tree, expand the **Services** node and select **Persistent Stores**.
3. In the **Summary of Persistent Stores** list, select the custom file store you want to modify.
4. On the **Configuration** tab for the custom file store, click **Advanced**.
5. Scroll down and deselect the **Enable File Locking** check box.
6. Click **Save**. If necessary, click **Activate Changes** in the Change Center.
7. If the custom file store was in use, you must restart the server for the changes to take effect.

The resulting `config.xml` entry looks like the following:

```
<file-store>
  <name>CustomFileStore-0</name>
  <directory>C:\custom-file-store</directory>
  <synchronous-write-policy>Direct-Write</synchronous-write-policy>
  <io-buffer-size>-1</io-buffer-size>
  <max-file-size>1342177280</max-file-size>
  <block-size>-1</block-size>
  <initial-size>0</initial-size>
  <file-locking-enabled>false</file-locking-enabled>
  <target>examplesServer</target>
</file-store>
```

Disabling File Locking for a JMS Paging File Store

To disable file locking for a JMS paging file store using the WebLogic Server Administration Console:

1. If necessary, click **Lock & Edit** in the Change Center (upper left corner) of the Administration Console to get an Edit lock for the domain.
2. In the **Domain Structure** tree, expand the **Services** node, expand the **Messaging** node, and select **JMS Servers**.
3. In the **Summary of JMS Servers** list, select the JMS server you want to modify.
4. On the **Configuration > General** tab for the JMS Server, scroll down and deselect the **Paging File Locking Enabled** check box.
5. Click **Save**. If necessary, click **Activate Changes** in the Change Center.
6. **Restart** the server you modified for the changes to take effect.

The resulting `config.xml` file entry will look like the following:

```
<jms-server>
  <name>examplesJMSServer</name>
```



```

<target>examplesServer</target>
<persistent-store>exampleJDBCStore</persistent-store>
...
<paging-file-locking-enabled>>false</paging-file-locking-enabled>
...
</jms-server>

```

Disabling File Locking for a Diagnostics File Store

To disable file locking for a Diagnostics file store using the WebLogic Server Administration Console:

1. If necessary, click **Lock & Edit** in the Change Center (upper left corner) of the Administration Console to get an Edit lock for the domain.
2. In the **Domain Structure** tree, expand the **Diagnostics** node and select **Archives**.
3. In the **Summary of Diagnostic Archives** list, select the server name of the archive that you want to modify.
4. On the **Settings for [server_name]** page, deselect the **Diagnostic Store File Locking Enabled** check box.
5. Click **Save**. If necessary, click **Activate Changes** in the Change Center.
6. **Restart** the server you modified for the changes to take effect.

The resulting `config.xml` file will look like this:

```

<server>
  <name>examplesServer</name>
  ...
  <server-diagnostic-config>
    <diagnostic-store-dir>data/store/diagnostics</diagnostic-store-dir>
    <diagnostic-store-file-locking-enabled>>false</diagnostic-store-file-locking-enabled>

<diagnostic-data-archive-type>FileStoreArchive</diagnostic-data-archive-type>
  <data-retirement-enabled>>true</data-retirement-enabled>
  <preferred-store-size-limit>100</preferred-store-size-limit>
  <store-size-check-period>1</store-size-check-period>
  </server-diagnostic-config>
</server>

```

3.11 Administration Server and Node Manager High Availability

The Administration Server is the WebLogic Server instance that configures and manages the WebLogic Server instances in its domain.

A domain can include multiple WebLogic Server clusters and non-clustered WebLogic Server instances. A domain can consist of only one WebLogic Server instance—however, in this case that sole server instance would be an Administration Server because each domain must have exactly one Administration Server.

There are a variety of ways to invoke the services of the Administration Server to accomplish configuration tasks. Whichever method is used, the Administration Server for a cluster must be running when you modify the configuration.

Note: Oracle recommends that you set the Administration Server listen address to the hostname that its clients need to access it on, particularly for systems with multiple Middleware homes or Oracle homes.

3.11.1 Administration Server Failure

The failure of an Administration Server for a domain does not affect the operation of managed servers in the domain. If an Administration Server for a domain becomes unavailable while the server instances it manages—clustered or otherwise—are up and running, those managed servers continue to run. If the domain contains clustered server instances, the load balancing and failover capabilities supported by the domain configuration remain available, even if the Administration Server fails.

Note: If an Administration Server fails because of a hardware or software failure on its host machine, other server instances on the same machine may be similarly affected. However, the failure of an Administration Server itself does not interrupt the operation of managed servers in the domain.

For instructions on re-starting an Administration Server, see the *Oracle Fusion Middleware Using Clusters for Oracle Server*.

3.11.2 Node Manager Failure

If Node Manager fails or is explicitly shut down, upon restart, it determines the server instances that were under its control when it exited. Node Manager can restart any failed server instances as needed.

Note: It is advisable to run Node Manager as an operating system service, so that it restarts automatically if its host machine is restarted.

3.12 Load Balancing

Load balancing configuration consists of three pieces of information: the load-balancing algorithm to be used, an indicator of whether local affinity should be applied, and weights that are assigned to each member of the topology to influence any routing algorithms that use weights.

The load-balancing algorithm specifies how requests are load balanced across components. Oracle Fusion Middleware uses the following three load-balancing methods:

- Round Robin - Requests are balanced across a list of available servers by selecting from the list sequentially.
- Random - Requests are balanced across a list of available servers by selecting a random server on each request.
- Weighted - Requests are balanced across a list of available servers using weights assigned to each server to determine the percentage of requests sent to each

Local affinity determines whether clients show a preference to servers that run on the same machine to avoid network latency. If the flag is set to true, then requests are

routed across the list of servers on the local machine using the load-balancing algorithm if any local servers are available. If no local servers are available, requests are routed to all available remote servers according to the load-balancing algorithm. If local affinity is set to false, requests are routed across all available servers (local and remote) based on the load-balancing algorithm.

You configure weights as single integer values that are associated with component instances. You can assign weights to components that are not currently in a group, however, the weight is not used unless you later configure the component as a member of a group and select the weighted load-balancing algorithm. The weight is a unitless number. The percentage of requests to be sent to each member is calculated by summing the weights of all available members and dividing the weight for each member by the sum of the weights.

3.13 GridLink Data Sources

A single GridLink data source provides connectivity between WebLogic Server and an Oracle Database Real Application Clusters (RAC). It uses the Oracle Notification Service (ONS) to respond adaptively to state changes in an Oracle RAC. It responds to FAN events to provide Fast Connection Failover (FCF), Runtime Connection Load-Balancing (RCLB), and RAC instance graceful shutdown. It also provides capabilities of Affinities.

See Using GridLink Data Sources in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

3.14 Multi Data Sources

A multi data source is an abstraction around a group of data sources that provides load balancing or failover processing at the time of connection requests, between the data sources associated with the multi data source. Multi data sources are bound to the JNDI tree or local application context just like data sources are bound to the JNDI tree. Applications look up a multi data source on the JNDI tree or in the local application context (java:comp/env) just as they do for data sources, and then request a database connection. The multi data source determines which data source to use to satisfy the request depending on the algorithm selected in the multi data source configuration: load balancing or failover.

A multi data source can be thought of as a pool of data sources. Multi data sources are best used for failover or load balancing between nodes of a highly available database system, such as redundant databases or Oracle Real Application Clusters (Oracle RAC).

3.15 Cluster Configuration and config.xml

The `config.xml` file is an XML document that describes the configuration of a WebLogic Server domain. The `domain` element in `config.xml` is the top-level element, and all elements in the domain descend from the `domain` element. The `domain` element includes child elements such as the `server`, `cluster`, and `application` elements. These child elements may have children of their own. For example, the `server` element can include the child elements `WebServer`, `SSL` and `Log`. The `Application` element includes the child elements `EJBComponent` and `WebAppComponent`.

Each element has one or more configurable attributes. An attribute defined in `config.dtd` has a corresponding attribute in the configuration API. For example, the

Server element has a `ListenPort` attribute, and likewise, the `Weblogic.management.configuration.ServerMBean` has a `ListenPort` attribute. Configurable attributes are readable and writable, that is, `ServerMBean` has a `getListenPort()` and a `setListenPort()` method.

To learn more about `config.xml`, see the *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

3.16 About Singleton Services

A singleton service is a service that must run on only a single server instance at any given time, such as JMS and the JTA transaction recovery system, when the hosting server instance fails. A managed server configured for automatic migration automatically migrates to another machine if a failure occurs.

3.17 WebLogic Server and LDAP High Availability

In a high availability environment, WebLogic Server must be able to access LDAP for these reasons:

- To access users and groups stored in LDAP for which WebLogic Server supports failover.

For information about configuring failover for LDAP authentication providers, see "Configuring Failover for LDAP Authentication Providers" in the *Oracle Fusion Middleware Securing Oracle WebLogic Server* manual.

- To access the LDAP-based policy store and credential store.

For information about configuring a domain to use an LDAP-based policy store, see "Configuring a Domain to Use an LDAP-Based Policy Store" in the *Oracle Fusion Middleware Application Security Guide* manual.

For information about configuring a domain to use an LDAP-based credential store, see "Configuring a Domain to Use an LDAP-Based Credential Store" in the *Oracle Fusion Middleware Application Security Guide* manual.

Considerations for High Availability Oracle Database Access

This chapter describes considerations for high availability Oracle database access.

The sections in this chapter are as follows:

- [Section 4.1, "Oracle Real Application Clusters and Fusion Middleware"](#)
- [Section 4.2, "Protecting Idle Connections from Firewall Timeouts"](#)
- [Section 4.3, "Troubleshooting"](#)
- [Section 4.4, "Oracle Fusion Middleware Products are Certified to be Used with 11.2 RDBMS Oracle RAC"](#)

4.1 Oracle Real Application Clusters and Fusion Middleware

Most Fusion Middleware components use a database as the persistent store for their data. You can configure the Oracle database back end in any number of high availability configurations, including Cold Failover Clusters, Oracle Real Application Clusters, Oracle Data Guard, or Oracle Streams. For more information, see the *Oracle Database High Availability Overview*. This chapter describes considerations for Oracle Fusion Middleware configured with a high availability Oracle database, Oracle Real Application Clusters.

Oracle Real Application Cluster (Oracle RAC) is a computing environment that harnesses the processing power of multiple, interconnected computers. Along with a collection of hardware (cluster), it unites the processing power of each component to become one robust computing environment. Oracle RAC simultaneously provides a highly scalable and highly available database for Oracle Fusion Middleware.

Every Oracle RAC instance in the cluster has equal access and authority, therefore, node and instance failure may affect performance, but doesn't result in downtime; the database service is available or can be made available on surviving server instances.

For more information on Oracle RAC, see the *Oracle Real Application Clusters Administration and Deployment Guide*.

Oracle Fusion Middleware provides the best integration with an Oracle database in a high availability environment. When Oracle Fusion Middleware behaves as a client for the database (either as a java or system client) it uses special communication and monitoring capabilities that provide fast failover and minimal middle tier disruption in reaction to database failure scenarios.

You can categorize Oracle Fusion Middleware components that access the database as follows:

- Java-based components deployed to Oracle WebLogic Server
- Java-based components that are standalone Java Clients
- Non-Java components

This chapter contains the following sections:

- [Section 4.1.1, "Java-Based Oracle Fusion Middleware Components Deployed to Oracle WebLogic Server"](#)
- [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC"](#)
- [Section 4.1.3, "Using Multi Data Sources with Oracle RAC"](#)
- [Section 4.1.4, "Configuring GridLink Data Sources with Oracle RAC"](#)
- [Section 4.1.5, "Configuring Multi Data Sources with Oracle RAC"](#)
- [Section 4.1.6, "JDBC Clients"](#)
- [Section 4.1.7, "System Clients"](#)

4.1.1 Java-Based Oracle Fusion Middleware Components Deployed to Oracle WebLogic Server

All Oracle Fusion Middleware components deployed to Oracle WebLogic Server support Oracle RAC. For establishing connection pools, Oracle Fusion Middleware supports GridLink data sources and multi data sources for the Oracle RAC back end for both XA and non-XA JDBC drivers. Oracle Fusion Middleware deployments do not support other connection failover features supported by Oracle JDBC drivers for Oracle RAC. See component specific guides for multi data source configuration details.

When an Oracle RAC node or instance fails, Oracle WebLogic Server or the Oracle Thin driver redirect session requests to another node in the cluster. There is no failover of existing connections, however, new connection requests from the application are managed using existing connections in the Oracle WebLogic pool or by new connections to the working Oracle RAC instance. When the database is the transaction manager, in-flight transactions typically roll back. When the WebLogic Server is the transaction manager, in-flight transactions fail over; they are driven to completion or rolled back based on the transaction state when failure occurs. If the application requires load balancing across Oracle RAC nodes, WebLogic Server supports this capability by using JDBC GridLink data sources or JDBC multi data sources configured for load balancing. See *Run Time Connection Load Balancing in the Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide for more information on GridLink load balancing.

4.1.2 Using GridLink Data Sources with Oracle RAC

A GridLink data source includes the features of generic data sources plus the following support for Oracle RAC:

- Fast Connection Failover
- Runtime Connection Load Balancing
- Graceful Handling for Oracle RAC Outages
- XA Affinity
- SCAN Addresses

- Secure Communication Using Oracle Wallet


See *Using GridLink Data Sources in the Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide for more information on the following topics:

- What is a GridLink Data Source?
- Using Socket Direct Protocol for a GridLink Data Source
- Configuring Connection Pool Features
- Configuring Oracle Parameters
- Configuring an ONS Client
- Tuning GridLink Data Source Connection Pools
- Setting Database Security Credentials
- Monitoring GridLink JDBC Resources

4.1.3 Using Multi Data Sources with Oracle RAC

When you deploy Oracle Fusion Middleware against Oracle RAC back ends it is configured out of the box with Multi Data Sources (MDS). The MDSs have constituent data sources for each RAC instance providing the database service. Oracle recommends that you add an additional data source to the MDS on the Fusion Middleware tier when you configure additional RAC instances that offer the database service. Ensure that each constituent data source that you create for the MDS are configured identically for properties in [Section 4.1.5, "Configuring Multi Data Sources with Oracle RAC."](#)

When you migrate the database from a non-RAC to a RAC database, you must create an equivalent, new MDS for each data source that is affected. The MDS that you create must have consistent data sources for each RAC instance. The data source values must be identical to the original single instance data source for the properties in [Section 4.1.5, "Configuring Multi Data Sources with Oracle RAC."](#) For example, if the single instance data source driver is `oracle.jdbc.xa.client.OracleXADataSource`, it must be `oracle.jdbc.xa.client.OracleXADataSource` for each constituent data source of the new Multi Data Source.

Figure 4–1 Multi Data Source Configuration

4.1.3.1 Configuring Multi Data Sources for MDS Repositories

Applications that use an MDS database-based repository can be configured for high availability Oracle database access. With this configuration, failure detection, recovery, and retry by MDS, as well as by the WebLogic infrastructure result in the application's read-only MDS operations being protected from Oracle RAC database planned and unplanned downtimes.

MDS multi data sources are exposed as MDS repositories in the Fusion Middleware Control navigation tree. These multi data sources can be selected during deployment plan customization of application deployment, and can be used with MDS WLST commands.

- **Configuring an application to retry read-only operations**

To configure an application to retry the connection, you can configure the `RetryConnection` attribute of the application's MDS `AppConfig` MBean. For information about MDS configuration, see the *Oracle Fusion Middleware Administrator's Guide*.

- **Registering an MDS multi data source**

In addition to the steps specified in [Section 4.1.5, "Configuring Multi Data Sources with Oracle RAC,"](#) consider the following for MDS:

- The child data sources that constitute a multi data source used for an MDS Repository must be configured as non-XA data sources.
- The multi data source's name must be pre-fixed with `mDS-`. This ensures that the multi data source can be recognized as an MDS repository that can be used

for MDS management functionality through Fusion Middleware Control, WLST, and JDeveloper.

Note: When an MDS data source is added as a child of a multi data source, this data source is no longer exposed as an MDS repository. For example, it does not appear under the Metadata Repositories folder in the Fusion Middleware Control navigation tree, you cannot perform MDS repository operations on it, and it does not appear in the list of selectable repositories during deployment.

- **Converting a data source to a multi data source**

There are two considerations when converting an MDS data source to a multi data source to make sure the application is configured correctly:

- To create a new multi data source with a new, unique name, redeploy the application and select this new multi data source as the MDS repository during deployment plan customization.
- To avoid redeploying the application, you can delete the data source and recreate the new multi data source using the same name and jndi-name attributes.

4.1.3.2 Oracle RAC Configuration Requirements

This section describes requirements for Oracle RAC configuration:

- **XA Requirements:** Many Oracle components participate in distributed transactions, or are part of container managed transactions. These components require the back-end database setup for XA recovery by Oracle WebLogic Transaction Manager. For repositories created using RCU, this is done automatically. For other databases participating in XA transactions, ensure that XA pre-requisites are met:

1. Log on to SQL*Plus as a system user, for example:

```
sqlplus "/ as sysdba"
```

2. Grant select on sys.dba_pending_transactions to public.
3. Grant execute on sys.dbms_xa to public.
4. Grant force any transaction to user.

Note: Ensure that the `distributed_lock_timeout` parameter for the Oracle database is set to a value higher than the JTA timeout. It should be higher than the highest value on the middle tier - between the default for WebLogic Server, a specific configuration for a data source, or one used by a component for a transaction.)

- **Server-side Load Balancing:** If the server-side load balancing feature is enabled for the Oracle RAC back end (using remote_listeners), the JDBC URL used in the data sources of a multi data source configuration should include the INSTANCE_NAME. For example, you can specify the URL in the following format:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=host-vip) (PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=dbservice) (INSTANCE_NAME=inst1)))
```

By default, the out-of-box installation assumes that `remote_listener` has been configured and creates the URL for data sources in a multi data source accordingly. Any multi data source created outside of the typical installation and configuration should follow the format described in this section.

If `remote_listeners` cannot be specified on the Oracle RAC side, and server side load balancing has been disabled, specifying the `INSTANCE_NAME` in the URL is not necessary. To disable remote listeners, delete any listed remote listeners in `spfile.ora` file on each Oracle RAC node. For example:

```
*.remote_listener=
```

In this case, the recommended URL that you use in the data sources of a multi data source configuration is:

```
jdbc:oracle:thin:@host-vip:port/dbservice
```

Or

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=host-vip) (PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=dbservice)))
```

- **Services:** When configuring Oracle Fusion Middleware for the Oracle database and specifically for Oracle RAC, Oracle recommends using the Oracle Services feature. Create the `service_name` provided as part of the database service location specifically for the application.

4.1.3.3 Configuring Schemas for Transactional Recovery Privileges

You need the appropriate database privileges to allow the Oracle WebLogic Server transaction manager to query for transaction state information and issue the appropriate commands, such as commit and rollback, during recovery of in-flight transactions after a WebLogic Server container crash.

To configure the schemas for transactional recovery privileges:

1. Log on to SQL*Plus as a user with sysdba privileges. For example:

```
sqlplus "/ as sysdba"
```

2. Grant select on `sys.dba_pending_transactions` to the `appropriate_user`.
3. Grant force any transaction to the `appropriate_user`.

Note: Grant these privileges to the soainfra schema owner, as determined by the RCU operations.

4.1.4 Configuring GridLink Data Sources with Oracle RAC

How you configure a GridLink data source depends on the Oracle component that you are working with and the domain you are creating. For detailed procedures, go to the section that describes the component type you are working with. For example, for SOA, see [Section 5.13.20.2, "Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console"](#).

Note: Oracle recommends that you use Oracle Single Client Access Name (SCAN) addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

For a generic overview of how to configure a GridLink data source, see Creating a GridLink Data Source in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

4.1.5 Configuring Multi Data Sources with Oracle RAC

You can configure multi data sources using the following:

- Oracle Fusion Middleware Configuration Wizard during WebLogic Server domain creation
- Oracle Universal Installer Java EE component configuration for Identity Management or Oracle Portal, Forms, Reports, and Discoverer.
- Oracle WebLogic Server Administration Console
- WLST Commands

Multi data sources support load balancing for both XA and non-XA data sources, including all Oracle database versions that Oracle Fusion Middleware components support.

Multi data sources encapsulate individual data sources that pool connections to specific instances of Oracle RAC. For multi data sources created manually, or modified after initial configuration, Oracle strongly recommends the following XA and Non-XA data source property values for optimal high availability behavior. Make changes only after careful consideration and testing if your environment requires that you do so:

Table 4–1 Recommended Multi Data Source Configuration

Property Name	Value
test-frequency-seconds	5
algorithm-type	Load-Balancing

For individual data sources, Oracle recommends the following for high availability environments. Oracle recommends that you set any other parameters according to application requirements.

Table 4–2 XA Data Source Configuration

Property Name	Value
Driver	oracle.jdbc.xa.client.OracleXADataSource
Property command	<pre><property> <name>oracle.net.CONNECT_TIMEOUT</name> <value>10000</value> </property></pre>
initial-capacity	0
connection-creation-retry-frequency-seconds	10
test-frequency-seconds	300
test-connections-on-reserve	true
test-table-name	SQL SELECT 1 FROM DUAL
seconds-to-trust-an-idle-pool-connection	0
global-transactions-protocol	TwoPhaseCommit
keep-xa-conn-till-tx-complete	true
xa-retry-duration-seconds	300
xa-retry-interval-seconds	60

Table 4–3 Non-XA Data Source Configuration

Property Name	Value
Driver	oracle.jdbc.OracleDriver
Property to set	<pre><property> <name>oracle.net.CONNECT_TIMEOUT</name> <value>10000</value> </property></pre>
initial-capacity	0
connection-creation-retry-frequency-seconds	10
test-frequency-seconds	300
test-connections-on-reserve	true
test-table-name	SQL SELECT 1 FROM DUAL
seconds-to-trust-an-idle-pool-connection	0
global-transactions-protocol	None

For examples of recommended multi data sources, see [Appendix B, "Recommended Multi Data Sources."](#)

Increasing Transaction Timeout for XA Data Sources

If you see WARNING messages in the server logs that include the following exception:

```
javax.transaction.SystemException: Timeout during commit processing
[ javax.transaction.SystemException: Timeout during commit processing
```

This message may indicate the XA timeout value you have in your setup must be increased. You can increase XA timeout for individual data sources when these warnings appear.

To increase this setting, use Administration Console:

1. Access the data source configuration.
2. Select the **Transaction** tab.
3. Set XA Transaction Timeout to a larger value, for example, **300**.
4. Select the **Set XA Transaction Timeout** checkbox. You *must* select this checkbox for the new XA transaction timeout value to take effect.
5. Click **Save**.

Repeat this configuration for all individual data sources of an XA multi data source.

4.1.6 JDBC Clients

Java J2SE-based Oracle Fusion Middleware components are optimized to work with the high availability features of Oracle RAC. You can deploy the components to use both the Oracle thin JDBC driver or the OCI based JDBC drivers.

The JDBC Thin client is a pure Java, Type IV driver. It is lightweight, easy to install and provides high performance, comparable to the performance of the JDBC Oracle Call Interface (OCI) driver. The JDBC Thin driver communicates with the server using TTC, a protocol developed by Oracle to access data from Oracle database. The driver allows a direct connection to the database by providing an implementation of TCP/IP that implements Oracle Net and TTC on top of Java sockets. The JDBC OCI client is a Type II driver and provides connections to JDBC clients over the Oracle Net. It uses the client side installation of Oracle Net and a deployment can customize behavior using Oracle Net configuration on the middle tier.

Note: These JDBC clients are used as part of standalone Java J2SE programs.

Oracle Virtual Directory

When used with database adapters, Oracle Virtual Directory connects to a database, and the connections are not pooled. For details about configuring database adapters for Oracle RAC, see "Creating Database Adapters" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

Database URL

To configure an Oracle Virtual Directory database adapter for an Oracle RAC database using the Oracle Directory Services Manager:

1. In the **Connection** screen, select **Use Custom URL** from the **URL Type** list.
2. In the **Database URL** field, enter the URL to connect to the Oracle RAC database, for example:

JDBC Thin

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(LOAD_
BALANCE=ON)(ADDRESS=(PROTOCOL=TCP)(HOST=host-name-1)(PORT=1521))(ADDRESS=
(PROTOCOL=TCP)(HOST=host-name-2)(PORT=1521)))(CONNECT_
DATA=(SERVER=DEDICATED)(SERVICE_NAME=database-service-name)))
```

JDBC OCI

```
jdbc:oracle:oci:@(DESCRIPTION=(ADDRESS_LIST=(LOAD_
BALANCE=ON) (ADDRESS=(PROTOCOL=TCP) (HOST=host-name-1) (PORT=1521)) (ADDRESS=
(PROTOCOL=TCP) (HOST=host-name-2) (PORT=1521))) (CONNECT_
DATA=(SERVER=DEDICATED) (SERVICE_NAME=database-service-name)))
```

Connection Timeout Configuration

To configure the connection timeout for an Oracle RAC database using the Oracle Directory Services Manager:

1. In the **Connection** screen, for **JDBC Thin**, specify the database adapter parameter **oracleNetConnectTimeout** for the timeout parameter in seconds.
2. For JDBC OCI, specify `TCP.CONNECT_TIMEOUT=n` in the `sqlnet.ora` in `ORACLE_INSTANCE/config` directory.

4.1.7 System Clients

Oracle Fusion Middleware 11g includes some non-Java components. These components are primarily C-based and include Oracle Internet Directory (OID), Oracle Forms, Oracle Reports, Oracle Discoverer, and Oracle Portal. These components use the Oracle Call Interface layer to interact with Oracle databases. For Oracle RAC-based systems, some components integrate with the Oracle high availability Event Notification database feature.

High availability Event Notification provides a signal to the non-Java application if database failure occurs. The applications can register a callback on the environment to monitor the database connection. When a database failure related to the non-Java client occurs, the callback is invoked. This callback contains information about the database failure, including the event payload, and a list of connections (server handles) that were disconnected as a result of the failure.

If another instance, for example, instance C, of the same database, goes down, the client is not notified, since it does not affect any of the client's connections.

High availability Event Notification improves the application response time if database failure occurs. Without Event Notification, database failure would result in the connection being broken only after the TCP time out expired, which could take minutes. With high availability Event Notification, OCI automatically breaks and cleans up standalone, connection pool, and session pool connections and the application callback is invoked within seconds of failure. If any server handles are TAF-enabled, OCI automatically engages failover.

The following section describes the recommended setting for non-Java client connections to Oracle RAC databases.

4.1.7.1 Oracle Internet Directory

Oracle Internet Directory integrates with high availability Event Notification. Oracle recommends using the Oracle Enterprise Manager Cluster Managed Services Page to create database services that client applications use to connect to the database.

You can also use SQL*Plus to configure your Oracle RAC database service.

To enable high availability event motivation for an Oracle RAC database connection:

1. Set the `AQ_HA_NOTIFICATIONS` attribute to **TRUE** and server-side Transparent Application Failover (TAF) settings are enabled. The failover retries and failover

delay can be adjusted based on the requirements of the deployment. So for the database service used by OID, Oracle recommends setting Oracle RAC DBMS_SERVICE property values according to [Table 4-4](#).

Table 4-4 *OID Database Services Property Settings*

Property Name	Value
AQ_HA_NOTIFICATIONS	TRUE
FAILOVER_METHOD	DBMS_SERVICE.FAILOVER_METHOD_BASIC
FAILOVER_TYPE	DBMS_SERVICE.FAILOVER_TYPE_SELECT
FAILOVER_RETRIES	5
FAILOVER_DELAY	5

2. Oracle also recommends setting TCP connect timeouts for the Oracle Net configuration. To configure this setting, specify `TCP.CONNECT_TIMEOUT=n` in the `sqlnet.ora` file in the `ORACLE_INSTANCE/config` directory.

4.1.7.2 Oracle Forms

Oracle Forms also integrates with high availability event notification. To enable this feature for Oracle Forms:

1. Use the Oracle Enterprise Manager Cluster Managed Services Page to create database services. For Oracle Forms, set the Oracle RAC DBMS_SERVICE property values according to [Table 4-5](#). The following is recommended to be set using the package of Oracle database.

Table 4-5 *Oracle Forms Database Services Property Settings*

Property Name	Value
AQ_HA_NOTIFICATIONS	TRUE
FAILOVER_METHOD	DBMS_SERVICE.FAILOVER_METHOD_NONE
FAILOVER_TYPE	DBMS_SERVICE.FAILOVER_TYPE_NONE

2. Oracle also recommends setting TCP connect timeouts for the Oracle Net configuration. To configure this setting, specify `TCP.CONNECT_TIMEOUT=n` in the `sqlnet.ora` file in the `ORACLE_INSTANCE/config` directory.

4.1.7.3 Oracle Portal

To configure Oracle Portal for optimal behavior in a high availability environment, set TCP connect timeouts for the Oracle Net configuration. To configure this setting, specify `TCP.CONNECT_TIMEOUT=n` in the `sqlnet.ora` file in the `ORACLE_INSTANCE/config` directory.

Oracle Portal also uses the death detection feature of `mod_plsql`.

`mod_plsql` maintains a pool of connections to the database, and reuses established database connections for subsequent requests. If there is no response from a database connection in a connection pool, `mod_plsql` detects this, discards the dead connection, and creates a fresh database connection for subsequent requests.

The dead database connection detection feature of `mod_plsql` eliminates the occurrence of random errors when a database node or instance goes down. This feature is also extremely useful in high availability configurations, such as Oracle

RAC. If a node in an Oracle RAC cluster goes down, `mod_plsql` detects this and immediately starts servicing requests using the other Oracle RAC nodes.

By default, when an Oracle RAC node or database instance goes down and `mod_plsql` had previously pooled connections to the node, the first `mod_plsql` request which uses a dead connection in its pool results in a failure response (HTTP-503) being sent back to the end-user. `mod_plsql` uses this failure to trigger the detection and removal of all dead connections in its pool. `mod_plsql` pings all connection pools that were created before the node failure. This ping operation is performed at the time of processing the next request that uses a pooled connection. If the ping operation fails, the database connection is discarded and a new connection is created and processed.

Note: If, after node failure, multiple `mod_plsql` requests come in concurrently and `mod_plsql` has not yet detected the first dead connection, there may be multiple failures at that instant.

`mod_plsql` provides two configuration options for tuning the dead database connection detection feature:

- Specifying the Option to Detect Dead Database Connections
- Specifying the Connection Validation and Timeout Period

Specifying the Option to Detect Dead Database Connections.

`mod_plsql` corrects connections after it detects a failure that could be caused by a database node going down. The `PlsqlConnectionValidation` parameter controls this.

Setting `PlsqlConnectionValidation` to **Automatic** (default) makes the `mod_plsql` module test all pooled database connections that were created before a failed request.

Setting `PlsqlConnectionValidation` to **AlwaysValidate** makes `mod_plsql` test all pooled database connections before issuing any request. Although the `AlwaysValidate` configuration option ensures greater availability, it also introduces additional performance overhead.

You can specify the timeout period for `mod_plsql` to test a bad database connection in a connection pool. The `PlsqlConnectionTimeout` parameter, which specifies the maximum time `mod_plsql` should wait for the test request to complete before it assumes that a connection is not usable. For more on this parameter, see the `mod_plsql` section in the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

Specifying the Connection Validation and Timeout Period

When the `PlsqlConnectionValidation` parameter is set to **Automatic** or **AlwaysValidate**, `mod_plsql` attempts to test pooled database connections.

You can specify the timeout period for `mod_plsql` to test a bad database connection in a connection pool. This is controlled by the `PlsqlConnectionTimeout` parameter, which specifies the maximum time `mod_plsql` should wait for the test request to complete before it assumes that a connection is not usable.

For details on the `PlsqlConnectionTimeout`, `PlsqlConnectionValidation`, and `PlsqlConnectionTimeout` parameters, refer to the "mod_plsql" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

4.1.7.4 Oracle Reports and Oracle Discoverer

To configure Oracle Reports and Oracle Discovery for optimal behavior in a high availability environment, set TCP connect timeouts for the Oracle Net configuration. To configure this setting, specify `TCP.CONNECT_TIMEOUT=n` in the `sqlnet.ora` file in the `ORACLE_INSTANCE/config` directory.

Oracle Discoverer also uses a TNS entry to connect to the Oracle RAC database:

```
frdisco = (DESCRIPTION = (LOAD_BALANCE = ON) (ADDRESS_LIST =
  (ADDRESS = (PROTOCOL = TCP) (HOST = stajo05-vip) (PORT = 1521))
  (ADDRESS = (PROTOCOL = TCP) (HOST = stajo06-vip) (PORT = 1521)))
(CONNECT_DATA = (SERVICE_NAME = orcl.us.oracle.com)))
```

Note: When Oracle Discoverer is configured to be connected to a customer Oracle RAC database, Oracle recommends using the following TNS connect string:

```
(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=host1-vip) (PORT=1521))
  (ADDRESS=(PROTOCOL=TCP) (HOST=host2-vip) (PORT=1521))) (CONNECT_
DATA=(SERVICE_NAM
E=orcl) (FAILOVER_
MODE=(TYPE=select) (METHOD=basic) (RETRIES=20) (DELAY=15))))
```

4.2 Protecting Idle Connections from Firewall Timeouts

Because most production deployments involve firewalls and database connections are made across firewalls, Oracle recommends configuring the firewall not to timeout the database connection. For Oracle RAC case, this specifically means not timing out the connections made on Oracle RAC VIPs and the database listener port.

If such a configuration is not possible, on the database server side, set `SQLNET.EXPIRE_TIME=n` in `ORACLE_HOME/network/admin/sqlnet.ora`. For Oracle RAC, this needs to be set on all the Oracle Homes. The `n` is in minutes. It should be set to less than the known value of the network device (firewall) timeout. Since the order of these times is normally more than ten minutes, and in some cases hours, the value should be set to the highest possible value.

4.3 Troubleshooting

If an Oracle RAC instances goes down, WebLogic Server attempts to determine the status of the of the database using the `SELECT 1 FROM DUAL` query. This query typically takes less than a few seconds to complete. However, if the database response is slow, WebLogic Server gives up and assumes the database is unavailable. The following is an example of the type of exception that results in the logs:

```
<Mar 30, 2009 2:14:37 PM CDT> <Error> <JDBC> <BEA-001112> <Test "SELECT 1 FROM
DUAL" set up for pool SOADDataSource-rac1" failed with exception:
oracle.jdbc.xa.OracleXAException".> [TopLink Warning]: 2009.03.30
14:14:37.890--UnitOfWork(14568040)--Exception [TOPLINK-4002] (Oracle TopLink -
11g Release 1 (11.1.1.1.0) (Build 090304)):
oracle.toplink.exceptions.DatabaseException Internal Exception:
java.sql.SQLException: Internal error: Cannot obtain XAConnection Creation of
XAConnection for pool SOADDataSource failed after waitSecs:30 :
weblogic.common.ResourceException: SOADDataSource(SOADDataSource-rac1): Pool
SOADDataSource-rac1 has been @ disabled because of hanging connection tests,
cannot allocate resources to applications. We waited 10938 milliseconds. A
```

typical test has been taking 16.

You can set the WebLogic Server parameter, `-Dweblogic.resourcepool.max_test_wait_secs=30` to increase the time WebLogic Server waits for a response from the database. This parameter is located in the `setDomainEnv.sh` file. By setting this parameter, WebLogic Server waits 30 seconds for the database to respond to the `SELECT 1 FROM DUAL` query before giving up.

4.4 Oracle Fusion Middleware Products are Certified to be Used with 11.2 RDBMS Oracle RAC

If your 11.2 RDBMS Oracle RAC database is not configured with SCAN, you can provide details of the Oracle RAC instances (instance address in terms of `host:port`) in the Configuration Wizard and Oracle Universal Installer, just as you entered them for previous database releases.

If your 11.2 RDBMS Oracle RAC database is configured with SCAN, provide details of the Oracle RAC instances in terms of the SCAN address. In Fusion Middleware wiring to an Oracle RAC instance, each Oracle RAC instance is uniquely identified using the service name, instance name, host, and port. This configuration information for each instance is provided when wiring to an Oracle RAC database. In the case of a SCAN-configured Oracle RAC database, since the `host:port` address of all such instances is the `SCANhost:port`, Oracle recommends that you use this same common address for all the instances.

Specifically, with Oracle Fusion Middleware configuration, use the following rules:

- In RCU installations, against an Oracle RAC database, the hostname specified can be `scan-hostname-address`.
- In Oracle Universal Installer based installation sessions, you can specify the following for Oracle RAC databases depending on the input format Oracle Universal Installer requires.

```
scan-address-hostname:port:instance1^scan-address-hostname:port:instance2@service-name
```

or:

```
scan-address-hostname:port^scan-address-hostname:port@service-name
```

- In Configuration Wizard-based installations that use a multi data source for the 11gR2 Oracle RAC database, the `scan-address-hostname`, `port`, `service-name` must be the same for each of the constituent data source and instance names must specific for each constituent data source, and targeted to the Oracle RAC end instance.

GridLink with SCAN does not use an instance name. A sample GridLink connection string (that does not use the RAC instance):

```
jdbc:oracle:thin:@(DESCRIPTION=(ENABLE=BROKEN)(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=<scan-host-name>)(PORT=<scan-port>)))(CONNECT_DATA=(SERVICE_NAME=rac-service-name)))
```

- If the connect string is specified explicitly, use the following base format:
(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=scan-hostname-address)(PORT=port)))(CONNECT_DATA=(SERVICE_NAME=service-name))) when the whole Oracle RAC database needs to be specified

```
(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=scan-hostname-address)
(PORT=port)))) (CONNECT_DATA=(SERVICE_NAME=service-name) (INSTANCE_NAME=inst1)))
```

when a specific Oracle RAC instance needs to be specified

4.4.1 SCAN Run Time Implications and Limitations

There are three possible supported scenarios when you configure against RAC:

Table 4–6 SCAN Run Time Implications and Limitations

Scenario	Description	High Availability Run time Outcome and Limitations
1. Non Scan	Multi data source with each subordinate multi data source pointing to a separate RAC instance	Gives run time High Availability of the database connections, which the WLS multi data source implementation manages.
2. SCAN	Multi data source with each subordinate data source pointing to the SCAN address	Gives run time High Availability of the database connections, which the WLS multi data source implementation manages. Limitation: Even if you reference a SCAN address, you are using the limited High Availability features of the WLS multi data source.
3. SCAN	A single data source pointing to the SCAN address.	Does not give runtime High Availability of the database Connections. Limitation: A SCAN address virtualizes the entry point to the RAC instances, however, if you specify a single Data Source on WLS, doing so does not provide High Availability. The reason for this is that each server is effectively bound to a single RAC instance.
4. SCAN	GridLink data source pointing to the SCAN address	You must have the correct database version that supports SCAN and set up ONS correctly, using the FAN enabled setting to receive the ONS status messages that the database sends.

Configuring High Availability for Oracle SOA Suite

Oracle SOA Suite provides a complete set of service infrastructure components for designing, deploying, and managing composite applications. Oracle SOA Suite enables services to be created, managed, and orchestrated into composite applications and business processes. Composites enable you to assemble multiple technology components into one SOA composite application. Oracle SOA Suite plugs in to heterogeneous IT infrastructures and enables enterprises to incrementally adopt SOA.

This chapter describes Oracle SOA Suite components from a high availability perspective. This chapter also includes sections that outline the single-instance concepts for SOA components that are important for designing a high availability deployment.


This chapter includes the following topics:

- [Section 5.1, "Introduction to Oracle SOA Suite"](#)
- [Section 5.2, "Oracle SOA Service Infrastructure High Availability"](#)
- [Section 5.3, "Oracle BPEL Process Manager and High Availability Concepts"](#)
- [Section 5.4, "Oracle BPM Suite and High Availability Concepts"](#)
- [Section 5.5, "Oracle Mediator and High Availability Concepts"](#)
- [Section 5.6, "Oracle Human Workflow and High Availability Concepts"](#)
- [Section 5.7, "Oracle B2B and High Availability Concepts"](#)
- [Section 5.8, "Oracle Web Services Manager and High Availability Concepts"](#)
- [Section 5.9, "Oracle User Messaging Service and High Availability Concepts"](#)
- [Section 5.10, "Oracle JCA Adapters and High Availability Concepts"](#)
- [Section 5.11, "Oracle Business Activity Monitoring and High Availability Concepts"](#)
- [Section 5.12, "Oracle Service Bus and High Availability Concepts"](#)
- [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines"](#)
- [Section 5.14, "Configuring High Availability for Oracle Service Bus, with SOA Service Infrastructure and Component Service Engines"](#)
- [Section 5.15, "Configuring High Availability for Oracle BAM"](#)


5.1 Introduction to Oracle SOA Suite

As shown in [Figure 5–1](#), Oracle SOA Suite provides a comprehensive suite of products for developing, securing, and monitoring service-oriented architecture (SOA). *Oracle SOA Suite 11g* provides a unified runtime based on the SCA standard. The runtime consists of service engines (BPEL Process Manager, Human Workflow, Mediator, Rules) and binding components (JCA Adapters, B2B) that a common service infrastructure manages and inter-connects. The service infrastructure also provides common services such as lifecycle management and deployment.

Figure 5–1 Oracle SOA Suite and Components



A SOA composite application is the basic unit of deployment to the SOA runtime. Service components (BPEL process, business rule, human task, and Mediator routing rule) are the building blocks of a SOA composite. Service components target service engines during deployment, while services and references are enabled using the binding components. At runtime, the binding component receives messages, which the service infrastructure then routes to the appropriate service engine(s).

Figure 5–2 Oracle SOA Infrastructure Stack Diagram

The SOA runtime runs within the context of an application server such as the Oracle WebLogic Application Server. It leverages the underlying application server capabilities for load balancing and high availability.

This guide provides high availability information for the following Oracle SOA Suite components:

- Oracle SOA Service Infrastructure
- Oracle BPEL Process Manager (Oracle BPEL PM)
- Oracle BPM Suite
- Oracle Mediator
- Oracle Human Workflow
- Oracle JCA Adapters
- Oracle B2B
- Oracle Web Services Manager
- Oracle User Messaging Service
- Oracle Business Activity Monitoring

5.2 Oracle SOA Service Infrastructure High Availability

Oracle SOA Service Infrastructure is a Java EE application that provides the foundation services for running Oracle SOA Suite. This Java EE application is a runtime engine that is automatically deployed when Oracle SOA Suite is installed. You deploy *composites*, the basic artifacts in a Software Component Architecture, to the Oracle SOA Infrastructure and it provides the required services for the composites to run. Oracle SOA Infrastructure provides deployment, wiring, and thread management services for the composites. These services sustain the composite's lifecycle and runtime operations.

This section guides you through the issues and considerations necessary for designing a SOA Service Infrastructure high availability cluster.

Backup and Recovery Considerations

For information on backing up SOA Service Infrastructure files, see the "Introducing Backup and Recovery" and "Backing Up Your Environment" chapters of the *Oracle Fusion Middleware Administrator's Guide*.

5.2.1 Oracle SOA Service Infrastructure Single-Instance Characteristics

Oracle SOA Service Infrastructure provides the required services for running composites. A composite is basic unit of deployment for Service Component Architectures (SCA). The SCA Assembly Model consists of a series of artifacts, which are defined by elements contained in XML files.


Composites are software packages made up of components, wires, services and references. For example, an Oracle BPEL process is a component, an inbound adapter is a service, and an outbound adapter is a reference. Wires provide connections between service engines.

Each component targets a specific service engine such as Oracle BPEL PM. Oracle SOA Service Infrastructure connects to a SOA database to maintain composite state and to the SOA Oracle Metadata Services (MDS) Repository to maintain composite metadata, such deployments, and version tracking. These two databases may be the same physical database but the schemas used for each purpose are different. SOA infra provides a servlet for remote deployment of composites. The metadata and artifacts for remote deployments are also stored in the MDS repository. For more information on the MDS repository, see [Section 4.1.3.1, "Configuring Multi Data Sources for MDS Repositories."](#)

The Oracle SOA Service Infrastructure application is also responsible for targeting the individual components to their specific engine and for instantiating these composites when requests reach the SOA system. After targeting and instantiation, the Oracle SOA Service Infrastructure controls thread and resource assignment. This happens in the JVM, where the composite runs.

As illustrated in [Figure 5-3](#), the Oracle SOA Service Infrastructure integrates SOA composite applications with UDDI registries. UDDI registries provide a standards-based foundation for locating published services, invoking services, and managing service metadata. The Oracle SOA Service Infrastructure is also the central hub used by the service engines to deliver messages through Oracle User Messaging Service infrastructure to communication channels, such as email and voice.

SOA Service Infrastructure provides the required services that sustain the different pieces in a Service Component Architecture, and enables the communication between them. For more details of the different components in an SCA system, refer to the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

Figure 5–3 Basic Single-Node SOA Service Infrastructure Architecture

5.2.1.1 Oracle SOA Service Infrastructure Application Characteristics

Oracle SOA Service Infrastructure Services are contained in the `soa-infra-wls.ear` file. None of the services provided by the Oracle SOA Service Infrastructure system are singletons, therefore, the Oracle SOA Service Infrastructure can run in full active-active mode. The SOA Service Infrastructure Java EE application contains a Web module that provides browsing of the deployed composites and links to the test pages for these composites. This Web module uses `/soa-infra` as the associated URL context. This Web module is stateless and does not have any specific session replication requirements.

Other modules in the Oracle SOA Service Infrastructure application provide task control for process instantiation and process tracking, and client services for accessing User Messaging System (UMS).

A task service controls instantiating and tracking processes asynchronously. In addition, there are multiple EJBs used by the Oracle SOA Service Infrastructure system. However, all of the EJBs are stateless, and there are no requirements for stateful session bean replication in an Oracle SOA cluster. The processing of transactions by these EJBs relies on Oracle WebLogic Server transaction control service. Configure the appropriate transaction stores as recommended in the basic Oracle WebLogic Server guidelines to guarantee recovery across failures in Oracle WebLogic Server container.

5.2.1.2 Oracle SOA Service Infrastructure Startup and Shutdown Lifecycle


An Oracle SOA composite consists of the following:

- Components such as a BPEL process, Human Workflow task, or Business Rules.
- Services and References for connecting Oracle SOA composite applications to external services, applications, and technologies.

These components are assembled together into an Oracle SOA composite application. This application is a single unit of deployment that simplifies the management and lifecycle of Oracle SOA applications.

When the Oracle SOA Service Infrastructure application starts, it initializes the different service engines and loads the present composites from the MDS repository. It targets the individual components to their specific engines. Once the composite loads, the system is available to receive requests. At runtime, the Oracle SOA Service Infrastructure manages all communication across service components. Calls between service engines are in-process calls. The following diagram shows the sequence for the Oracle SOA Service Infrastructure startup and processing of work:

Figure 5–4 Oracle SOA Service Infrastructure Application Startup and Shutdown Lifecycle



5.2.1.3 Oracle SOA Service Infrastructure External Dependencies

As the previous sections describe, the Oracle SOA Service Infrastructure system depends on the following components:

- Instance manager service depends on the runtime SOA database schema (soa-infra).
- Composite metadata is stored in the MDS database schema, which acts as a repository.
- In a clustered environment, the deployment coordinator service depends on the underlying Coherence cluster for signal propagation.

These three components must be available for the Oracle SOA Service Infrastructure to start and run properly.


5.2.1.4 Oracle SOA Service Infrastructure Startup and Shut Down of Processes

The Oracle SOA Service Infrastructure application starts default whenever any Oracle WebLogic Managed Server to which the Oracle SOA Service Infrastructure has been deployed is started. Normally, you should not need to stop the Oracle SOA Service Infrastructure or any of its components by themselves. Some operations may require

Oracle WebLogic Managed Server where the SOA Service Infrastructure runs to be rebooted. Only some patching scenarios could require stopping the application.

You can use Oracle WebLogic Server Administration Console to verify status and to start and stop Oracle WebLogic Server. You can also use the WebLogic Server WLST command line to control the application. Oracle Enterprise Manager Fusion Middleware Control also enables multiple operations and configuration of the Oracle SOA Service Infrastructure application as well as monitoring its status. See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* for details on monitoring and controlling the Oracle SOA Service Infrastructure application.

Figure 5–5 Monitoring and Controlling the Oracle SOA Service Infrastructure Application



5.2.1.5 Oracle SOA Service Infrastructure Configuration Artifacts

Starting with Oracle Fusion Middleware 11g Release 1 (11.1.1.2), the SOA MDS database stores configuration parameters for Oracle Service Infrastructure. You configure the main controls of the Oracle SOA Service Infrastructure using Oracle Enterprise Manager Fusion Middleware Control.

- The data source JNDI name for process dehydration.

Note: Once the JNDI names are read from this file, the databases used by the system are determined by the data sources that matched those JNDI names in WebLogic Server JDBS resources configuration.

- The server and callback URL

A Callback URL is the address that asynchronous services specify to be notified of a response to the service they invoked.

Note: If a request to an external or internal asynchronous service originates from SOA, the callback URL is determined using the following in decreasing order of preference:

- Use `callbackServerURL` specified as a binding property for the specific reference. (You can set this when modeling the composite or at runtime using the MBeans). This enables different service calls to have different callback URLs. At runtime this property is set using the System MBean Browser, through the corresponding binding Mbean. To add a specific URL add a `callbackServerURL` property to its `Properties` attribute, then invoke the save operation.
 - Use the callback URL as specified in the SOA Database. In this case, only one address can be specified and an address that works well for all possible services must be used.
 - Use the callback URL as the frontend host specified in WebLogic Server for the `SOA_Cluster`. In this case as well, only one address can be specified, and the recommendation is the same as for the SOA Database configuration option.
 - Use the local host name as provided by WebLogic Server MBean APIs. This is not recommended in high availability environments.
-

- The audit level of information to be collected by the message tracking infrastructure.

Other configuration options at the container level, such as data sources, JTA configuration, and persistent stores, are maintained as part of the WebLogic Server Domain configuration, and are synchronized across a cluster of Oracle WebLogic Servers by the Oracle WebLogic Server core infrastructure.

5.2.1.6 Oracle SOA Service Infrastructure Log File Locations

Oracle WebLogic Managed Server logs operations that Oracle SOA Service Infrastructure and its components perform where the SOA Service Infrastructure is running. You can find these logs at the following location:

`DOMAIN_HOME/servers/WLS_ServerName/logs/WLS_ServerName.log`

The log files for the different Oracle WebLogic Server managed servers are also available from Oracle WebLogic Server Administration Console. To verify the logs, access WebLogic Server Administration Console using the following URL: `admin_server_host:port/console`. Click **Diagnostics-Log Files**.

It is also important to verify the output of the Oracle WebLogic Managed Server where the Oracle SOA Service Infrastructure is running. This information is stored at the following location:

`DOMAIN_HOME/servers/WLS_ServerName/logs/WLS_ServerName.out`

Additionally, a diagnostic log is produced in the log directory for the managed server. This log's granularity and logging properties can be changed through the `domain_dir/config/fmwconfig/servers/WLS_ServerName/logging.xml` file. The properties in this file can also be modified from Oracle Enterprise Manager Fusion Middleware Control by selecting **Farm, SOA, SOA Server**. Right-click, and select **Logs**, and then **Log Configuration**.

Oracle Enterprise Manager Fusion Middleware Control enables performing selective searches in all the logs in the SOA domain. To do a selective search, access Oracle Fusion Middleware Control and click on **Farm-Logs** and enter the search criteria that pertain to `soa-infra` or deployed composites. See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* for more details on the logs and information reported for a SOA System in Oracle Enterprise Manager Fusion Middleware Control.

5.2.2 Oracle SOA Service Infrastructure High Availability Architecture and Failover Considerations

[Figure 5–6](#) illustrates a two-node Oracle SOA Service Infrastructure cluster running on two Oracle WebLogic Servers. The Oracle WebLogic Servers are frontended by two Oracle HTTP Server instances on web tier hosts, which receive requests from a load balancer in front of them.


Figure 5–6 Oracle SOA Service Infrastructure High Availability Architecture

Figure 5–6 illustrates the following main characteristics of this high availability configuration:

The SOA Service Infrastructure runs in Oracle WebLogic Server managed servers that are part of an Oracle WebLogic Server cluster. Oracle WebLogic Server cluster synchronizes the configuration for common artifacts of WebLogic Server cluster that Oracle SOA Service Infrastructure uses such as JTS configuration, data sources, and persistent store definitions.

The Oracle SOA Service Infrastructure uses the front end host and port information configured for the Oracle WebLogic Server cluster as the server and callback URL. You define these settings with Oracle WebLogic Server Administration Console. Select **Clusters**, **SOA_Cluster_Name**, **HTTP/HTTPS frontend host** and then **Port**. If there is no address for the Oracle WebLogic Server cluster where the Oracle SOA Service

Infrastructure is running, the system uses the physical host name as the server and callback URL.

For SOA high availability installations frontended by Oracle HTTP Server, monitoring should be done on the Listen ports of Oracle HTTP Server. This is the case when a deployment uses all components deployed to the SOA Managed Server. A simple HTTP monitor that pings the HTTP/HTTPS port and expects a pre-determined response in turn should suffice. If only a specific SOA component is being used (such as B2B), then a monitor that does a deeper level check all the way to the Managed server can be considered to validate the health of the component in use. Check with your load balancer vendor on setting up the HTTP monitors with your load balancer.

For more information about the server and callback URLs see the *Oracle Fusion Middleware Administrator's Guide*. Changing the HTTP front end address for a cluster requires a restart of the managed servers that are part of the cluster.

The deployment coordinator is configured and used for deployment and updates of composites. The deployment coordinator sends notifications to deployment coordinator cluster members to retrieve new artifacts from the MDS repository, when they are updated by the group leader. A leader node performs singleton operations for the cluster such as updating the MDS after deployments, or changes are made to the composites.

The Oracle SOA Service Infrastructure system uses the Oracle WebLogic Server cluster name as its key to confirm a deployment coordinator group. If all nodes in a WebLogic Server cluster can communicate (over multicast or unicast) the deployment coordinator cluster is the same as the WebLogic Server cluster in which the SOA Service Infrastructure runs.

The Administration Server runs in Active-Passive mode. Whenever a failure occurs in SOAHOST1, you can restart the Administration Server in SOAHOST2; it uses a virtual IP or virtual hostname as listen address.

For information about configuring virtual IPs for the Administration Server and configuring the Administration Server for high availability, see [Chapter 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

Oracle WebLogic Server Whole Server Migration

Although SOA Service Infrastructure can run in full active-active mode, the architecture uses the Oracle WebLogic Server Migration feature to protect some SOA components against failures. This implies that each of the WebLogic Managed Servers in which the SOA Service Infrastructure runs is listening on a Virtual IP that is migrated to another box upon failover. The Administration Server and Enterprise Manager run in AdminServer, not the managed server.

As shown in [Figure 5–6](#), WLS_SOA1 listens on VIP1, and WLS_SOA2 listens on VIP2. Each managed server uses the other node as a failover resource, therefore, the system is configured in a cross manner. In other words, WLS_SOA1 fails over to SOAHOST2, and WLS_SOA2 fails over to SOAHOST1. The appropriate capacity planning must be done to anticipate the scenario where the two SOA managed servers are running on the same node. For more information on Server Migration features, see [Chapter 3, "High Availability for WebLogic Server"](#) in this guide.

To resume transactions after a server migration, configure the transaction and JMS stores in a shared storage. In case of failure in one of the server infrastructure instances, other instances can resume transactions and JMS operation by reading the persistent stores from that shared storage.

The Metadata store is configured in an Oracle RAC database to protect from database failures. Similarly, the SOA process state information is also stored in an Oracle RAC database. In this example, both Oracle RAC databases are the same.

About Oracle SOA Service Infrastructure Components

These high availability characteristics apply to most of Oracle SOA components contained in the composite applications deployed across the cluster. For specific two-node high availability characteristics of the individual components, see the specific component sections that follow in this chapter.

5.2.2.1 Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior

This section describes how an Oracle SOA high availability cluster deployment protects components from failure and the expected behavior if component failure occurs.

The WebLogic Server infrastructure protects the SOA Service Infrastructure system from all process failures.

5.2.2.1.1 WebLogic Server Crash If a WLS_SOAx server crashes, Node Manager tries to restart it locally. If repeated restarts fail, the WebLogic Server infrastructure attempts to perform a server migration of the WLS_SOAx server to the other node in the cluster. While the failover takes place, the other SOA Service Infrastructure instance becomes the leader for deployments and composite updates and provides the basic services required by the service engines in the system.

Ongoing requests from the HTTP Server time out and new requests are directed to the other WLS_SOAx server. Once the server's restart completes on the other node, Oracle HTTP Server resumes routing any incoming requests to the server. The migrated server reads MDS repository for any updates that might have taken place during restart, and joins the deployment coordinator cluster to listen for new updates. The migrated server also resumes any pending transactions from the transaction logs in shared storage.

In the server migration scenario, the service engines, such as Oracle BPEL PM and Oracle Mediator, are failed over together with the SOA Service Infrastructure. They do not re-issue any requests to the other SOA Service Infrastructure instances by themselves. They resume operations together with the SOA Service Infrastructure once failover is complete.

The Oracle SOA Service Infrastructure application may be down due to failure in accessing resources, errors caused by the deployment coordinator infrastructure, or other issues unrelated to whether the managed server is running. Therefore, Oracle recommends administrators monitoring the soa-infra application for errors caused by the application in the managed server logs. For information about log file locations, see [Section 5.2.1.6, "Oracle SOA Service Infrastructure Log File Locations"](#).

5.2.2.1.2 Node Failure If node failure occurs, server migration is triggered after the available server verifies the time stamp in the database leasing system. If the failed server was the deployment coordinator cluster master, the available server becomes the new master and the SOA Service Infrastructure remains available for deployment and for composite lifecycle. After the time stamp for leasing is verified, the Node Manager in the node that still remains available attempts to migrate the VIP used by the failed managed server, and restarts it locally. This effectively results in the SOA Service Infrastructure application having two instances running in the same node. For more information on the failover process, see [Section 3.9, "Whole Server Migration"](#).

Service engines are deployed to the container as a part of the Service Infrastructure application. These service engines contain all of the ear files and library calls.

5.2.2.1.3 Database Failure GridLink data sources and multi data sources protect the SOA Service Infrastructure system against database failures. Typically you configure GridLink and multi data sources during the system setup (using the Configuration Wizard to define these multi-pools directly at installation time) and guarantee that when an Oracle RAC database instance fails, the connections reestablish with available database instances.


Note: You must disable the fast connection failover flag for SOA or a Null Pointer Exception occurs.

For more information about GridLink data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC."](#) For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

5.2.2.2 Oracle SOA Service Infrastructure Cluster-Wide Deployment

As previous sections explain, composite deployments are stored centrally by the SOA Service Infrastructure in the MDS repository. Each time the SOA Service Infrastructure is started, it synchronizes itself with the MDS repository and SOA store to get the deployment and process state. The deployment coordinator infrastructure orchestrates the notifications for composites deployments and updates. When a new deployment or update takes place, deployment coordinator notifies all members in the cluster. When all members in the cluster confirm that the deployment has succeeded, the master sends a notification to start the composite. If a deployment fails on any one of the nodes, it is rolled back to the rest of the cluster. An error message in the deployment coordinator master (WebLogic Server managed server), indicates the node on which the deployment failed. [Figure 5-7](#) explains this process:

Figure 5-7 Cluster-Wide Deployment of Oracle SOA Composites



5.2.2.3 Online Redeployment of Oracle SOA Service Infrastructure Composites in a Cluster

When the Oracle SOA Service Infrastructure performs an update or redeployment of a composite, it can overwrite an existing version (x) or create a new version (x+1). All composites are uniquely identified based on the composite name and revision. By default, clients accessing a composite use the version that is identified in the MDS

repository as the default version (also visible from Oracle Enterprise Manager Fusion Middleware Control). It is possible to manage the lifecycle of each version separately and perform online redeployments of composites, even with one single instance of SOA Service Infrastructure. To perform this operation:

1. Deploy version x.
2. Mark version x as default.
3. Deploy a new version of the composite (x+1) and mark version x as the default. This step stops people from accessing the new version in default access.
4. Test the new version by accessing from a test client by specifying the version of the composite (x+1)
5. When verifications are complete, mark x+1 as the default version. New clients are routed to version x+1, while old clients complete their work in version x.

It is possible to leave the previous version deployed, or undeploy it. If you undeploy a composite which has in-flight instances, they become stale and do not complete. The un-deployment of a composite removes the composite from the MDS repository.

5.2.2.4 Oracle SOA Service Infrastructure Cluster-Wide Configuration Changes

The standard Java EE artifacts that the Oracle SOA Service Infrastructure uses are configured as part of Oracle WebLogic Domain in which SOA is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts such as data sources, persistent stores, and JMS modules, across the WebLogic Server domain. At the same time, the WebLogic Server cluster controls synchronization of the deployments and libraries used by the SOA Service Infrastructure.

As explained in [Section 5.2.1, "Oracle SOA Service Infrastructure Single-Instance Characteristics,"](#) SOA Service Infrastructure-specific configuration is stored in the SOA database. Changes are applied once (per SOA server), but affect all SOA servers in the same SOA domain. For example, in the high availability topology described in this chapter, if you change the callback URL or audit level for server WLS_SOA1, the change also applies to WLS_SOA2.

Note: Changes to properties of a Composite from the Oracle Enterprise Manager Fusion Middleware Control using the MBean browser are also reflected in all the nodes of a SOA cluster. For example, if you change the Mbean by selecting **oracle.soa.config node, server:WLS_SOA1**(clustered node), **SCAComposite, My Composite**, this change propagates to the other servers in the cluster.

5.3 Oracle BPEL Process Manager and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle BPEL PM for high availability.

5.3.1 Oracle BPEL Process Manager Single-Instance Characteristics

Service engines are containers that host the business logic of service components in a SOA composite application. Each service component, such as Oracle BPEL PM, Oracle Human Workflow, Decision Service, or Oracle Mediator, is executed in its own service engine (Decision Service executes in the business rules service engine). A service engine plugs into the Oracle SOA Service Infrastructure. Oracle BPEL Process engine is

the service engine running in the Oracle SOA Service Infrastructure that enables the execution of BPEL Processes.


A BPEL process provides the standard for assembling a set of discrete services into an end-to-end process flow, and developing synchronous and asynchronous services into end-to-end BPEL process flows. It provides process orchestration and storage of long running, asynchronous processes.

Figure 5–8 illustrates that the BPEL engine is a stateless part of the Oracle SOA Service Infrastructure. It starts with the Oracle SOA Service Infrastructure application and contains different modules that take care of different aspects required to execute BPEL processes. If a deployed composite contains a BPEL process, SOA Service Infrastructure invokes BPEL Process Manager to get the component from the MetaDataStore (MDS).

BPEL Process Manager uses a Dispatcher Module that maintains an in-memory logical queue containing units of work to process the incoming messages from binding components (JMS, database, Web server).

The BPEL Process Manager engine saves processes' execution state in the SOA database through a persistence module based on Oracle TopLink. The Audit framework continuously audits the work being processed by storing process execution information in the SOA database.

Figure 5–8 Oracle BPEL PM Single-Instance Architecture



5.3.1.1 BPEL Process Manager Component Characteristics

The BPEL Service Engine runs inside the SOA Service Infrastructure Java EE application (soa-infra.ear). The BPEL Service Engine does not have any singleton services. All state associated with a BPEL process is stored in a database (dehydration store) and there is no in-memory state replication required.

The processing of work by SOA Service Infrastructure Java EE EJBs is transactional and relies on Oracle WebLogic Service transaction control service. You configure the appropriate transaction stores as recommended in the basic WebLogic Server guidelines to guarantee recovery across failures in the WebLogic Server container. Additionally the BPEL engine system does not contain any Web modules, therefore, session replication is not required in the servlet layer when running BPEL in active-active.

In this release of Oracle Fusion Middleware, BPEL Service Container does not rely on JMS for asynchronous message dispatching. Therefore, there is no dependency on a distributed JMS infra-structure.

External Dependencies

The BPEL engine system depends on the following components:


- SOA Service Infrastructure database for BPEL process state persistence
- MDS repository for BPEL process metadata store

Both components must be available for the BPEL engine system to start and run properly.

5.3.1.2 Oracle BPEL Process Manager Startup and Shutdown Lifecycle

As illustrated in [Figure 5–9](#), when the Oracle SOA Service Infrastructure application starts, it initializes the BPEL engine and loads the composites from the MDS repository. If the composite contains any BPEL processes, it targets those individual components to the BPEL engine. Once the process is loaded, the system is available to receive requests. At runtime, the BPEL engine waits for requests from different channels, such as JMS, the database, and HTTP.

Figure 5–9 Startup and Shutdown Lifecycle of Oracle BPEL PM



A detailed startup and shutdown lifecycle is as follows:

1. Start SOA Server.
2. Start BPEL Engine.
3. Composites are loaded from MDS repository by SOA Service Infrastructure.
4. BPEL components are dispatched to the BPEL engine to be loaded.
5. Composite-binding components are activated.
6. The BPEL engine services requests.
7. The shutdown signal is received by SOA Service Infrastructure.
8. SOA Service Infrastructure starts undeploying loaded composites.
9. Composite-binding components are disabled.
10. BPEL components are dispatched to the BPEL engine to be unloaded.
11. The BPEL engine shuts down.

5.3.1.3 Oracle BPEL Process Manager Request Flow and Recovery

Recoverable activities are activities that have failed and can be recovered. For example, if you are using the file adapter to initiate an asynchronous BPEL process and your system crashes while the instance is processing, you can manually perform recovery when the server restarts to ensure that all message records are recovered.

There are 2 types of BPEL Processes based on the invocation interface:


- **One-Way:** Most commonly asynchronous fire-and-forget pattern.
- **Two-Way:** Synchronous request-response pattern.

The recovery semantics after a server failure are based on whether the process is invoked synchronously and asynchronously. The following describes the behavior based on invocation and process type:

- **Synchronous Invocation (Request):** For synchronous requests, an error is thrown back to the client in the case of a server failure. It is the client's responsibility to handle the error message and take appropriate action such as retrying the request. This holds true for both Transient and Durable processes.

Note: In the case of durable processes, the message is persisted to the dehydration store at certain points. It is possible to recover the message from the dehydration store and replay the process using Enterprise Manager, however, the client cannot be notified of the response. Therefore, this is not a recommended option. It is preferable to handle all recovery from the client.

- **Asynchronous Invocation (Post):** There are two types of asynchronous invocations - one that starts a new process and one that is a *callback* to an existing process. In the case of a callback, the engine (after recognizing that it is a continue operation) attempts to resolve the subscribing process, first by conversation ID, and then by a correlation set if defined. Messages associated with asynchronous requests are persisted to the dehydration store as part of the client call. If there is a failure prior to the persistence, the client receives an error message and has to handle the error in the same way as an error for a synchronous invocation. If the persistence is successful, the client call returns and further processing is done outside the context of the client call. In the case of a server failure, one-way processes that were invoked asynchronously can be restarted using Enterprise Manager.

Figure 5–10 Enterprise Manager BPEL Engine Recovery

See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* and the *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite* for details on the synchronous and asynchronous models supported by BPEL engine.


5.3.1.4 Oracle BPEL Process Manager Configuration Artifacts

Starting with Oracle Fusion Middleware 11g Release 1 (11.1.1.2), the configuration parameters for Oracle BPEL PM are stored in the SOA database. These parameters are configurable from Oracle Enterprise Manager Fusion Middleware Control.

To configure these parameters, go to the Oracle Enterprise Manager Fusion Middleware Control, and in the navigation tree, select **SOA, soa-infra, (server_name)**. Right-click **SOA Administration**, and select **BPEL Properties**, or use the MBean Browser to navigate to the appropriate property.

These properties are specific to each WebLogic domain directory that contains WebLogic Servers running Oracle Fusion Middleware SOA. Other configuration options at the container level, such as data sources, JTA configuration, and persistent stores are maintained as part of the WebLogic Server Domain configuration, and are synchronized across a cluster of WebLogic Servers by the WebLogic Server core infrastructure. See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* for more details on the configuration for the BPEL engine.

Figure 5–11 shows the WebLogic Server Administration Console where you can configure Oracle BPEL PM properties

Figure 5–11 Oracle BPEL PM Configuration Properties

If you are using Oracle SOA Suite in a clustered environment, any configuration property changes you make in Oracle Enterprise Manager on one node must be made on all nodes. Configuration properties are set in Oracle Enterprise Manager through the following options of the SOA Infrastructure menu:

Administration > System MBean Browser

SOA Administration > any property selections.

5.3.2 Oracle BPEL Process Manager High Availability Architecture and Failover Considerations

Figure 5–6 describes an Oracle SOA Service Infrastructure two-node cluster running on two WebLogic Servers. Oracle BPEL PM is deployed as part of the Oracle SOA Service infrastructure.

5.3.2.1 Oracle BPEL Process Manager Protection from Failures and Expected Behavior

For information about Oracle SOA Service Infrastructure protection from failures and expected behavior, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#).

The BPEL engine system is protected from all process failures by the WebLogic Server infrastructure. The following process failure considerations apply to Oracle BPEL:

- If the managed servers crash, Node Manager attempts to restart them locally. If whole server migration is configured and repeated restarts fail, the WebLogic Server infrastructure attempts to perform server migration of the managed server to the other node in the cluster, if it is configured. Once the server on the other node is restarted, Oracle HTTP Server resumes routing any incoming requests to it. The migrated server reads the SOA database, resumes any pending processing, and resume transactions from the transaction logs in shared storage.
- The BPEL PM Service Engine or the entire SOA Service Infrastructure may be unavailable because of errors related to JDBC data sources, or Coherence

configuration, even after a successful managed server startup. Therefore, it is necessary to monitor SOA Service Infrastructure logs for errors. For the location of server logs, see [Section 5.2.1.6, "Oracle SOA Service Infrastructure Log File Locations"](#).

In addition, for the handling of failures in the BPEL Engine itself with the WebLogic Server infrastructure, you can define and perform fault recovery actions on BPEL process faults identified as **recoverable** in Oracle Enterprise Manager Fusion Middleware Control. The recovery actions you perform on faults are based on actions you defined in your fault recovery policy files for BPEL process service components.

Three types of faults can be displayed in Oracle Enterprise Manager Fusion Middleware Control:

- **Business:** Application-specific faults that are generated when there is a problem with the information being processed (for example, a social security number is not found in the database).
- **System:** Network and other types of errors, such as a database server being completely unavailable, or a Web service being unreachable.
- **Oracle Web Services Manager (OWSM):** Errors on policies attached to SOA composite applications, service components, or binding components.

Fault recovery policies are defined at design time. The fault policies files included with the composites define the actions to take should a failure occur. For more information on how to define and use Fault recovery, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

Missing Instances

The inbound payload for one-way invocation requests is stored in the dehydration store and committed with the current global transaction context. If the caller has already started a global transaction, the invocation message is saved when the caller commits the transaction; if no incoming transaction is present, a new transaction is started to commit the invocation message. If a BPEL instance cannot be found in Enterprise Manager, the invocation message may be present in the manual recovery console (the recovery console lists all of the incoming messages that have not been delivered yet). If the invocation message is not found in the recovery console, check the status of the transaction from the caller side.

Transactional Issues with Endpoints

BPEL asynchronously saves the audit trail for instances that cannot be dehydrated within the global transaction in which the request was started. If any transactional services that are referenced from a BPEL component happen to roll back the transaction, the BPEL instance audit trail should still appear in Enterprise Manager. If there is a problem with the dehydration store, the entire global transaction is rolled back and no audit trail is saved. In this case, a message or activity can be recovered to pick up processing from the last known dehydrated point.


If the dehydration point is an Oracle RAC database with multiple nodes and the Java EE server on which the BPEL engine is hosted loses its JDBC connection to the Oracle RAC node, the BPEL engine attempts to retry the transaction (thereby rolling back whatever changes were made in the current transaction). If a new connection cannot be established to a new Oracle RAC node, the BPEL message or activity can be recovered using the recovery console.

Logging

BPEL engine loggers are set to INFO level by default. When trying to track down a message or instance, you may find it helpful to set loggers to TRACE level to enable more output to the WebLogic Server logs. For dehydration log messages, enable `oracle.soa.bpel.engine.data`. For thread or dispatcher log messages, enable `oracle.soa.bpel.engine.dispatcher`. To capture all engine log messages, enable `oracle.soa.bpel.engine`.

5.3.2.1.1 Recovering Failed BPEL and Mediator Instances In the case of a server failure, in-flight one-way processes that were invoked asynchronously may require manual recovery. Such processes are marked as **Recoverable** and can be restarted using Enterprise Manager as shown in [Figure 5–12](#).

Figure 5–12 BPEL PM Instance Recovery using Oracle Enterprise Manager



5.3.2.2 Oracle BPEL Process Manager Cluster-Wide Configuration Changes

The standard Java EE artifacts that BPEL engine uses are configured as part of Oracle WebLogic Domain in which SOA is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts, such as data sources, persistent stores, and JMS modules across the WebLogic Server domain.

As explained in the [Section 5.3.1, "Oracle BPEL Process Manager Single-Instance Characteristics,"](#) BPEL engine-specific configuration is stored in the SOA database. Changes are applied once (per SOA server) but affect all BPEL instances in the same SOA domain. For example, If WLS_SOA1 and WLS_SOA2 are part of the same SOA cluster and the number of Dispatcher threads for BPEL in WLS_SOA1 is changed to 20, WLS_SOA2's dispatcher threads are set to 20 as well.


5.4 Oracle BPM Suite and High Availability Concepts

Oracle Fusion Middleware BPM Suite provides a complete set of components for designing, deploying, and managing Business Process Management (BPM) projects.

5.4.1 Oracle BPM Suite Single Instance Concepts

As shown in Figure 5–13, Oracle BPM Suite provides a comprehensive suite of products for developing, managing and monitoring BPM Projects. Based on Oracle SOA Suite 11g, the complete set of technologies and standards like SCA can be leveraged by the BPM Suite 11g. The BPM runtime consists of service engines (BPMN Service Engine, BPEL Process Manager, Human Workflow, Rules, Mediator) and binding components (JCA Adapters, B2B) that a common Service Infrastructure manages and interconnects. The Service Infrastructure also provides common services for lifecycle management and deployment of BPM projects.


Figure 5–13 BPM Suite Single-Instance Architecture



A BPM project is the basic unit of deployment to the BPM runtime. A BPM project comprises service components (such as BPMN Process, Business Rules, Human Workflow) and references organized as a SOA Composite, organizational data (such as Roles and Business Calendars), business indicator metadata and dashboard data.

Components are targeted to service engines during deployment while services and references are enabled using the binding components. The metadata for organizational data, business indicators and dashboards is persisted and evaluated by appropriate components at runtime. At runtime, the binding component or the BPM Workspace receives messages, which the Service Infrastructure routes to the appropriate service engine(s).


Figure 5–14 shows the Oracle BPM Suite Infrastructure stack diagram:

Figure 5–14 BPM Suite Infrastructure Stack Diagram

The BPM runtime executes within the context of an application server such as the Oracle WebLogic Server. It leverages the underlying application server capabilities for load balancing and high availability.

5.4.1.1 Oracle BPM Suite Component Characteristics

[Figure 5–15](#) shows the components in BPM Suite, their relationship, and the technologies they use to interface with each other.

Figure 5–15 BPM Suite Components and Interfaces

The BPM Suite components are:

- **BPMN Service Engine:**
BPMN is an abbreviation for Business Process Modeling Notation. The BPMN Service Engine is an extension of the existing BPEL Service Engine and as such it leverages the core infrastructure of the BPEL. The BPMN Service Engine leverages JPA/EclipseLink to store/recover the state of a process instance in the SOA Infrastructure dehydration store maintained by a database and to persist audit records that are created in the course of running a process. MDS APIs are used to retrieve metadata Information about the BPMN Process Model and other BPM project artifacts like the Business Catalog.
- **BPM Composer:**
The BPM Composer is a standard J2EE Web Application. End users access it from a browser using HTTP protocol. The BPM Composer leverages the BPM infrastructure libraries. It uses the Project Metadata Layer to create and retrieve artifacts of a BPM project (like a business rule or BPMN process) from MDS.
- **BPM Workspace:**
The BPM Workspace is a standard J2EE web application. End users access it from a browser using HTTP protocol. The BPM Workspace uses APIs of the Oracle Business Process Interface (OBPI) to retrieve information about BPMN process instances and user tasks for display in the Worklist. Furthermore, it leverages the Cube Persistence APIs for standard and custom dashboards.
- **BPM Infrastructure Libraries:**

The most important BPM Infrastructure Libraries are:

- OBPI:

The Oracle Business Process Interface provides functionality to access information about process instances, user tasks, and so on. It deploys as a library in WebLogic Server and exposes an EJB interface as well as Spring Beans for dependency injection. The OBPI leverages existing SOA Services such as Human Workflow Service, Identity Service, and Facade API. OBPI is the primary interface for clients that want to access the BPM Infrastructure.

- PML:

The Project Metadata Layer is an internal API for managing BPM projects. It leverages MDS for storage, retrieval and labeling of BPM projects and is used by BPM Composer and the BPMN Service Engine

- Analytics:

The Analytics library provides a set of functionality to manage Process Cubes. The library exposes a EJB3 (stateless) interface and uses JPA/EclipseLink for persistence. This library is used by BPM Workspace.

Note: Because Oracle BPMN Service Engine is the core piece of BPM Suite, the following sections describe it in detail.

5.4.1.2 Oracle BPM Suite Component Interaction

[Table 5–1](#) shows how the BPM Suite components interact and the tasks that can be performed by each component during the execution of a BPMN process.

After the heading row in [Table 5–1](#), each row in the table is for a particular BPM Suite component. After the first column in [Table 5–1](#), each column heading represents a point in time during the execution of the BPMN process; T1 is the first point in time, T2 is the next point in time, and so on. T8 is the last point in time for the process execution. The tasks that each component can perform are shown in the row for that component at the point in time that the task is performed during the BPMN process.

Table 5–1 Tasks Performed by BPM Suite Components During BPMN Process Execution

Component	T1	T2	T3	T4	T5	T6	T7	T8
BPM Workspace		Instantiate process instance using Initiate Task		Approve user task		Monitor task performance and workload		Monitor process performance
BPM Composer	Deploy BPM project							

Table 5–1 (Cont.) Tasks Performed by BPM Suite Components During BPMN Process Execution

Component	T1	T2	T3	T4	T5	T6	T7	T8
BPMN Service Engine			Execute process instance		Continue execution of process instance		Complete process instance	
BPM Studio	Deploy BPM project							
Web Service (WS) Client		Instantiate process instance invoking composite service						Return result

The following list provides additional details about the eight points in time (T1 through T8) during the BPMN process execution shown in [Table 5–1](#):

- T1:
The example in [Table 5–1](#) assumes that a BPM project is available either in BPM Studio or BPM Composer and is ready for deployment. As part of deploying a BPM project from BPM Studio or BPM Composer, an (SCA) archive is created and transferred to the standard SOA Composite Deployer Servlet running in the WebLogic Server. The composite deployer servlet stores the content of the BPM Project in MDS and notifies the appropriate service engines that a new BPM Project has been deployed. After that, the BPMN Service Engine is ready to receive requests for processing.
- T2:
There are two ways to initiate a BPMN process:
 - Using an Initiate Task:
If the user has modeled a Initiate Task in the BPMN process diagram to start the process instance, then after deployment the task would become visible in BPM Workspace and the process instance can be started from here.
 - Invocation of the Composite Service using any Web Service (WS client) Using an Initiate Task:
In the case where the BPMN Process exposes a service interface, that service is exposed as a composite service and as such any WS client can be used to instantiate the process instance.
- T3:
The BPMN Service Engine starts executing the process and continues executing it until the first User Task is executed. For a User Task, the BPMN Service Engine creates a Human Task using OBPI. As a result, a new task is created in the SOA Infrastructure dehydration store.
- T4:
BPM Workspace uses OBPI to query new tasks for display in the Worklist. The user logged into BPM Workspace is supposed to work on the task. Eventually the task is approved or rejected and work on it has completed. BPM Workspace uses OBPI to complete the task.

- T5:
When the task is completed, that also completes the User Task Activity in the BPMN Process, and the process is able to continue.
- T6:
At any time during the execution of the process instance, a user might use BPM Workspace to monitor the task performance and workload. For this, BPM Workspace leverages the Cube Persistence API to query task performance data from the Process Cubes stored in the database.
- T7:
Eventually the BPMN process is completed.
- T8:
Typically, the last step in a BPMN process is to return a result (via a Message End Event) to the client that invoked the process. When the process has completed, a user may use BPM Workspace to monitor the process performance.

5.4.1.3 Oracle BPM Suite Startup and Shutdown Lifecycle


For the BPMN Service Engine, the startup and shutdown lifecycle is the same as for the BPEL Service Engine, since the BPMN Service Engine is based on the BPEL Service Engine. For more information, see [Section 5.3.1.2, "Oracle BPEL Process Manager Startup and Shutdown Lifecycle."](#) and [Section 5.4.2.1.3, "Oracle BPMN Service Engine Startup and Shutdown Lifecycle."](#)

For BPM Composer and BPM Workspace, the lifecycle is the same as for any web application residing in Oracle WebLogic Server.


The lifecycle of the BPM Infrastructure Libraries is the lifecycle for the system that uses and makes invocations to the libraries.

5.4.1.4 Oracle BPM Suite Configuration Artifacts

To configure the BPMN Service Engine in Enterprise Manager, from the **SOA Infrastructure** list in soa-infra choose **SOA Administration > BPMN Properties**, as shown in [Figure 5-16](#).

Figure 5–16 Displaying the BPMN Service Engine Properties Page

This displays the BPMN Service Engine Properties page shown in [Figure 5–17](#). Additional parameters are available from the System MBean browser. Click **More BPMN Configuration Properties...** on the BPMN Service Engine Properties page to launch the BPMN MBean browser.

Figure 5–17 Launching the BPMN MBean Browser

For detailed information on configuring the Oracle BPMN service engine using Enterprise Manager, see *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

For detailed instructions on configuring high availability for the Oracle BPMN service engine, refer to [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines."](#)

5.4.2 Oracle BPMN Service Engine High Availability

This section provides single instance and high availability information for Oracle BPMN service engine.

5.4.2.1 Oracle BPMN Service Engine Single Instance Characteristics


Service engines are containers that host the business logic of service components in a BPM project. Each service component, such as Oracle BPMN Service Engine, Oracle BPEL Process Manager, Oracle Human Workflow, Oracle Business Rules, or Oracle Mediator is executed in its own service engine. A service engine plugs into the Oracle SOA Service Infrastructure.

Oracle BPMN service engine is the service engine running in SOA Service Infrastructure that enables the execution of BPMN processes.

A BPMN process provides the standard for assembling a business process using standardized activities, gateways and events in a well-defined process flow. The BPMN service engine provides functionality for the execution of potentially long-running BPMN process models. It leverages the core-infrastructure features of the Oracle BPEL Process Manager such as dehydration, dispatching, and service orchestration.

5.4.2.1.1 Oracle BPMN Service Engine Single Instance Architecture As shown in [Figure 5–18](#), The BPMN service engine is a stateless part of the Oracle SOA Service Infrastructure that builds on top of the Oracle BPEL Process Manager Service Engine. For detailed information about Oracle BPEL Process Manager and high availability see [Section 5.3](#), "Oracle BPEL Process Manager and High Availability Concepts."

Figure 5–18 Oracle BPMN Service Engine Single-Instance Architecture



The BPMN service engine leverages the BPEL PM dispatcher module to dispatch incoming messages from binding components (JMS, database, Web Services) for processing.

The state of process execution is saved in the SOA database through a persistence module based on Java Persistence Architecture (JPA). The auditing infrastructure of

the BPMN service engine continuously audits the work being processed by the engine and stores audit records in the BPM database. Those audit records are used as a source for measurements and integration with Oracle Business Activity Monitoring (BAM).

The BPMN service engine runs inside the SOA Service Infrastructure Java EE application (soa-infra.ear). Since it is built on top of the BPEL Process Manager, it has identical characteristics. See [Section 5.3.1.1, "BPEL Process Manager Component Characteristics"](#) for more information.

5.4.2.1.2 Oracle BPMN Service Engine External Dependencies The BPMN service engine depends on the following components:

- SOA Service Infrastructure database for BPMN process state persistence
- BPM database for persistence of analytics data

By default, the BPM database is collocated with the SOA Service Infrastructure database and requires no additional setup.

- MDS repository for BPMN process metadata store


The components in [Figure 5–18](#) must be available for the BPMN service engine to start and run properly.

Depending on the BPM project, the BPMN service engine might depend on the following additional components:

- BAM adapter
- User directory. For example, Oracle Internet Directory or another LDAP server that is configured to work with BPM.

5.4.2.1.3 Oracle BPMN Service Engine Startup and Shutdown Lifecycle As shown in [Figure 5–19](#), when the Oracle SOA Service Infrastructure application starts, it initializes the BPMN service engine and loads the composites from the MDS repository. If the composite contains any BPMN processes, it targets those individual components to the BPMN service engine. Once the process is loaded and its BPM specific metadata persisted in the database, the system is available to receive requests.

Figure 5–19 Startup and Shutdown Lifecycle of Oracle BPMN Service Engine



The detailed startup and shutdown lifecycle is:

1. Start BPM Server.

2. Start BPMN service engine.
3. BPM project composites are loaded from the MDS repository by the SOA Service Infrastructure.
4. BPMN components are dispatched to the BPMN service engine to be loaded.
5. The BPM project metadata, such as organization data and audit/measurement metadata, is persisted in the infrastructure database.
6. Composite binding components are activated.
7. The BPMN engine services requests.
8. The shutdown signal is received by the SOA Service Infrastructure.
9. The SOA Service Infrastructure starts unloading composites.
10. Composite binding components are disabled.
11. BPMN components are dispatched to the BPMN engine to be unloaded.
12. The BPMN service engine shuts down.

5.4.2.1.4 Oracle BPMN Service Engine Log Files BPMN service engine loggers are set to INFO by default. When trying to track down a message or instance, you may want to set the loggers to TRACE level to enable more output to the WebLogic Server log files.


The BPMN service engine logs trace information to the same files as the SOA Service Infrastructure. For the location of the server logs, see [Section 5.2.1.6, "Oracle SOA Service Infrastructure Log File Locations."](#)

5.4.2.2 Oracle BPMN Service Engine High Availability Considerations


This section guides you through the issues and considerations necessary for configuring Oracle BPMN Process Manager for high availability.

5.4.2.2.1 Oracle BPMN Service Engine High Availability Architecture and Failover Considerations The BPMN service engine leverages the functionality of BPEL PM for process instance recovery. See [Section 5.3.2.1, "Oracle BPEL Process Manager Protection from Failures and Expected Behavior"](#) for more details on instance recovery.

In Oracle Enterprise Manager Fusion Middleware Control, select **BPMN Engine (Service Engine)** for the administration of the BPMN service engine, as shown in [Figure 5-20](#).

Figure 5–20 BPMN Service Engine Home Page in Enterprise Manager

The home page of the BPMN service engine comprises several sub pages. The Dashboard page is the default page. To recover BPMN service instances, select the Recovery tab, which opens the BPMN Service Engine Recovery page; see [Figure 5–21](#).

Figure 5–21 BPMN Service Recovery Page in Enterprise Manager

5.4.2.2.2 Configuring Oracle BPMN Service Engine for High Availability To configure high availability for Oracle BPMN Service Engine, see [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines."](#)

5.4.2.2.3 Cluster-Wide Configuration Changes for Oracle BPMN Service Engine The standard Java EE artifacts that Oracle BPMN components use are configured as part of Oracle WebLogic's domain in which SOA is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts, such as data sources, persistent stores, and JMS modules across the WebLogic Server domain. At the same time, the WebLogic Server cluster is in charge of synchronizing the deployments and libraries used by the Oracle BPMN components. As explained in [Section 5.4.1.4, "Oracle BPM](#)

[Suite Configuration Artifacts](#)," the Oracle BPMN components' configuration is stored in the SOA MDS database. Changes are applied once (per SOA server), but affect all Oracle BPMN instances in the same SOA domain.

5.4.3 Oracle Business Process Web Applications High Availability


This section provides single instance and high availability information for Oracle Business Process web applications.

5.4.3.1 Oracle Business Process Web Applications Single Instance Characteristics

This section provides single instance information for Oracle Business Process web applications.

5.4.3.1.1 Oracle Business Process Web Applications Single Instance Architecture [Figure 5–22](#) shows Oracle BPM Suite web applications and their interactions with other BPMN components.

Figure 5–22 Oracle BPM Suite Web Applications



The web applications are stateless. The state is maintained in the Dehydration Service and the data is exposed through the Operational APIs or in MDS.

5.4.3.1.2 Oracle Business Process Web Applications External Dependencies Both BPM Workspace and BPM Composer are deployed separately from the BPMN Service Engine, but depend on it completely. If the BPMN Service Engine is stopped (soa-infra system), neither BPM Workspace nor BPM Composer will be able to access information and metadata about deployed projects. Given the dependency of soa-infra on the SOA MDS database, by extension the Business Process web applications depend on the MDS database to be available.

5.4.3.1.3 Oracle Business Process Web Applications Startup and Shutdown Lifecycle BPM Workspace and BPM Composer are standard J2EE applications. They are started when the WebLogic Server where BPM has been deployed is started. They can be controlled

from the WebLogic Server Administration Console and can be stopped with a forced shutdown or a graceful shutdown.

5.4.3.1.4 Oracle Business Process Web Applications Log Files The Workspace and Composer web applications write to the SOA WebLogic Server log file and SOA WebLogic Server output. For the location of the server logs, refer to [Section 5.2.1.6, "Oracle SOA Service Infrastructure Log File Locations."](#)

Oracle Enterprise Manager Fusion Middleware Control can be used for diagnosing log file messages.

5.4.3.2 Oracle Business Process Web Applications High Availability Considerations

This section describes high availability considerations for Oracle Business Process web applications.

5.4.3.2.1 Oracle Business Process Web Applications High Availability Architecture and Failover Considerations Both BPM Workspace and BPM Composer are stateless web applications. When the WebLogic Servers they reside on are deployed behind a load balancer or an HTTP Server, the front end devices route requests indistinctly to either node on which the applications are running. When a node failure occurs, requests are redirected to the other available WebLogic Server and work on the user interface can continue without interruptions.

5.4.3.2.2 Configuring Oracle Business Process Web Applications for High Availability For detailed instructions on configuring high availability for Oracle Business Process web applications, refer to [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines."](#) These web applications are clustered in a stateless cluster as part of the configuration for Oracle SOA Service infrastructure.

5.4.3.2.3 Cluster-Wide Configuration Changes for Oracle Business Process Web Applications There are no specific configuration files for the web applications that reside locally on the file system for the servers. Some property files, like the Workspace property file, `workspace.properties`, are part of the `OracleBPMWorkspace.ear` file but do not include any instance-specific settings. They are deployed and available wherever Oracle BPM Workspace is deployed.

5.4.4 Oracle Business Process Analytics High Availability


This section provides single instance and high availability information for Oracle Business Process Analytics.

5.4.4.1 Oracle Business Process Analytics Single Instance Characteristics

The following section describes Oracle Business Process Analytics single instance characteristics.

5.4.4.1.1 Oracle Business Process Analytics Single Instance Architecture The Oracle BPM Suite has built-in analytical capabilities for support of business-friendly process dashboards and real-time monitoring of business processes using Oracle Business Activity Monitoring (BAM).

[Figure 5–23](#) shows the Oracle BPM Suite Analytical Infrastructure.

Figure 5–23 Oracle BPM Suite Analytical Infrastructure

From a technical perspective, the Oracle BPM Suite Analytical Infrastructure is used for the support of process analytics as follows:

- **Audit Persistence**

The BPMN service engine is continuously generating audit events comprised of activity runtime data. The data of those audit events is persisted in audit tables of the service engine dehydration store. The audit data is the source of all analytical data.

- **JMS Topic**

To de-couple process execution from the preparation and publishing of analytical data, a JMS topic configured as part of the SOA Service Infrastructure is used.

- **Cube Action MDB**

A message driven bean (MDB) is used to trigger aggregation and persistence of the analytical data to the BPM cube schema stored in the SOA Service Infrastructure database.

- **BAM Action MDB**

A message driven bean (MDB) is used to publish analytical data towards the BAM adapter installed as part of the SOA Service Infrastructure.

- **Process Persistence**

For the persistence of audit events and analytical data to the SOA Infrastructure database, the Oracle BPM Suite leverages the Java Persistence API (JPA) Infrastructure.

The persistence unit is configured to use the JTA data source
 jdbc/SOADatasource and the provider
 org.eclipse.persistence.jpa.PersistenceProvider.

5.4.4.1.2 Oracle Business Process Analytics External Dependencies The Oracle Business Process Analytics libraries and components run as part of the BPMN Service Engine infrastructure. They rely heavily on the BPMN Service Engine artifacts (queues, stores,

and JDBC resources) to function properly. As is the case for the service infrastructure, the SOA database must be available for the Analytics framework to work. Additionally, when feeding data to BAM, the required BAM system must be up and running to process the analytical information.

5.4.4.1.3 Oracle Business Process Analytics Startup and Shutdown Lifecycle Since Oracle Business Process Analytics is part of the BPMN Service Engine, its lifecycle is the same as the BPMN Service Engine's lifecycle.

5.4.4.1.4 Oracle Business Process Analytics Log Files The following loggers are available for tracing of the Oracle Business Process Analytics components:

- oracle.bpm.analytics.measurement
- oracle.bpm.analytics.cube
- oracle.bpm.analytics.bam

To enable logging for one of the analytical components, set the log level to TRACE.

5.4.4.2 Oracle Business Process Analytics High Availability Considerations

This section provides information about high availability considerations for Oracle Business Process Analytics.

5.4.4.2.1 Oracle Business Process Analytics High Availability Architecture and Failover Considerations Because Oracle Business Process Analytics is a subsystem in charge of feeding information to other systems (mainly read-oriented information), no special considerations for failover are required. A number of JMS queues used by Oracle Business Process Analytics are configured (as part of the SOA system high availability setup) as uniform distributed destinations for optimum availability and load balancing. All BPM Cube data persisted uses the SOA JDBC services, which use multi datasources configured for Oracle RAC databases.

5.4.4.2.2 Configuring Oracle Business Process Analytics for High Availability For detailed instructions on configuring high availability for Oracle Business Process Analytics, refer to [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines."](#)

5.4.4.2.3 Cluster-Wide Configuration Changes for Oracle Business Process Analytics Oracle Business Process Analytics instances do not store any local instance-specific properties. Cluster configuration changes are applied to each Oracle Business Process Analytics instance.

5.5 Oracle Mediator and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle Mediator for high availability.

5.5.1 Oracle Mediator Single-Instance Characteristics

Oracle Mediator is a service engine within the Oracle SOA Service Infrastructure. Oracle Mediator provides the framework to mediate between various providers and consumers of services and events. The Oracle Mediator service engine runs in-place with the SOA Service Infrastructure Java EE application. The runtime state for execution started by asynchronous interactions or involving parallel routing rules is

maintained in the SOA runtime database. For details about administering Oracle Mediator, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

- Synchronous and Asynchronous Interactions

Oracle Mediator provides support for synchronous and asynchronous request-response interactions. In a synchronous interaction, the client making the request remains blocked, awaiting the response. In an asynchronous interaction, the client invokes the service but does not wait for the response. You can specify a time out period for an asynchronous interaction.

- Sequential and Parallel Message Routing

Oracle Mediator can route messages to their destinations either sequentially, or in parallel.

- In a sequential routing, processing of data takes place in one single transaction.
- In a parallel routing scenario, one transaction is used for en-queueing information, and another one for de-queueing it.

5.5.1.1 Oracle Mediator Component Characteristics

If a composite contains an Oracle Mediator component, SOA Service Infrastructure targets the component to the Oracle Mediator engine for deployment. None of the services provided by the Oracle Mediator engine system are singletons, therefore, Oracle Mediator engines can run in full active-active mode. The processing of messages by the worker threads in Oracle Mediator is transactional and relies on Oracle WebLogic Server transaction control service. Configure the appropriate transaction stores as recommended by WebLogic Server guidelines to guarantee recovery across failures in the WebLogic Server container. Additionally, Oracle Mediator's engine does not contain any stateful Web modules or stateful session beans, therefore you are not required to configure any sort of session replication when running Oracle Mediator in active-active mode. The state of work and work-to-be processed is maintained by Oracle Mediator in the database. Therefore, it is critical that Oracle Mediator's database be highly available. This requires configuring multi data sources for the SOA data source as described in [Section 5.13.4, "Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA Domain"](#). For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

External Dependencies

Oracle Mediator depends on the following components:

- SOA database for Mediator message and message state persistence
- MDS repository for composite metadata store


Both components must be available for Oracle Mediator to start or run properly.

5.5.1.2 Oracle Mediator Startup and Shutdown Lifecycle

When the Oracle SOA Service Infrastructure application starts, it initializes the Oracle Mediator engine and loads the composites from the MDS repository. If the composite contains any Oracle Mediator components, it targets them to the Oracle Mediator engine. At runtime, Oracle Mediator routing rules can be invoked through an inbound binding component or by another service engine. Graceful shutdown of the Oracle Mediator engine is initiated by the SOA Service Infrastructure and involves sending signals to in-flight instances and unloading of loaded components.

Figure 5–24 illustrates Oracle Mediator Startup lifecycle.

Figure 5–24 Oracle Mediator Startup Lifecycle



5.5.1.3 Oracle Mediator Request Flow

The recovery semantics of Oracle Mediator after a server failure are based on client interaction type and routing rule type.

Oracle Mediator provides support for synchronous and asynchronous request-response interactions. The following describes the behavior based on interaction type:

- **Synchronous Interaction:** For synchronous interactions, an error is thrown back to the client in the case of a server failure. It is the client's responsibility to handle the error message and take appropriate action such as retrying the request.
- **Asynchronous Interaction:** There are two types of asynchronous invocations - one that starts a new routing rule execution and one that is a callback to an existing rule execution. In the case of a callback, the engine attempts to resolve the subscribing instance through a correlation id. If there is a failure in handling the callback, the client receives an error message and has to handle the error appropriately. The client invocation must be transactional in order to guarantee reliable handling of a callback in the case of a server failure.

Oracle Mediator can route messages either sequentially or in parallel. Only messages processed in parallel need manual recovery in the case of a server failure. The following describes the behavior based on routing rule type:

- **Sequential Routing Rule:** For sequential routing rule, complete rule is executed in a single transaction. In the case of server failure, Oracle Mediator relies on the underlying transaction manager for recovery.
- **Parallel Routing Rule:** For parallel routing rule, Oracle Mediator uses a store-and-forward paradigm that involves two separate transactions - one transaction for persisting the message and a second one for executing the routing rule. If there is a failure prior to the persistence, the client receives an error message and has to handle the error in the same way as in a sequential routing rule. If the persistence is successful, the client call returns and further processing is done outside the context of the client call. In the case of a server failure, Oracle Mediator will initiate recovery upon server restart after a configurable time interval specified by `ContainerIdLeaseRefresh`.

5.5.1.4 Oracle Mediator Configuration Artifacts

Starting with Oracle Fusion Middleware 11g Release 1 (11.1.1.2), the configuration parameters for Oracle Mediator are stored in the SOA MDS database. You can configure these parameters using Oracle Enterprise Manager Fusion Middleware Control.

To configure these parameters, go to the Oracle Enterprise Manager Fusion Middleware Control, and in the navigation tree, select **SOA, soa-infra, (server_name)**. Right-click **SOA Administration**, and select **Mediator Properties**, or use the MBean Browser to navigate to the appropriate property.

5.5.2 Oracle Mediator High Availability Architecture and Failover Considerations

Oracle Mediator is an embedded service engine that leverages the same high availability architecture as the Oracle SOA Service Infrastructure. [Figure 5–6](#) describes a two-node Oracle SOA Service Infrastructure cluster running on two WebLogic Servers. Oracle Mediator is deployed as part of the Oracle SOA Service infrastructure composite application.

In a clustered configuration, Oracle Mediator can run in an active-active mode as there are no singleton services and all state is stored in the SOA runtime database.

Note: In a single-node environment, Oracle Mediator performance is slow if there is a high latency between the middle tier and database, for example, if the middle tier and database are in different subnets.

5.5.2.1 Oracle Mediator Protection from Failures and Expected Behavior

For information about Oracle SOA Service Infrastructure protection from failures and expected behavior, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior."](#)

During the execution of a Parallel Routing Rule, Oracle Mediator obtains a logical lock on a batch of messages in the SOA runtime database. This lock contains a reference to the container ID that uniquely identifies the Oracle Mediator engine. The container ID is assigned at startup time. The user can configure the batch size using `DeferredMaxRowsRetrieved`. A smaller batch size ensures that there are a lower number of locked rows requiring recovery (as explained above) in the case of a server failure.

In case of unplanned outages, you must wait as much time as specified in the `ContainerIdLeaseRefresh` interval after restarting the server. This enables the server to complete the instances still in the running state.

In a multi-node cluster environment, if Oracle Mediator is used for an asynchronous message exchange pattern, there could be a possibility that the callback will not be handled by Oracle Mediator if it arrives before the request has completed. This could happen in a scenario where the request initiated by Oracle Mediator to the target service takes longer to complete, and callback from the target service arrives before the request completes.

Process Failure

Oracle Mediator is protected from all process failures by the WebLogic Server infrastructure. The following process failure considerations apply to Oracle Mediator:

- If the managed servers crash, Node Manager attempts to restart them locally. If whole server migration is configured and repeated restarts fail, the WebLogic Server infrastructure attempts to perform server migration of the managed server to the other node in the cluster, if it is configured. Once the server on the other node is restarted, Oracle HTTP Server resumes routing any incoming requests to it. The migrated server reads the SOA database, resumes any pending processing, and resumes transactions from the transaction logs in shared storage.

- The SOA Service Infrastructure application where the Oracle Mediator engine runs may be down due to failure in accessing resources, errors caused by the coherence infrastructure, or other issues unrelated to the status of the managed server where it is located. Therefore, monitor the Oracle SOA Service Infrastructure application to watch for errors caused by the application in the managed server logs. For the location of server logs, see [Section 5.2.1.6, "Oracle SOA Service Infrastructure Log File Locations"](#)).
- During recovery after a server crash, Oracle Mediator attempts to redeliver messages that were partially processed. If the automatic retry fails, Oracle Mediator enqueues the message to the error hospital for manual recovery.

Node Failure


In case of a node failure, server migration is triggered after the available server verifies the time stamp in the database leasing system. If the crashed server was the coherence cluster master, the available server become the new master, and the Oracle Mediator engine remains available for processing messages from binding components. After the time stamp for leasing is verified, the Node Manager in the remaining node attempts to migrate the VIP used by the failed managed server, and restarts it locally. This effectively results in the Oracle Mediator engine having two instances running in the same node.

Cluster-Wide Configuration Changes

The standard Java EE artifacts that Oracle Mediator engine uses are configured as part of Oracle WebLogic's domain in which SOA is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts, such as data sources, persistent stores, and JMS modules across the WebLogic Server domain. At the same time, the WebLogic Server cluster is in charge of synchronizing the deployments and libraries used by the Oracle Mediator engine.

As explained in [Section 5.5.1, "Oracle Mediator Single-Instance Characteristics"](#), Oracle Mediator engine-specific configuration is stored in the SOA MDS database. Changes are applied once (per SOA server), but affect all Oracle Mediator instances in the same SOA domain. For example, if WLS_SOA1 and WLS_SOA2 are part of the same SOA cluster and the heartbeat period for Oracle Mediator in WLS_SOA1 is changed to 20, WLS_SOA2's heartbeat period is also set to 20.

5.5.2.1.1 Recovering Failed Mediator Instances During recovery from a server failure, under certain circumstances, messages with in-flight routing rules may be placed in the error hospital for manual recovery. These messages can be recovered using Enterprise Manager as shown in [Figure 5–25](#).

Figure 5–25 Mediator Instance Recovery using Oracle Enterprise Manager

5.5.2.1.2 Oracle Mediator's Resequencing in a Cluster Oracle Mediator's resequencer enables rearranging a stream of related (but not necessarily ordered) messages back into order. It sequences the incoming messages that arrive in a random order, and then sends them to the target services in an orderly manner.

The resequencer works with two central concepts: groups and sequence IDs. The sequence ID is an identifier part of the message, based on which messages are rearranged. Messages arriving for resequencing are split into groups, and the messages within a group are sequenced according to the sequence ID. The sequencing is based on the sequencing strategy selected (the Standard, FIFO and Best Effort strategies are available in the Oracle JDeveloper Composite Designer).

Message processing is based on assigning a group id to batches of messages in Oracle Mediator's queue system. The resequencer dequeue system uses an instance identifier (ID) to lock a group, then processes the group messages. The heartbeat infrastructure creates and maintains the instance ID in Oracle Mediator. There is one instance ID for every Oracle Mediator instance. The heartbeat infrastructure inserts the instance ID and current time into the table `MEDIATOR_CONTAINERID_LEASE`. The heartbeat infrastructure has one heartbeat thread. This thread periodically updates the time associated with the instance ID to announce its presence to other Oracle Mediator instances. The Oracle Mediator configuration parameter `ContainerIdLeaseRefresh`, specified in minutes (default: 1), is used for this purpose. The heartbeat thread also looks for instance IDs that have not been updated within a configurable time period. The Oracle Mediator configuration parameter `ContainerIdLeaseTimeout`, specified in minutes (default: 5), is used for this purpose. The thread releases the locks held by these instance IDs. After the groups are unlocked, the dequeue system of other Oracle Mediator instances can pick up the groups for processing.

The possible states for a group stored in the `MEDIATOR_GROUP_STATUS` table are:

- Group Ready for processing - 0
- Group Locked - 1
- Group Processing Completed - 2
- Error in processing group - 3

You can configure both the `ContainerIdLeaseTimeout` and `ContainerIdLeaseRefresh` parameters from the Oracle Enterprise Manager Fusion Middleware Control configuration screens for Oracle Mediator. These properties determine the behavior of a cluster of Oracle Mediator Instances configured for re-sequencing. Each group is processed by one single server until its lease expires (the instance ID is not updated). If a managed server goes down while it was processing a group, the other instances in the cluster will proceed to its ordering once the lease has expired (that is, once the unique ID assigned to the group has not been updated during a lease timeout period).

5.5.2.2 Troubleshooting Oracle Mediator High Availability

To debug Oracle Mediator failures, check the database tables to determine which container failed. To identify requests that were in progress when Oracle Mediator failed, find the rows that are still locked, and unlock them. You may also view the payload, as it is stored as a blob.

The poll interval for Oracle Mediator Instance Manager should be the same across all the nodes. The time stamp used is the database time stamp.

The possible states of messages in the database are:

- Ready
- Locked
- Completed
- Error

5.6 Oracle Human Workflow and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle Human Workflow for high availability.

5.6.1 Oracle Human Workflow Single-Instance Characteristics

Oracle Human Workflow is a service engine running in the Oracle SOA Service Infrastructure that enables the execution of interactive human driven processes. A human workflow provides the human interaction support such as approve, reject, and reassign actions within a process or outside of any process. The Human Workflow service consists of a number of services that handle various aspects of human interaction with a business process.

All human task metadata is stored and managed in the Metadata Service (MDS) repository. The Human Workflow engine consists of a Service Engine running within the SOA Service Infrastructure and additional Java EE applications for `DefaultToDoTaskFlow` and `Worklist` applications. A human workflow leverages an internal JMS queue for notifications related to a human task.

For details about administering Oracle Human Workflow, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

5.6.1.1 Oracle Human Workflow Startup and Shutdown Lifecycle

The Human Workflow startup consists of two phases: loading of the Java EE applications, and initialization of the service engine. The Java EE applications are loaded as part of the application server startup. The service engine initialization and shutdown is controlled by SOA Service Infrastructure. Post initialization, composites that contain Human Workflow components are targeted to the Human Workflow service engine by SOA Service Infrastructure.

5.6.1.2 Oracle Human Workflow Request Processing

A human workflow can be initiated by an invocation from another SOA engine (such as Oracle BPEL PM). The message is routed to the engine by SOA Service Infrastructure and is persisted by the Workflow engine within its runtime schema. The message becomes available for human actions through the browser based UI after the client transaction that is associated with the invocation is committed. Each update on the message or the runtime state through a user action is a separate transaction.

As soon as Workflow commits its transaction, the control passes back to BPEL which almost instantaneously commits its transaction. Between this window, if the Oracle RAC instance goes down, on failover, the message is retried and can cause duplicate tasks.

5.6.1.3 Oracle Human Workflow Configuration Artifacts

Starting with Oracle Fusion Middleware 11g Release 1 (11.1.1.2), the configuration parameters for Oracle Human Workflow are stored in the SOA MDS database. You can configure these parameters using Oracle Enterprise Manager Fusion Middleware Control.

To configure these parameters, go to the Oracle Enterprise Manager Fusion Middleware Control, and in the navigation tree, select **SOA, soa-infra, (server_name)**. Right-click **SOA Administration, Workflow Notification/Task Service Properties**, or use the MBean Browser to navigate to the appropriate property.

5.6.1.3.1 Managing the URI of the Human Task Service Component Task Details Application You can add or remove the URI of the task details application used in Oracle Human Workflow.

To manage the URI of the human task service component task details application:

1. Access this page through one of the following options:

From the SOA Infrastructure Menu...	From the SOA Folder in the Navigator...
1. Select Home. 2. Select the Deployed Composites tab. 3. In the Composite section, select a specific SOA composite application.	1. Under soa-infra, select a specific SOA composite application.

2. Select the human task service component in the **Component Metrics** table.
3. Click **Administration**.

The Administration page shows the URI for the task details application.

testall [3.0] ⓘ Logged in as weblogic | Page Refreshed Apr 17, 2009 6:33:20 PM PDT ↺

SOA Composite ▼

testall [3.0] > Task2

Task2 (Human Workflow Component) ⓘ Related Links ▼

Dashboard Instances Faults Policies **Administration**

Add or remove the URI for the user defined task details application.

View ▼ + Add URI ✗ Remove

Application Name	Host Name	HTTP Port	HTTPS Port	URI
worklist	myhost39.us.oracle.com	8001	0	/workflow/testallTaskFlow/fi

Apply Revert

4. Click the **Add** icon to specify the following details for the URI:

- Application name
 - Host name
 - HTTP port
 - HTTPS port (optional)
 - URI
5. Click **Apply**.

5.6.2 Oracle Human Workflow High Availability Architecture and Failover Considerations

Figure 5–6 describes a two-node Oracle SOA Service Infrastructure cluster running on two WebLogic Servers. Oracle Human Workflow is deployed as part of the Oracle SOA Service infrastructure composite application.

5.6.2.1 Oracle Human Workflow Protection from Failures and Expected Behavior

Oracle Human Workflow's engine uses transactional EJBs for persistence and JMS queues for user notification. All state is stored in the database and the JMS queue, and there is no in-memory session state to be replicated for recovery. Therefore, Oracle Human Workflow's service engine and the associated Java EE applications are run in an active-active topology on a WebLogic cluster. In the case of a server or hardware crash, whole server migration must be configured for recovering pending transactions and JMS messages stored on the local queue. Notifications are not sent out until the server is restarted.

5.6.2.2 Manual Recovery Required for Human Workflow Task in Rejected MSG Table

The `FabricInstanceManager.persistCompositeInstanceBean` API (or other `InstanceManager` APIs) does not have built-in retry logic. When there are transient failures caused by Oracle RAC failover, the WS invocation is rejected directly. This manual recovery of these rejected messages is done using Oracle Enterprise Manager Fusion Middleware Control.

To manually recover the rejected messages, in the Oracle Enterprise Manager Fusion Middleware Control, select **SOA, soa-infra(soa_server1), <composite_name>, Faults and Rejected Messages**, select the message, and choose **Recover With Options**. This causes the workflow request to be retried.

5.6.3 Troubleshooting Oracle Human Workflow High Availability

Human Workflow works like a standard Java EE application with a browser based client. Once a message is persisted to its runtime schema, it can be processed using the workflow UIs. If a server crashes in the middle of the transaction, you must recover the server for successful recovery or rollback. The following logs are relevant for debugging errors associated with Oracle Human Workflow:

- `oracle.soa.services.common`
- `oracle.soa.services.notification`
- `oracle.soa.services.workflow`
- `oracle.soa.services.workflow.common`
- `oracle.soa.services.workflow.evidence`

- oracle.soa.services.workflow.metadata
- oracle.soa.services.workflow.persistency
- oracle.soa.services.workflow.query
- oracle.soa.services.workflow.runtimeconfig
- oracle.soa.services.workflow.soa
- oracle.soa.services.workflow.task
- oracle.soa.services.workflow.task.dispatch
- oracle.soa.services.workflow.task.routing

5.7 Oracle B2B and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle B2B for high availability.

5.7.1 Oracle B2B Single-Instance Characteristics

Oracle B2B connects SOA composite applications to external services, applications, and technologies. Oracle B2B offers a multi-protocol gateway that supports industry-recognized B2B standards. Oracle B2B extends Oracle SOA Suite with business protocol standards, such as electronic data interchange (EDI), ebXML, HL7, and RosettaNet. For details about Oracle B2B's functionality, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

Oracle B2B is a binding component within SOA Service Infrastructure. B2B metadata consists of trading partner details along with their supported documents and delivery channels which is stored within the MDS repository. The SOA composite has a reference to this metadata.

Oracle B2B receives messages from different channels. A listener thread logs the messages in a table in the SOA database and sends the corresponding events to a JMS queue. Event handlers listen for events to process the requests in the JMS queues.

Oracle B2B supports the following transport protocols for communication with business entities:

- HTTP(S)
- Oracle Advanced Queue
- Email (SMTP 1.0, IMAP 1.0, POP3)
- File
- FTP and SFTP (SSH FTP)
- TCP/IP
- JMS

For HTTP protocol, Oracle B2B uses the system's HTTP listener. For FTP and Email, Oracle B2B uses the external FTP and email server configured in the system. Even in high availability mode, only one B2B instance polls for incoming message for FTP, File, and email protocols.

Note: MLLP & TCP/IP protocols are not supported in a clustered environment.

5.7.1.1 Oracle B2B Component Characteristics

Oracle B2B runs in-place with SOA Service Infrastructure Java EE application. The Oracle B2B user interface is a Web application that is deployed as a standalone war file on the same managed server as the SOA Service Infrastructure. Oracle B2B UI application is stateful and stores information in the HTTP Session.

Oracle B2B stores state information within JMS queues and the SOA run-time database. Whole server migration is required for automatic JMS message and transaction recovery after a server failure.

Oracle B2B uses JMS intensively. Oracle B2B uses the Oracle WebLogic Server migration feature for protecting Oracle B2B JMS resources against failures. Oracle B2B's engine uses EJBs and Servlets, which are stateless. All state information is persisted in JMS queues and the database.

High availability of B2B depends on the high availability of the run-time database. Hence, the SOA runtime data sources should be configured as GridLink data sources or multi data sources for high availability. See [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC"](#) or [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

External Dependencies

Oracle B2B depends on the following components:

- Oracle SOA database for Oracle B2B message and message state persistence
- MDS repository for Oracle B2B metadata store
- FTP and email servers if the corresponding adapters are used

The SOA database and the MDS repository must be available for Oracle B2B to start and run properly.


5.7.1.2 Oracle B2B Startup and Shutdown Lifecycle

When the Oracle SOA Service Infrastructure application starts, it initializes Oracle B2B's engine. Oracle B2B metadata deployment should occur before composites are deployed to the Oracle SOA Service Infrastructure. Oracle B2B end-points are defined as channels which are started as part of the engine initialization. Based on the protocol, each configured channel has external dependencies, such as:

- If FTP is used, the appropriate FTP server must be started.
- If email is used, the appropriate email server must be started.
- If HTTP/HTTPS is used, the HTTP server front ending the B2B system must be started

5.7.1.3 Oracle B2B Request Flow

In an Oracle B2B system all processing is done asynchronously. When a request arrives, the corresponding message is inserted in the SOA database. A notification is inserted in a JMS queue as a placeholder for the pending work. The event handler threads listen for events in these queues and process the work sequentially.

Figure 5–26 Oracle B2B Request Flow

5.7.1.4 Oracle B2B Configuration Artifacts

You can enable and disable metrics for Oracle B2B's Server using Oracle Enterprise Manager Fusion Middleware Control. This is the only property directly exposed by Oracle Enterprise Manager Fusion Middleware Control. The configuration for Oracle B2B's Server is maintained in the SOA database and most of it is exposed by the Oracle B2B user interface application. For details about Oracle B2B configuration, see the *Oracle Fusion Middleware User's Guide for Oracle B2B*.

5.7.2 Oracle B2B High Availability Architecture and Failover Considerations

[Figure 5–6](#) describes a two-node Oracle SOA Service Infrastructure cluster running on two WebLogic Servers. Oracle B2B is deployed as part of the Oracle SOA Service infrastructure composite application. The following high availability characteristics are specific to an Oracle B2B high availability deployment:

- The Oracle B2B user interface application runs inside each one of Oracle WebLogic Server servers, and as part of the same cluster as the Oracle SOA Service Infrastructure application.
- Oracle B2B's server maintains the partners, documents, and channels definitions in the SOA database using an Oracle RAC database.

5.7.2.1 Oracle B2B Protection from Failures and Expected Behavior

This section describes how an Oracle B2B high availability cluster deployment protects components from failure and the expected behavior if component failure occurs.

Oracle B2B UI Failure

The Oracle B2B user interface application maintains some navigation information in memory. When a failure occurs in one of the managed servers running the Oracle B2B user interface, users' requests are redirected to another active WebLogic Server running the application. Ongoing requests from Oracle HTTP Server time out according the time out setting in Oracle HTTP Server's configuration. Failover requires those users accessing the failed instance to log in again, since the application is not enabled for serializing the session information.

The Oracle B2B user interface in-memory data is non-transactional, and some steps may need to be revisited in the case of failover.

For general information about Oracle SOA Service Infrastructure process failure, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#)

Node Failure

If a node failure occurs, or if the local Oracle WebLogic Server Node Manager reaches the maximum restart tries on the failed managed server, whole server migration is triggered after the available server verifies the time stamp in the database leasing system. The other Oracle B2B engine remains available for processing new messages from the different channels. At the same time, the remaining Oracle B2B server should resume the pending operations for singleton channels, such as FTP, email, or file, after the default timeout period is reached in the time stamp the failed node sets in the database (two minutes by default). For Oracle B2B's User Interface application, ongoing requests from Oracle HTTP Server time out according to Oracle HTTP Server configuration.

Database Failure

For information about Oracle B2B database failure, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#)

5.7.2.2 Oracle B2B Cluster-Wide Configuration Changes

The standard Java EE artifacts that the Oracle B2B engine uses are configured as part of the Oracle WebLogic domain in which the Oracle SOA Service Infrastructure is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts, such as data sources, persistent stores, and JMS modules across the WebLogic Server domain. At the same time, the WebLogic Server cluster controls synchronization of deployments and libraries used by Oracle B2B's engine.

As explained in [Section 5.7.1, "Oracle B2B Single-Instance Characteristics,"](#) all Oracle B2B server-specific configuration is maintained in the database, and configuration changes are applied to all the SOA Servers running in a WebLogic Server domain. Therefore, configuration properties, such as Payload Size, and Outbound Dispatcher Counts are applied cluster-wide, meaning they are used by all instances in the Oracle SOA cluster.

To set up File, FTP, or Email transports in a high availability environment, specify a unique name for each instance by using `b2b.HAInstanceName unique_instance_name`. If you use `#ServerName#` for the value, Oracle B2B retrieves the WebLogic Server name as the `HAInstanceName`.

"Properties to Set in Fusion Middleware Control" in the *Oracle Fusion Middleware User's Guide for Oracle B2B* also describes setting up File, FTP, or Email transports in a high availability environment.

5.7.2.3 Oracle B2B Deployments in a Cluster

Using the command line utility for deploying, purging, or importing metadata in B2B may cause inconsistencies and errors in the B2B system. For B2B, deploy agreements and purge or import metadata **ONLY** from the GUI available in the B2B console, instead of using the command line utility.

5.7.2.4 Troubleshooting Oracle B2B Active-Active Configuration

This section provides troubleshooting tips and possible resolutions for Oracle B2B active-active configurations.

5.7.2.4.1 Purge, Import, or Deployment of B2B Metadata When performing purge, import, or deployment of B2B metadata in a cluster, error timing and load balancing may cause exceptions that are unlikely to happen if a retry of the operation is performed.

There is no clean up nor other additional steps required. If errors, such as [java] MDS-02202: Content of the metadata object appear for deployment or "postTransfer: MDS-00521: error while reading document... appear for purge or import, retry the operation.

5.7.2.4.2 Error While Retrieving Oracle B2B Document Definitions Problem: Error happens when trying to retrieve a document definition XSD from Oracle B2B. B2B resides in a cluster and is accessed through a load balancer. B2B console report the following:

```
An error occurred while loading the document definitions.
java.lang.IllegalArgumentException: Cluster address must be set when clustering
is enabled.
```

Solution: This occurs if you do not set the front end HTTP host and port for the Oracle WebLogic cluster where Oracle B2B resides. To eliminate this error, set the front end address for the SOA Cluster:

1. In the WebLogic Server Administration Console, in the Change Center section, click **Lock & Edit**.
2. In the left pane, choose the **Environment in the Domain Structure** window and then choose **Clusters**. The Summary of Clusters page appears.
3. Select the WLS_SOA cluster.
4. Select **HTTP**.
5. Set the values for the following:
 - **Frontend Host:** soa.mycompany.com
 - **Frontend HTTPS Port:** 443
 - **Frontend HTTP Port:** 80
6. Click **Save**.
7. To activate the changes, click **Activate Changes** in the Change Center section of the Administration Console.
8. Restart the servers to make the Frontend Host directive in the cluster effective.

5.8 Oracle Web Services Manager and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle WSM for high availability.

5.8.1 Oracle WSM Single-Instance Characteristics

Oracle Web Services Manager (Oracle WSM) provides a policy framework to manage and secure Web services consistently across your organization. It provides capabilities to build, enforce, execute and monitor Web service policies including security, WSRM, MTOM and addressing policies. It typically gets deployed along with SOA Service Infrastructure.

Oracle Web Services Manager is made up of the following components:

- **Policy Manager** reads and writes security and management policies including predefined and custom policies from the MDS repository. It is typically deployed on the Oracle SOA Service Infrastructure managed servers. However, you can deploy Policy Manager on separate managed servers.

Policy Manager is a stateless Java EE application. It exposes its capabilities through stateless session beans. Policy Manager does not cache any data, (all accesses go to the MDS repository).


- **Agent** is responsible for policy enforcement, execution and gathering of runtime statistics. The Oracle WSM Agent is available on all Oracle Fusion Middleware managed servers. It is configured on the same server as the application which it protects.

The Oracle WSM Agent is made up of a set of jar files, which are a part of underlying Web service stack. It does not have any session state. The Agent maintains an in-memory policy cache, which is populated at the Agent startup time. It does not use any JTA or JMS.

The Oracle WSM Agent is made up of the following two pieces:

- **Policy Access Point (PAP)** communicates with Policy Manager. The Agent communicates with the Policy Manager through EJB invocations over T3 (or T3s if SSL is enabled) protocol.
 - **Policy Interceptor** is generated when a Web service is deployed and activated, or when a policy is attached to a Web service using Enterprise Manager. If new web services are protected using Oracle WSM, an additional instance of the interceptor is generated for each new Web service. Interceptor is responsible for policy enforcement.
- **Metadata Store** Policies are stored in the MDS repository. It is typically backed by an Oracle database. For high availability purposes, Oracle recommends using an Oracle RAC database as the back end for MDS repository.
 - **Enterprise Manager** is used to configure Oracle WSM. It also displays different web services metrics gathered by Oracle WSM.

For more details of Oracle WSM architecture, please refer to the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

Figure 5–27 Oracle WSM Single-Instance Architecture

5.8.1.1 Oracle WSM Component Characteristics

The Oracle WSM Agent is a set of .jar files available on every Oracle Fusion Middleware managed server in a Web services stack.

Policy Manager is contained in the `wsm-pm.ear` file. None of the services provided by Oracle WSM are singletons, therefore, it can run in full active-active mode. Oracle WSM services can be validated by `http://SOAHOSTx:port/wsm-pm/validator`. This validator displays Oracle WSM policies.

The Oracle WSM Agent and Oracle Enterprise Manager interact with Policy Manager using the EJB interfaces. The EJBs used in Oracle WSM are stateless and can be deployed in a clustered environment. Therefore, there is no requirement to enable state replication in the cluster.

The Oracle WSM Agent and Policy Manager need not be co-located. However, the Agent expects Policy Manager to be deployed on at least one node of the domain. The Oracle WSM Agent has capabilities to auto-discover Policy Managers deployed in the domain.

Neither the Oracle WSM Agent nor Policy Manager use JTA or JMS messaging. The MDS-based policy store also does not support JTA.

External Dependencies

Oracle WSM Policy Manager depends on the following components:

- MDS repository for storing the policies
- Oracle WSM Agent depends only on Oracle WSM Policy Manager.

Both components must be available for Oracle WSM to start and run properly.

5.8.1.2 Oracle WSM Startup and Shutdown Lifecycle

The following key Oracle WSM components are involved in the startup lifecycle for Oracle WSM:

- Oracle WSM Policy Manager
- Oracle WSM Agent

Policy Manager is a stateless application which does not perform any caching. There is no special application level startup sequence performed when the managed server where Policy Manager is deployed starts up. Policy Manager communicates with the MDS repository to retrieve policies. The MDS repository can be stored in an Oracle RAC database to provide MDS high availability.

When a managed server on which an Agent is configured comes up, the Agent connects to Policy Manager to get latest revision of policies. If it succeeds, the changes to the policies are downloaded and cached. Once the Agent is up and running, it periodically attempts a cache refresh at a configurable interval. The default time is every one minute.

Oracle WSM Agent communicates with the Policy Manager through EJB invocations over T3 (or T3s if SSL is enabled) protocol. If Policy Manager is deployed on different nodes and some of them have SSL enabled and others don't, Agent communicates only with the nodes with SSL connections.

If the Policy Manager to which the Oracle WSM Agent is connected becomes unavailable, the underlying infrastructure automatically connects to another Policy Manager instance running elsewhere in the cluster. This is achieved through Oracle WebLogic EJB clustering.

For high availability scenarios, if an Oracle WSM application is targeted to multiple nodes, it should be targeted to a cluster rather than to individual managed servers.

If a managed server has Web services deployed which are protected by Oracle WSM, and the Oracle WSM Agent is not able to communicate with any of the Policy Managers at startup time, Web service invocation fails.

5.8.1.3 Oracle WSM Request Flow

When a protected Web service is accessed by a client application, the Oracle WSM Agent queries the policy cache and enforces the applicable policies. Based on the policies, the request is authenticated, encrypted, decrypted, authorized or logged. It does not connect to Policy Manager for any of these operations. Runtime availability of Policy Manager does not affect the functioning of the Oracle WSM Agent, unless there is a configuration change, such as new Web services, which are protected by Oracle WSM, being deployed, or new policies attached to existing Web services. If there is such a configuration change, then the Oracle WSM Agent must connect to Policy Manager to get the applicable policies. If it cannot connect after initial startup, it continues to operate based on the cached policies.

5.8.1.4 Oracle WSM Configuration Artifacts

Oracle WSM Agent configuration resides in `policy-accessor-config.xml`. This file is located in the `DOMAIN_HOME/config/fmwconfig` directory. With this configuration file you can specify:

- Policy Manager URL (if configured)
- Cache Refresh Interval
- Clock skew, to allow for difference in system clock of the client and servers

These options are also available from Oracle Enterprise Manager Fusion Middleware Control and are specific to each Oracle WSM Agent installation.

Other configuration options at the container level, such as data sources for MDS repository location, and application targeting, are maintained as part of Oracle WebLogic Server Domain configuration, and are synchronized across a cluster of Oracle WebLogic Servers by Oracle WebLogic Server core infrastructure.

5.8.2 Oracle WSM High Availability Architecture and Failover Considerations

[Figure 5–6](#) describes an Oracle SOA Service Infrastructure two-node cluster running on two WebLogic Servers. Oracle WSM is deployed as part of the Oracle SOA Service infrastructure composite application. The following high availability characteristics are specific to an Oracle WSM high availability deployment:

- Oracle WSM Policy Manager and Oracle WSM Agents are deployed on WLS_SOA_INFRA. Oracle WSM Agents are available on any custom WLS cluster in case there is a need to protect any custom web services deployed on them.
- Oracle WSM Policy Manager and Oracle WSM Agents run in WebLogic Server Infrastructure managed servers that are part of a WebLogic Server cluster. The WebLogic Server cluster synchronizes configuration for common artifacts of WebLogic Server used by Oracle SOA Service Infrastructure (JDBC).
- Oracle WSM Policy Manager EJBs leverage clustering and high availability capabilities of the WebLogic Server cluster.
- All Oracle WSM Policy Manager instances in the cluster point to the same MDS repository.
- The MDS repository where Oracle WSM policies are stored is configured in an Oracle RAC database to protect from database failure.

5.8.2.1 Oracle WSM Protection from Failures and Expected Behavior

Since the Oracle WSM Agent is deployed on the same managed server as the application is deployed, it will be available again as soon as the application becomes available due to server restart/migration. The following two sections describe the failover for Policy Manager.

Process Failure

Oracle WSM components are protected from process failures by the WebLogic Server infrastructure:

- If the managed servers crash, Node Manager attempts to restart them locally. If whole server migration is configured and repeated restarts fail, the WebLogic Server infrastructure attempts to perform server migration of the managed server to another node in the cluster. Once the server on the other node is restarted, Oracle HTTP Server resumes routing any incoming requests to it. At startup time, Oracle WSM Agent and Policy Manager go through the startup lifecycle as described in [Section 5.8.1.2, "Oracle WSM Startup and Shutdown Lifecycle"](#).
- If the managed server running Policy Manager is restarted or migrated, the failover is transparent to the agents connected to it. Policy Manager leverages underlying EJB clustering and the failover infrastructure of Oracle WebLogic Server.

Node Failure

If node failure occurs and whole server migration is not configured, the Oracle WSM Agent fails over transparently to the other Policy Manager instance.

If whole server migration is configured, server migration is triggered after the available server verifies the time stamp in the database leasing system. After the time stamp for leasing is verified, the Node Manager in the node that still remains available attempts to migrate the VIP used by the failed managed server, and restarts it locally. This effectively results in the Oracle WSM Policy Manager application having two instances running in the same node. Refer to the [Section 3.9, "Whole Server Migration"](#) for more details on the failover process. In this situation, the Oracle WSM Agents load balance against both Policy Manager instances running on the same node.

5.8.2.2 Oracle WSM Cluster-Wide Configuration Changes

The standard Java EE artifacts that Policy Manager uses are configured as part of the Oracle WebLogic domain in which the Oracle SOA Service Infrastructure is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts, such as data sources, across the WebLogic Server domain. At the same time, the WebLogic Server cluster controls synchronizing the deployments and libraries used by different Oracle WSM components.

As explained in the single-instance section, Oracle WSM instance specific aspects are configured individually per WebLogic Server domain directory (typically one per node) through the `policy-accessor-config.xml` file. This file is included in the .jar file that is created when using the Oracle WebLogic Server Pack utility, therefore it is propagated to other nodes when you run `pack/unpack` to set up high availability for the Oracle SOA Service Infrastructure. However, ongoing changes to `policy-accessor-config.xml` are not replicated automatically to the other nodes. This implies that for updates, you must modify each node individually, for example, if you updated the Policy Manager URL or Cache Refresh Interval. These modifications must be manually performed in all the domain directories across a high availability topology. You can edit the `policy-accessor -config.xml` file directly, or use Oracle Enterprise Manager Fusion Middleware Control to make the configuration change on each managed server.

Oracle WSM Agents are capable of automatically discovering the deployed Oracle WSM Policy Manager deployments. If you want to override the behavior and some point to a specific Policy Manager instance, you can do so by editing the `policy-accessor-config.xml` file.

5.8.2.3 Configuring the Java Object Cache for Oracle WSM

The Java Object Cache (JOC) should be configured among all the servers running Oracle WSM. This local cache is provided to increase the performance of Oracle WSM.

For instructions on configuring the Java Object Cache among all the servers running Oracle WSM, see [Section F.1, "Configuring the Java Object Cache."](#)

5.8.2.4 Configuring Distributed Notifications for the MDS Repository

In high availability environments, Oracle recommends that you configure distributed notifications for the MDS repository.

For instructions on configuring distributed notifications for the MDS repository, see [Appendix G, "Configuring Distributed Notifications for MDS."](#)

5.9 Oracle User Messaging Service and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle User Messaging Service for high availability.

5.9.1 Oracle User Messaging Service Single-Instance Characteristics

Oracle User Messaging Service (Oracle UMS) enables two way communication between users and deployed applications. It has support for a variety of channels, such as email, IM, SMS, and text-to-voice messages. Oracle UMS is integrated with Oracle Fusion Middleware components such as Oracle BPEL PM and Oracle WebCenter Portal. It is typically deployed along with the Oracle SOA Service Infrastructure.

Oracle UMS comprises the following components:


- **UMS Server** is a Java EE application that runs on Oracle WebLogic Server. The UMS Server orchestrates message flows between applications and users. The server routes outbound messages from a client application to the appropriate driver, and routes inbound messages to the correct client application. The server also maintains a repository of previously sent messages in a persistent store, and correlates delivery status information with previously sent messages.
- **UMS Drivers** connect UMS to the messaging gateways, adapting content to the various protocols supported by UMS. Drivers can be deployed or undeployed independently of one another depending on the messaging channels available in a given installation.

UMS Drivers adapt sent and received content to and from external protocols, such as email, XMPP, and SMPP. UMS Drivers are also Java EE applications deployed to an Oracle WebLogic Server.

- **UMS Client applications** implement the business logic of sending and receiving messages. Examples of a UMS client application include an Oracle SOA application that sends messages as one step of an Oracle BPEL PM workflow, or an Oracle WebCenter Portal: Spaces application that can send messages from a Web interface.

UMS client applications have either a UMS-specific EJB module embedded in them, or interact as standard Web service clients.

[Figure 5–28](#) illustrates the services and dependencies of an Oracle UMS single-instance architecture.

Figure 5–28 Oracle User Messaging Service Single-Instance Architecture

5.9.1.1 Oracle User Messaging Service Component Characteristics

This section describes the characteristics of Oracle UMS components.

UMS Server is made up of the following:

- Message driven beans (MDBs) that dequeue messages from JMS queues
- Stateless session beans to implement messaging business logic
- JAX-WS servlets to implement the messaging web services
- A simple Oracle ADF Faces user interface component for managing end user messaging preferences.

UMS drivers typically contain the following:

- JCA resource adapter (embedded within the EAR) in order to interface with external protocol gateways

- Two MDBs, one for inbound, and one for outbound, that interface between the resource adapter and the JMS queues
- Some UMS drivers that implement an HTTP-based protocol, such as VoiceXML, have a servlet module as well.

UMS client applications that use the EJB interface have a stateless session bean that provides the API used by client business logic, and an MDB to asynchronously receive inbound messages.

UMS depends heavily on JMS. JMS queues are also used to buffer content between clients and servers, and between servers and drivers.

Each typical messaging operation (sending an outbound message or receiving an inbound message) involves two JMS queues, one between the client and server, and one between the server and driver.

UMS messaging state is stored in the database and in persistent JMS queues.

External Dependencies

UMS depends on an external database repository to maintain message and configuration state. It shares this state between clustered instances. The UMS Server accesses the database using a Java EE JDBC data source provisioned at installation time. This data source is a non-XA data source. Therefore, UMS does not depend on the JTA capabilities of the underlying infrastructure.

UMS uses JMS to deliver messages among messaging applications. By default it is configured to use a file-based persistent JMS store, therefore it depends on the storage device where those files are located.

5.9.1.2 Oracle User Messaging Service Startup and Shutdown Lifecycle

As Java EE applications, all UMS components are started by Oracle WebLogic Server container.

UMS Server Startup

At server startup time, UMS Server initializes a TopLink session that creates a connection to the database repository. UMS Server then begins listening on Web service endpoints. EJBs also become available for invocation of functionality such as sending messages or retrieving delivery status.

UMS Driver Startup

When a UMS driver is deployed to a cluster, it sends a registration message to the local UMS server. The UMS server records the registration information in the database, making it available to all UMS servers in the cluster. Once this occurs, any UMS server can route messages to any UMS driver in the cluster. This happens when new drivers are deployed, or when existing drivers are restarted following a configuration change.

At the time of server startup, UMS drivers typically establish a connection to the external gateway for which they are configured. Some of these are persistent socket-level connections, such as SMPP and XMPP. Some connections are established and torn down for each request, such as SMTP and IMAP connections. Drivers then make a remote EJB call to the UMS server to register themselves.

UMS Client Application Startup

When the managed server where the clients are deployed starts up, EJB clients typically make a remote EJB call to register themselves with the UMS server. Web service clients do not have explicit registration mechanisms.

5.9.1.3 Oracle User Messaging Service Request Flow

UMS clients make a series of short-lived requests to the UMS server. Web service clients make requests using SOAP or HTTP. EJB clients make remote EJB calls to the server for initialization and configuration management, in addition to interacting using JMS.


After initial startup and registration, subsequent UMS request flows can be categorized as follows:

- **Outbound messaging:** For an outbound message, a client application wishes to send a message to a user. The client application either makes a Web service request to the UMS Server, or enqueues a message in the UMS Server's sending JMS queue. The UMS Server receives the request, records the messaging state, and selects an appropriate driver. The UMS Server enqueues the message in a different outbound JMS queue. The corresponding driver, which is listening to the outbound queue, dequeues the message, converts it to the appropriate format for the given protocol, and delivers it to an external gateway (such as an SMTP server or IM gateway.)
- **Inbound messaging:** For an inbound message, the end user wishes to send a message back to the client application. The user sends a message from a device, such as an IM client or phone. The message passes from the external gateway to the driver. The driver adapts the message to the UMS format, and enqueues it in a JMS queue. The server dequeues the message from the JMS queue, records the message state, and enqueues the message in the appropriate application's inbound JMS queue. The client application then dequeues the message and processes it as needed.

5.9.1.4 Oracle User Messaging Service Configuration Artifacts

The UMS Server and drivers each have an XML configuration file that is the artifact of a configuration Mbean, implemented using Oracle Fusion Middleware JMX Framework. You can perform configuration changes using Oracle Enterprise Manager Fusion Middleware Control.

Figure 5–29 Configuring Oracle User Messaging Service using Oracle Enterprise Manager Fusion Middleware Control



All UMS components use standard Java EE deployment descriptors and Oracle WebLogic proprietary descriptors for configuring Java EE components. You can configure these descriptors using standard Oracle WebLogic tools, such as WebLogic

Server Administration Console or WLST. For more information, see UMS administration topics in *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

If you are using Oracle SOA Suite in a clustered environment, any configuration property changes you make in Oracle Enterprise Manager on one node must be made on all nodes. Configuration properties are set in Oracle Enterprise Manager through the following options of the SOA Infrastructure menu:

Administration > System MBean Browser

SOA Administration > any property selections.

5.9.2 Oracle User Messaging Service High Availability Architecture and Failover Considerations

See [Section 5.2.2, "Oracle SOA Service Infrastructure High Availability Architecture and Failover Considerations"](#) for a description of the Oracle SOA Service Infrastructure two-node high availability characteristics. [Figure 5–6](#) illustrates an Oracle SOA Service Infrastructure cluster running on two WebLogic Servers. Oracle User Messaging Service is deployed as part of the Oracle SOA Service infrastructure composite application. The following high availability characteristics are specific to an Oracle User Messaging Service high availability deployment:

- All UMS components are deployed to Oracle WebLogic Service Infrastructure managed servers that are part of an Oracle WebLogic Server cluster. An Oracle WebLogic Server cluster synchronizes configuration for common artifacts of Oracle WebLogic Server used by UMS, such as JDBC data sources.
- All Oracle UMS components are stateless.
- UMS Server and Client stateless EJBs leverage clustering and high availability capabilities of Oracle WebLogic Server cluster
- UMS Server relies on a shared database repository for persistent storage.
- UMS relies on JMS distributed destinations for load balancing and availability across cluster nodes. It also relies on the JMS connection factory's capability to failover to a different JMS server if a failure occurs.
- The user messaging preferences user interface does not require session stickiness. It remains available through the use of a basic load balancing. There are no sticky session routing requirements, as all session state is persisted in the database and shared across the clustered instances.
- UMS does not participate in any global transactions.
- UMS uses Oracle WebLogic Server GridLink data source or multi data source to connect to the back-end Oracle RAC database.

5.9.2.1 Oracle User Messaging Service Protection from Failures and Expected Behavior

Oracle UMS is typically deployed on the same managed server as the Oracle SOA Service Infrastructure. The Oracle SOA Service Infrastructure is protected from process and node failures using Oracle WebLogic Server whole server migration. Whole server migration also provides failover capabilities for JMS usage.

For information about Oracle UMS protection from failures and expected behavior, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#).

Process Failure

This section describes specific considerations for process failure in an Oracle User Messaging Service high availability configuration:

- Node Manager attempts to restart the managed servers locally if a crash occurs.
- If whole server migration is configured for the Oracle WebLogic Server managed server to which Oracle UMS components are deployed, and the restart count threshold is exceeded, Oracle WebLogic Server infrastructure attempts to perform a server migration of the managed server to another node in the cluster. After the server migration completes successfully, at the startup time, UMS Server and Drivers go through the startup cycle as previously described, including driver registration. Driver registration is an independent operation and does not have any affect on other available instances. See [Chapter 3, "High Availability for WebLogic Server"](#) for more information about server migration.
- At restart (on the same or different node), the UMS JMS server in the managed server starts producing and consuming messages from its JMS store.
- If a managed server running UMS Server and Drivers is restarted or migrated, the failover is transparent to the connected UMS Clients. The failover is transparent because UMS components are stateless. Once the server's restart is finished, the Web server starts routing requests to it for Web service clients. Similarly, EJB clients become aware of the server availability and start routing requests to it. This is made possible by Oracle WebLogic clustering infrastructure.

Node Failure

For information about Oracle UMS node failure, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#)

Database Failure

For information about Oracle UMS database failure, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#)

Protection From External Messaging Gateway Failures

Before attempting a message delivery, UMS first persists the message to the database. If an external Messaging Gateway becomes unavailable, the corresponding UMS driver periodically attempts to reconnect to the gateway and deliver any undelivered messages persisted in the database. Alternatively, if the messages are not delivered, administrators can manually resend the messages using the UMS server's Message Status Page in Oracle Fusion Middleware Enterprise Manager.

5.9.2.2 Oracle User Messaging Service Cluster-Wide Configuration Changes

UMS configuration is file-based, standard Java EE deployment descriptors and JMX configuration Mbeans, using a standard JMX framework. Changes are propagated using standard Oracle WebLogic Server Mbean server mechanisms. There are no cluster-wide configuration capabilities. As a result, configuration changes must be repeated on every member of a cluster

UMS uses standard Java EE artifacts, configured as part of Oracle WebLogic's domain in which Oracle UMS is installed. Oracle WebLogic Server clusters provide automatic configuration synchronization for artifacts, such as data sources, across Oracle WebLogic Server domain. At the same time, Oracle WebLogic Server cluster controls synchronization of deployments and libraries used by different Oracle UMS components.

5.10 Oracle JCA Adapters and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle JCA Adapters for high availability.


5.10.1 Oracle JCA Adapters Single-Instance Characteristics

Oracle JCA Adapters are JCA binding components that allow the Service Infrastructure to communicate to endpoints using different protocols. Oracle JCA Adapters are deployed as a JCA resource (RAR) and are not part of the Oracle SOA Service Infrastructure.

The run-time component of Oracle JCA Adapters is the J2CA 1.5 resource adapter for the specific back-end application. Oracle JCA Adapters are deployed in J2CA container of Oracle WebLogic Server. Oracle Fusion Middleware integrates with these J2CA 1.5 adapters through the JCA Binding Component, which converts Web service messages into J2CA interactions and back. Adapters can, therefore, fully leverage the scalability and high availability of the underlying Oracle WebLogic Server platform.

Figure 5–30 illustrates the services and dependencies of Oracle JCA Adapters single-instance architecture.

Figure 5–30 Oracle JCA Adapters Single Instance Architecture



5.10.1.1 Oracle JCA Adapters Component Lifecycle

The lifecycle of an adapter is controlled by the Oracle SOA Service Infrastructure. There are three high-level steps in the initialization of an Oracle JCA Adapter:

- **Deployment:** Oracle JCA Adapters are deployed during installation as J2CA 1.5 resource adapters (RAR files) within the same Oracle WebLogic Server container as Oracle Fusion Middleware. The physical deployment of adapters involves using the RAR file and registering the adapters as connectors with the underlying Oracle WebLogic Server, or the middle tier platform. The RAR file contains the `ra.xml` file, which contains declarative information about the contract between Oracle WebLogic Server and the resource adapter. In addition, the RAR file contains `weblogic-ra.xml` which is the deployment descriptor file containing deployment-specific information about the resource adapter.
- **Loading:** Loading refers to the process of creating in-memory resources and creating connections to the configured end-point. Although Oracle JCA Adapters are physically deployed as J2CA 1.5 resource adapters, their logical deployment involves creating the Connection Factory entries for the J2CA 1.5 resource adapter by editing the `weblogic-ra.xml` file. This file is Oracle WebLogic Server-specific deployment descriptor for a resource adapter. It contains configurations for deploying resource adapters to Oracle WebLogic Server, which includes the back-end application connection information as specified in the deployment descriptor of the resource adapter, Java Naming and Directory Interface (JNDI) name to be used, connection pooling parameters, resource principal mapping mechanism, and configurations.
- **Activation:** Activation refers to initiation of a JCA binding component (Service and Reference) within a Composite. Listeners are started for the endpoint referenced by the Adapter configuration within the Composite.

Property Changes During Oracle JCA Adapters Runtime

Oracle JCA Adapters expose the underlying back-end operation-specific properties as header properties and allow the manipulation of these elements within a business process. The underlying properties are as follows:

- **interactionspec or activationspec Properties** - These properties require the adapter endpoint to be recycled (re-activated).
- **Endpoint Properties** - Changes to these properties are notified to the adapter without requiring the endpoint to be restarted.

5.10.1.2 Oracle JCA Adapters Reliability and Transactional Behavior

Oracle JCA Adapters support global transactions based on the JCA 1.5 XA contracts that leverage the underlying application server transaction manager. Adapters supporting XA transactions include Oracle Adapters for Oracle Applications, database, Advanced Queuing, JMS and MQSeries. Non-transactional adapters include Oracle File Adapter and Oracle FTP Adapter.

Inbound Transactions

For a synchronous process, the global transaction initiated by the adapter spans message delivery and composite execution.

For an asynchronous service entry point, a transactional adapter initiates a global JTA transaction before sending an inbound message to a composite. When control returns to the adapter, it commits the JTA transaction, executing the following set of actions as an atomic unit of work:

1. Commit the removal of the message from the inbound adapter endpoint, for example, table and queue.
2. Commit the execution of the composite instance.

If anything fails during this process, both of these actions are rolled back based on XA guarantees.

Outbound Transactions

For transactional adapters, outbound JCA interactions (the invoke activities) are scoped with the global JTA transaction of the Composite instance. This means that all composite activities, including Oracle JCA adapter invocations, are part of a global transaction, and as such all activities are either committed or rolled back if an error occurs. Therefore, one can guarantee exactly-once message delivery when both inbound and outbound adapters are transactional and the connection factories have been configured to support XA global transactions.

Nontransactional

The Oracle File Adapter picks up a file from an inbound directory, processes the file, and sends the processed file to an output directory. However, during this process if a failover occurs in the Oracle RAC backend or in an SOA managed server, then the file is processed twice because of the nontransactional nature of Oracle File Adapter. As a result, there can be duplicate files in the output directory.

5.10.1.3 Oracle JCA Adapters - Rejected Message Handling

The messages that error out before being posted to the Oracle SOA Service Infrastructure mesh are referred to as rejected messages. For example, Oracle File Adapter picks a file having data in CSV format and tries to translate it to the XML format using NXSD. Now, if there is any errors in the translation, the message is rejected and is not posted to the target composite.

All rejected messages are stored in the database with payload. There are two destinations for rejected messages:

- Adapter Rejection Table - Messages are stored in the Adapter Rejection table when there is a binding error related to a corrupt message.
In case the database hosting the Rejected Message Table is offline, rejected messages are temporarily stored in the local file system (as a backup) and eventually loaded back into the database when it is detected as being online again.
- Composite Instance Failure Table - This table is part of instance tracking, populated by each component detecting a failure.

In the case of errors being thrown by the SOA Service Infrastructure layer, the Adapter framework behavior depends on whether or not the error is marked as retrievable. If the error is not retrievable, the message is treated as a rejected message and persisted to the Rejected Message Table.

Configuring Rejection Handlers

Rejection handlers are defined and configured by using fault policies as described in the following steps:

1. Define a fault policy for the rejected messages in the `fault-policies.xml` file, which is stored along with `composite.xml` in the Oracle JDeveloper project directory.

2. Associate the fault policy with a service endpoint of the composite in the `fault-bindings.xml` file.
3. Copy the `fault-policies.xml` file and the `fault-bindings.xml` file to Oracle SOA Composite project directory.
4. Deploy the Oracle SOA Composite project.

Retry frequency and retry interval are both configurable. Service engines can mark an invocation as retrievable. In a clustered environment, if rollback occurs, it may be picked up by another instance. If another node picks up the message, it generates a unique ID. The number of retries are configured within the composite, under the `binding.jca` element. For more information, see the *Oracle Fusion Middleware User's Guide for Technology Adapters*.

System-Defined Rejected Message Handlers

The following system-defined error handlers can be configured in fault policies:

- **Web Service Handler** - A predefined WSDL interface must be implemented by the target service
- **Custom Java Handler** - A predefined Java interface must be implemented by the target class.
- **JMS Queue** - The rejected message is enqueued to a queue as a JMS message with the appropriate context and payload.
- **File** - The rejected message is stored in the file system with the proper context and payload.

Payload Persistence

For resubmitting rejected messages, payload persistence is imperative. Payloads are stored in the database and be viewed through Oracle Enterprise Manager Fusion Middleware Control. Error handlers invoked during automatic error handling also receive the payload during their invocation. Error payload persistence in the database is enabled by default.

5.10.2 Oracle JCA Adapters High Availability Architecture and Failover Considerations

This section describes Oracle JCA Adapters high availability considerations for configuring Oracle JCA Adapters to run on an Oracle SOA cluster.

5.10.2.1 Oracle JCA Adapters High Availability Error Handling

This section describes Oracle JCA High Availability error handling.

Connection Errors

Oracle JCA Adapters and Oracle Adapters for Oracle Applications handle connection errors for the following interactions:

- **Outbound Interaction:** Oracle JCA Adapters and Oracle Adapter for Oracle Applications raise the `oracle.tip.adapter.api.exception.PCRetriableResourceException` exception for transient connection errors, which are recoverable connection errors. For example, a database listener may not have started, resulting in connection errors. You can define the maximum number of attempts to reconnect made using the `fault-policy.xml` file. The parameters for attempts to reconnect can be specified using this file. After the configured number of retries is reached, the `fabricInvocationException` exception is thrown.

Fault handling and retry properties are defined as part of the binding level configuration. When the binding level retry has expired, the fault is handled based on policies specified within the `fault-policy.xml`. For more information on binding properties, see the *Oracle Fusion Middleware User's Guide for Technology Adapters*.

- **Inbound Interaction:** Oracle JCA Adapters, Oracle Adapter for Oracle Applications, and legacy adapters support a poll model for connecting to the back-end application for receiving events. In case of unrecoverable connection failures, the adapters recycle old connections, and send out alerts or notifications to the Oracle SOA Service Infrastructure. The inbound interaction connection errors are written to a log, and can be viewed through Oracle WebLogic Server Administration Console.

5.10.2.2 Oracle File and FTP Adapters High Availability

The Oracle File and FTP Adapters can be configured for high availability using an active-active topology. A cluster-based high availability configuration is supported for both inbound and outbound operations. File and FTP Adapters do not support XA and hence can only guarantee at-least-once delivery of messages. Thus, it is possible to have duplicate messages after recovery from a server crash.

Prerequisites for High Availability

The following list describes prerequisites for Oracle File and FTP Adapters high availability:

- In a clustered configuration, inbound adapters across managed servers must point to the same physical directory, a directory on a shared drive.
- Connection-factories must specify the same shared folder as the control directory, and their names must match. For example, if the deployment descriptor for one connection-factory has `/shared/control_dir` as the value for `controlDir`, the other deployment descriptor must also have the same value.

The Oracle File and FTP Adapters must ensure that only one node processes a particular file in a distributed topology. You can use the database table as a coordinator to ensure that Oracle File and FTP Adapters are highly available for inbound operations. See the *Oracle Fusion Middleware User's Guide for Technology Adapters* for details on configuring a database table as a coordinator.

The Oracle File and FTP Adapters must ensure that if multiple references write to the same directory, then these do not overwrite each other. The following locking capabilities can be used to make Oracle File and FTP Adapters highly available for outbound operations:

- Database mutex
- User-defined mutex

See the *Oracle Fusion Middleware User's Guide for Technology Adapters* for details on configuring a database mutex.

Oracle File and FTP Adapters High Availability Configuration

A cluster-based high availability configuration is supported for both inbound and outbound operations. However, consider the following:

- **Inbound Operations** - Oracle File and FTP Adapters must ensure that only one node processes a particular file in a distributed topology.
- **Outbound Operations** - Oracle File and FTP Adapters must ensure that if multiple references write to the same directory, these do not overwrite each other.

Database-based mutexes are used as coordinators to ensure these behaviors in a clustered topology.

Configuring a Database Mutex

Configure a database table as a coordinator by modifying Oracle File and FTP Adapter deployment descriptor for the connection-instance corresponding to `eis/HFileAdapter` from Oracle WebLogic Server Administration Console:

1. Click **FileAdapter** under **Summary of Deployments** on Oracle WebLogic Server Administration Console.
2. Click the **Outbound Connection Pools** tab, and expand **javax.resource.cci.ConnectionFactory** to see the configured connection factories.
3. Click **eis/HFileAdapter**. The **Outbound Connection Properties** for the connection factory corresponding to high availability is displayed.
4. Update the connection factory properties.

Update the Adapter configuration to use the connection factory as shown in the following example:

```
<adapter-config name="FlatStructureOut" adapter="File Adapter"
xmlns="http://platform.integration.oracle/blocks/adapter/fw/metadata">
    <connection-factory location="eis/HFileAdapter" adapterRef=""/>
    <endpoint-interaction portType="Write_ptt" operation="Write">
        <interaction-spec
className="oracle.tip.adapter.file.outbound.FileInteractionSpec">
            <property../>
            <property../>
        </interaction-spec>
    </endpoint-interaction>
</adapter-config>
```

Note: The location attribute is set to `eis/HFileAdapter` for the connection factory.

The new parameters in connection factory for Oracle File and FTP Adapters are as follows:

- **controlDir** - Set this parameter to the directory structure where you want the control files to be stored. Set it to a shared location if multiple Oracle WebLogic Server instances are running in a cluster.
- **inboundDataSource** - Set this parameter to `jdbc/SOADDataSource`. This is the data source where the schemas corresponding to high availability are pre-created. The pre-created schemas are located under `MW_HOME/AS11gR1SOA/rcu/integration/soainfra/sql/adapter/createschema_adapter_oracle.sql`. To create the schemas elsewhere, use this script. Set the `inboundDataSource` property accordingly if you choose a different schema.
- **outboundDataSource** - Set this parameter to `jdbc/SOADDataSource`. This is the data source where the schemas corresponding to high availability are pre-created. The pre-created schemas can be found under `MW_HOME/AS11gR1SOA/rcu/integration/soainfra/sql/adapter/createschema_adapter_oracle.sql`. To create the schemas elsewhere, use this script. Set the `outboundDataSource` property if you do create the schemas elsewhere.

- **outboundLockTypeForWrite** - Set this parameter to `oracle` if you are using Oracle database. By default Oracle File and FTP Adapters use an in-memory mutex to lock outbound write operations. Choose one of the following values for synchronizing write operations:
 - **memory** - Oracle File and FTP Adapters use an in-memory mutex to synchronize access to the file system.
 - **oracle** - The adapter uses Oracle database sequence.
 - **db** - The adapter uses a pre-created database table (`FILEADAPTER_MUTEX`) as the locking mechanism. Use this option only if you are using a schema other than the Oracle database schema.
 - **user-defined** - The adapter uses a user-defined mutex. In order to configure the user-defined mutex, implement the mutex interface: `oracle.tip.adapter.file.Mutex` and then configure a new binding-property with the name `oracle.tip.adapter.file.mutex` and value as the fully qualified class name for the mutex for the outbound reference.

Note: For large payloads, increase transaction timeout for the SOA DataSource by adding the following:

```
<xa-set-transaction-timeout>true</xa-set-transaction-timeout>
```

```
<xa-transaction-timeout>1000</xa-transaction-timeout>
```

Increase global transaction timeouts if a database is used as the coordinator.

5.10.2.3 Oracle Database Adapters High Availability

The Oracle Database Adapter supports high availability in an active-active setup. In an active-active setup, distributed polling techniques can be used for inbound Database Adapters to ensure that the same data is not retrieved more than once. For more information on best practices for distributed polling, see the "Oracle JCA Adapter for Database" chapter in *Oracle Fusion Middleware User's Guide for Technology Adapters*. Similar to other adapters, an Oracle Database Adapter can also be configured for singleton behavior within an active-passive setup. This enables a high performance multi-threaded inbound Oracle Database Adapter instance running in an active-passive setup, to follow a fan out pattern and invoke multiple composite instances across a cluster. The Oracle Database Adapter also supports the high availability feature when there is a database failure or restart. The DB adapter picks up again without any message loss.

For information about how an inbound rejected message is handled by using fault policy, see the *Oracle Fusion Middleware User's Guide for Technology Adapters*.

Surviving Database Restart

The Oracle Database Adapter can survive a database down scenario and pick up again without any message loss. For avoiding data loss, the data source needs to be XA enabled and configured for Oracle RAC (multi data source). For information on configuring a datasource for high availability, see [Appendix B, "Recommended Multi Data Sources."](#)

5.10.2.4 Oracle JMS Adapters High Availability

The Oracle JMS Adapter supports multiple providers including WLS JMS, MQ Series, and Oracle AQ. For each provider, the high availability aspects may vary and depend on the underlying messaging infrastructure capabilities. For example, Oracle AQ requires configuring the appropriate data source for XA and Oracle RAC. See [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC"](#) to configure a GridLink data source for high availability, or [Section 4.1.3, "Using Multi Data Sources with Oracle RAC"](#) to configure a multi data source

Transaction control is very relevant for consistency and reliability purposes when using the Oracle JMS Adapter. The Oracle JMS Adapter supports global transactions based on the JCA 1.5 XA contracts that leverage the underlying application server transaction manager. In context of the JMS Adapter, a transaction enables an application to coordinate a group of messages for production and consumption, treating messages sent or received as a single unit. When an application commits a transaction, all messages it received within the transaction are removed by the JMS provider. The messages it sent within the transaction are delivered as one unit to all JMS consumers. If the application rolls back the transaction, then the messages it received within the transaction are returned to the messaging system and the messages it sent are discarded.

When using the WLS JMS provider, the Adapter supports XA transactions by using an XA-enabled connection factory (`weblogic.jms.XAConnectionFactory`, which is the default). For this provider, the high availability behavior is tied to the configuration of the queues and topics that the adapter uses and to the JMS Infrastructure that Oracle WebLogic Server provides. There are two main aspects that affect the behavior of the JMS Adapter when using Oracle WLS JMS as the provider:

- Message redelivery
- Use of WebLogic Server distributed destinations

The next two sections describe these topics. For information on configuring high availability for JMS Adapter, see *Enabling High Availability for Oracle JMS Adapter in the Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

5.10.2.4.1 Message Redelivery Message redelivery plays an important role in a highly available configuration. In WLS JMS, you can delay the redelivery of messages when a temporary, external condition prevents an application from properly handling a message. This enables an application to temporarily inhibit the receipt of "poison" messages that it cannot currently handle. When a message is rolled back or recovered, the redelivery delay is the amount of time a message is put aside before an attempt is made to redeliver the message. If JMS immediately redelivers the message, the error condition may not be resolved and the application may still not be able to handle the message. However, if an application is configured for a redelivery delay, then when it rolls back or recovers a message, the message is set aside until the redelivery delay has passed, at which point the messages are made available for redelivery. All messages consumed and subsequently rolled back or recovered by a session receive the redelivery delay for that session at the time of rollback or recovery. Messages consumed by multiple sessions as part of a single user transaction will receive different redelivery delays as a function of the session that consumed the individual messages. Messages that are left unacknowledged or uncommitted by a client, either intentionally or as a result of a failure, are not assigned a redelivery delay.

A session in JMS context inherits the redelivery delay from its connection factory when the session is created. The `RedeliveryDelay` attribute of a connection factory is configured using the Administration Console (Configuration->Default Delivery

screen). WebLogic JMS defines two default connection factories, which you can look up using the following JNDI names:

- `weblogic.jms.ConnectionFactory`
- `weblogic.jms.XAConnectionFactory`

The JMS Adapter uses `weblogic.jms.XAConnectionFactory` by default when configured for the WLS JMS Provider, therefore using the same redelivery settings as the `XAConnectionFactory`. For details on configuring connection factories for different adapters, refer to "Oracle JMS Adapters Use Cases" in *Oracle Fusion Middleware User's Guide for Technology Adapters*.

5.10.2.4.2 Use of WebLogic Server Distributed Destinations Previous releases of Oracle JMS Adapter could not guarantee exactly-once message delivery when using some types of Oracle WebLogic Uniform Distributed Destinations. Oracle Fusion Middleware 11gR1 includes several enhancements in the WLS JMS layer that have been consumed by the JMS Adapter to better support Uniform Distributed Destinations. Shared ClientID, Partitioned Distributed Topics, and Shared Subscriptions are the main enhancements. These enhancements facilitate scalability, high availability, and consistent message delivery in different scenarios:

- **Shared ClientID:** Enables durable subscribers to consume from distributed destinations using the same ClientID. The JMS standard requires that topic applications that consume from durable subscriptions must specify a ClientID and that the same ClientID should not be used by more than one running connection at any given time. This feature enables "overriding" this limitation.
- **Partitioned Distributed Topics:** A partitioned distributed topic is a blend of uniform distributed queue and uniform distributed topic behavior. It enables messages to be load balanced among distributed topic members as in a distributed queue, and then allows for multiple receivers of the message (subscribers) on the single distributed topic member. One message published to a partitioned distributed topic land on only one of its members, not all of them as in previous implementations.
- **Shared Subscription:** Multiple subscribers can optionally share the same subscription, durable or non-durable, by setting the connection's `SubscriptionSharingPolicy` to `SHARABLE`. When you use this feature with a partitioned topic and an `UNRESTRICTED` connection client-id policy, it guarantees that, among all subscribers who share the same distributed subscription, only one subscriber receives a particular message sent to the topic. All subscribers that share the same distributed subscription will together process each of the messages that are published to the topic.

Also, the WLS JMS infrastructure provides notification helpers in 11gR1 that allow for notifications of member runtime availabilities of specific distributed destinations.

From the JMS Adapter perspective, these JMS features have the following impact when using distributed queues:

Inbound - Distributed Queues

The JMS Adapter automatically registers a notification listener with WLS JMS on endpoint activation. If an "unavailable" notification is received, the corresponding poller threads are stopped and the necessary cleanup is done. Similarly, if an "available" notification is received, the poller thread is started and the adapter begins consumption from the new member.

Note: Oracle BAM's Enterprise Message Source (EMS) cannot be configured with distributed queues fed by producers load balancing to the UDD. UDD's behavior pins each consumer to one single physical destination, and this could cause missed messages on the EMS side. As an alternative, configure EMS with the standard destination cluster wide JNDI name, but force producers to pin to one single physical JMS queue (by using that particular server host rather than using a cluster-wide JNDI name).

Outbound - Distributed Queues

There is no change in the JMS Adapter behavior when producing a message to a distributed queue. Available/Unavailable/Failure notification has no impact on the working of outbound adapter reference since messages are produced by creating a MessageProducer for the distributed destination and not for a specific member.

These JMS features have the following impact when using distributed topics:

Inbound - Distributed Topics

The JMS Adapter automatically registers a notification listener with WLS JMS on endpoint activation. If an "unavailable" notification is received, the corresponding poller threads are stopped and the necessary cleanup is done. Similarly, if an "available" notification is received, the poller thread is started and the adapter begins consumption from the new member. The durable subscription is maintained in a similar fashion as in a non-distributed topic scenario. The adapter will unregister the notification listener on endpoint deactivation. Any message arriving at a distributed topic will be processed based on the various settings used and the type of distributed destination in use, as detailed below:

- One copy of message per application: Each message is processed exactly once (no duplicate processing). The use of partitioned distributed topics and configuring the JMS adapter to use unrestricted clientid and shared subscription policy (these are the default settings) is recommended. In this scenario, the client id and subscription name are the same for everyone, and each adapter instance creates subscriptions on every member. The name is unique and immutable across server restarts. The specific clientid configuration in the connection instance (weblogic-ra.xml for a local cluster) for the adapter is required:

```
<property>
<name>FactoryProperties</name>
  <value>ClientID=SOAClient1;</value>
</property>
```

When using replicated distributed topics, it is also possible to force one copy of message per application. For this, it will be required to configure JMS adapter to use unrestricted clientid and shared subscription policy (the default settings), but in addition, it will require specifying the following message selector "(NOT JMS_WL_DDForwarded)" when defining an activation spec. Note that this type of configuration is highly inefficient and not recommended, because it distributes messages to all consumers but then does the rejection on the consumer end.

- One message per adapter endpoint instance in a cluster: In this scenario either the clientid or subscription name are unique for each adapter instance. Also, the unique part of the name should be immutable across server restarts. This type of configuration is possible using both Replicated and Partitioned Topics.

For Partitioned Distributed Topics, configure the JMS adapter to use unrestricted clientid and a shared subscription policy (that is, set "SubscriptionSharingPolicy" property to SHARABLE), which are the default settings. Also, to ensure subscription name uniqueness, set the TopicMessageDistributionAll property to true (default value is false). This property can be defined by setting the "FactoryProperties" property of the connection instance in weblogic-ra.xml. An example usage (snippet of connection instance from weblogic-ra.xml is:

```
<property>
<name>FactoryProperties</name>
  <value>ClientID=SOAClient2; TopicMessageDistributionAll=true</value>
</property>
```

For Replicated Distributed Topics, configure the JMS adapter to use unrestricted clientid and shared subscription policy as well (that is, set "SubscriptionSharingPolicy" property to SHARABLE), which are the default settings. As with the Partitioned Distributed Topic case, to achieve subscription name uniqueness, the adapter requires that the property TopicMessageDistributionAll be set to true (the default value is false). The property can be defined by setting the "FactoryProperties" property of the connection instance in weblogic-ra.xml. An example usage (snippet of connection instance from weblogic-ra.xml for a local cluster) is shown above.

For Replicated Distributed Topics to guarantee one copy per adapter endpoint, the message selector "NOT JMS_WL_DDForwarded" needs to be used when defining an activation spec, as shown in the following .jca file example:

```
<activation-spec
className="oracle.tip.adapter.jms.inbound.JmsConsumeActivationSpec">
<property name="DestinationName" value="jms/DemoInTopic"/>
<property name="UseMessageListener" value="false"/>
<property name="DurableSubscriber" value="dsub1"/>
<property name="MessageSelector" value="NOT JMS_WL_DDForwarded"/>
<property name="PayloadType" value="TextMessage"/>
</activation-spec>
```

Note, however, that this configuration is highly inefficient and not recommended; it distributes messages to all consumers but does the rejection on the consumer end. For better performance, use Partitioned Distributed Topics.

Outbound - Distributed Topics

There is no change to adapter behavior when producing message to a distributed topic. Available/Unavailable/Failure notification does not affect the working of outbound adapter reference. The message is produced by creating a MessageProducer for the Distributed Destination and not for a specific member.

Note: Oracle BAM does not support durable, distributed topics but it does support non-durable (transient) distributed topics.

5.10.2.5 Oracle JCA Adapters Log File Locations

You can view the logs for Oracle JCA Adapters as follows:

Oracle JCA Adapters and Oracle Adapter for Oracle Applications: For both outbound and inbound interactions, the log files are redirected to the soa-diagnostic.log file. The log files for Oracle SOA Suite that is deployed to the server-soa managed server are located in:

```
MW_HOME/user_projects/domains/domain_  
name/servers/server-soa/logs/soa-diagnostic.log
```

Packaged-application adapters: These adapters do not implement the LogManager interface because it is not part of the J2CA 1.5 standard. Therefore, for OPMN-managed components the log outputs are redirected to

```
ORACLE_INSTANCE\diagnostics\logs\component_type\component_name
```

For outbound interactions, the logs are directed to the same location. On the other hand, for inbound interactions, logs are redirected to soa-diagnostic.log.

Legacy adapters: In addition to the J2CA resource adapter, legacy adapters consists of Oracle Connect, which consists of native adapters for communicating with the mainframe application and data stores. Oracle Connect logs can be viewed using Oracle Studio, which is the mainframe adapter design-time tool and Oracle Connect management tool. Oracle Connect generates various types of logs, such as the daemon log, workspace log, and server process log.


5.11 Oracle Business Activity Monitoring and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for configuring Oracle BAM for high availability.

5.11.1 Oracle Business Activity Monitoring Single-Instance Characteristics

Oracle BAM provides the tools for monitoring business services and processes in the enterprise. It allows correlating of market indicators to the actual business process and to changing business processes quickly or taking corrective actions if the business environment changes. Oracle Business Activity Monitoring (Oracle BAM) provides the necessary tools and runtime services for creating dashboards that display real-time data inflow and define rules to send alerts under specified conditions.

[Figure 5–31](#) illustrates the main services and dependencies that characterize an Oracle BAM instance.

Figure 5–31 Oracle Business Activity Monitoring Single-Instance Architecture

5.11.1.1 Oracle Business Activity Monitoring Component Characteristics

Oracle BAM is made up of the following components:

- **Oracle BAM Server** is a set of runtime components that handle incoming data from different data sources. Oracle BAM components are also used to evaluate conditions for sending alerts to users and triggering the required actions configured for these alerts. The following are the main components in an Oracle BAM Server:
 - **Active Data Cache** is designed and optimized to handle large amounts of data in a real-time solution. To make data readily accessible and deliverable, it maintains real-time views of the data. The data feed to the Active Data Cache is a combination of business data sources, from data warehouse information to transactional feeds and other enterprise sources. The various data streaming technologies integrated with Oracle BAM send this information to the Active Data Cache in a continuous stream as data changes occur.
- **Event Engine:** Event Engine monitors complex data conditions and implements specified rules. Rules can include a series of conditions and actions attached to an event. The Event Engine continuously monitors the information in the Active Data Cache for certain conditions and executes the related actions defined in associated rules.

The Event Engine is responsible for tracking events based on date, time or data changes. The design of the Event Engine uses a satellite concept, in which

there are four different systems (satellites), with which event clauses can be registered and in which they can be tracked.

- **Report Cache** off loads the burden of maintaining the view set snapshot in memory from the Active Data Cache. The Report Cache opens view sets and active view sets in the Active Data Cache for the Report Server in Oracle BAM Web Applications set of components. It then caches the snapshot (in chunks) and the active data before sending it to the Report Server. This allows for random access into the snapshot and recovery from losing internet connectivity. The Report Cache also allows the Report Server to be stateless, and with the Active Data Cache it supports view set sharing.
- **Real Time Data Streaming Clients:** In an Oracle BAM system clients can feed data into Oracle BAM server by using different types of protocols and APIs. The following are the available mechanisms to send data to the Oracle BAM server:
 - **Oracle BAM Adapter** is a JCA-compliant Adapter, which Java EE applications can use to send data to the Oracle BAM Server.
 - **Enterprise message sources** are used by applications to provide direct Java Message Service (JMS) connectivity to the Oracle BAM server by mapping messages directly to Oracle BAM data objects. The Oracle BAM server can read data directly from any JMS based message queue or topic. This option offers guaranteed messaging.

In an Oracle BAM high availability environment, involving Enterprise Message Sources, the Enterprise Message Source property **Start when BAM Server starts** should be set to **Yes**.

Durable Subscriber Name must be set in the Enterprise Message Source to ensure that no messages are lost during the failover.

- **Oracle Data Integrator** is the ETL tool that is used with Oracle BAM to perform rigorous data transformations. The Oracle BAM Server has been implemented as an ODI Technology (for example, DB2, SQL Server are ODI Technologies) and Oracle BAM has ODI Knowledge Modules which let ODI perform all of the operations on the Oracle BAM Server to facilitate reading and writing data in various ways, including Change Data Capture.
- **Web Services API** interacts directly with Oracle BAM data objects from a remote client.
- **Oracle BAM Web Applications** are the Web user interfaces (Java EE Web applications) for viewing Oracle BAM data. Oracle BAM Web Applications also allow for building data models and creating dashboards and alerts. Oracle BAM Web Applications are the browser based interfaces provided for building reports and administrating the users that have access to the BAM User Interface system. As part of the Web applications, a Report Server applies the report definitions to the data sets retrieved from the Active Data Cache for presentation in a browser.
- **ICommand** is a command line interface for administrating the data in the Active Data Cache. It provides a set of commands to export, import, rename, clear, and delete items from Active Data Cache.

The BAM application runs separately from the rest of Oracle SOA Suite components. Oracle BAM Web Applications are deployed together with the BAM Server. Oracle BAM Server uses EJBs and DAOs for implementing its services and persisting information to the database. All of the EJBs are stateless. However, the Oracle BAM Server is a singleton: only one active Oracle BAM server is used for maintaining fed data and for pushing it into the Web Application Reports. Oracle WebLogic Server Migration feature is used to protect Oracle BAM server from failures. Oracle BAM

server uses JMS intensively. However, all Oracle BAM JMS queues are internal and do not require to be configured in a high availability configuration. When failover occurs, the queues are recreated on the new active Oracle BAM instance.

Oracle BAM Web Application Modules can run in full active-active mode connecting to the available Oracle BAM server. The information is retrieved using RMI protocol to access the EJBs that provide the information to remote consumers. Oracle BAM Web Applications are stateless, except for the Reports Server, which maintains a session ID for reports or users accessing the systems. This requires enabling session replication for the WebLogic Managed Server where Oracle BAM Web Applications run.

The processing of work by the EJBs and threads in Oracle BAM is non transactional and does not require Oracle WebLogic Server special transaction logs configuration. At the same time, Oracle BAM uses the database intensively, therefore it is important that Oracle BAM's database access is prepared for failure in the database. This requires configuring GridLink data sources or multi data sources for Oracle BAM data source as [Section 5.15, "Configuring High Availability for Oracle BAM"](#) describes.


External Dependencies

Oracle BAM engine system depends only on the Oracle BAM database (which contains Oracle BAM schemas) for Oracle BAM data and report metadata. The database must be available for Oracle BAM system to start and run properly.

5.11.1.2 Oracle Business Activity Monitoring Startup/Shutdown Lifecycle

As shown in [Figure 5–32](#), when the WebLogic Server is started, the Oracle BAM Server application is initialized. During startup, the Active Data Cache loads all the data from the repository. The Reports Cache is initialized with the data required for the available systems. The Event Engine starts analyzing the data and preparing notifications. Oracle BAM Web Applications are configured at boot time with an Oracle BAM Server's address. When Oracle BAM Web Applications are started, they connect to the Active Data Cache in Oracle BAM server and retrieve the corresponding viewsets through the Reports Server. Once initialization is done, the system is ready to receive data from clients and raise events and update reports dynamically.


Figure 5–32 Oracle Business Activity Monitoring Startup/Shutdown Lifecycle



5.11.1.3 Oracle Business Activity Monitoring Startup and Shutdown of Processes


Oracle BAM Server is an application, `oracle-bam`, that operates independently from the rest of the SOA Service Infrastructure. It runs in the managed server where Oracle BAM is installed. It gets started by default with Oracle WebLogic Managed Server to which Oracle BAM is deployed. Use WebLogic Server Administration Console to verify Oracle BAM Server's status and to start and stop it. You can also use WebLogic Server WLST command line to control the application.

Figure 5–33 Startup and Shutdown of Oracle Business Activity Monitoring using WebLogic Server



Oracle Enterprise Manager Fusion Middleware Control allows multiple operations and configuration of the BAM Server as well as monitoring its status. For details about monitoring and controlling the Oracle BAM engine, *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

Figure 5–34 Startup and Shutdown of Oracle Business Activity Monitoring using Fusion Middleware Control




The BAM Web Applications run collocated by default in the same managed server as BAM Server. For example, the BAM server and BAM Web Applications are part of the same deployment in WebLogic Server Administration Console. They can't be stopped and managed separately from BAM Server by default. It is possible though, as explained in later sections, to manipulate WebLogic Server application targeting to

separate both components. The Oracle WebLogic Administration Console allows starting and stopping Oracle BAM. This can also be done using WLST commands.

Oracle Enterprise Manager Fusion Middleware Control allows multiple operations and configuration of the BAM Web Applications.

Figure 5–35 Oracle Enterprise Manager Fusion Middleware Control




Note: Although Enterprise Manager offers separate start and stop for Oracle BAM Server and Oracle BAM Web Applications, this is in reality a stop operations that affects both components, meaning if Oracle BAM Server is stopped from Enterprise Manager, the corresponding Oracle BAM Web Applications are also stopped, and the reverse is also true. This also applies for start operations.

5.11.1.4 Oracle Business Activity Monitoring Configuration Artifacts

Oracle Enterprise Manager Fusion Middleware Control exposes some of the configuration options for Oracle BAM Server.

Figure 5–36 Configuring Oracle Business Activity Monitoring




The properties exposed by Oracle Enterprise Manager are a mix of the information contained in two different files: `BAMServerConfig.xml` and

BAMCommonConfig.xml. These files are located under the `DOMAIN_HOME/servers/BAM_Server_Name/tmp/_WL_user/oracle-bam-11.1.1/yrhyfp/APP-INF/classes/config` directory (notice that yrhyfp is the random directory generated at installation time for deploying BAM applications). This is the random directory generated at installation time for deploying Oracle BAM applications. The properties in these files are modifiable by using the Mbeans exposed in Oracle Enterprise Manager Fusion Middleware Control. For details on the configuration options for Oracle BAM Servers, refer to the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite* and the *Oracle Fusion Middleware User's Guide for Oracle Business Activity Monitoring*.

Similarly, some properties are exposed for Oracle BAM Web Applications by Oracle Enterprise Manager Fusion Middleware Control as shown in [Figure 5–37](#).

Figure 5–37 Oracle Business Activity Monitoring Configuration Properties



Configuration options at the container level, such as data sources and persistent stores, are maintained as part of Oracle WebLogic Server Domain configuration and are synchronized across a cluster of Oracle WebLogic Servers by Oracle WebLogic Server core infrastructure. See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* for details on configuring Oracle BAM.

5.11.2 Oracle Business Activity Monitoring High Availability Architecture and Failover Considerations


[Figure 5–38](#) describes a two-node Oracle BAM cluster running on two Oracle WebLogic Servers. The Oracle WebLogic Servers are frontended by Oracle HTTP Servers which load balance incoming requests to them. The following are the main characteristics of this configuration:

- Oracle BAM Web Applications run on two clustered WebLogic Server managed servers. The WebLogic Server Cluster synchronizes configuration for common artifacts of WebLogic Server used by Oracle BAM Web Applications, such as data sources, persistent store, and definitions. Oracle BAM Server is targeted by any of the servers where BAM Web applications are running. Only one WebLogic server runs Oracle BAM Server.
- The architecture uses Oracle WebLogic Server Migration feature to protect Oracle WebLogic Server that runs both Oracle BAM Web Applications and Oracle BAM Server (in BAMHOST1) against failures. This provides protection for Oracle BAM Server which is a singleton. The WebLogic Managed Server in which Oracle BAM Server runs is listening on a Virtual IP that gets migrated to another node when

the failover occurs. This is the address that Oracle BAM Web Applications in BAMHOST2 use to connect to an Oracle BAM Server. Plan appropriately to account for the scenario where Oracle BAM Server and two instances of Oracle BAM Web Applications are running on BAMHOST2. For more information on Server Migration features, see [Chapter 3, "High Availability for WebLogic Server."](#)

- Oracle BAM's database is configured with Oracle Real Application Clusters (Oracle RAC) to protect from database failures. Oracle BAM Server performs the appropriate reconnection and operations retries if database instance failure occurs.

Figure 5–38 Oracle BAM High Availability Architecture



5.11.2.1 Oracle Business Activity Monitoring Protection from Failures and Expected Behavior

Oracle BAM Server and Oracle BAM Web applications are protected from all process failures by the WebLogic Server infrastructure. This section also describes expected behavior if component failure occurs.

Process Failure

Oracle BAM Server and BAM Web Applications are protected from all process failures by the WLS infrastructure. This section describes process failure considerations for Oracle BAM.

- If the WLS_BAMx server crashes, Node Manager attempts to restart it locally. It attempts to restart the servers according to the configured restart count threshold.
 - For failures related to the WebLogic Server on which the BAM Server is running, if repeated restarts fail, the WebLogic Server infrastructure attempts to perform a server migration to the other node in the cluster. Ongoing requests from the clients time out during failover. Once the server's restart completes on the other node, the clients should be restarted to continue routing to it. For opened reports running in the Web browser, if a request is delivered though the BAM WebApps running in the same managed server as BAM Server, a **Reconnecting** message appears in the BAM WebApps report. If the request is being delivered through the BAM WebApps running in the other node, no message appears. In this situation, and during failover, you may be viewing stale data (while a failover takes place or while connection from webapps to servers is reestablished). The BAMWebApps Servlet remains available, and reconnects as soon as BAM Server comes up. For new reports or requests, BAM Web Applications inBAMHOST2 reconnect once the VIP is migrated and the managed server is restarted. BAM Web Applications listening on VIP1 become functional once the server migration completes. At this point, the HTTP Server restarts routing HTTP requests to the Managed Server.
 - Failures affecting the WebLogic Server where only the BAM Web Applications are running, do not affect the BAM system. Other BAM Web Applications (running in BAMHOST1) remain available and maintain session information by using the WebLogic Server Session replication cluster. Failover should be transparent.
- The `oracle-bam` application where BAM Server and BAM Web Applications run may be down due to failure in accessing resources, errors in reading the database, or other issues unrelated to the status of the managed server where it is located. Therefore, you should monitor the Oracle SOA Service Infrastructure application and watch for errors caused by the application in the managed server logs, for log file locations, see [Section 5.2.1.6, "Oracle SOA Service Infrastructure Log File Locations"](#).
- Failover may not succeed for an open report if when the report is opened, only one BAM managed server in the cluster was up, and after starting an additional managed server in the cluster, no other operations were performed. One trigger for session state replication is opening a report. Session state replication is not triggered by active data updates.

In an Oracle BAM high availability environment, when the Oracle BAM Active Viewer is running on an Oracle BAM server that is shut down and then restarted, a "viewset not found" error will be written to the server's log. This error does not affect the functionality of the Oracle BAM Active Viewer and can be ignored.

Node Failure

For node failures in SOAHOST2, the behavior in case of a node failure is equivalent to the process failure scenario: Oracle BAM Web Applications in the other node remain available and can serve requests. Sessions are preserved by the session replication framework provided by WebLogic Server and failover to the other node should be transparent.

For node failures in SOAHOST1, server migration is triggered after the available server verifies the time stamp in the database leasing system. While failover occurs, clients are unable to feed data into the system and retry appropriately. Oracle BAM Web Applications connecting to the server attempt to reconnect until the VIP is migrated and the server is restarted.

Database Failure

For information about Oracle BAM database failure, see [Section 5.2.2.1, "Oracle SOA Service Infrastructure Protection from Failures and Expected Behavior"](#)

5.11.2.2 Oracle Business Activity Monitoring Cluster-Wide Configuration Changes

The standard Java EE artifacts that Oracle BAM Server and Oracle BAM Web Application use are configured as part of Oracle WebLogic Domain in which Oracle BAM is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts such as data sources, persistent stores, and JMS modules, across the WebLogic Server domain. At the same time, the WebLogic Server cluster is in charge of synchronizing the deployments, and libraries used by Oracle BAM Web Applications and Oracle BAM Server.

As explained in the single instance section Oracle BAM Server's and Oracle BAM Web Applications configuration are maintained in the `DOMAIN_HOME/servers/BAM_Server_Name/tmp/_WL_user/oracle-bam_11.1.1/yhryfp/APP-INF/classes/config` directory (notice that `yhryfp` is the random directory generated at installation time for deploying BAM applications). The properties in these files can be modified by using the Mbeans exposed in Oracle Enterprise Manager Fusion Middleware Control. The properties exposed through MBeans are specific to each server. The properties exposed through Enterprise Manager-specific screens are cluster-wide and are only modified on one server. All properties, whether applied in Enterprise Manager or in an MBean browser require a restart of Oracle WebLogic Servers where Oracle BAM runs. For details on the configuration options for BAM Server and BAM Web Applications see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* and *Oracle Business Process Management Suite* and the *Oracle Fusion Middleware User's Guide for Oracle Business Activity Monitoring*.

One of the configuration options related to high availability environments is the Application URL which is used to determine the front end host used by the system in a cluster configuration. This option is used to produce the copy shortcut URL for reports and alerts. The other relevant parameter for Oracle BAM high availability configuration is the Server Name in OracleBAM Web configuration screen. This parameter is used by Oracle Web Application to determine Oracle BAM server to which it connects for accessing the Active Data Cache.

For SOA high availability installations frontended by Oracle HTTP Server, monitoring should be done on the Oracle HTTP Server ports of the real backend servers. This is the case when a deployment is using all the components deployed to the SOA Managed Server. A simple HTTP monitor that pings the HTTP/HTTPS port and expects a pre-determined response in turn should suffice. If only a specific SOA component is being used (such as B2B), then a monitor that does a deeper level check

all the way to the Managed server can be considered to validate the health of the component in use. Please check with your load balancer vendor on setting up the HTTP monitors with your load balancer.

5.11.2.3 Considerations for BAM Client Retries

Oracle BAM clients communicating with the BAM Server will retry "retryable" methods if the failure of the method execution was due to a "retryable" exception. A typical "retryable" exception is one that indicates a failure to communicate with the BAM Server (like a network connection exception). Some of the BAM Server's methods are automatically retried by the BAM client when a failure occurs. Additionally the list of exceptions to be retried can be extended. This can be done by adding a comma-separated list of exception class names to be treated as retryable inside the following configuration fields in the clients' configuration files:

```
<BamEjbRetryableExceptions></BamEjbRetryableExceptions>
```

Other configuration parameters for retryable exceptions are:

- `<BamEjbRetryCount>180</BamEjbRetryCount>`: This parameter configures the number of retries to be attempted if getSession or retryable method invocations fail due to a retryable exception. The default is 180.
- `<BamEjbRetryInterval>10000</BamEjbRetryInterval>`: This parameter configures the number of milliseconds to pause between retry attempts. The default is 10000 milliseconds or 10 seconds.

These parameters can be added in the configuration files for the following clients accessing the BAM Server:


- BAM Server, BAM WEB, BAM Data Control: `BamCommonConfig.xml` configuration file
- ICommand: `BAMICCommandConfig.xml` configuration file
- ODI: `BAMODIConfig.xml` configuration file
- BAM Adapter: `BAMAdapterConfig.xml` configuration file

5.12 Oracle Service Bus and High Availability Concepts

Oracle Service Bus is a proven Enterprise Service Bus (ESB) built from the ground up for SOA lifecycle management that provides foundation capabilities for service discovery and intermediation, rapid service provisioning and deployment, and governance.

[Figure 5–39](#) shows the high level functional features of Oracle Service Bus.

Figure 5–39 Oracle Service Bus Functional Features



For more detailed information about the functional capabilities of Oracle Service Bus, see *Oracle Fusion Middleware Concepts and Architecture for Oracle Service Bus*.

5.12.1 Oracle Service Bus Single Node Characteristics

Oracle Service Bus is a stateless Java EE application that runs on Oracle WebLogic Server.


Oracle Service Bus uses the following Java EE technologies:

- Servlets: Oracle Service Bus Web Console, Service Bus resource servlet, WSIL servlet, and HTTP Proxy Services.
- EJBs: Are used for:
 - The Java EJB protocol proxy service. This service, also known as JEJB, is used to expose proxy services as EJBs.
 - Service Bus transport. Also known as SB, this protocol is used for communication between two Oracle Service Bus services.
 - The test console for testing web services.
- JMS: Is used for inbound message processing for the FTP, SFTP, Email, and File transports. It is also used for the Oracle Service Bus reporting framework and for Web Services Reliable Messaging (WS-RM).
- JDBC Data Source: Oracle Service Bus defines a data source for the database tables associated with the JMS Reporting provider.
- JCA Adapters: Oracle Service Bus supports the DB, Oracle Applications, AQ, File, BAM, FTP, SAP, Siebel, PeopleSoft, JD Edwards adapters using JCA transport. The adapter implementation for these adapters is same as the one used in SOA Suite.

Oracle Service Bus also uses the following non-Java EE components:

- **Configuration Framework:** The configuration framework component provides CRUD (Create, Read, Update, Delete) capabilities for Oracle Service Bus resources along with resource and file integrity protection, session management, caching and indexing, referential integrity, and configuration propagation. The Oracle WebLogic Server Administration Server has the master configuration in files and these are propagated to the managed servers using Weblogic Deployment Service.
- **Monitoring and Alert Framework:** The monitoring framework defines statistics based on counters (message, error counts), intervals (min/max/avg. response or execution time) and status (end point up/down). The monitoring framework has the following sub-components:
 - **Collector:** collects statistics at one minute intervals called checkpoints, stores collected statistics to Weblogic Persistent Store, and sends the collected statistics to the Aggregator. Collector runs on each managed server in the cluster.
 - **Aggregator:** aggregates statistics received from each of the managed servers in the cluster. Aggregator is a singleton service that runs on only one managed server in the cluster.
 - **Retriever:** retrieves the statistics that are stored in the memory. This is present only in the managed server, which contains the Aggregator.
 - **Alert Manager:** The Alert Manager fires alerts based on the aggregated statistics. This is present only in the managed server that contains the Aggregator. The Aggregator pushes aggregated statistics to the Alert Manager when there is an aggregation interval (defined for a service or an alert rule) for which rules need evaluation. The Alert Manager is triggered in a separate thread than the Aggregator.
- **Transport SDK:** Oracle Service Bus processes messages independently of how they flow into or out of the system. The Transport SDK provides a layer of abstraction between Oracle Service Bus and components that deal with the flow of data in and out of Oracle Service Bus. This layer of abstraction allows users to develop new transport providers to handle unique transport protocols. A transport provider implements the interfaces of the Transport SDK and provides a bridge between Oracle Service Bus and mechanisms by which messages are sent or received. A transport provider manages the life cycle and runtime behavior of transport endpoints. An endpoint is a resource where messages originate or are targeted.
- **Message Flow Engine:** The Message Flow Engine provides execution of the message flow logic using actions configured by the message flow designer.
- **Split-Join Engine:** The Split-Join Engine provides execution of logic specified in Split-Join, which is used for performing high-performance, in-memory processing of messages in parallel.

[Figure 5–40](#) is a high-level architecture diagram that illustrates Oracle Service Bus and its functional subsystems.

Figure 5–40 Oracle Service Bus Single Instance Architecture

When Oracle Service Bus is deployed in a high availability configuration, the Aggregator and Alert Manager components are singletons that should be deployed only on one host in the cluster.

Customers can write custom client applications using the Oracle Service Bus extensibility framework, such as the Transport SDK and reporting framework.

For more information about the Oracle Service Bus architecture, see the "Oracle Service Bus Architecture" chapter in *Oracle Fusion Middleware Concepts and Architecture for Oracle Service Bus*.

5.12.1.1 Oracle Service Bus Session State

Oracle Service Bus is optimized for short-lived message processing/mediation use cases. Thus, all message context state is kept in memory. For synchronous transports such as HTTP, EJB, and SB, the client must retry if the Oracle Service Bus server fails in the middle of processing the synchronous request. Oracle Service Bus supports request/response pattern for asynchronous transports such as JMS and WS (WS-RM). In these cases correlation information, which maps the request to its corresponding response destination, is kept in memory. If the Oracle Service Bus server fails before the response is sent, the correlation information is lost. Therefore, Oracle Service Bus will fail to send the response after server restart. For more details about JMS transports and WS transports, see the "JMS Transport" chapter and the "WS Transport" chapter in the *Oracle Fusion Middleware Developer's Guide for Oracle Service Bus*.

5.12.1.2 Oracle Service Bus External Dependencies

Oracle Service Bus has the following external dependencies:

- The JMS Reporting Provider uses a database to store the report data.
- UDDI (Oracle Service Registry) is used to store Oracle Service Bus proxy services and to browse services to create Oracle Service Bus business services.
- Oracle Enterprise Repository is used to browse and consume services. The result is generation of Oracle Service Bus business services. Oracle Service Bus proxy services are harvested into Oracle Enterprise Repository.
- Oracle Service Bus implements an Oracle Web Services Manager (OWSM) agent that enforces OWSM security policies attached to Oracle Service Bus proxy and business services.

- Oracle Service Bus depends on JCA Adapter implementations (JCA Adapters for HTTP, FTP, and so on) and the JCA framework from SOA Suite.
- Oracle Service Bus depends on the Coherence infrastructure for Service Result Caching (caching of results from services invoked by Oracle Service Bus).

Oracle Service Bus message flows are short lived in general. Message flow state is always in memory. Oracle Service Bus provides various out of box transports. For example, the HTTP transport supports HTTP 1.1 persistent connections.

Oracle Service Bus supports the following transports: HTTP, JMS, Email, File, FTP, SFTP, JCA - DB, AQ, Oracle Applications, Peoplesoft, Siebel, SAP, JD Edwards, File, FTP, JEJB, EJB (outbound only), Local (within same JVM), MQ, SB (RMI based, between two OSB servers/domains), Tuxedo, Web Services Reliable Messaging (WS-RM), SOA-Direct (RMI based), BPEL 10g (RMI based).

Oracle Service Bus uses JTA for some transport types, including JMS, File, FTP, Email, SFTP, WS-RM, JCA - DB, AQ, SB, JEJB, EJB, SOA-DIRECT, and BPEL 10g.

5.12.1.3 Oracle Service Bus Configuration Artifacts

The main configuration files for Oracle Service Bus are:

- `DOMAIN_HOME/config/config.xml`: has all applications and libraries, JMS system resources, JDBC system resources, work manager, startup/shutdown classes, SAF agent, security configuration.
- `DOMAIN_HOME/config/osb/coherence/osb-coherence-override.xml`: This Oracle Coherence override file specifies Oracle Coherence unicast/multicast listener information. This file is propagated by the Administration Server from the Administration Server domain directory to the domain directory for other managed servers.
- `DOMAIN_HOME/config/osb/coherence/osb-coherence-cache-config.xml`: This Coherence cache configuration file defines the cache used by the Oracle Service Bus service result caching feature.
- `DOMAIN_HOME/config/osb/transport/sftp/known_hosts`: This file is used by the SFTP transport. It contains information about remote SFTP servers. See *Oracle Fusion Middleware Developer's Guide for Oracle Service Bus* for information about creating the `known_hosts` file. Note that virtual host names can be used in the `known_hosts` file.

5.12.1.4 Oracle Service Bus Deployment Artifacts

You can use Eclipse IDE or the Oracle Service Bus web console to export your Oracle Service Bus artifacts (resources), such as proxy service, business service, wsdl, xsd, and xslt resources, to a jar file with a name that you choose. The jar file contains all the Oracle Service Bus resources that were exported. You have two ways to import the configuration jar file (and the Oracle Service Bus resources inside it) into an Oracle Service Bus server:

- WLST/ANT scripts using public Oracle Service Bus APIs
- Oracle Service Bus web console: Using the **Import from Configuration Jar** option

Oracle Service Bus resources are managed by the configuration framework component. The configuration framework stores all its data in `DOMAIN_HOME/osb/config` (`CONFIG_HOME`) on each managed server (and admin server). Each resource is stored in its own file in `CONFIG_HOME/core`. The Oracle Service

Bus configuration framework automatically assigns unique file names to the file used for each resource to ensure that no name collisions occur among these files. The *CONFIG_HOME/core* directory is the master view of all the resources and configuration data. The runtime works off of this data. It changes only when changes made in a session are activated. Each managed server has a copy of the resources in *CONFIG_HOME/core*. Depending on the location of the managed server, this directory could be mounted on a separate system.

5.12.1.5 Oracle Service Bus Startup and Shutdown

You can use Node Manager or the *startWebLogic/stopWebLogic* scripts to start and stop the Oracle Server Bus server. Oracle Enterprise Manager Grid Control supports management and monitoring of Oracle Service Bus services.

Oracle Service Bus is initialized through application listeners. The ALSB Framework Starter Application does the main framework initialization through an application listener. This application is deployed to the Administration Server and the cluster. The prestart for this application listener initializes Logging, ALSBConfigService, Security service, cluster timer service, and ALSBStatisticsManager. XBus Kernel is the next important application deployed that performs the rest of the initialization and start procedures for the Oracle Service Bus components. This application is targeted to the Administration Server and cluster, and it initializes upgraders, resources, split-join, Coherence, and so on.

Death detection capabilities are provided by Oracle WebLogic Server.

5.12.1.6 Oracle Service Bus Log File Locations


Oracle Service Bus operations are logged to the Oracle WebLogic Managed Server where the application is running:

DOMAIN_HOME/servers/WLS_ServerName/logs/WLS_ServerName.log

Oracle Service Bus also supports Debug logging, which can be enabled or disabled by updating *DOMAIN_HOME/alsbdebug.xml* and *DOMAIN_HOME/configfwkdebug.xml*.

5.12.2 Oracle Service Bus High Availability Architecture and Failover Considerations

[Figure 5–41](#) shows a two-node Oracle Service Bus cluster running on two Oracle WebLogic Servers. The Oracle WebLogic Servers are frontended by two Oracle HTTP Servers on web tier hosts, which receive requests from a load balancer in front of them.

Figure 5–41 Oracle Service Bus High Availability Architecture

The main characteristics of the high availability configuration in [Figure 5–41](#) are:

- The Oracle WebLogic Server server migration feature is used to protect some of the components against failure. This implies that each of the WebLogic managed

servers in which Oracle Service Bus runs is listening on a virtual IP that is migrated to another box upon failover. Enterprise Manager is deployed to the Administration Server in [Figure 5-41](#).

As shown in [Figure 5-41](#), the WLS_OSB1 managed server listens on VIP1, and the WLS_OSB2 managed server listens on VIP2. Each managed server uses the other node as a failover resource, therefore, the system is configured in a cross manner. In other words, WLS_OSB1 fails over to OSBHOST2, and WLS_OSB2 fails over to OSBHOST1. The appropriate capacity planning must be done to anticipate the scenario where the two Oracle Service Bus managed servers are running on the same node. For more information on server migration features, see [Section 3.9, "Whole Server Migration."](#)

To resume transactions after a server migration, configure the transaction and JMS stores in a shared storage. In case of failure in one of the server infrastructure instances, other instances can resume transactions and JMS operation by reading the persistent stores from that shared storage.

- The Administration Server runs in Active-Passive mode. Whenever a failure occurs in OSBHOST1, the Administration Server can be restarted in OSBHOST2. Therefore it uses a virtual IP or virtual hostname as listen address. For information about configuring virtual IPs for the Administration Server and configuring the Administration Server for high availability, see [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability."](#)
- An Oracle RAC database is recommended to protect from database failures. In this Oracle Service Bus high availability configuration, an Oracle RAC database is used for the JMS Reporting Provider and Oracle Web Services Manager policies.

Note: To keep the Oracle Service Bus high availability configuration figure above readable, the Oracle SOA Service Infrastructure high availability configuration is not shown. See [Figure 5-6](#) to view the Oracle SOA Service Infrastructure high availability figure.

The ALSB Cluster Singleton Marker Application determines the managed server in the cluster that runs the singleton monitoring Aggregator and Alert Manager. Similarly the Message Reporting Purger MDB, which helps purging reporting data from the database (for the JMS Reporting Provider), is deployed on only one managed server in the cluster. The remaining Oracle Service Bus components are all either targeted to the cluster, or to the cluster and Administration Server. Only the ServiceBus_Console, which provides Web console support for Oracle Service Bus, and the UDDI Manager, which registers resource types for UDDI import with the configuration framework are targeted to the Administration Server only. These components are not defined as Weblogic Singleton services. Their high availability is assured by Weblogic whole server migration.

All Oracle Service Bus services and resources are deployed uniformly to the cluster (to all managed servers in the cluster). The only exceptions to this are:

- Inbound (Proxy Service) endpoints for the poller transports: Email, File, FTP, and SFTP. The poller polls for messages on only one managed server in the cluster. The user configuration determines the managed server that polls.
- Inbound (Proxy Service) endpoints for JMS transport where the JMS transport is listening to a Weblogic JMS Topic destination (using compatibility Topic Message Distribution Mode) and is configured to process the message from the Topic only once (in the cluster) on the managed server configured by the user.

The basic highly available topology is a Weblogic Server homogenous cluster (except for the Oracle Service Bus singletons services mentioned above) with one Administration Server and two managed servers running on different systems.

Only one cluster of Oracle Service Bus is supported in a Weblogic domain. Local data should be stored on shared storage, such as a SAN storage system, multipointed disk, or NAS storage.

The local data (accessed as local files, private to a managed server or Administration Server) includes:

- System files like Weblogic configuration files and server logs. Optionally, you can store WLS JMS data in an Oracle RAC database. JMS is used both internally in Service Bus and as transports to external and proxy services.
- Oracle Service Bus data such as configuration files and logs
- User-defined Oracle Service Bus configuration data
- User files, such as data files, that are read or written by a proxy service with a File/FTP transport
- The JMS persistent store, which stores alert logs and aggregated performance metrics

The Oracle RAC database is used for high availability of the reporting provider and also for the leasing datasource used for server migration.

Optionally, a web server farm may be used as a front end to an Oracle Service Bus cluster. A hardware load balancer can load balance web servers or application servers directly. If web servers are used, WebLogic plugins should be configured to load balance HTTP traffic among the Oracle Service Bus managed servers.

Email, FTP, or NFS servers can exist in the network, and third party JMS servers or server clusters can also exist in the network.

5.12.2.1 Oracle Service Bus Protection from Failures and Expected Behavior

This section describes how an Oracle Service Bus high availability cluster deployment protects components from failure. This section also describes expected behavior if component failure occurs.

Oracle Service Bus is protected from all process failures by the WebLogic Server infrastructure.

5.12.2.1.1 WebLogic Server Failure Oracle Service Bus does not maintain any state, nor does it support the concept of user sessions. Therefore, Oracle Service Bus does not implement session state replication and failover. For synchronous inbound transports such as HTTP, SB, and EJB, if the managed server processing the request goes down in the middle of request processing, the client receives a connection exception and needs to retry.

Node Manager should be configured to perform automatic WebLogic server migration when a managed server fails. For details on WebLogic server migration, see [Section 3.9, "Whole Server Migration."](#)

Use WebLogic server migration for JMS failover.

There is no automatic failover for Oracle Service Bus singleton components such as the Aggregator, Alert Manager, and Reporting Message Purger. Use WebLogic server migration to perform the failover of singleton components.

5.12.2.1.2 Node Failure The external load balancer can be used to directly load balance HTTP requests to Oracle Service Bus HTTP Proxy Services on the managed servers in the cluster or to a Web Server/Oracle HTTP Server (OHS) cluster front-ending the Oracle Service Bus cluster. There are no sticky session routing requirements.

Node Manager can be configured to failover failed Oracle Service Bus nodes.

If an Oracle Service Bus managed server fails in the middle of processing HTTP requests, clients need to submit those requests again. For JEJB and SB transports, a client gets a connection error when the server goes down. For all poller transports, the integrity of message processing is handled by transactional semantics as the message is dequeued using XA connection factory. Poller transports provide "at least once" semantics. However, after whole server migration (of the failed node), the poller needs to have access to the resources it is polling. JTA transactions will also preserve transactional integrity when using JMS transport. JTA TLogs need to be recovered to recover in-flight JTA transactions.

Oracle Service Bus guarantee "Exactly Once" message delivery semantics. This behavior is controlled by the QoS property in the \$outbound context variable configured in the proxy service. For more detail on the different QoS settings, see *Oracle Fusion Middleware Administrator's Guide for Oracle Service Bus*.

When a managed server starts up, any in-flight global transactions are recovered. For example, recovery of global transactions may be required when the proxy service is bridging between two JMS providers using XA connection factories, and QoS of exactly once. Another example is the JMS reporting provider. With the default JMS Reporting Provider, the report message is first written to the JMS queue. The MDB dequeues the report message from the JMS queue and then writes it to the database. The (database) data source is configured with transaction semantics of LLR (Last Logging Resource optimization). In this case, the database must be running during recovery, otherwise the server will not start. The Tlog is used even when LLR is in effect. However, this happens only on a per transaction basis. Therefore the Tlog must still be made highly available even when all transactions are LLR transactions. Transaction Manager persists checkpoint TLog records that are unrelated to specific transactions, but that are still required to provide full transactional safety.

There are two cases where failure of the managed server causes certain state to be lost, which prevents correct processing of messages:

- JMS request/response business service: Maintains in-memory table mapping correlation information. This is lost when the managed server goes down. Thus, the response from the JMS service cannot be sent back to the original client of the proxy service, which routed to the JMS request/response business service.
- WS (Reliable Messaging) business service: Similar to the JMS request/response business service, the WS business service keeps an in-memory table mapping correlation information. This is lost when the managed server goes down, and the response cannot be handled.

5.12.2.1.3 Database Failure Beyond the implications for database leasing (used for server migration) database failures are relevant for Oracle Service Bus only in the context of the JMS reporting provider functionality. When there is report action in the proxy service, the report data is enqueued into the reporting JMS queue. The reporting MDB dequeues the message from the JMS queue and inserts the data into the database using the data source configured with the Logging Last Resource (LLR) global transaction protocol. The reporting JMS queue is configured with a redelivery limit of 2 and an error queue. When a database failure occurs, report messages are moved to the error queue after the redelivery limit is reached. When the database is running again, you can move these report messages back to the JMS reporting queue so that

they are inserted into the database. If the database fails in the middle of LLR transaction processing, transaction recovery is performed as described in the chapter on Logging Last Resource Transaction Optimization in *Oracle Fusion Middleware Programming JTA for Oracle WebLogic Server*.

5.12.2.2 Oracle Service Bus Cluster-Wide Deployment

Oracle Service Bus resources are updated/deployed onto the WebLogic Administration Server. The Administration Server uses the WebLogic deployment service to propagate the Oracle Service Bus resources to all the managed servers in the cluster.

5.12.2.3 Online Redeployment of Oracle Service Bus in a Cluster

Oracle Service Bus resources do not have the notion of versions. When a resource is updated, it is as if a new version is deployed and the older version is overwritten. The configuration framework makes changes to resources in a session, which is a copy of the core (that all managed servers are currently using) when the session is created. When the session is committed, the core is updated with the changes. The changes are then propagated in a transactional manner to the managed servers in the cluster. In-flight requests (requests already executing the Proxy Service Message flow) continue to use the old copy of the resources when the update is in progress. Requests after the changes are added to the core use the updated resources.

5.12.2.4 Oracle Service Bus Cluster-Wide Configuration Changes

Oracle Service Bus resources are updated at the cluster level. The WebLogic Administration Server propagates these changes to all the managed servers in the cluster. You can make configuration changes using the Oracle Service Bus web console, the WebLogic Server Administration Console, and public MBeans.

The following files are excluded from the automatic propagation to the managed servers in the cluster:

- Files in the `DOMAIN_HOME/config/osb/coherence/` directory
- `DOMAIN_HOME/config/osb/transport/sftp/known_hosts`

Each managed server must have access to the `DOMAIN_HOME` to access files in these directories.

Note: For detailed instructions on installing and configuring Oracle Service Bus in a high availability configuration, see [Section 5.14, "Configuring High Availability for Oracle Service Bus, with SOA Service Infrastructure and Component Service Engines."](#)

5.13 Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines

The procedures described in this section include setting up the Service engines contained in the Oracle SOA Service Infrastructure system, such as Oracle BPEL/Oracle PM, Oracle Mediator, Oracle Human Workflow and Oracle Decision Services, as well as Oracle B2B and Oracle User Messaging Service.

This section includes these topics:

- [Section 5.13.1, "Preparing the Environment: Prerequisite Steps Before Setting up a SOA High Availability Configuration"](#)

- Section 5.13.2, "Installing Oracle Fusion Middleware Home"
- Section 5.13.3, "Enabling VIP1 in SOAHOST1 and VIP2 in SOAHOST2"
- Section 5.13.4, "Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA Domain"
- Section 5.13.5, "Creating boot.properties for the Administration Server on SOAHOST1"
- Section 5.13.6, "Starting and Validating the Administration Server in SOAHOST1"
- Section 5.13.7, "Disabling Host Name Verification for the Administration Server and the WLS_SOAn Managed Servers"
- Section 5.13.8, "Configuring Oracle Coherence for Deploying Composites"
- Section 5.13.9, "Starting the System in SOAHOST1"
- Section 5.13.10, "Propagating the Domain Configuration to SOAHOST2 with pack/unpack Utilities"
- Section 5.13.11, "Extracting XEngine Files in the Second Node"
- Section 5.13.12, "Starting the System in SOAHOST2"
- Section 5.13.13, "Configuring Oracle HTTP Servers for the Administration Server and the WLS_SOAn Managed Servers"
- Section 5.13.14, "Validating Access Through Oracle HTTP Server"
- Section 5.13.15, "Configuring JMS Persistence Store as Shared Across the Servers"
- Section 5.13.16, "Configuring a Default Persistent Store for Transaction Recovery"
- Section 5.13.17, "Setting the Front End HTTP Host and Port"
- Section 5.13.18, "Setting the WLS Cluster Address for Direct Binding/RMI Invocations to Composites"
- Section 5.13.19, "Deploying Applications"
- Section 5.13.20, "Configuring Server Migration for the WLS_SOA Servers"
- Section 5.13.21, "Scaling the Topology"

Note: Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

Note: The architectures and deployment procedures defined in this guide enable simple clustered deployments. The procedures described in these chapters can be used as a building block to enable this and other similar high availability topologies for these Fusion Middleware components. It is also expected that production deployments will use other required procedures, such as associating security policies with a centralized LDAP server. For complete details of secured, multi-tiered architectures, and deployment procedures, please refer to the Enterprise Deployment Guide for the component you are configuring.

Figure 5–42 represents the example architecture that the configuration steps in this section create.


Figure 5–42 Oracle SOA Service Infrastructure High Availability Architecture

Figure 5–42 describes a two-node SOA cluster running on two Oracle WebLogic Servers. The Oracle WebLogic Servers are frontended by Oracle HTTP Servers, which load balance incoming requests. A load balancer front ends the system and distributes incoming requests from clients to the two Oracle HTTP Servers. A separate Oracle WebLogic Server (not shown in the figure) is typically used for custom logic and application deployment. This configuration uses an Oracle RAC database for storing metadata and SOA schemas, and shared storage for transaction and JMS stores. Virtual IP addresses (VIPs) provide manual failover for the Administration Server and for Oracle SOA Servers (for Server Migration). For more details about the components contained in this architecture, see the individual component sections in this chapter.

For information about configuring virtual IPs for the Administration Server and configuring the Administration Server for high availability, see [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

5.13.1 Preparing the Environment: Prerequisite Steps Before Setting up a SOA High Availability Configuration

The following sections provide prerequisite steps before setting up an Oracle SOA Service Infrastructure high availability configuration.

For information about platform-specific commands, see the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

- [Section 5.13.1.1, "Database Prerequisites"](#)
- [Section 5.13.1.2, "VIP and IP Prerequisites"](#)
- [Section 5.13.1.3, "Shared Storage Prerequisites"](#)
- [Section 5.13.1.4, "Installing and Configuring an LDAP Provider"](#)
- [Section 5.13.1.5, "Synchronizing System Clocks"](#)
- [Section 5.13.1.6, "Terminology for Directories and Directory Environment Variables"](#)
- [Section 5.13.1.7, "Installing and Configuring the Database Repository"](#)
- [Section 5.13.1.8, "Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database"](#)
- [Section 5.13.1.9, "Configuring Virtual Server Names and Ports for the Load Balancer"](#)
- [Section 5.13.1.10, "Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2"](#)

5.13.1.1 Database Prerequisites

Oracle SOA Suite requires the presence of a supported database and schemas.

To check if your database is certified or to see all certified databases, refer to the "Certified Databases" section in the Certification Document:

http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

5.13.1.2 VIP and IP Prerequisites

As shown in [Table 5–2](#), you configure the Administration Server and the SOA managed servers to listen on different virtual IPs. This requires the provisioning of the corresponding VIP in the node and related host names in the DNS system in your network. Ensure that the different VIPs are available and are reachable from SOAHOST1, SOAHOST2, and the client systems before running the installation.

Table 5–2 Virtual Hosts

Virtual IP	VIP Maps to...	Description
VIP0	SOAHOST1VHN0	SOAHOST1VHN0 is the virtual host name that is the listen address for the Administration Server and fails over with manual failover of the Administration Server. It is enabled on the node where the Administration Server process is running (SOAHOST1 by default).

Table 5–2 (Cont.) Virtual Hosts

Virtual IP	VIP Maps to...	Description
VIP1	SOAHOST1VHN1	SOAHOST1VHN1 is the virtual host name that maps to the listen address for WLS_SOA1 and fails over with server migration of this managed server. It is enabled on the node where WLS_SOA1 process is running (SOAHOST1 by default).
VIP2	SOAHOST2VHN1	SOAHOST2VHN1 is the virtual host name that maps to the listen address for WLS_SOA2 and fails over with server migration of this managed server. It is enabled on the node where WLS_SOA2 process is running (SOAHOST2 by default).

5.13.1.3 Shared Storage Prerequisites

For proper recovery in case of failure, store both JMS and transaction logs in a location that is accessible to all the nodes that can resume the operations after a failure in a managed server. This requires a shared storage location that can be referenced by multiple nodes. [Table 5–3](#) lists the contents of shared storage.

Table 5–3 Contents of Shared Storage

Server	Type of Data	Vol in Shared Storage	Directory	Files
WLS_SOA1	Tx Logs	VOL1	ORACLE_ BASE/admin/domain _name/soa_cluster_ name/tlogs	Common location (stores decided by WebLogic Server)
WLS_SOA2	Tx Logs	VOL1	ORACLE_ BASE/admin/domain _name/soa_cluster_ name/tlogs	Common location (stores decided by WebLogic Server)
WLS_SOA1	JMS Stores	VOL1	ORACLE_ BASE/admin/domain _name/soa_cluster_ name/jms	Common location but Individual store per server (for example: SOAJMSStore1, UMSJMSStore1)
WLS_SOA2	JMS Stores	VOL1	ORACLE_ BASE/admin/domain _name/soa_cluster_ name/jms	Common location but Individual store per server (for example: SOAJMSStore2, UMSJMSStore2)

The shared storage can be a NAS or SAN device. Specifically for NFS mounted systems, different issues related to file locking and abrupt node failures have been detected. Check the *Oracle Fusion Middleware Release Notes* and with your storage vendor for the main recommended parameters to be used as mount options. The following is an example command based on a NAS device. Your options may be different from the ones specified in this section:

```
SOAHOST1> mount nasfiler:/vol/volX/FMWshared
MW_HOME -t nfs -o
rw,bg,hard,nointr,tcp,vers=3,timeo=300,rsize=32768,wsiz=32768
```

5.13.1.4 Installing and Configuring an LDAP Provider

For production environments, it is a mandatory requirement for Oracle SOA Suite high availability topologies to have an external LDAP policy store. For more

information on the supported policy stores and instructions on configuring LDAP, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

5.13.1.5 Synchronizing System Clocks

Synchronizing system clocks on each of the cluster nodes for high availability SOA deployments is required.

5.13.1.6 Terminology for Directories and Directory Environment Variables

The following list describes the directories and variables used in this chapter:

- **ORACLE_BASE:** This environment variable and related directory path refers to the base directory under which Oracle products are installed.
- **MW_HOME:** This environment variable and related directory path refers to the location where Fusion Middleware (FMW) resides.
- **WL_HOME:** This environment variable and related directory path contains installed files necessary to host a WebLogic Server.
- **ORACLE_HOME:** This environment variable and related directory path refers to the location where Oracle FMW SOA Suite is installed.
- **ORACLE_COMMON_HOME:** This environment variable and related directory path refers to the Oracle home that contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF).

There can be only one Oracle Common home within each Middleware home.

- **DOMAIN Directory:** This directory path refers to the location where the Oracle WebLogic Domain information (configuration artifacts) is stored.
- **ORACLE_INSTANCE:** An Oracle instance contains one or more system components, such as Oracle Web Cache, Oracle HTTP Server, or Oracle Internet Directory. An Oracle instance directory contains updateable files, such as configuration files, log files, and temporary files.

The values used and recommended for consistency for this directories are:

ORACLE_BASE:

`/u01/app/oracle`

MW HOME (Application Tier):

`ORACLE_BASE/product/fmw`

WL_HOME:

`MW_HOME/wlserver_10.3`

ORACLE_HOME:

`MW_HOME/soa`

ORACLE_COMMON_HOME:

`MW_HOME/oracle_common`

Location for JMS file based stores and Tlogs:

`ORACLE_BASE/admin/domain_name/soa_cluster_name/jms`

`ORACLE_BASE/admin/domain_name/soa_cluster_name/tlogs`

Mount point is:

`ORACLE_BASE/admin/domain_name/soa_cluster_name/`

5.13.1.7 Installing and Configuring the Database Repository

This section describes how to install and configure the database repository.

Oracle Clusterware

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.
- When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

Oracle Real Application Clusters

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

You must install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database before you install the Oracle Fusion Middleware components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. You install RCU in its own, separate Middleware home.

Use the latest version of RCU to install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

Database Initialization Parameters

Ensure that the following initialization parameter is set to the required value. It is checked by Repository Creation Assistant.

Table 5–4 Required Initialization Parameters

Parameter	Required Value	Parameter Class
PROCESSES	300 or greater	Static

To check the value of the initialization parameter using SQL*Plus, you can use the SHOW PARAMETER command.

As the SYS user, issue the SHOW PARAMETER command as follows:

```
SQL> SHOW PARAMETER processes
```

Set the initialization parameter using the following command:

```
SQL> ALTER SYSTEM SET processes=300 SCOPE=SPFILE
```

Restart the database.

Note: The method that you use to change a parameter's value depends on whether the parameter is static or dynamic, and on whether your database uses a parameter file or a server parameter file. See the *Oracle Database Administrator's Guide* for details on parameter files, server parameter files, and how to change parameter values.

5.13.1.8 Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database

Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and Oracle SOA schemas into a Real Application Clusters database before you install Oracle Fusion Middleware SOA components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about installing RCU.

5.13.1.8.1 Running RCU Run RCU to install the required metadata for Oracle Fusion Middleware 11g.

1. For Linux, go to `RCU_Home/bin` and use the following command:

```
./rcu
```

2. In the Welcome screen, click **Next**.
3. In the Create Repository screen, select **Create** to load component schemas into a database, click **Next**.
4. In the Database Connection Details screen, enter connection information for your database:
 - Database Type: Select **Oracle Database**.
 - Host Name: Enter the name of the node that is running the database. For an Oracle RAC database, specify the VIP name, or one of the node names as the host name: **SOADBHOST1VIRTUAL**.
 - Port: The port number for the database: **1521**
 - Service Name: Enter the service name of the database: **soaha.mycompany.com**
 - Username: **SYS**
 - Password: Enter the password for the SYS user.
 - Role: **SYSDBA**
5. Click **Next**.
6. If you receive the following warning message:

The database you are connecting is with non-UTF8 charset, if you are going to use this DB for multilingual support, you may have data loss. If you are not using for multilingual support you can continue, otherwise we strongly recommend using UTF-8 database.

7. Click **Ignore** or **Stop**.
8. In the Select Components screen, select **Create a New Prefix**, and enter a prefix to use for the database schemas, for example, **SOAHA**
Write down the schema names so they are available in later procedures.
 - Select the following schemas:
 - Under AS Common Schemas, select **Metadata Services**.
 - Under SOA Service Infrastructure, select **SOA Service Infrastructure** and **User Messaging**.
9. Click **Next**.
10. In the Schema Passwords screen, enter passwords for the main and additional (auxiliary) schema users, and click **Next**.
11. In the Map Tablespaces screen, choose the tablespaces for the selected components, and click **Next**.
12. In the Summary screen, click **Create**.
13. In the Completion Summary screen, click **Close**.

See the *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about using RCU.

5.13.1.8.2 Configuring SOA Schemas for Transactional Recovery Privileges You need the appropriate database privileges to allow the Oracle WebLogic Server transaction manager to query for transaction state information and issue the appropriate commands, such as commit and rollback, during recovery of in-flight transactions after a WebLogic Server container crash.

To configure the SOA schemas for transactional recovery privileges:

1. Log on to SQL*Plus as a user with sysdba privileges. For example:

```
sqlplus "/ as sysdba"
```
2. Grant select on sys.dba_pending_transactions to **soaha_soainfra**.
3. Grant force any transaction to **soaha_soainfra**.

Note: These privileges should be granted to the owner of the soainfra schema, as determined by the RCU operations.

5.13.1.9 Configuring Virtual Server Names and Ports for the Load Balancer

This section describes the load balancer prerequisites for deploying an Oracle SOA Suite high availability environment.

Load Balancers

Oracle SOA Suite uses a hardware load balancer when deployed in a high availability configuration with two Oracle HTTP Servers as web tier components. The hardware load balancer should have the following features:

- Ability to load-balance traffic to a pool of real servers through a virtual host name: Clients access services using the virtual host name (instead of using actual host

names). The load balancer can then load balance requests to the servers in the pool.

- Port translation configuration: The load balancer should have the ability to perform port translation, where it allows incoming requests received on one port to be routed to a server process running on a different port. For example, a request received on port 80 can be routed to port 7777.
- Monitoring of ports (HTTP, HTTPS)
- Virtual servers and port configuration: Ability to configure virtual server names and ports on your external load balancer, and the virtual server names and ports must meet the following requirements:
 - The load balancer should allow configuration of multiple virtual servers. For each virtual server, the load balancer should allow configuration of traffic management on more than one port. For example, the load balancer typically needs to be configured with a virtual server and ports for HTTP and HTTPS traffic.
 - The virtual server names must be associated with IP addresses and be part of your DNS. Clients must be able to access the load balancer through the virtual server names.
- Ability to detect node failures and immediately stop routing traffic to the failed node.
- Resource monitoring/port monitoring/process failure detection: The load balancer must be able to detect service and node failures (through notification or some other means) and to stop directing non-Oracle Net traffic to the failed node. If your load balancer has the ability to automatically detect failures, you should use it.
- Fault-tolerant mode: It is highly recommended that you configure the load balancer to be in fault-tolerant mode.
- Other: It is highly recommended that you configure the load balancer virtual server to return immediately to the calling client when the back-end services to which it forwards traffic are unavailable. This is preferred over the client disconnecting on its own after a timeout based on the TCP/IP settings on the client system.
- Sticky routing capability: Ability to maintain sticky connections to components based on cookies or URL.
- SSL acceleration: SSL acceleration is a method of offloading the processor-intensive public key encryption algorithms involved in SSL transactions to a hardware accelerator.

This feature is recommended but not required.

[Table 5–5](#) shows some example virtual server names to use for the external load balancer in the Oracle SOA Suite high availability environment.

Table 5–5 Virtual Server Names for the External Load Balancer

Component	Virtual Server Name
Oracle SOA Suite	soa.mycompany.com
WebLogic Server Administration Console	admin.mycompany.com
Oracle Enterprise Manager Fusion Middleware Control	admin.mycompany.com

Virtual Server Names

This section describes the virtual server names that should be set up for the high availability deployments described in this chapter.

soa.mycompany.com

`soa.mycompany.com` is a virtual server name that acts as the access point for all HTTP traffic to the runtime SOA components, such as `soa-infra` and `workflow`. Traffic to the both SSL and non-SSL ports is configured and typically non-SSL is redirected to SSL. Clients access this service using the address `soa.mycompany.com:443`. This virtual server is defined on the load balancer.

admin.mycompany.com

This virtual server acts as the access point for all internal HTTP traffic that gets directed to the administration services. The incoming traffic from clients is non-SSL enabled. Thus, the clients access this service using the address `admin.mycompany.com:80` and in turn forward these to ports `7777` on `WEBHOST1` and `WEBHOST2`. The services accessed on this virtual host include the WebLogic Administration Server Console and Oracle Enterprise Manager.

In addition, ensure that the virtual server names are associated with IP addresses and are part of your Domain Name System (DNS). The computers on which Oracle Fusion Middleware is running must be able to resolve these virtual server names.

5.13.1.10 Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2

To install Oracle HTTP Server on `WEBHOST1`:

1. Verify that the servers meet the following requirements:
 - The system, patch, kernel, and other requirements meet the requirements specified in the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite*.
 - Port `7777` is not used by any service on the nodes. You can verify this by running the following command:

Unix:

```
netstat -an | grep LISTEN | grep "7777"
```

Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7777"
```

If the ports are in use, make them available.

2. On UNIX platforms, if the `/etc/oraInst.loc` or `/var/opt/oracle/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct, and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, skip this step.

3. Start Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:

For UNIX, run this command: `runInstaller`.

For Windows, double-click **setup.exe**.

4. In the Welcome screen, click **Next**.
5. In the Select Installation Type screen, select **Install and Configure**, and click **Next**.

6. In the Prerequisite Checks screen, ensure that all the prerequisites are met, and click **Next**.
7. In the Specify Installation Location screen, set the location to:
`/u01/app/oracle/product/11.1.1/ohs_1`
8. Click **Next**.
9. In the Configure Components screen:
 - Select **Oracle HTTP Server**.
 - Do not select **Associate Selected Components with WebLogic Domain**.
 - Click **Next**.
10. In the Specify Component Details screen, enter the following values:
 - Instance Home Location: Instance Home
 - Instance Home Location: `/u01/app/oracle/product/11.1.1/ohs_1/instances/ohs_instance1`
 - Instance Name: `ohs_instance1`
 - OHS Component Name: `ohs1`
11. Click **Next**.
12. In the Specify Web Tier Port Details screen:
 - Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File**, and use the **Browse** function to select the file.
 - Enter the **Oracle HTTP Server port**, for example, `7777`.
13. Click **Next**.
14. In the Configuration Summary screen, ensure that the selections are correct, and click **Install**.
15. In the Installation Progress screen:

For UNIX systems, a dialog box appears prompting you to run the `oracleRoot.sh` script. Open a command window and run the script, following the prompts.

Click **Next**.
16. In the Configuration screen, several configuration assistants are launched in succession. When the configuration assistants are finished, the Configuration Completed screen appears.
17. In the Configuration Completed screen, click **Finish** to exit.

Repeat the steps for WEBHOST2 and configure your LBR with a pool containing both the WEBHOST1 and WEBHOST2 addresses.

5.13.1.10.1 Validating Oracle HTTP Server To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URL in a browser:

HTTP://WEBHOST1:7777/

If Oracle HTTP Server is set up correctly, the Oracle FMW Welcome screen appears in the browser.

5.13.2 Installing Oracle Fusion Middleware Home

This section describes the procedure for installing Oracle Fusion Middleware on all nodes in the application tier that run Oracle WebLogic Server and Oracle SOA Suite. Repeat the procedures (described below for SOAHOST1) for installing WebLogic Server and Oracle SOA in SOAHOST2. The directory paths for binary files and domains used when installing new nodes must be exactly the same as those used for first node. If these paths and domains are not exactly the same as those used for the first node, failover will not work properly.

Install the following Oracle Fusion Middleware components:

- Oracle WebLogic Server (see [Section 5.13.2.1, "Installing Oracle WebLogic Server"](#))
- Oracle SOA Suite (see [Section 5.13.2.2, "Installing Oracle Fusion Middleware for Oracle SOA"](#))

5.13.2.1 Installing Oracle WebLogic Server

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

To install Oracle WebLogic Server on all nodes in the application tier:

1. Start Oracle WebLogic Server Installer.
2. In the Welcome screen, click **Next**.
3. In the Choose Middleware Home Directory screen:
 - Select **Create a New Middleware Home**.
 - For the **Middleware Home Directory** field, enter **MW_HOME**.
 - Click **Next**.
4. In the Register for Security Updates screen, enter your contact information for security update notifications, and click **Next**.
5. In the Choose Install Type screen, select **Custom**, and click **Next**.
6. In the Choose Products and Components screen, click **Next**.
7. In the JDK Selection screen, select only **Oracle JRockit 160_20_D1.0.1-2124 SDK**, and click **Next**.
8. In the Choose Product Installation Directories screen, accept the following directory:
`WL_HOME`
Click **Next**.
9. In the Installation Summary screen, click **Next**.
10. In the Installation Complete screen, deselect **Run QuickStart**, and click **Done**.

5.13.2.2 Installing Oracle Fusion Middleware for Oracle SOA

On Linux platforms, if the `/etc/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, you can skip this step.

1. Start Oracle Fusion Middleware 11g Oracle SOA Suite Installer:

On UNIX:

```
SOAHOST1> runInstaller
```

On Windows:

```
SOAHOST1> setup.exe
```

When Oracle Fusion Middleware 11g Oracle SOA Suite Installer prompts you for a **JRE/JDK location** enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `MW_HOME/jrockit_160_<version>`.

2. In the Specify Inventory Directory screen, do the following:
 - a. Enter `HOME/oraInventory`, where HOME is the home directory of the user performing the installation (this is the recommended location).
 - b. Enter the OS group for the user performing the installation.
 - c. Click **Next**.
 - d. Follow the instructions on the screen to execute `/createCentralInventory.sh` as root.
 - e. Click **OK**.
3. In the Welcome screen, click **Next**.
4. In the Prerequisite Checks screen, verify that the checks complete successfully, and click **Next**.
5. In the Specify Installation Location screen:
 - For Middleware Home, enter **MW_HOME**.
 - For Oracle Home Directory, enter **soa**.Click **Next**.
6. In the Installation Summary screen, click **Install**.
7. In the Installation Complete screen, click **Finish**.

Note: Before you run the Configuration Wizard by following the instructions in [Section 5.13.4, "Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA Domain,"](#) ensure that you applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home so that you have the latest version of Oracle Fusion Middleware.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

5.13.3 Enabling VIP1 in SOAHOST1 and VIP2 in SOAHOST2

The SOA domain uses virtual hostnames as the listen addresses for the SOA managed servers. You must enable VIP mapping for each of these hostnames on the two SOA Machines, (VIP1 on SOAHOST1 and VIP2 on SOAHOST2), and must correctly resolve the virtual hostnames in the network system used by the topology (either by DNS Server or by hosts resolution).

For information about configuring virtual IPs, see [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster,"](#) where an example is provided for

the Administration Server. Refer to [Section 5.13.20, "Configuring Server Migration for the WLS_SOA Servers"](#) for more details on configuring server migration for the SOA servers.

5.13.4 Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA Domain

Run Oracle Fusion Middleware Configuration Wizard from the SOA home directory to create a domain containing the Administration Server and Oracle SOA components. Ensure that the database where you installed the repository is running. For Oracle RAC databases, all the instances must be running.

Note: Multiple SOA clusters are not allowed in the same WebLogic Server domain.

Note: Oracle BPM requires the WL_HOME and ORACLE_HOME homes to be patched to the appropriate Oracle Fusion Middleware patch set level before you perform Oracle Fusion Middleware Configuration Wizard steps to extend a domain.

1. Change the directory to the location of the Oracle Fusion Middleware Configuration Wizard, located in the Middleware home, SOA directory:

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
```
2. Start Oracle Fusion Middleware Configuration Wizard:
For Linux:

```
SOAHOST1> ./config.sh
```


For Windows:

```
SOAHOST1> config.cmd
```
3. In the Welcome screen, select **Create a New WebLogic Domain**. Click **Next**.
4. In the Select Domain Source screen:
 - Select **Generate a domain configured automatically to support the following products**.
 - Select the following products:
 - **Basic Weblogic Server Domain - 10.3.6.0 [wlserver_10.3]** (Selected by default and grayed out)
 - **Oracle BPM Suite - 11.1.1.0 [Oracle_SOA1]** (ONLY FOR BPM Systems)
 - **Oracle SOA Suite - 11.1.1.0 [Oracle_SOA1]** (Selected by default when selecting BPM Suite)
 - **Oracle Service Bus OWSM Extension - 11.1.1.6**
 - **Oracle Enterprise Manager - 11.1.1.0 [oracle_common]** (Selected by default)
 - **Oracle Service Bus - 11.1.1.6**
 - **WebLogic Advanced Web Services for JAX-RPC Extension - 10.3.6.0**

- **Oracle WSM Policy Manager - 11.1.1.0 [oracle_common]** (Selected by default when selecting SOA/BPM Suite)
- **Oracle JRF - 11.1.1.0 [oracle_common]** (Selected by default when selecting SOA/BPM Suite)

If you accidentally deselect some of the targets, ensure that the following selections are made in this screen:

- Oracle BPM Suite (ONLY FOR BPM Systems)
- Oracle SOA Suite
- Oracle Enterprise Manager
- Oracle WSM Policy Manager
- Oracle JRF

Click **Next**.

5. In the Specify Domain Name and Location screen, make the following entries:

- Domain Name: **soadomain**
- Domain Location: accept the default
- Application Location: accept the default

Click **Next**.

6. In the Configure Administrator User Name and Password screen, enter the username and password to use for the domain's administrator, and click **Next**.
7. In the Configure Server Start Mode and JDK screen, select the following selections:
 - WebLogic Domain Startup Mode: select **Production Mode**
 - JDK Selection: select **Oracle JRockit 160_20_D1.0.1-2124 SDK**.

Click **Next**.

8. In the Configure JDBC Component Schema screen:

- a. Select all component schemas that appear in the table at the bottom: **SOA Infrastructure, User Messaging Service, OWSM MDS Schema and SOA MDS Schema**.
- b. For the RAC configuration for component schemas, select **Convert to GridLink** or **Convert to RAC multi data source**. For single database configuration, select **Don't convert**.
- c. Ensure that the following data source appears on the screen. The user names in [Table 5–6](#) assume that you used **soaha** as the prefix for schema creation from RCU.

For information about GridLink data source configuration with Oracle RAC, see [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC."](#)

For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

Table 5–6 Configuring Values for Data Sources

Component Schema	Schema Owner
SOA Infrastructure	SOAHA_SOAINFRA
User Messaging Service	SOAHA_ORASDPM
OWSM MDS Schema	SOAHA_MDS
SOA MDS Schema	SOAHA_MDS

- d. Click Next.
9. Enter values for the following fields, specifying connect information for the Oracle RAC database that was seeded with RCU.

For multi data sources, enter values for these fields:

Figure 5–43 Configure RAC Multi Data Source Component Schema Screen

Note: Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: *Oracle's Driver (Thin) for RAC Service-Instance |

Service Name: soaha.mycompany.com

Username: Varies among component schemas

Password: *****

Host Name	Instance Name	Port
soadbhost1.vir	soahadb1	1521
soadbhost2.vir	soahadb2	1521

Add Delete

	RAC Component Schema	Service Name	Schema Owner	Schema Password
☒	SOA Infrastructure	soaha.mycompany.com	SOAHA_SOAINFRA	*****
☒	User Messaging Service	soaha.mycompany.com	SOAHA_ORASDPM	*****
☒	OWSM MDS Schema	soaha.mycompany.com	SOAHA_MDS	*****
☒	SOA MDS Schema	soaha.mycompany.com	SOAHA_MDS	*****

Exit Help Previous Next

For GridLink data sources, enter values for these fields:

Figure 5–44 Configure GridLink RAC Component Schema

Fusion Middleware Configuration Wizard

Configure GridLink RAC Component Schema

Note: Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: *Oracle's Driver (Thin) for GridLink Connections; Version: [dropdown]

Service Name: [Enter a value]

Username: [Varies among component schemas]

Password: [Enter a value]

Enable FAN: ☒

Enable SSL: ☐

Wallet File: [Enter a value] [Browse]

Wallet Password: [Enter a value]

Service Listener	Port	Protocol
		TCP

ONS Host	Port

[Add] [Delete]

RAC Component Schema	Service Name	Schema Owner	Schema Password
☒ SOA Infrastructure		DEV_SOAINFRA	
☒ User Messaging Service		DEV_ORASDPM	
☒ OWSM MDS Schema		DEV_MDS	
☒ SOA MDS Schema		DEV_MDS	

[Exit] [Help] [Previous] [Next]

- Driver: For RAC: Select **Oracle driver (Thin) for RAC Service-Instance connections, Versions:10, 11**. For GridLink: **Oracle Driver (Thin) for GridLink Connections, Versions:10 and later**
- Service Name: Enter the service name of the database. **For a GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example, `<mydbservice>.us.oracle.com`

Note: The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example, `us.oracle.com`

- Username prefix: Enter the prefix for the schemas. The user names in [Table 5–6](#) assume that **soaha** was used as prefix for schema creation from RCU.
- Password and Confirm Password: Enter the password for access to the schemas.
- For a GridLink data source**, enter the SCAN address in the Listener Address field and SCAN port in the Port field. Enter the ONS host and port information in the ONS Host and Port fields, respectively.

Note: Oracle recommends that you use SCAN addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

For a multi data source, click **Add**. Enter details for Oracle RAC instances.

- f. Update each schema: select one data source at a time and add the appropriate details.

Ensure that the information is entered for all schemas: **SOA Infrastructure**, **User Messaging Service**, **OWSM MDS Schema**, and **SOA MDS Schema**.

- g. Click **Next**.

10. In the Test JDBC Data Sources screen, the connections are tested automatically. The **Status** column shows the results. Ensure that all connections were successful. If not, click **Previous** to correct your entries.

Click **Next** when all the connections are successful.

11. In the Select Optional Configuration screen, select the following:

- **Administration Server**
- **JMS Distributed Destinations**
- **Managed Servers, Clusters and Machines**

12. In the Configure the Administration Server screen, enter the following values:

- Name: **AdminServer**
- Listen Address: Enter the hostname for the VIP0 virtual IP
- Listen Port: **7001**
- SSL listen port: N/A
- SSL enabled: leave unchecked

Click **Next**.

13. Verify that the UDD settings are **UDD for UMSJMSSystemResource**, **SOAJMSModule**, and **BPMJMSModule**.

14. In the Configure Managed Servers screen, add the following managed servers:

Table 5–7 Configuring Managed Servers

Name	Listen Address	Listen Port	SSL Listen Port	SSL Enabled
WLS_SOA1	SOAHOST1VHN1	8001	n/a	No
WLS_SOA2	SOAHOST2VHN1	8001	n/a	No

Do not delete any server that appears. You can modify the servers. If you delete a server and add a new one, targeting fails.

Note: Although the standard recommendation is to run custom applications and other systems in a separate WebLogic Managed Server, the creation of the custom WLS managed servers described in [Figure 5-42](#) is not addressed here.

Click **Next**.

15. In the Configure Clusters screen, add the following cluster:

- Name: **SOA_Cluster**
- Cluster Messaging Mode: **unicast**
- Multicast Address: N/A
- Multicast Port: N/A
- Cluster Address: Leave empty

Click **Next**.

16. In the Assign Servers to Clusters screen, assign the following servers to SOA_Cluster:

- WLS_SOA1
- WLS_SOA2

Click **Next**.

17. In the Configure Machines screen:


- Delete the **LocalMachine** that appears by default.
- Click the **Unix Machine** tab, and add the following machines:

Table 5-8 Configuring Machines

Name	Node Manager Listen Address
SOAHOST1	Hostname of SOAHOST1
SOAHOST2	Hostname of SOAHOST2

Leave all other fields as their default values, and click **Next**.

18. In the Assign Servers to Machines screen, assign servers to machines as follows:

Figure 5–45 Assign Servers to Machines Screen

- SOAHOST1: AdminServer, WLS_SOA1
- SOAHOST2: WLS_SOA2

Click **Next**.

19. In the Review WebLogic Domain screen, click **Next**.
20. In the Configuration Summary screen, click **Create**.
21. In the Creating Domain screen, click **Done**.

Note: The multicast and unicast addresses are different from the ones used by the WebLogic Server cluster for cluster communication. SOA guarantees that composites are deployed to members of a single WebLogic Server cluster even though the communication protocol for the two entities (the WebLogic Server cluster and the groups to which composites are deployed) are different.

5.13.5 Creating boot.properties for the Administration Server on SOAHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a boot.properties file for the Administration Server on SOAHOST1.

For the Administration Server:

1. Create the following directories:

```
SOAHOST1> mkdir -p ORACLE_BASE/product/fmw/user_
projects/domains/soadomain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` in the security directory created in the previous step, and enter the following lines in the file:

```
username=adminuser  
password=password
```

Note: When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time that the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible for the entries to be encrypted.

5.13.6 Starting and Validating the Administration Server in SOAHOST1

This section describes procedures for starting and validating the Administrator Server in SOAHOST1.

5.13.6.1 Starting the Administration Server on SOAHOST1

To start the Administration Server on SOAHOST1, run the following commands:

```
SOAHOST1> cd MW_HOME/user_projects/domains/soadomain/bin
```

```
SOAHOST1> ./startWebLogic.sh
```

5.13.6.2 Validating the Administration Server

To verify that the Administration Server is properly configured:

1. In a browser, go to `http://vip0:7001/console`.
2. Log in as the administrator.
3. Verify that the WLS_SOA1 and WLS_SOA2 managed servers are listed.
4. Verify that the SOA_Cluster cluster is listed.
5. Verify that you can access Enterprise Manager at `http://vip0:7001/em`.

5.13.7 Disabling Host Name Verification for the Administration Server and the WLS_SOA Managed Servers

This step is required if you have not set up the appropriate certificates for hostname verification between the Administration Server and Node Manager. If SSL is not set up, you receive an error message unless you disable host name verification.

You can re-enable host name verification when you have set up SSL communication between the Administration Server and the Node Manager.

To disable host name verification:

1. In Oracle WebLogic Server Administration Console, select **Servers**, and then **AdminServer**.
2. Select **SSL**, and then **Advanced**.
3. In the Change Center, click **Lock & Edit**.
4. Set **Hostname Verification** to **None**.

5. Save and activate the changes.
6. Repeat these steps for the WLS_SOA1 and WLS_SOA2 servers.
7. Restart the Administration Server.

5.13.8 Configuring Oracle Coherence for Deploying Composites

Although deploying composites uses multicast communication by default, Oracle recommends using unicast communication for SOA high availability. Use unicast if you disable multicast communication for security reasons.


Note: An incorrect configuration of the Oracle Coherence framework that is used for deployment may prevent the SOA system from starting. The deployment framework must be properly customized for the network environment on which the SOA system runs. Oracle recommends the following configuration described in this section.

Enabling Communication within Clusters Using Unicast Communication

Multicast communication enables Oracle Fusion Middleware SOA to discover all of the members of a cluster to which it deploys composites dynamically. However, unicast communication does not enable nodes to discover other cluster members in this way. Consequently, you must specify the nodes that belong to the cluster. You do not need to specify all of the nodes of a cluster, however. You need only specify enough nodes so that a new node added to the cluster can discover one of the existing nodes. As a result, when a new node has joined the cluster, it is able to discover all of the other nodes in the cluster. Additionally, in a configuration where multiple IPs are available in the same node, you must configure Oracle Coherence to use a specific hostname to create Oracle Coherence cluster.

Tip: To guarantee high availability during deployments of SOA composites, specify enough nodes so that at least one of them is running at any given time.

Specify the nodes using the `tangosol.coherence.wka<n>` system property, where `<n>` is a number between 1 and 9. You can specify up to 9 nodes. Start the numbering at 1. This numbering must be sequential and must not contain gaps. In addition, specify the hostname used by Oracle Coherence to create a cluster through the `tangosol.coherence.localhost` system property. This hostname should be the virtual hostname used by the SOA servers that maps the corresponding listener addresses (VIP1 and VIP2). Set this property by adding the `-Dtangosol.coherence.localhost` parameters to the Arguments field of Oracle WebLogic Server Administration Console's Server Start tab ([Figure 5-46](#)).

Figure 5–46 Setting the Hostname Using the Start Server Tab of Oracle WebLogic Server Administration Console

Specifying the hostname

To add the hostname used by Oracle Coherence:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node.
3. Click **Servers**. The Summary of Servers page appears.
4. Click the name of the server (represented as a hyperlink) in Name column of the table. The settings page for the selected server appears.
5. Click the **Server Start** tab (illustrated in Figure 5–46).
6. Enter the following for WLS_SOA1 and WLS_SOA2 into the Arguments field.

For WLS_SOA1, enter the following (on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1 -Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost1vhn1
```

For WLS_SOA2, enter the following (on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1 -Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost2vhn1
```

Note: The Coherence cluster used for deployment uses port 8088 by default. This port can be changed by specifying a different port (for example, 8089) with the `-Dtangosol.coherence.wka1.port` and `-Dtangosol.coherence.localport` startup parameters. For example:

WLS_SOA1 (enter the following into the Arguments field on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1
-Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost1vhn1
-Dtangosol.coherence.localport=8089
-Dtangosol.coherence.wka1.port=8089
-Dtangosol.coherence.wka2.port=8089
```

WLS_SOA2 (enter the following into the Arguments field on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1
-Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost2vhn1
-Dtangosol.coherence.localport=8089
-Dtangosol.coherence.wka1.port=8089
-Dtangosol.coherence.wka2.port=8089
```

7. Click **Save** and activate the changes.
8. This change requires the SOA servers to be restarted.

Note: The multicast and unicast addresses are different from the ones used by the WebLogic Server cluster for cluster communication. SOA guarantees that composites are deployed to members of a single WebLogic Server cluster even though the communication protocol for the two entities (the WebLogic Server cluster and the groups to which composites are deployed) are different.

Do not copy the text from above to your Administration Console's arguments text field, which can result in HTML tags being inserted in the Java arguments. The text should not include any text or characters other than the ones shown above.

5.13.9 Starting the System in SOAHOST1

This section describes how to start Node Manager on SOAHOST1 and how to start and validate the WLS_SOA1 managed server on SOAHOST1.

5.13.9.1 Starting Node Manager on SOAHOST1

To start Node Manager on SOAHOST1:

1. Run the `setNMProps.sh` script, which is located in the `ORACLE_COMMON_HOME/common/bin` directory, to set the `StartScriptEnabled` property to `true` before starting Node Manager

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
SOAHOST1> ./setNMProps.sh
```

Note: You must use the `StartScriptEnabled` property to avoid class loading failures and other problems.

2. Start Node Manager:

```
SOAHOST1> cd WL_HOME/server/bin
SOAHOST1> ./startNodeManager.sh
```

5.13.9.2 Starting and Validating the WLS_SOA1 Managed Server

To start the WLS_SOA1 managed server and check that it is configured correctly:

1. Start the WLS_SOA1 managed server using Oracle WebLogic Server Administration Console.
2. When WLS_SOA1 starts, the following URLs become available:
 - `http://SOAHOST1VHN1:8001/b2bconsole`
 - Verify login to B2B console
 - `http://SOAHOST1VHN1:8001/integration/worklistapp`
 - Verify login to worklist console
 - `http://SOAHOST1VHN1:8001/wsm-pm`
 - Verify the policy validator link.
 - `http://SOAHOST1VHN1:8001/soa/composer`
 - `http://SOAHOST1VHN1:8001/soa-infra`
 - `http://SOAHOST1VHN1:8001/bpm/composer` (ONLY FOR BPM Systems)
 - `http://SOAHOST1VHN1:8001/bpm/workspace` (ONLY FOR BPM Systems)

5.13.10 Propagating the Domain Configuration to SOAHOST2 with pack/unpack Utilities

Follow these steps to propagate the domain configuration to SOAHOST2 using pack/unpack utilities:

1. Run the following pack command on SOAHOST1 to create a template pack:

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
SOAHOST1> ./pack.sh -managed=true
-domain=ORACLE_BASE/product/fmw/user_projects/domains/soadomain/
-template=soadomaintemplate.jar
-template_name=soa_domain_template
```

2. Run the following command on SOAHOST1 to copy the template file created in the previous step to SOAHOST2:

```
SOAHOST1> scp soadomaintemplate.jar
oracle@node2:ORACLE_COMMON_HOME/common/bin
```

3. Run the unpack command on SOAHOST2 to unpack the propagated template:

```
SOAHOST2> cd ORACLE_COMMON_HOME/common/bin
SOAHOST2> ./unpack.sh
-domain=ORACLE_BASE/product/fmw/user_projects/domains/soadomain/
-template=soadomaintemplate.jar
```

Note: When you run the Fusion Middleware Configuration Wizard after a pack/unpack procedure, products that already exist in the original domain and appear in the Select Extension Source screen, are not selected or greyed out.

Templates created with the pack command are considered user-created templates. They are treated differently from default templates. User-created templates are self-contained and do not contain any information related to the creation of the original domain.

5.13.11 Extracting XEngine Files in the Second Node

To enable B2B's XEngine in the second node, it is required to extract the content of the ZEngine tar manually:

```
SOAHOST2>cd ORACLE_HOME/soa/thirdparty/edifecs
SOAHOST2>tar xzvf XEngine.tar.gz
```

5.13.12 Starting the System in SOAHOST2

This section describes procedures for starting the system in SOAHOST2.

5.13.12.1 Starting Node Manager on SOAHOST2

To start the Node Manager on SOAHOST2, repeat the steps from [Section 5.13.9.1, "Starting Node Manager on SOAHOST1"](#) on SOAHOST2.

5.13.12.2 Starting and Validating the WLS_SOA2 Managed Server

To start the WLS_SOA2 managed server and verify that it is configured correctly:

1. Start the WLS_SOA2 managed server using Oracle WebLogic Server Administration Console
2. When WLS_SOA2 is started, the following URLs become available:
 - <http://SOAHOST2VHN1:8001/b2bconsole>
Verify login to B2B console
 - <http://SOAHOST2VHN1:8001/integration/worklistapp>
Verify login to worklist console
 - <http://SOAHOST2VHN1:8001/wsm-pm>
Verify the policy validator link
 - <http://SOAHOST2VHN1:8001/soa/composer>
 - <http://SOAHOST2VHN1:8001/soa-infra>
 - <http://SOAHOST2VHN1:8001/bpm/composer> (ONLY FOR BPM Systems)
 - <http://SOAHOST2VHN1:8001/bpm/workspace> (ONLY FOR BPM Systems)

5.13.13 Configuring Oracle HTTP Servers for the Administration Server and the WLS_SoAn Managed Servers

To enable Oracle HTTP Server to route to the SOA Cluster, which contains the WLS_SoAn managed servers, you must set the WebLogicCluster parameter to the list of nodes in the cluster.

1. On WEBHOST1 and WEBHOST2, add the following lines to the ORACLE_BASE/admin/<instance_name>/config/OHS/<component_name>/mod_wl_ohs.conf file:

```
# WSM-PM
<Location /wsm-pm>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# SOA soa-infra app
<Location /soa-infra>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# SOA inspection.wsil
<Location /inspection.wsil>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# Worklist
<Location /integration/>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# B2B
<Location /b2b>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

<Location /b2bconsole>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# UMS
<Location /sdpmessaging/ >
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# UMS WS
<Location /ucs/messaging/webservice >
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

# Default to-do taskflow
<Location /DefaultToDoTaskFlow/>
    SetHandler weblogic-handler
```

```
        WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
    </Location>

    # Workflow
    <Location /workflow>
        SetHandler weblogic-handler
        WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
    </Location>

    #Required if attachments are added for workflow tasks
    <Location /ADFAttachmentHelper>
        SetHandler weblogic-handler
        WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
    </Location>

    # SOA composer application
    <Location /soa/composer>
        SetHandler weblogic-handler
        WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
    </Location>

    # BPM composer (ONLY FOR BPM Systems)
    <Location /bpm/composer >
        SetHandler weblogic-handler
        WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
    </Location>

    # BPM workspace (ONLY FOR BPM Systems)
    <Location /bpm/workspace >
        SetHandler weblogic-handler
        WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
    </Location>
```

2. Make sure the httpd.conf file located in the same directory as the mod_wl_ohs file contains the following lines:

```
NameVirtualHost *:7777
<VirtualHost *:7777>
    ServerName https://soa.mycompany.com:443
    ServerAdmin you@your.address
    RewriteEngine On
    RewriteOptions inherit
</VirtualHost>

NameVirtualHost *:7777
<VirtualHost *:7777>
    ServerName admin.mycompany.com:80
    ServerAdmin you@your.address
    RewriteEngine On
    RewriteOptions inherit
</VirtualHost>
```

Note: Values such as soa.mycompany.com:443, 7777, admin.mycompany.com:80, and you@youraddress that are noted in this document serve as examples only. Enter values based on the actual environment.

3. Perform the same steps for the Oracle HTTP Server on WEBHOST2.

4. Restart Oracle HTTP Server on both WEBHOST1 and WEBHOST2:

```
WEBHOST1> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs1
```

```
WEBHOST2> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs2
```

5.13.14 Validating Access Through Oracle HTTP Server

Verify that the SOA Servers status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming," wait for the server status to change to "Started." If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors. Verify that you can access these URLs:

- <http://WEBHOST1:7777/wsm-pm>
- <http://WEBHOST2:7777/wsm-pm>
- <http://WEBHOST1:7777/soa-infra>
- <http://WEBHOST2:7777/soa-infra>
- <http://WEBHOST1:7777/soa/composer>
- <http://WEBHOST2:7777/soa/composer>
- <http://WEBHOST1:7777/integration/worklistapp>
- <http://WEBHOST2:7777/integration/worklistapp>
- <http://WEBHOST1:7777/sdpMessaging/userprefs-ui>
- <http://WEBHOST2:7777/sdpMessaging/userprefs-ui>
- <http://WEBHOST1:7777/b2bconsole>
- <http://WEBHOST2:7777/b2bconsole>
- <http://WEBHOST1:7777/bpm/composer> (ONLY FOR BPM Systems)
- <http://WEBHOST2:7777/bpm/composer> (ONLY FOR BPM Systems)
- <http://WEBHOST1:7777/bpm/workspace> (ONLY FOR BPM Systems)
- <http://WEBHOST2:7777/bpm/workspace> (ONLY FOR BPM Systems)

Verify these URLs also using your load balancing router address:

- <http://soa.mycompany.com:80/wsm-pm>
- <http://soa.mycompany.com:80/soa-infra>
- <http://soa.mycompany.com:80/soa/composer>
- <http://soa.mycompany.com:80/integration/worklistapp>
- <http://soa.mycompany.com:80/sdpMessaging/userprefs-ui>
- <http://soa.mycompany.com:80/b2bconsole>
- <http://soa.mycompany.com:80/bpm/composer> (ONLY FOR BPM Systems)
- <http://soa.mycompany.com:80/bpm/workspace> (ONLY FOR BPM Systems)

Follow these instructions to ensure that routing and failover from the HTTP Server to the SOA_Cluster is working correctly:

1. While WLS_SOA2 is running, stop WLS_SOA1 from Oracle WebLogic Server Administration Console.
2. Access the following URLs and verify the appropriate functionality:
 - WEBHOST1:7777/wsm-pm
 - WEBHOST1:7777/soa-infra
 - WEBHOST1:7777/soa/composer
 - WEBHOST1:7777/integration/worklistapp
 - WEBHOST1:7777/sdpMessaging/userprefs-ui
 - WEBHOST1:7777/b2bconsole
 - WEBHOST1:7777/bpm/composer (ONLY FOR BPM Systems)
 - WEBHOST1:7777/bpm/workspace (ONLY FOR BPM Systems)
3. Start WLS_SOA1 from Oracle WebLogic Server Administration Console.
4. Stop WLS_SOA2.
5. Access the URLs in Step 2 above again and verify the appropriate functionality:

5.13.15 Configuring JMS Persistence Store as Shared Across the Servers

Configure the location for all persistence stores to a directory visible from both nodes. Change all persistent stores to use this shared base directory.

From the Fusion Middleware Administration Console, select **Services, Persistence Store, Persistence_Store_Name**, and then **Directory**.

To enable resume of pending JMS messages, you must specify a location on a persistent storage solution (NAS, SAN) that is available to other servers in the cluster. Therefore, the directory that you enter must be accessible by both WLS_SOA1 and WLS_SOA2. This directory must exist before the server is restarted.

Note: See [Considerations for Using File Stores on NFS](#) for important information about JMS messages and transaction logs, as well as releasing locks on file stores.

5.13.16 Configuring a Default Persistent Store for Transaction Recovery

Each server has a transaction log, which stores information about committed transactions coordinated by the server that may not have been completed. WebLogic Server uses the transaction log when recovering from system crashes or network failures. To take advantage of the migration capability of the Transaction Recovery Service for servers in a cluster, you must store the transaction log in a location that is available to a server and its backup servers, preferably on a dual-ported SCSI disk, or a Storage Area Network (SAN). To configure the default persistent store:

1. In the **Domain Structure** tree, expand **Environment** and select **Servers**.
2. Select the server you want to modify.
3. Select the **Configuration, Services** tab.
4. Under **Default Store**, in the **Directory** field, enter the path to the folder where you want the default persistent store to store its data files.

To enable migration of the Transaction Recovery Service, you must specify a location on a persistent storage solution that is available to other servers in the cluster. Therefore, the directory that you enter must be accessible by both WLS_SOA1 and WLS_SOA2. This directory must exist before the server is restarted.

5.13.17 Setting the Front End HTTP Host and Port

You must set the front end HTTP host and port for Oracle WebLogic Server cluster:

1. In the WebLogic Server Administration Console, in the Change Center section, click **Lock & Edit**.
2. In the left pane, select **Environment > Clusters**.
3. Select the SOA_Cluster.
4. Select **HTTP**.
5. Set the values for the following:
 - **Frontend Host:** soa.mycompany.com
Frontend HTTP Port: 80
 - **Frontend HTTPS Port:** leave it blank

Note: Make sure this address is correct and available (the load balancing router is up). An incorrect value, for example, **http://** in the address, or trailing **/** in the hostname, may prevent the SOA system from being accessible even when using the virtual IPs to access it.

6. Click **Save**.
7. To activate the changes, click **Activate Changes** in the Change Center section of the Administration Console.
8. If you have started the server before, notice this change requires a restart of the managed servers in the cluster.

Note: The SOA system calculates the callback URL as follows:

- If a request to SOA originates from an external or internal service is originating a request to SOA, then SOA uses the callback URL specified by the client.
- If a request to an external or internal service originates from SOA, the callbackURL cannot be populated in the SOA request dynamically because SOA is the originator. Instead, callbackServerURL is used if it is specified as a binding property for the specific reference. (You can set this when modeling the composite or at runtime using the MBeans accessed through Oracle Enterprise Manager console.) This allows different service calls to have different callback URLs. That is, a callback URL from an external service is different than one to an internal service.

However, if the callbackServer URL is not specified as a binding property for that reference, then the system uses the callback URL as specified in `soa-infra-config`. If the callback URL is not specified in `soa-infra-config`, then the system uses the callback URL as the front end host specified in WLS. If the front end host is not specified in WLS, the system uses the callback URL as the local host name as provided by WLS MBean APIs.

For SOA high availability installations frontended by Oracle HTTP Server, monitoring should be done on the Oracle HTTP Server ports of the real backend servers. This is the case when a deployment is using all the components deployed to the SOA Managed Server. A simple HTTP monitor that pings the HTTP/HTTPS port and expects a pre-determined response in turn should suffice. If only a specific SOA component is being used (such as B2B), then a monitor that does a deeper level check all the way to the managed server can be considered to validate the health of the component in use. Please check with your load balancer vendor on setting up the HTTP monitors with your load balancer.

If you do not set the front end HTTP host and port, you get the following message when trying to retrieve a document definition XSD from Oracle B2B:

```
An error occurred while loading the document definitions.  
java.lang.IllegalArgumentException: Cluster address must be set when  
clustering is enabled.
```

5.13.18 Setting the WLS Cluster Address for Direct Binding/RMI Invocations to Composites

You must set the WLS Cluster address for the SOA_Cluster if you are going to use direct binding to composites. To do this follow these steps:

1. In the WebLogic Server Administration Console, in the Change Center section, click **Lock & Edit**.
2. In the left pane, choose **Environment** in the Domain Structure window and then choose **Clusters**. The Summary of Clusters page appears.
3. Select the **SOA_Cluster** cluster.
4. In the **Configuration > General** tab, enter the following in the **Cluster Address** field: `SOAHOST1VHN1:8001, SOAHOST2VHN1:8001`

5. Click **Save**.
6. To activate the changes, click **Activate Changes** in the Change Center section of the Administration Console.
7. Restart the servers to make the Frontend Host directive in the cluster effective.

Note: For asynchronous request/response interactions over direct binding, the SOA composites need to provide the JNDI provider URL for the invoked service to look up the beans for callback.

If soa-infra config properties are not specified but the WebLogic Server cluster address is specified, the cluster address will be used to form the JNDI provider URL. This cluster address can be a single DNS name which maps to the clustered servers' IP-addresses or a comma separated list of server *ip:port*. Alternatively, if the soa-infra config property `JndiProviderURL/SecureJndiProviderURL` is explicitly set by users, it can be used also for the same purpose.

5.13.19 Deploying Applications

You can deploy SOA composite applications from Oracle Enterprise Manager Fusion Middleware Control Console with the Deploy SOA Composite wizard. Use the Deploy SOA Composite wizard to deploy any of the following:

- A new SOA composite application for the first time
- A new revision (for example, 2.0) alongside an older revision (for example, 1.0) without impacting the latter. The revision deployed last becomes the new default revision of that composite (unless you specify otherwise at a later step during deployment).
- A bundle (ZIP file) containing multiple SOA composite application revisions (for example, revisions 2.0, 3.0, and 4.0) of a SOA composite application that already has a different revision that is currently deployed (for example, 1.0). This option enables you to deploy revisions 1.0, 2.0, 3.0, and 4.0 at the same time. The bundle can also contain revisions of different composites. There is no restriction that all revisions must be of the same composite application.

Deployment extracts and activates the composite application in SOA Service Infrastructure. Once an application is deployed, you can perform administration tasks, such as creating instances, configuring properties, monitoring performance, managing instance life cycles, and managing policies and faults.

Note: To redeploy an *existing* revision of an application, do *not* use this wizard. Instead, use the Redeploy SOA Composite wizard.

To deploy applications:

1. Access the Deploy SOA Composite wizard through one of the following options:

From SOA Infrastructure Menu...	From the SOA Folder in the Navigator...	From the SOA Infrastructure Home Page...	From the SOA Composite Menu...
1. Select SOA Deployment > Deploy .	1. Right-click soa-infra . 2. Select SOA Deployment > Deploy .	1. Click the Deployed Composites tab. 2. Above the Composite table, click Deploy .	1. Select SOA Deployment > Deploy Another Composite .

The Select Archive page appears.

- In the **Archive or Exploded Directory** section, specify the archive of the SOA composite application to deploy. The archive contains the project files of the composite to be deployed (for example, **HelloWorld_rev1.0.jar** for a single archive or **OrderBooking_rev1.0.zip** for multiple archives).
- In the **Configuration Plan** section, optionally specify the configuration plan to include with the archive. The configuration plan enables you to define the URL and property values to use in different environments. During process deployment, the configuration plan is used to search the SOA project for values that must be replaced to adapt the project to the next target environment.
- Click **Next**.
The Select Target page appears.
- Select the WebLogic Server or cluster to which to deploy the SOA composite application archive. You can deploy to multiple servers and clusters.
- Click **Next** and review your selections.
- Select whether or not to deploy the SOA composite application as the default revision. The default revision instantiates when a new request comes in.
- Click **Deploy**.
- When deployment completes, click **Close**.

See Also: See *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite* for instructions on creating configuration plans and deploying applications from Oracle JDeveloper

5.13.20 Configuring Server Migration for the WLS_SOA Servers

The high availability architecture for a SOA system uses server migration to protect some singleton services against failures. For more information on whole server migration, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

The WLS_SOA1 managed server is configured to restart on SOAHOST2 in case of failure, and the WLS_SOA2 managed server is configured to restart on SOAHOST1 in case of failure. For this configuration the WLS_SOA1 and WLS_SOA2 servers listen on specific floating IPs that are failed over by WLS Server Migration.

The following steps enable server migration, enabling a managed server to failover to another node in the case of server or process failure.

Configuring server migration for the WebLogic managed servers consists of these steps:

- Step 1: [Setting Up a User and Tablespace for the Server Migration Leasing Table](#)
- Step 2: [Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console](#)
- Step 3: [Editing Node Manager's Properties File](#)
- Step 4: [Setting Environment and Superuser Privileges for the wlsifconfig.sh Script](#)
- Step 5: [Configuring Server Migration Targets](#)
- Step 6: [Testing the Server Migration](#)

5.13.20.1 Setting Up a User and Tablespace for the Server Migration Leasing Table

The first step to set up a user and tablespace for the server migration leasing table:

Note: If other servers in the same domain are already configured with server migration, you can use the same tablespace and data sources. In this case, you don't need to recreate the data sources and multi data source for database leasing, however, you must retarget them to the clusters you're configuring for server migration.

1. Create a tablespace named `leasing`. For example, log on to SQL*Plus as the `sysdba` user and run the following command:

```
SQL> create tablespace leasing logging datafile 'DB_
HOME/oradata/orcl/leasing.dbf' size 32m autoextend on next 32m maxsize 2048m
extent management local;
```

2. Create a user named `leasing` and assign to it the `leasing` tablespace:

```
SQL> create user leasing identified by welcome1;
SQL> grant create table to leasing;
SQL> grant create session to leasing;
SQL> alter user leasing default tablespace leasing;
SQL> alter user leasing quota unlimited on LEASING;
```

3. Create the leasing table using the `leasing.ddl` script:

- a. Copy the `leasing.ddl` file located in either the `WL_HOME/server/db/oracle/817` or the `WL_HOME/server/db/oracle/920` directory to your database node.
- b. Connect to the database as the `leasing` user.

- c. Run the leasing.ddl script in SQL*Plus:

```
SQL> @Copy_Location/leasing.ddl;
```

5.13.20.2 Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console

You create a data source to each of the Oracle RAC database instances during the process of setting up the multi data source, both for these data sources and the global leasing multi data source. When you create a data source:

- Ensure that this is a non-XA data source.
- Use Oracle's Driver (Thin) Version 9.0.1, 9.2.0, 10, 11.
- Data sources do not require support for global transactions. Therefore, do *not* use any type of distributed transaction emulation/participation algorithm for the data source (do not choose the **Supports Global Transactions** option, or the **Logging Last Resource, Emulate Two-Phase Commit**, or **One-Phase Commit** options of the **Supports Global Transactions** option), and specify a service name for your database.
- Target these data sources to the cluster(s).
- Make sure the data source's connection pool initial capacity is set to 0 (zero). To do this, select **Services**, **JDBC** and then **Datasources**. In the Datasources screen, click the **Datasource Name**, click the **Connection Pool** tab, and enter 0 (zero) in the **Initial Capacity** field.
- Ensure that an ONS daemon is running on your database servers at all times. To start the ONS daemon on a database server, run the `onsctl` command: `start`

Creating a GridLink Data Source

To create a GridLink Data Source:

1. Log in to the Oracle WebLogic Server Administration Console.
2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the **Domain Structure** tree, expand **Services**, then select **Data Sources**.
4. On the **Summary of Data Sources** page, click **New** and select **GridLink Data Source**.
5. Enter the following information:
 - a. Enter a logical name for the datasource in the Name field. For example, **gridlink**
 - b. Enter a name for JNDI. For example, **jdbc/gridlink**.
 - c. Click **Next**.
6. In the **Transaction Options** page, de-select **Supports Global Transactions**, and click **Next**.
7. Select **Enter individual listener information** and click **Next**.
8. Enter the following connection properties:
 - d. **Service Name**: Enter the service name of the database. For a **GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example, `<mydbservice>.us.oracle.com`

Note: The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example, `us.oracle.com`

- **Host Name** - The DNS name or IP address of the server that hosts the database. For an Oracle GridLink service-instance connection, this must be the same for each data source in a given multi data source.
- **Port** - The port on which the database server listens for connections requests.
- **Database User Name:** The database user name. For example, `myDataBase`.
- **Password:** The password. For example, `myPassword1`.
- Confirm Password and click **Next**.

Tip: For more information, see the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

The console automatically generates the complete JDBC URL. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=left) (PORT=1234)) (ADDRES
S=(PROTOCOL=TCP) (HOST=right) (PORT=1234)) (ADDRESS=(PROTOCOL
=TCP) (HOST=center) (PORT=1234))) (CONNECT_DATA=(SERVICE_
NAME=myService)))
```

9. On the **Test GridLink Database Connection** page, review the connection parameters and click **Test All Listeners**.

Oracle WebLogic attempts to create a connection from the Administration Server to the database. Results from the connection test appear at the top of the page. If the test is unsuccessful, correct any configuration errors and retry the test.

Click **Next**.

10. In the **ONS Client Configuration** page, do the following:

- Select **Fan Enabled** to subscribe to and process Oracle FAN events.
- In ONS host and port, enter a comma-separate list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.
- Click **Next**.

11. On the **Test ONS client configuration** page, review the connection parameters and click **Test All ONS Nodes**.

Click **Next**.

12. In the **Select Targets** page, select **Dept1_Cluster1** as the target and **All Servers in the cluster**.

13. Click **Finish**.

14. Click **Activate Changes**.

Creating a Multi Data Source

To create a multi data source:

1. Log into the WebLogic Server Administration Console using the Admin credentials.
2. In the Domain Structure window, expand the **Services** node then expand the **DataSource** node.
3. Click **Lock and Edit** in the Change Center.
4. Click **New** then click **Multi Data Sources**.
5. Enter `leasing` as the name.
6. Enter `jdbc/leasing` as the JNDI name.
7. Select **Failover as algorithm** (default).
8. Click **Next**.
9. Select the target cluster(s).
10. Click **Next**.
11. Select **non-XA driver** (the default).
12. Click **Next**.
13. Click **Create New Data Source**.
14. Enter `leasing-rac0` as the name. Enter `jdbc/leasing-rac0` as the JNDI name. Enter `oracle` as the database type. For the driver type, select Oracle Driver (Thin) for RAC server-Instance connection Version 10,11.

Note: When you create multi data sources for the leasing table, enter names in the format `<MultiDS>-rac0`, `<MultiDS>-rac1`, and so on.

15. Click **Next**.
16. Deselect **Supports Global Transactions**.
17. Click **Next**.
18. Enter the service name, database name (that is, the RAC Node instance name such as `racdb1`, `racdb2`), host port, and password for your leasing schema.
19. Click **Next**.
20. Click **Test Configuration** and verify that the connection works.
21. Click **Next**.
22. Target the data source to the cluster(s).
23. Select the data source and add it to the right screen.
24. Click **Create a New Data Source** for the second instance of your Oracle RAC database, target it to the cluster(s), repeating the steps for the second instance of your Oracle RAC database.
25. Add the second data source to your multi data source.
26. Save and click **Activate Changes**.

5.13.20.3 Editing Node Manager's Properties File

You must edit the `nodemanager.properties` file to add the following properties for each node where you configure server migration:

```
Interface=eth0
NetMask=255.255.255.0
UseMACBroadcast=true
```

- **Interface:** Specifies the interface name for the floating IP (such as `eth0`).

Note: Do not specify the sub interface, such as `eth0:1` or `eth0:2`. This interface is to be used without the `:0`, or `:1`. The Node Manager's scripts traverse the different `:X` enabled IPs to determine which to add or remove. For example, the valid values in Linux environments are `eth0`, `eth1`, or, `eth2`, `eth3`, `ethn`, depending on the number of interfaces configured.

Note: In Windows, the Interface should be set to the Network Interface Name. For example: `Interface="Local Area Connection"`.

- **NetMask:** Net mask for the interface for the floating IP. The net mask should be the same as the net mask on the interface; `255.255.255.0` is an example. The actual value depends on your network.
- **UseMACBroadcast:** Specifies whether or not to use a node's MAC address when sending ARP packets, that is, whether or not to use the `-b` flag in the `arping` command.

Verify in Node Manager's output (shell where Node Manager starts) that these properties are being used or problems may arise during migration. You should see an entry similar to the following in Node Manager's output:

```
...
StateCheckInterval=500
Interface=eth0
NetMask=255.255.255.0
...
```

Note: The following steps are not required if the server properties (start properties) are properly set and Node Manager can start the servers remotely.

1. Set the following property in the `nodemanager.properties` file:
 - **StartScriptEnabled:** Set this property to `true` to enable Node Manager to start the managed servers.
2. Start Node Manager on `OIMHOST1` and `OIMHOST2` by running the `startNodeManager.sh` script in the `WL_HOME/server/bin` directory.

Note: When running Node Manager from a shared storage installation, multiple nodes start using the same `nodemanager.properties` file. However, each node may require different `NetMask` or `Interface` properties. In this case, specify individual parameters on a per-node basis using environment variables. For example, to use a different interface (`eth3`) in `HOSTn`, use the `Interface` environment variable as follows: `HOSTn> export JAVA_OPTIONS=-DInterface=eth3` and start Node Manager after the variable is set in the shell.

5.13.20.4 Setting Environment and Superuser Privileges for the `wlsifconfig.sh` Script

To set environment and superuser privileges for the `wlsifconfig.sh` script:

Note: On Windows, the script is `wlsifconfig.cmd` and users with the Administrator privilege can run it.

1. Ensure that your `PATH` environment variable includes these files:

Table 5–9 Files Required for the `PATH` Environment Variable

File	Located in this directory
<code>wlsifconfig.sh</code>	<code>DOMAIN_HOME/bin/server_migration</code>
<code>wlscontrol.sh</code>	<code>WL_HOME/common/bin</code>
<code>nodemanager.domains</code>	<code>WL_HOME/common</code>

2. Grant sudo configuration for the `wlsifconfig.sh` script.
 - Configure sudo to work without a password prompt.
 - For security reasons, Oracle recommends restricting to the subset of commands required to run the `wlsifconfig.sh` script. For example, perform the following steps to set the environment and superuser privileges for the `wlsifconfig.sh` script:
 - Grant sudo privilege to the WebLogic user (`oracle`) with no password restriction, and grant execute privilege on the `/sbin/ifconfig` and `/sbin/arping` binaries.
 - Ensure that the script is executable by the WebLogic user. The following is an example of an entry inside `/etc/sudoers` granting sudo execution privilege for `oracle` and also over `ifconfig` and `arping`:

```
oracle ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
```

Note: Ask the system administrator for the sudo and system rights as appropriate to this step.

5.13.20.5 Configuring Server Migration Targets

You first assign all the available nodes for the cluster's members and then specify candidate machines (in order of preference) for each server that is configured with server migration. To configure cluster migration in a migration in a cluster:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand **Environment** and select **Clusters**.
3. Click the cluster you want to configure migration for in the Name column.
4. Click the **Migration** tab.
5. Click **Lock and Edit**.
6. In the **Available** field, select the machine to which to allow migration and click the right arrow.
7. Select the data source to use for automatic migration. In this case, select the leasing data source.
8. Click **Save**.
9. Click **Activate Changes**.
10. Set the candidate machines for server migration. You must perform this task for all managed servers as follows:

- a. In the Domain Structure window of the Oracle WebLogic Server Administration Console, expand **Environment** and select **Servers**.

Tip: Click **Customize this table** in the Summary of Servers page and move Current Machine from the Available window to the Chosen window to view the machine that the server runs on. This will be different from the configuration if the server migrates automatically.

- b. Select the server that you want to configure migration for.
- c. Click the **Migration** tab.
- d. In the **Available** field, located in the Migration Configuration section, select the machines you want to allow migration to and click the right arrow.
- e. Select **Automatic Server Migration Enabled**. This enables Node Manager to start a failed server on the target node automatically.
- f. Click **Save** then Click **Activate Changes**.
- g. Repeat the steps above for any additional managed servers.
- h. Restart the administration server, Node Managers, and the servers for which server migration has been configured.

5.13.20.6 Testing the Server Migration

To verify that server migration is working properly, follow these steps:

From SOAHOST1:

1. Force stop the WLS_SOA1 managed server.

To do this, run this command:

```
SOAHOST1> kill -9 <pid>
```

pid specifies the process ID of the managed server. You can identify the pid in the node by running this command:

```
SOAHOST1> ps -ef | grep WLS_SOA1
```

Note: For Windows, you can terminate the Managed Server by using the `taskkill` command. For example:

```
taskkill /f /pid <pid>
```

Where *<pid>* is the process Id of the Managed Server.

To determine the process Id of the Managed Server, run the following command and identify the pid of the WLS_SOA1 Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

-
2. Watch the Node Manager console: you should see a message indicating that WLS_SOA1's floating IP has been disabled.
 3. Wait for the Node Manager to try a second restart of WLS_SOA1. Node Manager waits for a fence period of 30 seconds before trying this restart.
 4. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.


From SOAHOST2:

1. Watch the local Node Manager console. After 30 seconds since the last try to restart WLS_SOA1 on Node 1, Node Manager on Node 2 should prompt that the floating IP for WLS_SOA1 is being brought up and that the server is being restarted in this node.
2. Access the soa-Service Infrastructure console in the same IP.

To verify migration in the Oracle WebLogic Server Administration Console:

1. Log into Oracle WebLogic Server Administration Console.
2. Click on **Domain** on the left pane.
3. Click on the **Monitoring** tab and then on the **Migration** sub-tab.

The Migration Status table provides information on migration status.

Figure 5–47 Migration Status Screen in Oracle WebLogic Server Administration Console

Note: On Windows, when you manually shut down multiple servers at the same time on the same machine and then, on another machine, attempt to start one of the servers that was shut down, the IP bind may not work. This happens because the original machine still has claim to the IP address even though `netsh` reports that the IP address has been removed.

To resolve this, you must check the network configuration either by using the `ipconfig` utility or Windows Network Configuration. Either of these may show that one of the virtual/floating IP addresses is still configured even though the servers are shut down. You can then use Windows Network Configuration to remove the IP address using the following procedure:

1. From Windows Control Panel, select **Network Connections**.
2. Select the appropriate network interface, right-click, and select **Properties**.
3. Select **Internet Protocol (TCP/IP)** and click the **Properties** button.
4. Select **Advanced**.
5. Select the appropriate IP address and click the **Remove** button.

Note: After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

5.13.21 Scaling the Topology

You can scale out or scale up this high availability topology. When you "scale up" the topology, you add new managed servers to nodes that are already running one or more managed servers. When you scale out the topology, you add new managed servers to new nodes.

5.13.21.1 Scaling Up the Topology (Adding Managed Servers to Existing Nodes)

In this case, you already have a node that runs a managed server configured with SOA components. The node contains a Middleware home, an Oracle HOME (SOA) and a domain directory for existing managed servers.

You can use the existing installations (the Middleware home, and domain directories) for creating new WLS_SOA servers. There is no need to install SOA binaries in a new location, or to run pack and unpack.

Follow these steps for scaling up the topology:

1. Using the Administration Console, clone WLS_SOA1 into a new managed server. The source managed server to clone should be one that already exists on the node where you want to run the new managed server.

To clone a managed server:

- a. Select **Environment** -> **Servers** from the Administration Console.
- b. Select the managed server that you want to clone (for example, WLS_SOA1).
- c. Select **Clone**.

Name the new managed server WLS_SOA n , where n is a number to identify the new managed server.

The rest of the steps assume that you are adding a new server to SOAHOST1, which is already running WLS_SOA1.

2. For the listen address, assign the virtual host name to use for this new managed server. If you are planning to use server migration as recommended for this server, this virtual host name will enable it to move to another node. The virtual host name should be different from the one used by the managed server that is already running.

Perform these steps to set the managed server listen address:

- a. Log into the Oracle WebLogic Server Administration Console.
- b. In the Change Center, click **Lock & Edit**.
- c. Expand the **Environment** node in the Domain Structure window.
- d. Click **Servers**. The Summary of Servers page is displayed.
- e. Select the managed server whose listen address you want to update in the **Names** column of the table. The Setting page for that managed server is displayed.
- f. Set the Listen Address to SOAHOST1VHN1.
- g. Click **Save**.
- h. Save and activate the changes.

The changes will not take effect until the managed server is restarted.

3. Update the cluster address to include the new server:

- a. Select **Environment -> Cluster** in the Domain Structure window in the Administration Console.
- b. Click on the SOA_Cluster server. The Settings Screen for the SOA_Cluster appears.
- c. In the Change Center, click **Lock & Edit**.
- d. Add the new server's address and port to the **Cluster Address** field. For example:

SOAHOST1VHN1:8001, SOAHOST2VHN1:8001, SOAHOST1VHN2:8001

4. Create JMS Servers for SOA, BPM (if applicable) and UMS on the new managed server.
 - a. Use the WebLogic Server Administration Console to create a new persistent store for the new SOAJMSJMServer and name it, for example, SOAJMSFileStore_N. Specify the path for the store. This should be a directory on shared storage, as recommended in [Section 5.13.1.3, "Shared Storage Prerequisites"](#):

Note: This directory must exist before the managed server is started or the start operation fails.

ORACLE_BASE/admin/DOMAIN_NAME/cluster_name/jms/soajmsfilestore_n

- b. Create a new JMS Server for SOA, for example, SOAJMSJMServer_N. Use the SOAJMSFileStore_N for this JMSServer. Target the SOAJMSJMServer_N Server to the recently created Managed Server (WLS_SOAn).
- c. Create a new persistence store for the new UMSJMSServer, for example, UMSJMSFileStore_N. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.13.1.3, "Shared Storage Prerequisites"](#):

ORACLE_BASE/admin/domain_name/cluster_name/jms/umsjmsfilestore_n

Note: This directory must exist before the managed server is started or the start operation fails.

You can also assign SOAJMSFileStore_N as store for the new UMS JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

- d. Create a new JMS Server for UMS, for example, UMSJMSServer_N. Use the UMSJMSFileStore_N for this JMSServer. Target the UMSJMSServer_N Server to the recently created Managed Server (WLS_SOAn).
- e. (ONLY FOR BPM Systems) Create a new persistence store for the new BPMJMSServer, for example, BPMJMSFileStore_N. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.13.1.3, "Shared Storage Prerequisites"](#):

ORACLE_BASE/admin/domain_name/cluster_name/jms/bpmjmsfilestore_n

Note: This directory must exist before the managed server is started or the start operation fails.

You can also assign SOAJMSFileStore_N as store for the new BPM JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

- f. (ONLY FOR BPM Systems) Create a new JMS Server for BPM, for example, BPMJMSServer_N. Use the BPMJMSFileStore_N for this JMSServer. Target the BPMJMSServer_N Server to the recently created Managed Server (WLS_SOAn).
- g. Update the SubDeployment targets for the SOA JMS Module to include the recently created SOA JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **SOAJMSModule** (represented as a hyperlink in the **Names** column of the table). The Settings page for SOAJMSModule appears. Click the **SubDeployments** tab. The subdeployment module for SOAJMS appears.

Note: This subdeployment module name is a random name in the form of SOAJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS_SOA1 and WLS_SOA2).

Click the **SOAJMSServerXXXXXX** subdeployment. Add the new JMS Server for SOA called SOAJMSServer_N to this subdeployment. Click **Save**.

- h. Update the SubDeployment targets for the UMSJMSSystemResource to include the recently created UMS JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **UMSJMSSystemResource** (represented as a hyperlink in the **Names** column of the table). The Settings page for UMSJMSSystemResource appears. Click the **SubDeployments** tab. The subdeployment module for UMSJMS appears.

Note: This subdeployment module name is a random name in the form of UCMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS_SOA1 and WLS_SOA2).

Click the **UMSJMSServerXXXXXX** subdeployment. Add the new JMS Server for UMS called UMSJMSServer_N to this subdeployment. Click **Save**.

- i. (ONLY FOR BPM Systems) Update the SubDeployment targets for BPMJMSModule to include the recently created BPM JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **BPMJMSSystemResource** (represented as a hyperlink in the **Names** column

of the table). The Settings page for BPMJMSSystemResource appears. Click the **SubDeployments** tab. The subdeployment module for BPMJMS appears.

Note: This subdeployment module name is a random name in the form of BPMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS_SOA1 and WLS_SOA2).

Click the **BPMJMSServerXXXXXX** subdeployment. Add the new JMS Server for BPM called BPMJMSServer_N to this subdeployment. Click **Save**.

5. Configure Oracle Coherence for deploying composites. See [Section 5.13.8, "Configuring Oracle Coherence for Deploying Composites."](#)

Note: Only the localhost field needs to be changed for the server. Replace the localhost with the listen address of the new server added, for example:

```
Dtangosol.coherence.localhost=SOAHOST1VHNn
```

6. Configure TX persistent store for the new server. This should be a location visible from other nodes as indicated in the recommendations about shared storage.

From the Administration Console, select **Server_name > Services** tab. Under **Default Store**, in **Directory**, enter the path to the folder where you want the default persistent store to store its data files.

7. Disable host name verification for the new managed server. Before starting and verifying the WLS_SOA_n managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in SOAHOST_n. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required (the hostname verification settings is propagated to the cloned server).

To disable host name verification:

- a. In the Oracle Enterprise Manager Console, select Oracle WebLogic Server Administration Console.
- b. Expand the **Environment** node in the Domain Structure window.
- c. Click **Servers**. The Summary of Servers page appears.
- d. Select WLS_SOA_n in the **Names** column of the table. The Settings page for the server appears.
- e. Click the **SSL** tab.
- f. Click **Advanced**.
- g. Set Hostname Verification to **None**.
- h. Click **Save**.
8. Start and test the new managed server from the Administration Console.
 - a. Shut down the existing managed servers in the cluster.
 - b. Ensure that the newly created managed server, WLS_SOA_n, is up.

- c. Access the application on the newly created managed server (`http://vip:port/soa-infra`). The application should be functional.
9. Configure Server Migration for the new managed server.

Note: Since this is a scale-up operation, the node should already contain a Node Manager and environment configured for server migration. The floating IP for the new SOA managed server should also be already present.

Configure server migration following these steps:

- a. Log into the Administration Console.
 - b. In the left pane, expand **Environment** and select **Servers**.
 - c. Select the name of the new managed server for which you want to configure migration.
 - d. Click the Migration tab.
 - e. In the **Available** field, in the "Migration Configuration" section, select the machines to which to allow migration and click the right arrow. Select the same migration targets as for the servers that already exist on the node.

For example, for new managed servers on SOAHOST1, which is already running WLS_SOA1, select SOAHOST2. For new managed servers on SOAHOST2, which is already running WLS_SOA2, select SOAHOST1.

Make sure the appropriate resources are available to run the managed servers concurrently during migration.
 - f. Select the "Automatic Server Migration Enabled" option. This enables the Node Manager to start a failed server on the target node automatically.
 - g. Click **Save**.
 - h. Restart the Administration Server, managed servers, and Node Manager.
10. Test server migration for this new server. Follow these steps from the node where you added the new server:
 - a. Stop the WLS_SOAn managed server.

To do this, run "kill -9 <pid>" on the PID of the managed server. You can identify the PID of the node using "ps -ef | grep WLS_SOAn".

Note: For Windows, the Managed Server can be terminated by using the taskkill command. For example:

```
taskkill /f /pid <pid>
```

Where <pid> is the process Id of the Managed Server.

To determine the process Id of the Managed Server, run the following command and identify the pid of the WLS_SOAn Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

- b. Watch the Node Manager Console: you should see a message indicating that WLS_SOAn's floating IP has been disabled.

- c. Wait for the Node Manager to try a second restart of WLS_SOAn. Node Manager waits for a fence period of 30 seconds before trying this restart.
- d. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

Note: After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

5.13.21.2 Scaling Out the Topology (Adding Managed Servers to New Nodes)

When you scale out the topology, you add new managed servers configured with SOA to new nodes.

Before performing the steps in this section, check that you meet these requirements:

- There must be existing nodes running managed servers configured with SOA within the topology.
- The new node can access the existing home directories for WebLogic Server and SOA. (Use the existing installations in shared storage for creating a new WLS_SOAn or WLS_WSM managed server. You do not need to install WebLogic Server or SOA binaries in a new location but you do need to run pack and unpack to bootstrap the domain configuration in the new node.)

Note: If there is no existing installation in shared storage, installing WebLogic Server and SOA in the new nodes is required. See [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines."](#)

Note: When an ORACLE_HOME or WL_HOME is shared by multiple servers in different nodes, Oracle recommends keeping the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. To update the oraInventory in a node and "attach" an installation in a shared storage to it, use `ORACLE_HOME/oui/bin/attachHome.sh`. To update the Middleware home list to add or remove a WL_HOME, edit the `user_home/boa/beahomelist` file. See the following steps.

Follow these steps for scaling out the topology:

1. On the new node, mount the existing Middleware home, which should include the SOA installation and the domain directory, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To attach ORACLE_HOME in shared storage to the local Oracle Inventory, execute the following command:

```
SOAHOSTn>cd ORACLE_BASE/product/fmw/soa/  
SOAHOSTn>./attachHome.sh -jreLoc ORACLE_BASE/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the *MW_HOME/boa/beahomelist* file and add *ORACLE_BASE/product/fmw* to it.

3. Log in to the Oracle WebLogic Administration Console.
4. Create a new machine for the new node that will be used, and add the machine to the domain.
5. Update the machine's Node Manager's address to map the IP of the node that is being used for scale out.
6. Use the WebLogic Server Administration Console to clone WLS_SOA1 into a new managed server. Name it WLS_SOA n , where n is a number and assign it to the new machine.

Note: These steps assume that you are adding a new server to node n , where no managed server was running previously.

7. For the listen address, assign the virtual host name to use for this new managed server. If you are planning to use server migration as recommended for this server, this virtual host name will enable it to move to another node. The virtual host name should be different from the one used by the managed server that is already running.

Perform these steps to set the managed server listen address:

- a. Log into the Oracle WebLogic Server Administration Console.
- b. In the Change Center, click **Lock & Edit**.
- c. Expand the **Environment** node in the Domain Structure window.
- d. Click **Servers**. The Summary of Servers page is displayed.
- e. Select the managed server whose listen address you want to update in the **Names** column of the table. The Setting page for that managed server is displayed.
- f. Set the Listen Address to SOAHOST n VHN1.
- g. Click **Save**.
- h. Save and activate the changes.

The changes will not take effect until the managed server is restarted.

8. Update the cluster address to include the new server:
 - a. Select **Environment -> Cluster** from the Administration Console.
 - b. Click on the SOA_Cluster server. The Settings Screen for the SOA_Cluster appears.
 - c. In the Change Center, click **Lock & Edit**.
 - d. Add the new server's address and port to the **Cluster Address** field. For example:

SOAHOST1VHN1:8001, SOAHOST2VHN1:8001, SOAHOST n VHN1:8001

9. Create JMS servers for SOA, BPM (if applicable), and UMS on the new managed server.

- a. Use the WebLogic Server Administration Console to create a new persistent store for the new SOAJMSJMServer and name it, for example, SOAJMSFileStore_N. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.13.1.3, "Shared Storage Prerequisites."](#)

ORACLE_BASE/admin/domain_name/cluster_name/jms/SOAJMSFileStore_N

Note: This directory must exist before the managed server is started or the start operation fails.

- b. Create a new JMS Server for SOA, for example, SOAJMSJMServer_N. Use the SOAJMSFileStore_N for this JMSJMServer. Target the SOAJMSJMServer_N Server to the recently created managed server (WLS_SOAn).
- c. Create a new persistence store for the new UMSJMSJMServer, and name it, for example, UMSJMSFileStore_N. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.13.1.3, "Shared Storage Prerequisites."](#)

ORACLE_BASE/admin/domain_name/cluster_name/jms/UMSJMSFileStore_N

Note: This directory must exist before the managed server is started or the start operation fails.

Note: It is also possible to assign SOAJMSFileStore_N as the store for the new UMS JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

- d. Create a new JMS Server for UMS: for example, UMSJMSJMServer_N. Use the UMSJMSFileStore_N for this JMS Server. Target the UMSJMSJMServer_N Server to the recently created managed server (WLS_SOAn).
- e. (ONLY FOR BPM Systems) Create a new persistence store for the new BPMJMSJMServer, and name it, for example, BPMJMSFileStore_N. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.13.1.3, "Shared Storage Prerequisites."](#)

ORACLE_BASE/admin/domain_name/cluster_name/jms/BPMJMSFileStore_N

Note: This directory must exist before the managed server is started or the start operation fails.

Note: It is also possible to assign SOAJMSFileStore_N as the store for the new BPM JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

- f. (ONLY FOR BPM Systems) Create a new JMS Server for BPM: for example, BPMJMSJMServer_N. Use the BPMJMSFileStore_N for this JMS Server. Target the

BPMJMSServer_N Server to the recently created managed server (WLS_SOAn).

- g. Update the SubDeployment targets for the SOA JMS Module to include the recently created SOA JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose JMS Modules from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **SOAJMSModule** (represented as a hyperlink in the Names column of the table). The Settings page for SOAJMSModule appears. Open the **SubDeployments** tab. The subdeployment module for SOAJMS appears.

Note: This subdeployment module name is a random name in the form of SOAJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS_SOA1 and WLS_SOA2).

Click the **SOAJMSServerXXXXXX** subdeployment. Add the new JMS Server for SOA called SOAJMSServer_N to this subdeployment. Click **Save**.

- h. Update the SubDeployment targets for UMSJMSSystemResource to include the recently created UMS JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose JMS Modules from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **UMSJMSSystemResource** (represented as a hyperlink in the Names column of the table). The Settings page for UMSJMSSystemResource appears. Open the **SubDeployments** tab. The subdeployment module for UMSJMS appears

Note: This subdeployment module is a random name in the form of UMSJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS_SOA1 and WLS_SOA2).

Click the **UMSJMSServerXXXXXX** subdeployment. Add the new JMS Server for UMS called UMSJMSServer_N to this subdeployment. Click **Save**.

- i. (ONLY FOR BPM Systems) Update the SubDeployment Targets for BPMJMSSModule to include the recently created BPM JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose JMS Modules from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **BPMJMSSystemResource** (represented as a hyperlink in the Names column of the table). The Settings page for BPMJMSSystemResource appears. Click the **SubDeployments** tab. The subdeployment module for BPMJMS appears.

Note: This subdeployment module is a random name in the form of BPMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS_SOA1 and WLS_SOA2).

Click the **BPMJMSXXXXXX** subdeployment. Add the new JMS Server for BPM called BPMJMSServer_N to this subdeployment. Click **Save**.

10. Run the pack command on SOAHOST1 to create a template pack as follows:

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
```

```
SOAHOST1> ./pack.sh -managed=true -domain=MW_HOME/user_
projects/domains/soadomain/
-template=soadomaintemplateScale.jar -template_name=soa_domain_templateScale
```

Run the following command on SOAHOST1 to copy the template file created to SOAHOSTN:

```
SOAHOST1> scp soadomaintemplateScale.jar oracle@SOAHOSTN:/ ORACLE_
BASE/product/fmw/soa/common/bin
```

Run the unpack command on SOAHOSTN to unpack the template in the managed server domain directory as follows:

```
SOAHOSTN> cd ORACLE_BASE/product/fmw/soa/common/bin
```

```
SOAHOSTN> ./unpack.sh -domain=ORACLE_BASE/product/fmw/user_
projects/domains/soadomain/ -template=soadomaintemplateScale.jar
```

11. Configure Oracle Coherence for deploying composites, as described in [Section 5.13.8, "Configuring Oracle Coherence for Deploying Composites."](#)

Note: Only the localhost field needs to be changed for the server. Replace the localhost with the listen address of the new server added, for example:

```
Dtangosol.coherence.localhost=SOAHOSTnVHN1
```

12. Configure TX persistent store for the new server. This should be a location visible from other nodes as indicated in the recommendations about shared storage.

From the Administration Console, select **Server_name > Services** tab. Under **Default Store**, in **Directory**, enter the path to the folder where you want the default persistent store to store its data files.

13. Disable host name verification for the new managed server. Before starting and verifying the WLS_SOAN managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in SOAHOSTn. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required (the hostname verification setting is propagated to the cloned server).

To disable host name verification:

- a. In the Oracle Enterprise Manager Console, select **Oracle WebLogic Server Administration Console**.
- b. Expand the **Environment** node in the Domain Structure window.
- c. Click **Servers**. The Summary of Servers page appears.
- d. Select **WLS_SOAn** in the Names column of the table.

The Settings page for server appears.

- e. Click the **SSL** tab.
- f. Click **Advanced**.

- g. Set Hostname Verification to **None**.
 - h. Click **Save**.
14. Start the Node Manager on the new node. To start the Node Manager, use the installation in shared storage from the existing nodes, and start Node Manager by passing the host name of the new node as a parameter as follows:
- ```
SOAHOSTN> WL_HOME/server/bin/startNodeManager <new_node_ip>
```
15. Start and test the new managed server from the Oracle WebLogic Server Administration Console:
- a. Shut down all the existing managed servers in the cluster.
  - b. Ensure that the newly created managed server, WLS\_SOAn, is running.
  - c. Access the application on the newly created managed server (<http://vip:port/soa-infra>). The application should be functional.
16. Configure server migration for the new managed server.

---

**Note:** Since this new node is using an existing shared storage installation, the node is already using a Node Manager and an environment configured for server migration that includes netmask, interface, wlsifconfig script superuser privileges. The floating IP for the new SOA Managed Server is already present in the new node.

---

Configure server migration following these steps:

- a. Log into the Administration Console.
- b. In the left pane, expand **Environment** and select **Servers**.
- c. Select the server (represented as hyperlink) for which you want to configure migration from the Names column of the table. The Setting page for that server appears.
- d. Click the **Migration** tab.
- e. In the Available field, in the Migration Configuration section, select the machines to which to allow migration and click the right arrow.

---

**Note:** Specify the least-loaded machine as the migration target for the new server. The required capacity planning must be completed so that this node has enough available resources to sustain an additional managed server.

---

- f. Select the **Automatic Server Migration Enabled** option. This enables the Node Manager to start a failed server on the target node automatically.
- g. Click **Save**.
- h. Restart the Administration Server, managed servers, and Node Manager.
- i. Test server migration for this new server. Follow these steps from the node where you added the new server:
  - 1. Abruptly stop the WLS\_SOAn managed server.

2. To do this, run "kill -9 <pid>" on the PID of the managed server. You can identify the PID of the node using "ps -ef | grep WLS\_SOAn".

---

**Note:** For Windows, the Managed Server can be terminated by using the taskkill command. For example:

```
taskkill /f /pid <pid>
```

Where <pid> is the process Id of the Managed Server.

To determine the process Id of the Managed Server run the following command and identify the pid of the WLS\_SOAn Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

3. Watch the Node Manager Console: you should see a message indicating that WLS\_SOAn's floating IP has been disabled.

4. Wait for the Node Manager to try a second restart of WLS\_SOAn. Node Manager waits for a fence period of 30 seconds before trying this restart.

5. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

---

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

---

## 5.14 Configuring High Availability for Oracle Service Bus, with SOA Service Infrastructure and Component Service Engines

The procedures described in this section include setting up Oracle Service Bus, and also other Service engines contained in the Oracle SOA Service Infrastructure system (such as Oracle BPEL/Oracle PM, Oracle Mediator, Oracle Human Workflow and Oracle Decision Services, as well as Oracle B2B and Oracle User Messaging Service).

This section includes these topics:

- [Section 5.14.1, "Preparing the Environment: Prerequisite Steps Before Setting up a SOA High Availability Configuration"](#)
- [Section 5.14.2, "Installing Oracle Fusion Middleware Home"](#)
- [Section 5.14.3, "Enabling VIP1 and VIP3 in SOAHOST1 and VIP2 and VIP4 in SOAHOST2"](#)
- [Section 5.14.4, "Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA and OSB Domain"](#)
- [Section 5.14.5, "Creating boot.properties for the Administration Server on SOAHOST1"](#)
- [Section 5.14.6, "Starting and Validating the Administration Server in SOAHOST1"](#)
- [Section 5.14.7, "Disabling Host Name Verification for the Administration Server and the WLS\\_SOAn/WLS\\_OSBN Managed Servers"](#)

- Section 5.14.8, "Configuring Oracle Coherence for Deploying Composites"
- Section 5.14.9, "Configure Oracle Coherence for the Oracle Service Bus Result Cache"
- Section 5.14.10, "Setting Connection Destination Identifiers for B2B Queues"
- Section 5.14.11, "Starting the System in SOAHOST1"
- Section 5.14.12, "Propagating the Domain Configuration to SOAHOST2, OSBHOST1, and OSBHOST2 with pack/unpack Utilities"
- Section 5.14.13, "Extracting XEngine Files in the Second Node"
- Section 5.14.14, "Starting the System in SOAHOST2, OSBHOST1, and OSBHOST2"
- Section 5.14.15, "Configuring Oracle HTTP Servers for the Administration Server, WLS\_SOAn, and WLS\_OSBN Managed Servers"
- Section 5.14.16, "Validating Access Through Oracle HTTP Server"
- Section 5.14.17, "Setting the Front End HTTP Host and Port"
- Section 5.14.18, "Configuring a Default Persistent Store for Transaction Recovery"
- Section 5.14.19, "Deploying Applications"
- Section 5.14.20, "Configuring Server Migration for the WLS\_SOA Servers"
- Section 5.14.21, "Configuring Server Migration for the WLS\_OSBN Servers"
- Section 5.14.22, "Scaling the Topology"

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---



---

**Note:** The architectures and deployment procedures defined in this guide enable simple clustered deployments. The procedures described in these chapters can be used as a building block to enable this and other similar high availability topologies for these Fusion Middleware components. It is also expected that production deployments will use other required procedures, such as associating security policies with a centralized LDAP server. For complete details of secured, multi-tiered architectures, and deployment procedures, please refer to the Enterprise Deployment Guide for the component you are configuring.

---

Figure 5–48 represents the example architecture that the configuration steps in this section create.



**Figure 5–48 Oracle SOA Service Infrastructure, Component Services Engines, and Oracle Service Bus High Availability Architecture**




Figure 5–48 shows a two-node Oracle Service Bus cluster running on two Oracle WebLogic Servers. The Oracle WebLogic Servers are frontended by Oracle HTTP

Servers, which load balance incoming requests. A load balancer front ends the system and distributes incoming requests from clients to the two Oracle HTTP Servers. This configuration uses an Oracle RAC database to store reporting data from the JMS reporting provider, and two phase commit logs using LLR with the JTA transaction manager (only used by the reporting data provider). It may also contain heartbeat data from servers in the cluster (used in the WebLogic Server whole server migration feature). The persistent store for transaction and JMS stores is on shared storage. Virtual IP addresses (VIPs) provide manual failover for the Administration Server and for Oracle Service Bus managed servers (for server migration).

---

**Note:** To keep the Oracle Service Bus high availability configuration figure above readable, the Oracle SOA Service Infrastructure high availability configuration is not shown. See [Figure 5–6](#) to view the Oracle SOA Service Infrastructure high availability figure.

---

For information about configuring virtual IPs for the Administration Server and configuring the Administration Server for high availability, see [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

## 5.14.1 Preparing the Environment: Prerequisite Steps Before Setting up a SOA High Availability Configuration

The following sections provide prerequisite steps before setting up an Oracle SOA Service Infrastructure high availability configuration.

For information about platform-specific commands, see the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

- [Section 5.14.1.1, "Database Prerequisites"](#)
- [Section 5.14.1.2, "VIP and IP Prerequisites"](#)
- [Section 5.14.1.3, "Shared Storage Prerequisites"](#)
- [Section 5.14.1.4, "Installing and Configuring an LDAP Provider"](#)
- [Section 5.14.1.5, "Synchronizing System Clocks"](#)
- [Section 5.14.1.6, "Terminology for Directories and Directory Environment Variables"](#)
- [Section 5.14.1.7, "Installing and Configuring the Database Repository"](#)
- [Section 5.14.1.8, "Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database"](#)
- [Section 5.14.1.9, "Configuring Virtual Server Names and Ports for the Load Balancer"](#)
- [Section 5.14.1.10, "Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2"](#)

### 5.14.1.1 Database Prerequisites

Oracle SOA Suite requires the presence of a supported database and schemas. Oracle Service Bus depends upon the database only for the JMS Reporting Provider and the OWSM Policy Manager.

To check if your database is certified or to see all certified databases, refer to the "Certified Databases" section in the Certification Document:

[http://www.oracle.com/technology/software/products/ias/files/fusion\\_certification.html](http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html)

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

#### 5.14.1.2 VIP and IP Prerequisites

As shown in [Table 5–10](#), you configure the Administration Server and the SOA and OSB managed servers to listen on different virtual IPs. This requires the provisioning of the corresponding VIP in the node and related host names in the DNS system in your network. Ensure that the different VIPs are available and are reachable from SOAHOST1, SOAHOST2, and the client machines before running the installation.

**Table 5–10 Virtual Hosts**

| Virtual IP | VIP Maps to... | Description                                                                                                                                                                                                                                                          |
|------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VIP0       | SOAHOST1VHN0   | SOAHOST1VHN0 is the virtual host name that is the listen address for the Administration Server and fails over with manual failover of the Administration Server. It is enabled on the node where the Administration Server process is running (SOAHOST1 by default). |
| VIP1       | SOAHOST1VHN1   | SOAHOST1VHN1 is the virtual host name that maps to the listen address for WLS_SOA1 and fails over with server migration of this managed server. It is enabled on the node where the WLS_SOA1 process is running (SOAHOST1 by default).                               |
| VIP2       | SOAHOST2VHN1   | SOAHOST2VHN1 is the virtual host name that maps to the listen address for WLS_SOA2 and fails over with server migration of this managed server. It is enabled on the node where the WLS_SOA2 process is running (SOAHOST2 by default).                               |
| VIP3       | OSBHOST1VHN1   | OSBHOST1VHN1 is the virtual host name that maps to the listen address for WLS_OSB1 and fails over with server migration of this managed server. It is enabled on the node where the WLS_OSB1 process is running (SOAHOST1 by default).                               |
| VIP4       | OSBHOST2VHN1   | OSBHOST2VHN1 is the virtual host name that maps to the listen address for WLS_OSB2 and fails over with server migration of this managed server. It is enabled on the node where the WLS_OSB2 process is running (SOAHOST2 by default).                               |

#### 5.14.1.3 Shared Storage Prerequisites

For proper recovery in case of failure, store both JMS and transaction logs in a location that is accessible to all the nodes that can resume the operations after a failure in a managed server. This requires a shared storage location that can be referenced by multiple nodes. [Table 5–11](#) lists the contents of shared storage.

**Table 5–11 Contents of Shared Storage**

| Server   | Type of Data | Vol in Shared Storage | Directory                                                        | Files                                                                                                 |
|----------|--------------|-----------------------|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| WLS_SOA1 | Tx Logs      | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/soa_cluster_<br>name/tlogs | Common location<br>(stores decided by<br>WebLogic Server)                                             |
| WLS_SOA2 | Tx Logs      | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/soa_cluster_<br>name/tlogs | Common location<br>(stores decided by<br>WebLogic Server)                                             |
| WLS_OSB1 | Tx Logs      | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/osb_cluster_<br>name/tlogs | Common location<br>(stores decided by<br>WebLogic Server)                                             |
| WLS_OSB2 | Tx Logs      | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/osb_cluster_<br>name/tlogs | Common location<br>(stores decided by<br>WebLogic Server)                                             |
| WLS_SOA1 | JMS Stores   | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/soa_cluster_<br>name/jms   | Common location but<br>Individual store per<br>server (for example:<br>SOAJMSStore1,<br>UMSJMSStore1) |
| WLS_SOA2 | JMS Stores   | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/soa_cluster_<br>name/jms   | Common location but<br>Individual store per<br>server (for example:<br>SOAJMSStore2,<br>UMSJMSStore2) |
| WLS_OSB1 | JMS Stores   | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/osb_cluster_<br>name/jms   | Common location but<br>Individual store per<br>server (for example:<br>WseeFileStore_auto_1)          |
| WLS_OSB2 | JMS Stores   | VOL1                  | ORACLE_<br>BASE/admin/domain<br>_name/osb_cluster_<br>name/jms   | Common location but<br>Individual store per<br>server (for example:<br>WseeFileStore_auto_2)          |

The shared storage can be a NAS or SAN device. Specifically for NFS mounted systems, different issues related to file locking and abrupt node failures have been detected. Check the *Oracle Fusion Middleware Release Notes* and with your storage vendor for the main recommended parameters to be used as mount options. The following is an example command based on a NAS device. Your options may be different from the ones specified in this section:

```
SOAHOST1> mount nasfiler:/vol/volX/FMWshared
MW_HOME -t nfs -o
rw,bg,hard,nointr,tcp,vers=3,timeo=300,rsz=32768,wsz=32768
```

#### 5.14.1.4 Installing and Configuring an LDAP Provider

For production environments, it is a mandatory requirement for Oracle SOA Suite high availability topologies to have an external LDAP policy store. For more information on the supported policy stores as well as instructions on configuring LDAP, see the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite*.

### 5.14.1.5 Synchronizing System Clocks

Synchronizing system clocks on each of the cluster nodes for high availability SOA deployments is required.

### 5.14.1.6 Terminology for Directories and Directory Environment Variables

The following list describes the directories and variables used in this chapter:

- **ORACLE\_BASE:** This environment variable and related directory path refers to the base directory under which Oracle products are installed.
- **MW\_HOME:** This environment variable and related directory path refers to the location where Fusion Middleware (FMW) resides.
- **WL\_HOME:** This environment variable and related directory path contains installed files necessary to host a WebLogic Server.
- **ORACLE\_HOME:** This environment variable and related directory path refers to the location where Oracle FMW SOA Suite or Oracle Service Bus is installed.
- **ORACLE\_COMMON\_HOME:** This environment variable and related directory path refers to the Oracle home that contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF).

There can be only one Oracle Common home within each Middleware home.

- **DOMAIN Directory:** This directory path refers to the location where the Oracle WebLogic Domain information (configuration artifacts) is stored.
- **ORACLE\_INSTANCE:** An Oracle instance contains one or more system components, such as Oracle Web Cache, Oracle HTTP Server, or Oracle Internet Directory. An Oracle instance directory contains updateable files, such as configuration files, log files, and temporary files.

The values used and recommended for consistency for this directories are:

**ORACLE\_BASE:**

*/u01/app/oracle*

**MW HOME (Application Tier):**

*ORACLE\_BASE/product/fmw*

**WL\_HOME:**

*MW\_HOME/wlserver\_10.3*

**ORACLE\_HOME:**

*MW\_HOME/soa*

**ORACLE\_COMMON\_HOME:**

*MW\_HOME/oracle\_common*

**Location for JMS file based stores and Tlogs:**

*ORACLE\_BASE/admin/domain\_name/soa\_cluster\_name/jms*

*ORACLE\_BASE/admin/domain\_name/soa\_cluster\_name/tlogs*

**Mount point is:**

*ORACLE\_BASE/admin/domain\_name/soa\_cluster\_name/*

Shared Storage location:

Shared storage location: `ORACLE_BASE/admin/domain_name/aserver`

### 5.14.1.7 Installing and Configuring the Database Repository

This section describes how to install and configure the database repository.

#### Oracle Clusterware

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

#### Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.
- When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

#### Oracle Real Application Clusters

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

You must install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database before you install the Oracle Fusion Middleware components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. You install RCU in its own, separate Middleware home.

Use the latest version of RCU to install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

#### Database Initialization Parameters

Ensure that the following initialization parameter is set to the required value. It is checked by Repository Creation Assistant.

**Table 5–12 Required Initialization Parameters**

| Parameter | Required Value | Parameter Class |
|-----------|----------------|-----------------|
| PROCESSES | 300 or greater | Static          |

To check the value of the initialization parameter using SQL\*Plus, you can use the SHOW PARAMETER command.

As the SYS user, issue the SHOW PARAMETER command as follows:

```
SQL> SHOW PARAMETER processes
```

Set the initialization parameter using the following command:

```
SQL> ALTER SYSTEM SET processes=300 SCOPE=SPFILE
```

Restart the database.

---

**Note:** The method that you use to change a parameter's value depends on whether the parameter is static or dynamic, and on whether your database uses a parameter file or a server parameter file. See the *Oracle Database Administrator's Guide* for details on parameter files, server parameter files, and how to change parameter values.

---

#### 5.14.1.8 Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database

Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and Oracle SOA schemas into a Real Application Clusters database before you install Oracle Fusion Middleware SOA components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about installing RCU.

**5.14.1.8.1 Running RCU** Run RCU to install the required metadata for Oracle Fusion Middleware 11g.

1. For Linux, go to `RCU_Home/bin` and use the following command:

```
./rcu
```

2. In the Welcome screen, click **Next**.
3. In the Create Repository screen, select **Create** to load component schemas into a database, click **Next**.
4. In the Database Connection Details screen, enter connection information for your database:
  - Database Type: Select **Oracle Database**.
  - Host Name: Enter the name of the node that is running the database. For an Oracle RAC database, specify the VIP name, or one of the node names as the host name: **SOADBHOST1VIRTUAL**.
  - Port: The port number for the database: **1521**
  - Service Name: Enter the service name of the database: **soaha.mycompany.com**
  - Username: **SYS**
  - Password: Enter the password for the SYS user.
  - Role: **SYSDBA**
5. Click **Next**.
6. If you receive the following warning message:

The database you are connecting is with non-UTF8 charset, if you are going to use this DB for multilingual support, you may have data loss. If you are not using for multilingual support you can continue, otherwise we strongly recommend using UTF-8 database.

7. Click **Ignore** or **Stop**.
8. In the Select Components screen, select **Create a New Prefix**, and enter a prefix to use for the database schemas, for example, **SOAHA**  
Write down the schema names so they are available in later procedures.
  - Select the following schemas:
    - Under AS Common Schemas, select **Metadata Services**.
    - Under SOA Service Infrastructure, select **SOA Service Infrastructure** and **User Messaging** (the OSB Reporting Provider schema is part of SOA Service Infrastructure).
9. Click **Next**.
10. In the Schema Passwords screen, enter passwords for the main and additional (auxiliary) schema users, and click **Next**.
11. In the Map Tablespaces screen, choose the tablespaces for the selected components, and click **Next**.
12. In the Summary screen, click **Create**.
13. In the Completion Summary screen, click **Close**.

See the *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about using RCU.

**5.14.1.8.2 Configuring SOA Schemas for Transactional Recovery Privileges** You need the appropriate database privileges to allow the Oracle WebLogic Server transaction manager to query for transaction state information and issue the appropriate commands, such as commit and rollback, during recovery of in-flight transactions after a WebLogic Server container crash.

To configure the SOA schemas for transactional recovery privileges:

1. Log on to SQL\*Plus as a user with sysdba privileges. For example:  

```
sqlplus "/ as sysdba"
```
2. Grant select on sys.dba\_pending\_transactions to **soaha\_soainfra**.
3. Grant force any transaction to **soaha\_soainfra**.

---

**Note:** These privileges should be granted to the owner of the soainfra schema, as determined by the RCU operations.

---

#### 5.14.1.9 Configuring Virtual Server Names and Ports for the Load Balancer

This section describes the load balancer prerequisites for deploying an Oracle SOA Suite high availability environment.

##### Load Balancers

Oracle SOA Suite uses a hardware load balancer when deployed in a high availability configuration with two Oracle HTTP Servers as web tier components. The hardware load balancer should have the following features:

- Ability to load-balance traffic to a pool of real servers through a virtual host name: Clients access services using the virtual host name (instead of using actual host



names). The load balancer can then load balance requests to the servers in the pool.

- Port translation configuration: The load balancer should have the ability to perform port translation, where it allows incoming requests received on one port to be routed to a server process running on a different port. For example, a request received on port 80 can be routed to port 7777.
- Monitoring of ports (HTTP, HTTPS)
- Virtual servers and port configuration: Ability to configure virtual server names and ports on your external load balancer, and the virtual server names and ports must meet the following requirements:
  - The load balancer should allow configuration of multiple virtual servers. For each virtual server, the load balancer should allow configuration of traffic management on more than one port. For example, the load balancer typically needs to be configured with a virtual server and ports for HTTP and HTTPS traffic.
  - The virtual server names must be associated with IP addresses and be part of your DNS. Clients must be able to access the load balancer through the virtual server names.
- Ability to detect node failures and immediately stop routing traffic to the failed node.
- Resource monitoring/port monitoring/process failure detection: The load balancer must be able to detect service and node failures (through notification or some other means) and to stop directing non-Oracle Net traffic to the failed node. If your load balancer has the ability to automatically detect failures, you should use it.
- Fault-tolerant mode: It is highly recommended that you configure the load balancer to be in fault-tolerant mode.
- Other: It is highly recommended that you configure the load balancer virtual server to return immediately to the calling client when the back-end services to which it forwards traffic are unavailable. This is preferred over the client disconnecting on its own after a timeout based on the TCP/IP settings on the client machine.
- Sticky routing capability: Ability to maintain sticky connections to components based on cookies or URL.
- SSL acceleration: SSL acceleration is a method of offloading the processor-intensive public key encryption algorithms involved in SSL transactions to a hardware accelerator.

This feature is recommended but not required.

[Table 5–13](#) shows some example virtual server names to use for the external load balancer in the Oracle SOA Suite high availability environment.

**Table 5–13 Virtual Server Names for the External Load Balancer**

| Component                                           | Virtual Server Name |
|-----------------------------------------------------|---------------------|
| Oracle SOA Suite                                    | soa.mycompany.com   |
| WebLogic Server Administration Console              | admin.mycompany.com |
| Oracle Enterprise Manager Fusion Middleware Control | admin.mycompany.com |

**Table 5–13 (Cont.) Virtual Server Names for the External Load Balancer**

| Component                      | Virtual Server Name |
|--------------------------------|---------------------|
| Oracle Service Bus Web Console | admin.mycompany.com |

**Virtual Server Names**

This section describes the virtual server names that should be set up for the high availability deployments described in this chapter.

**soa.mycompany.com**

soa.mycompany.com is a virtual server name that acts as the access point for all HTTP traffic to the runtime SOA components, such as soa-infra and workflow. Traffic to the both SSL and non-SSL ports is configured and typically non-SSL is redirected to SSL. Clients access this service using the address soa.mycompany.com:443. This virtual server is defined on the load balancer.

**admin.mycompany.com**

This virtual server acts as the access point for all internal HTTP traffic that gets directed to the administration services. The incoming traffic from clients is non-SSL enabled. Thus, the clients access this service using the address admin.mycompany.com:80 and in turn forward these to ports 7777 on WEBHOST1 and WEBHOST2. The services accessed on this virtual host include the WebLogic Administration Server Console, Oracle Enterprise Manager, and the Oracle Service Bus Web Console.

In addition, ensure that the virtual server names are associated with IP addresses and are part of your Domain Name System (DNS). The computers on which Oracle Fusion Middleware is running must be able to resolve these virtual server names.

**5.14.1.10 Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2**

To install Oracle HTTP Server on WEBHOST1:

1. Verify that the servers meet the following requirements:
  - The system, patch, kernel, and other requirements meet the requirements specified in the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite*.
  - Port 7777 is not used by any service on the nodes. You can verify this by running the following command:

Unix:

```
netstat -an | grep LISTEN | grep "7777"
```

Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7777"
```

If the ports are in use, make them available.

2. On UNIX platforms, if the /etc/oraInst.loc or /var/opt/oracle/oraInst.loc file exists, check that its contents are correct. Specifically, check that the inventory directory is correct, and that you have write permissions for that directory.

If the /etc/oraInst.loc file does not exist, skip this step.

3. Start Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:

For UNIX, run this command: `runInstaller`.

For Windows, double-click **setup.exe**.

4. In the Welcome screen, click **Next**.
5. In the Select Installation Type screen, select **Install and Configure**, and click **Next**.
6. In the Prerequisite Checks screen, ensure that all the prerequisites are met, and click **Next**.
7. In the Specify Installation Location screen, set the location to:  
`/u01/app/oracle/product/11.1.1/ohs_1`
8. Click **Next**.
9. In the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Do not select **Associate Selected Components with WebLogic Domain**.
  - Click **Next**.
10. In the Specify Component Details screen, enter the following values:
  - Instance Home Location: Instance Home
  - Instance Home Location: `/u01/app/oracle/product/11.1.1/ohs_1/instances/ohs_instance1`
  - Instance Name: `ohs_instance1`
  - OHS Component Name: `ohs1`
11. Click **Next**.
12. In the Specify Web Tier Port Details screen:
  - Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File**, and use the **Browse** function to select the file.
  - Enter the **Oracle HTTP Server port**, for example, `7777`.
13. Click **Next**.
14. In the Configuration Summary screen, ensure that the selections are correct and click **Install**.
15. In the Installation Progress screen:

For UNIX systems, a dialog box appears prompting you to run the `oracleRoot.sh` script. Open a command window and run the script, following the prompts.

Click **Next**.

In the Configuration screen, several configuration assistants launch in succession. When the configuration assistants are finished, the Configuration Completed screen appears.
16. In the Configuration Completed screen, click **Finish** to exit.

Repeat the steps for WEBHOST2 and configure your LBR with a pool containing both the WEBHOST1 and WEBHOST2 addresses.

**5.14.1.10.1 Validating Oracle HTTP Server** To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URL in a browser:

**HTTP://WEBHOST1:7777/**

If Oracle HTTP Server is set up correctly, the Oracle FMW Welcome screen appears in the browser.

## 5.14.2 Installing Oracle Fusion Middleware Home

This section describes the procedure for installing Oracle Fusion Middleware on all nodes in the application tier that run Oracle WebLogic Server and Oracle SOA Suite. Repeat the procedures (described below for SOAHOST1) for installing WebLogic Server and Oracle SOA in SOAHOST2. The directory paths for binary files and domains used when installing new nodes must be exactly the same as those used for first node. If these paths and domains are not exactly the same as those used for the first node, failover will not work properly.

Install the following Oracle Fusion Middleware components:

- Oracle WebLogic Server (see [Section 5.14.2.1, "Installing Oracle WebLogic Server"](#))
- Oracle SOA Suite (see [Section 5.14.2.2, "Installing Oracle Fusion Middleware for Oracle SOA"](#))
- Oracle Service Bus (see [Section 5.14.2.3, "Installing Oracle Service Bus"](#))

### 5.14.2.1 Installing Oracle WebLogic Server

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

To install Oracle WebLogic Server on all nodes in the application tier:

1. Start Oracle WebLogic Server Installer.
2. In the Welcome screen, click **Next**.
3. In the Choose Middleware Home Directory screen:
  - Select **Create a New Middleware Home**.
  - For the **Middleware Home Directory** field, enter **MW\_HOME**.
  - Click **Next**.
4. In the Register for Security Updates screen, enter your contact information for security update notifications, and click **Next**.
5. In the Choose Install Type screen, select **Custom**, and click **Next**.
6. In the Choose Products and Components screen, click **Next**.
7. In the JDK Selection screen, select only **Oracle JRockit 160\_20\_D1.0.1-2124 SDK**, and click **Next**.
8. In the Choose Product Installation Directories screen, accept the following directory:  
  
WL\_HOME  
  
Click **Next**.
9. In the Installation Summary screen, click **Next**.

10. In the Installation Complete screen, deselect **Run QuickStart**, and click **Done**.

#### 5.14.2.2 Installing Oracle Fusion Middleware for Oracle SOA

On Linux platforms, if the `/etc/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, you can skip this step.

1. Start Oracle Fusion Middleware 11g Oracle SOA Suite Installer:

On UNIX:

```
SOAHOST1> runInstaller
```

On Windows:

```
SOAHOST1> setup.exe
```

When Oracle Fusion Middleware 11g Oracle SOA Suite Installer prompts you for a **JRE/JDK location** enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `MW_HOME/jrockit_160_<version>`.

2. In the Specify Inventory Directory screen, do the following:
  - a. Enter `HOME/oraInventory`, where `HOME` is the home directory of the user performing the installation (this is the recommended location).
  - b. Enter the OS group for the user performing the installation.
  - c. Click **Next**.
  - d. Follow the instructions on the screen to execute `/createCentralInventory.sh` as root.
  - e. Click **OK**.
3. In the Welcome screen, click **Next**.
4. In the Prerequisite Checks screen, verify that the checks complete successfully, and click **Next**.
5. In the Specify Installation Location screen:
  - For Middleware Home, enter `MW_HOME`.
  - For Oracle Home Directory, enter `soa`.Click **Next**.
6. In the Installation Summary screen, click **Install**.
7. In the Installation Complete screen, click **Finish**.

#### 5.14.2.3 Installing Oracle Service Bus

To install Oracle Service Bus, follow these steps:

1. Start the Oracle Fusion Middleware 11g Oracle Service Bus Installer:

On UNIX:

```
SOAHOST1> runinstaller
```

On Windows:

```
SOAHOST1> setup.exe
```

When Oracle Fusion Middleware 11g Oracle Service Bus Installer prompts you for a JRE/JDK location enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `MW_HOME/jrockit_160_<version>`.

2. In the Welcome screen, click **Next**.
3. In the Install Software Updates screen, enter your preferences for updates.
4. In the Specify Installation Location screen:
  - For **Middleware Home**, enter `MW_HOME`.
  - For **Oracle Home Directory**, enter `osb`.Click **Next**.
5. In the Installation Type screen, select **Custom** and click **Next**.
6. In the Components to Install screen, deselect **Oracle Service Bus IDE** and **Oracle Service Bus Examples**.
7. In the Prerequisite Checks screen, verify that the checks complete successfully, and click **Next**.
8. In the Product Home Location screen, accept the default and click **Next**.
9. In the Installation Summary screen, click **Install**.
10. In the Installation Complete screen, click **Finish**.

---

**Note:** Before you run the Configuration Wizard by following the instructions in [Section 5.14.4, "Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA and OSB Domain,"](#) ensure that you applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home, so that you have the latest version of Oracle Fusion Middleware.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

---

### 5.14.3 Enabling VIP1 and VIP3 in SOAHOST1 and VIP2 and VIP4 in SOAHOST2

The SOA domain uses virtual hostnames as the listen addresses for the SOA and OSB managed servers. You must enable VIP mapping for each of these hostnames on the two SOA Machines, (VIP1 and VIP3 on SOAHOST1 and VIP2 and VIP4 on SOAHOST2), and must correctly resolve the virtual hostnames in the network system used by the topology (either by DNS Server or by hosts resolution).

For information about configuring virtual IPs, see [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster,"](#) where an example is provided for the Administration Server. Refer to [Section 5.14.20, "Configuring Server Migration for the WLS\\_SOA Servers"](#) for more details on configuring server migration for the SOA servers.

### 5.14.4 Running Oracle Fusion Middleware Configuration Wizard on SOAHOST1 to Create the SOA and OSB Domain

Run Oracle Fusion Middleware Configuration Wizard from the `MW_HOME/oracle_common/common/bin` directory to create a domain containing the Administration Server and Oracle SOA and OSB components. Ensure that the database where you installed the repository is running. For Oracle RAC databases, all instances must be running.

---

**Note:** Multiple SOA clusters are not allowed in the same WebLogic Server domain.

Multiple OSB clusters are not allowed in the same WebLogic Server domain.

One single cluster holding both SOA and OSB is not allowed in a single WebLogic Server domain.

---

1. Change the directory to the location of the Oracle Fusion Middleware Configuration Wizard, located in the Middleware home, `oracle_common/common/bin` directory:

```
SOAHOST1> cd ORACLE_HOME/oracle_common/common/bin
```

2. Start Oracle Fusion Middleware Configuration Wizard:

For Linux:

```
SOAHOST1> ./config.sh
```

For Windows:

```
SOAHOST1> config.cmd
```

3. In the Welcome screen, select **Create a New WebLogic Domain**, and click **Next**.
4. In the Select Domain Source screen:
  - Select **Generate a domain configured automatically to support the following products**.
  - Select the following products:
    - **Basic Weblogic Server Domain - 10.3.6.0 [wlserver\_10.3]** (Selected by default and grayed out)
    - **Oracle SOA Suite - 11.1.1.0 [soa]** (Selected by default)
    - **Oracle Service Bus OWSM Extension - 11.1.1.6 [osb]**
    - **Oracle Enterprise Manager - 11.1.1.0 [oracle\_common]** (Selected by default)
    - **Oracle Service Bus - 11.1.1.6 [osb]**
    - **WebLogic Advanced WebService for JAX-RPC Extension - 10.3.6.0 [wlserver\_10.3]**
    - **Oracle WSM Policy Manager - 11.1.1.0 [oracle\_common]** (Selected by default when selecting SOA Suite)
    - **Oracle JRF Web Services - 11.1.1.0 [oracle\_common]** (Selected by default when selecting SOA Suite)

If you accidentally deselect some of the targets, ensure that the selections above are made in this screen.

Click **Next**.

5. In the Specify Domain Name and Location screen, make the following entries:

- Domain Name: **soadomain**
- Domain Location: accept the default
- Application Location: accept the default

Click **Next**.

6. In the Configure Administrator Username and Password screen, enter the username and password to be used for the domain's administrator, and click **Next**.
7. In the Configure Server Start Mode and JDK screen, make the following selections:

- WebLogic Domain Startup Mode: select **Production Mode**
- JDK Selection: select **Oracle JRockit 160\_20\_D1.0.1-2124 SDK**.

Click **Next**.

8. In the Configure JDBC Component Schema screen:

- a. Select all the component schemas that appear in the table at the bottom: **SOA Infrastructure**, **User Messaging Service**, **OWSM MDS Schema**, **SOA MDS Schema**, and **OSB JMS Reporting Provider**.
- b. For the RAC configuration for component schemas, select one of these options: **Convert to GridLink**, **Convert to RAC multi data source**, or **Don't convert**.
- c. Ensure that the following data source appears on the screen. The user names shown in [Table 5–14](#) assume that **soaha** was used as the prefix for schema creation from RCU.

For more information about Gridlink Data Sources configuration with Oracle RAC and the MDS repository, see [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC"](#).

For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

**Table 5–14 Configuring Values for Data Sources**

| <b>Data Source Schema</b>  | <b>Schema Owner</b> |
|----------------------------|---------------------|
| SOA Infrastructure         | SOAHA_SOAINFRA      |
| User Messaging Service     | SOAHA_ORASDPM       |
| OWSM MDS Schema            | SOAHA_MDS           |
| SOA MDS Schema             | SOAHA_MDS           |
| OSB JMS Reporting Provider | SOAHA_SOAINFRA      |

- d. Click **Next**.

9. In the Configure JDBC Component Schema screen, enter values for the following fields, specifying the connect information for the Oracle RAC database that was seeded with RCU.

**For multi data sources**, enter values for these fields:



**Figure 5–49 Configure RAC Multi Data Source Component Schema Screen**

**Fusion Middleware Configuration Wizard**

### Configure RAC Multi Data Source Component Schema

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for RAC Service-Instance c...

Service Name: soaha.mycompany.com

Username: Varies among component schemas

Password: \*\*\*\*\*

| Host Name               | Instance Name | Port |
|-------------------------|---------------|------|
| soadbhost1-virtual.myco | soahadb1      | 1521 |
| soadbhost2-virtual.myco | soahadb2      | 1521 |

Add Delete

| Multi Data Source Schema                                   | Service Name          | Schema Owner   | Schema Password |
|------------------------------------------------------------|-----------------------|----------------|-----------------|
| <input type="checkbox"/> SOA Infrastructure                | soahadb.us.oracle.com | SOAHA_SOAINFRA | *****           |
| <input checked="" type="checkbox"/> User Messaging Service | soahadb.us.oracle.com | SOAHA_ORASDPM  | *****           |
| <input type="checkbox"/> OWSM MDS Schema                   | soahadb.us.oracle.com | SOAHA_MDS      | *****           |
| <input type="checkbox"/> SOA MDS Schema                    | soahadb.us.oracle.com | SOAHA_MDS      | *****           |
| <input type="checkbox"/> OSB JMS Reporting Provider        | soahadb.us.oracle.com | SOAHA_SOAINFRA | *****           |

Exit Help Previous Next

For GridLink data sources, enter values for these fields:

**Figure 5–50 Configure GridLink RAC Component Schema Screen**

**Fusion Middleware Configuration Wizard**

### Configure GridLink RAC Component Schema

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for GridLink Connections; Versio...

Service Name: soaha.mycompany.com

Username: SOAHA\_ORASDPM

Password: \*\*\*\*\*

Enable FAN: ☒

Enable SSL: ☐

Wallet File: Enter a value Browse

Wallet Password: Enter a value

| Service Listener            | Port | Protocol |
|-----------------------------|------|----------|
| soadbhost12-scn.mycompany.c | 1521 | TCP      |

| ONS Host                      | Port |
|-------------------------------|------|
| soadbhost12-scn.mycompany.com | 6200 |

Add Delete

| RAC Component Schema                                       | Service Name        | Schema Owner   | Schema Password |
|------------------------------------------------------------|---------------------|----------------|-----------------|
| <input type="checkbox"/> SOA Infrastructure                | soaha.mycompany.com | SOAHA_SOAINFRA | *****           |
| <input checked="" type="checkbox"/> User Messaging Service | soaha.mycompany.com | SOAHA_ORASDPM  | *****           |
| <input type="checkbox"/> OWSM MDS Schema                   | soaha.mycompany.com | SOAHA_MDS      | *****           |
| <input type="checkbox"/> SOA MDS Schema                    | soaha.mycompany.com | SOAHA_MDS      | *****           |
| <input type="checkbox"/> OSB JMS Reporting Provider        | soaha.mycompany.com | SOAHA_SOAINFRA | *****           |

Exit Help Previous Next

- a. Driver: For RAC: Select Oracle driver (Thin) for RAC Service-Instance connections, Versions:10, 11. For GridLink: Oracle Driver (Thin) for GridLink Connections, Versions:10 and later

- b. Service Name: Enter the service name of the database. **For a GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example,  
`<mydbservice>.us.oracle.com`

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example,  
`us.oracle.com`

---

- c. Username prefix: Enter the prefix for the schemas. The user names shown in [Table 5–14](#) assume that **soaha** was used as prefix for schema creation from RCU.
- d. Password and Confirm Password: Enter the password for access to the schemas.
- e. **For a GridLink data source**, enter the SCAN address in the Listener Address field and SCAN port in the Port field. Enter the ONS host and port information in the ONS Host and Port fields, respectively.

---

**Note:** Oracle recommends that you use SCAN addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

---

**For a multi data source**, click **Add**. Enter the details for Oracle RAC instances.

- f. Update each schema: select one data source at a time and add the appropriate details.

Ensure that the information is entered for all the multi data source schemas: **SOA Infrastructure**, **User Messaging Service**, **OWSM MDS Schema**, **SOA MDS Schema**, and **OSB JMS Reporting Provider**.

- g. Click **Next**.
10. In the Test JDBC Data Sources screen, the connections are tested automatically. The **Status** column displays the results. Ensure that all connections were successful. If not, click **Previous** to correct your entries.

Click **Next** when all the connections are successful.

11. In the Select Optional Configuration screen, select the following:

- **Administration Server**
- **JMS Distributed Destinations**
- **Managed Servers, Clusters and Machines**
- **JMS File Store**

12. In the Configure the Administration Server screen, enter the following values:

- Name: **AdminServer**

- Listen Address: Enter the hostname for the VIP0 virtual IP
- Listen Port: **7001**
- SSL listen port: N/A
- SSL enabled: leave unchecked

Click **Next**.

13. Verify that the JMS Distributed Destination Type is set to **UDD**.

14. In the Configure Managed Servers screen, add the following managed servers:

**Table 5–15 Configuring Managed Servers**

| Name     | Listen Address | Listen Port | SSL Listen Port | SSL Enabled |
|----------|----------------|-------------|-----------------|-------------|
| WLS_SOA1 | SOAHOST1VHN1   | 8001        | n/a             | No          |
| WLS_SOA2 | SOAHOST2VHN1   | 8001        | n/a             | No          |
| WLS_OSB1 | OSBHOST1VHN1   | 8011        | n/a             | No          |
| WLS_OSB2 | OSBHOST2VHN1   | 8011        | n/a             | No          |

Do not delete any server that appears. You can modify the servers. If you delete a server and add a new one, targeting fails.

---

**Note:** The standard recommendation is to run custom applications and other systems in a separate WebLogic Managed Server. However, this section does not cover creating custom WLS managed servers as [Figure 5–48](#) describes.

---

Click **Next**.

15. In the Configure Clusters screen, add the following clusters:

**Table 5–16 Adding the SOA and OSB Clusters**

| Name        | Cluster Messaging Mode | Multicast Address | Multicast Port | Cluster Address                         |
|-------------|------------------------|-------------------|----------------|-----------------------------------------|
| SOA_Cluster | unicast                | N/A               | N/A            | SOAHOST1VHN1:8001,<br>SOAHOST2VHN1:8001 |
| OSB_Cluster | unicast                | N/A               | N/A            | OSBHOST1VHN1:8011,<br>OSBHOST2VHN1:8011 |

Click **Next**.

16. In the Assign Servers to Clusters screen, assign the following servers to SOA\_Cluster:

- WLS\_SOA1
- WLS\_SOA2

Assign the following servers to the OSB\_Cluster:

- WLS\_OSB1
- WLS\_OSB2

Click **Next**.

17. In the Configure Machines screen:

- Delete the **LocalMachine** that appears by default.
- Click the **Unix Machine** tab, and add the following machines:


**Table 5–17 Configuring Machines**

| Name     | Node Manager Listen Address |
|----------|-----------------------------|
| SOAHOST1 | Hostname of SOAHOST1        |
| SOAHOST2 | Hostname of SOAHOST2        |
| OSBHOST1 | Hostname of OSBHOST1        |
| OSBHOST2 | Hostname of OSBHOST2        |

Leave all other fields as their default values, and click **Next**.

18. In the Assign Servers to Machines screen, assign servers to machines as follows:

**Figure 5–51 Assign Servers to Machines Screen**



- SOAHOST1: **AdminServer**, **WLS\_SOA1**
- SOAHOST2: **WLS\_SOA2**
- OSBHOST1: **WLS\_OSB1**
- OSBHOST2: **WLS\_OSB2**

Click **Next**.

19. In the Configure JMS Files Stores screen, enter a directory location for each store that is accessible from all nodes that the servers can be failed over to. As recommended in [Section 5.14.1.3, "Shared Storage Prerequisites,"](#) use a directory like the following on a shared device:

**Table 5–18 Configuring JMS File Stores**

| <b>Name</b>            | <b>Directory</b>                                                      | <b>Synchronous Write Policy*</b> |
|------------------------|-----------------------------------------------------------------------|----------------------------------|
| FileStore_auto_1       | /u01/app/oracle/admin/soadomain<br>/soacluster/FileStore_auto_1       | Direct-Write                     |
| FileStore_auto_2       | /u01/app/oracle/admin/soadomain<br>/soacluster/FileStore_auto_2       | Direct-Write                     |
| SOAJMSFileStore_auto_1 | /u01/app/oracle/admin/soadomain<br>/soacluster/SOAJMSFileStore_auto_1 | Direct-Write                     |
| SOAJMSFileStore_auto_2 | /u01/app/oracle/admin/soadomain<br>/soacluster/SOAJMSFileStore_auto_2 | Direct-Write                     |
| UMSJMSFileStore_auto_1 | /u01/app/oracle/admin/soadomain<br>/soacluster/UMSJMSFileStore_auto_1 | Direct-Write                     |
| UMSJMSFileStore_auto_2 | /u01/app/oracle/admin/soadomain<br>/soacluster/UMSJMSFileStore_auto_2 | Direct-Write                     |
| WseeFileStore_auto_1   | /u01/app/oracle/admin/soadomain<br>/soacluster/WseeFileStore_auto_1   | Direct-Write                     |
| WseeFileStore_auto_2   | /u01/app/oracle/admin/soadomain<br>/soacluster/WseeFileStore_auto_2   | Direct-Write                     |

\* The default synchronous write policy (when none is specified) is Direct-Write.

Click **Next**.

**20.** In the Configuration Summary screen, click **Create**.

**21.** In the Creating Domain screen, click **Done**.

### 5.14.5 Creating boot.properties for the Administration Server on SOAHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a boot.properties file for the Administration Server on SOAHOST1.

For the Administration Server:

1. Create the following directories:

```
SOAHOST1> mkdir -p ORACLE_BASE/product/fmw/user_
projects/domains/soadomain/servers/AdminServer/security
```

2. Use a text editor to create a file called boot.properties in the security directory created in the previous step and enter the following lines in the file:

```
username=adminuser
password=password
```

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible in order for the entries to be encrypted.

---

## 5.14.6 Starting and Validating the Administration Server in SOAHOST1

This section describes procedures for starting and validating the Administrator Server in SOAHOST1

### 5.14.6.1 Starting the Administration Server on SOAHOST1

To start the Administration Server on SOAHOST1, run the following commands:

```
SOAHOST1> cd MW_HOME/user_projects/domains/soadomain/bin
```

```
SOAHOST1> ./startWebLogic.sh
```

### 5.14.6.2 Validating the Administration Server

To verify that the Administration Server is properly configured:

1. In a browser, go to <http://vip0:7001/console>.
2. Log in as the administrator.
3. Verify that the WLS\_SOA1 and WLS\_SOA2 managed servers are listed. Also verify that the WLS\_OSB1 and WLS\_OSB2 managed servers are listed.
4. Verify that the SOA\_Cluster cluster is listed. Also verify that the OSB\_Cluster cluster is listed.
5. Verify that you can access Enterprise Manager at <http://vip0:7001/em>.
6. Verify that you can access the Oracle Service Bus Console at <http://vip0:7001/sbconsole>.

## 5.14.7 Disabling Host Name Verification for the Administration Server and the WLS\_SOAn/WLS\_OSBn Managed Servers

This step is required if you have not set up the appropriate certificates for hostname verification between the Administration Server and Node Manager. If SSL is not set up, you receive an error message unless you disable host name verification.

You can re-enable host name verification when you have set up SSL communication between the Administration Server and the Node Manager.

To disable host name verification:

1. In Oracle WebLogic Server Administration Console, select **Servers**, and then **AdminServer**.
2. Select **SSL**, and then **Advanced**.
3. In the Change Center, click **Lock & Edit**.
4. Set **Hostname Verification** to **None**.
5. Save and activate the changes.

6. Repeat these steps for the WLS\_SOA1, WLS\_SOA2, WLS\_OSB1, and WLS\_OSB2 servers.
7. Restart the Administration Server.

### 5.14.8 Configuring Oracle Coherence for Deploying Composites

Although deploying composites uses multicast communication by default, Oracle recommends using unicast communication for SOA high availability. Use unicast if you disable multicast communication for security reasons.

---

---

**Note:** An incorrect configuration of the Oracle Coherence framework that is used for deployment may prevent the SOA system from starting. The deployment framework must be properly customized for the network environment on which the SOA system runs. Oracle recommends the following configuration described in this section.

---


---

#### Enabling Communication within Clusters Using Unicast Communication

Multicast communication enables Oracle Fusion Middleware SOA to discover all of the members of a cluster to which it deploys composites dynamically. However, unicast communication does not enable nodes to discover other cluster members in this way. Consequently, you must specify the nodes that belong to the cluster. You do not need to specify all of the nodes of a cluster, however. You need only specify enough nodes so that a new node added to the cluster can discover one of the existing nodes. As a result, when a new node has joined the cluster, it is able to discover all of the other nodes in the cluster. Additionally, in a configuration where multiple IPs are available in the same node, you must configure Oracle Coherence to use a specific hostname to create Oracle Coherence cluster.

**Tip:** To guarantee high availability during deployments of SOA composites, specify enough nodes so that at least one of them is running at any given time.

Specify the nodes using the `tangosol.coherence.wka<n>` system property, where `<n>` is a number between 1 and 9. You can specify up to 9 nodes. Start the numbering at 1. This numbering must be sequential and must not contain gaps. In addition, specify the hostname used by Oracle Coherence to create a cluster through the `tangosol.coherence.localhost` system property. This hostname should be the virtual hostname used by the SOA servers that maps the corresponding listener addresses (VIP1 and VIP2). Set this property by adding the `-Dtangosol.coherence.localhost` parameters to the Arguments field of Oracle WebLogic Server Administration Console's Server Start tab (Figure 5-52).

**Figure 5–52** *Setting the Hostname Using the Start Server Tab of Oracle WebLogic Server Administration Console*

### Specifying the hostname

To add the hostname used by Oracle Coherence:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node.
3. Click **Servers**. The Summary of Servers page appears.
4. Click the name of the server (represented as a hyperlink) in Name column of the table. The settings page for the selected server appears.
5. Click the **Server Start** tab (illustrated in [Figure 5–52](#)).
6. Enter the following for WLS\_SOA1 and WLS\_SOA2 into the Arguments field.

For WLS\_SOA1, enter the following (on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1 -Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost1vhn1
```

For WLS\_SOA2, enter the following (on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1 -Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost2vhn1
```



---

**Note:** The Coherence cluster used for deployment uses port 8088 by default. This port can be changed by specifying a different port (for example, 8089) with the `-Dtangosol.coherence.wka1.port` and `-Dtangosol.coherence.localport` startup parameters. For example:

WLS\_SOA1 (enter the following into the Arguments field on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1
-Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost1vhn1
-Dtangosol.coherence.localport=8089
-Dtangosol.coherence.wka1.port=8089
-Dtangosol.coherence.wka2.port=8089
```

WLS\_SOA2 (enter the following into the Arguments field on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1
-Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost2vhn1
-Dtangosol.coherence.localport=8089
-Dtangosol.coherence.wka1.port=8089
-Dtangosol.coherence.wka2.port=8089
```

---

7. Click **Save** and activate the changes.
8. This change requires the SOA servers to be restarted.

---

**Note:** The multicast and unicast addresses are different from the ones used by the WebLogic Server cluster for cluster communication. SOA guarantees that composites are deployed to members of a single WebLogic Server cluster even though the communication protocol for the two entities (the WebLogic Server cluster and the groups to which composites are deployed) are different.

Do not copy the text from above to your Administration Console's arguments text field, which can result in HTML tags being inserted in the Java arguments. The text should not include any text or characters other than the ones shown above.

---

### 5.14.9 Configure Oracle Coherence for the Oracle Service Bus Result Cache

The default configuration for the Coherence infrastructure that Oracle Service Bus result cache uses is multicast. Oracle recommends using unicast communication for the Oracle Service Bus result cache.

To enable unicast for the Oracle Service Bus result cache Coherence infrastructure:

1. Log into Oracle WebLogic Server Administration Console. In the Change Center, click **Lock & Edit**.
2. In the Domain Structure window, expand the **Environment** node.
3. Click **Servers**.

4. Click the name of the server (represented as a hyperlink) in the **Name** column of the table. The settings page for the selected server appears.
5. Click the **Server Start** tab.
6. Enter the following for WLS\_OSB1 on a single line, no carriage returns:  

```
-DOSB.coherence.localhost=osbhost1vhn1 -DOSB.coherence.localport=7890
-DOSB.coherence.wka1= osbhost1vhn1 -DOSB.coherence.wka1.port=7890
-DOSB.coherence.wka2= osbhost2vhn1 -DOSB.coherence.wka2.port=7890
```

  
For WLS\_OSB2, enter the following on a single line, no carriage returns:  

```
-DOSB.coherence.localhost=osbhost2vhn1 -DOSB.coherence.localport=7890
-DOSB.coherence.wka1= osbhost1vhn1 -DOSB.coherence.wka1.port=7890
-DOSB.coherence.wka2= osbhost2vhn1 -DOSB.coherence.wka2.port=7890
```
7. Save and activate the changes. You must restart Oracle Service Bus servers for these changes take effect.

#### 5.14.10 Setting Connection Destination Identifiers for B2B Queues

Oracle B2B uses specific JMS Destination Member calls and requires setting the Create Destination Identifier (CDI) for these calls to succeed. To set up the CDI:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Services** node and the **Messaging** node.
3. Click **JMS Modules** and then **SOAJMSModule**.
4. In the Change Center, click **Lock & Edit**.
5. Click the **dist\_B2BEventQueue\_auto**, **Configuration**, and the General tab, and then click **Advanced**.
6. In the **Create Destination Identifier** field, add the following jndi name for the queue:  
`jms/b2b/B2BEventQueue`
7. Repeat these steps, creating the following Create Destination Identifiers for the queues listed below:
  - **B2B\_OUT\_QUEUE**: `jms/b2b/B2B_OUT_QUEUE`
  - **B2B\_IN\_QUEUE**: `jms/b2b/B2B_IN_QUEUE`
  - **B2BBroadcastTopic**: `jms/b2b/B2BBroadcastTopic`
  - **XmlSchemaChangeNotificationTopic**:  
`jms/fabric/XmlSchemaChangeNotificationTopic`
8. Click **Save** and **Activate Changes**.

#### 5.14.11 Starting the System in SOAHOST1

This section describes how to start Node Manager on SOAHOST1 and how to start and validate the WLS\_SOA1 managed server on SOAHOST1.

##### 5.14.11.1 Starting Node Manager on SOAHOST1

Perform these steps to start Node Manager on SOAHOST1:

1. Run the `setNMProps.sh` script, which is located in the `ORACLE_COMMON_HOME/common/bin` directory, to set the `StartScriptEnabled` property to `true` before starting Node Manager

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
SOAHOST1> ./setNMProps.sh
```

---

**Note:** You must use the `StartScriptEnabled` property to avoid class loading failures and other problems.

---

2. Start Node Manager:

```
SOAHOST1> cd WL_HOME/server/bin
SOAHOST1> ./startNodeManager.sh
```

#### 5.14.11.2 Starting and Validating the WLS\_SOA1 Managed Server

To start the WLS\_SOA1 managed server and check that it is configured correctly:

1. Start the WLS\_SOA1 managed server using Oracle WebLogic Server Administration Console.
2. When WLS\_SOA1 starts, the following URLs are available:
  - `http://SOAHOST1VHN1:8001/b2b`
    - Verify login to B2B console
  - `http://SOAHOST1VHN1:8001/integration/worklistapp`
    - Verify login to worklist console
  - `http://SOAHOST1VHN1:8001/wsm-pm`
    - Verify the policy validator link.
  - `http://SOAHOST1VHN1:8001/soa/composer`
  - `http://SOAHOST1VHN1:8001/soa-infra`

#### 5.14.12 Propagating the Domain Configuration to SOAHOST2, OSBHOST1, and OSBHOST2 with pack/unpack Utilities

To propagate the domain configuration to SOAHOST2 using pack/unpack utilities:

1. Run the following pack command on SOAHOST1 to create a template pack:

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
SOAHOST1> ./pack.sh -managed=true
-domain=ORACLE_BASE/product/fmw/user_projects/domains/soadomain/
-template=soadomaintemplate.jar
-template_name=soa_domain_template
```

2. Run the following command on SOAHOST1 to copy the template file created in the previous step to SOAHOST2:

```
SOAHOST1> scp soadomaintemplate.jar
oracle@node2:ORACLE_COMMON_HOME/common/bin
```

3. Run the unpack command on SOAHOST2 to unpack the propagated template:

```
SOAHOST2> cd ORACLE_COMMON_HOME/common/bin
```

```
SOAHOST2> ./unpack.sh
-domain=ORACLE_BASE/product/fmw/user_projects/domains/soadomain/
-template=soadomaintemplate.jar
```

---

**Note:** When you run the Fusion Middleware Configuration Wizard after a pack/unpack procedure, products already in the original domain that appear in the Select Extension Source screen are not selected or greyed out.

Templates that you create with the pack command are considered user-created templates and treated differently than default templates. User-created templates are self-contained and do not have any information related to the creation of the original domain.

---

4. Repeat the steps above for OSBHOST1 and OSBHOST2.

### 5.14.13 Extracting XEngine Files in the Second Node

To enable B2B's XEngine in the second node, you must extract ZEngine tar content manually:

```
SOAHOST2>cd ORACLE_HOME/soa/thirdparty/edifecs
SOAHOST2>tar xzvf XEngine.tar.gz
```

### 5.14.14 Starting the System in SOAHOST2, OSBHOST1, and OSBHOST2

This section describes procedures for starting the system in SOAHOST2.

#### 5.14.14.1 Starting Node Manager on SOAHOST2, OSBHOST1, and OSBHOST2

To start the Node Manager on SOAHOST2, OSBHOST1, and OSBHOST2, repeat the steps from [Section 5.14.11.1, "Starting Node Manager on SOAHOST1"](#) on those nodes.

#### 5.14.14.2 Starting and Validating the WLS\_SOA2, WLS\_OSB1, and WLS\_OSB2 Managed Server

To start the WLS\_SOA2 managed server and verify that it is configured correctly:

1. Start the WLS\_SOA2 managed server using Oracle WebLogic Server Administration Console. Watch for errors in the server's output file while it starts.
2. When WLS\_SOA2 starts, the following URLs become available:
  - <http://SOAHOST2VHN1:8001/b2b>  
Verify login to B2B console
  - <http://SOAHOST2VHN1:8001/integration/worklistapp>  
Verify login to worklist console
  - <http://SOAHOST2VHN1:8001/wsm-pm>  
Verify the policy validator link
  - <http://SOAHOST2VHN1:8001/soa/composer>
  - <http://SOAHOST2VHN1:8001/soa-infra>

To start the WLS\_OSB1 managed server and verify that it is configured correctly:

1. Start the WLS\_OSB1managed server using Oracle WebLogic Server Administration Console. Watch for errors in the server's output file while it starts.
2. When WLS\_OSB1 is started, the following URL becomes available:

`http://OSBHOST1VHN1:8011/sbinspection.wsil/`

An HTTP reply like the following should be seen:

```
<?xml version="1.0" encoding="UTF-8" ?>
<ins:inspection xmlns:ins="http://schemas.xmlsoap.org/ws/2001/10/inspection/">
<ins:link
referencedNamespace="http://schemas.xmlsoap.org/ws/2001/10/inspection/"
location="http://OSBHOST1VHN1:8011/sbinspection.wsil?refpath=default">
<ins:abstract>default</ins:abstract>
<ins:abstract>LinkType: Project</ins:abstract>
</ins:link>
</ins:inspection>
```

Repeat these steps for WLS\_OSB2.

### 5.14.15 Configuring Oracle HTTP Servers for the Administration Server, WLS\_SOAn, and WLS\_OSBn Managed Servers

To enable Oracle HTTP Server to route to the SOA Cluster (contains the WLS\_SOAn managed servers), set the WebLogicCluster parameter to the list of nodes in the cluster:

1. On WEBHOST1 and WEBHOST2, add the following lines to the ORACLE\_BASE/admin/<instance\_name>/config/OHS/<component\_name>/mod\_wl\_ohs.conf file:

```
WSM-PM
<Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001, SOAHOST2VHN1:8001
</Location>

SOA soa-infra app
<Location /soa-infra>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001, SOAHOST2VHN1:8001
</Location>

SOA inspection.wsil
<Location /inspection.wsil>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001, SOAHOST2VHN1:8001
</Location>

Worklist
<Location /integration/>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001, SOAHOST2VHN1:8001
</Location>

B2B
<Location /b2b>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001, SOAHOST2VHN1:8001
</Location>
```

```
<Location /b2bconsole>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

UMS
<Location /sdpmessaging/ >
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

UMS WS
<Location /ucs/messaging/webservice >
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

Default to-do taskflow
<Location /DefaultToDoTaskFlow/>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

Workflow
<Location /workflow>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

#Required if attachments are added for workflow tasks
<Location /ADFAttachmentHelper>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

SOA composer application
<Location /soa/composer>
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

OSB wsil
<Location /sbinspection.wsil*>
 SetHandler weblogic-handler
 WebLogicCluster OSBHOST1VHN1:8011,OSBHOST2VHN2:8011
</Location>

OSB resource
<Location /sbresource*>
 SetHandler weblogic-handler
 WebLogicCluster OSBHOST1VHN1:8011,OSBHOST2VHN2:8011
</Location>
```

2. Make sure the httpd.conf file located in the same directory as the mod\_wl\_ohs file contains the following lines:

```
NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName https://soa.mycompany.com:443
 ServerAdmin you@your.address
```

```

RewriteEngine On
RewriteOptions inherit
</VirtualHost>

NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName admin.mycompany.com:80
 ServerAdmin you@your.address
 RewriteEngine On
 RewriteOptions inherit
</VirtualHost>

```

---

**Note:** Values such as `soa.mycompany.com:443`, `7777`, `admin.mycompany:80`, and `you@youraddress` are examples only. Enter values based on the actual environment.

---



---

**Note:** Each HTTP proxy service in Oracle Service Bus exposes its own context root, which is the full path of the proxy service (default). For example, `ordersubmit.proxy` in the project `OrderProcessing` and in folder `submission` exposes `/OrderProcessing/submission/ordersubmit` as context path. You must add these url contexts in Oracle HTTP Server for each service.

---

3. Perform the same steps for the Oracle HTTP Server on WEBHOST2.
4. Restart Oracle HTTP Server on both WEBHOST1 and WEBHOST2:

```

WEBHOST1> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs1

```

```

WEBHOST2> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs2

```

### 5.14.16 Validating Access Through Oracle HTTP Server

Verify that the SOA Servers status is Running in the Administration Console. If the server status is Starting or Resuming, wait for the server status to change to Started. If another status is reported Admin, Failed), check the server output log files for errors. Verify that you can access these URLs:

- `http://WEBHOST1:7777/wsm-pm`
- `http://WEBHOST2:7777/wsm-pm`
- `http://WEBHOST1:7777/soa-infra`
- `http://WEBHOST2:7777/soa-infra`
- `http://WEBHOST1:7777/soa/composer`
- `http://WEBHOST2:7777/soa/composer`
- `http://WEBHOST1:7777/integration/worklistapp`
- `http://WEBHOST2:7777/integration/worklistapp`

- <http://WEBHOST1:7777/sdpMessaging/userprefs-ui>
- <http://WEBHOST2:7777/sdpMessaging/userprefs-ui>
- <http://WEBHOST1:7777/b2bconsole>
- <http://WEBHOST2:7777/b2bconsole>
- <http://WEBHOST1:7777/sbinspection.wsil>
- <http://WEBHOST2:7777/sbinspection.wsil>

Verify these URLs also using your load balancing router address:

- <http://soa.mycompany.com:80/wsm-pm>
- <http://soa.mycompany.com:80/soa-infra>
- <http://soa.mycompany.com:80/soa/composer>
- <http://soa.mycompany.com:80/integration/worklistapp>
- <http://soa.mycompany.com:80/sdpMessaging/userprefs-ui>
- <http://soa.mycompany.com:80/b2bconsole>
- <http://soa.mycompany.com:80/sbinspection.wsil>

To ensure that routing and failover from the HTTP Server to the SOA\_Cluster is working correctly:

1. While WLS\_SOA2 is running, stop WLS\_SOA1 from Oracle WebLogic Server Administration Console.
2. Access the following URLs and verify the appropriate functionality:
  - [WEBHOST1:7777/wsm-pm](http://WEBHOST1:7777/wsm-pm)
  - [WEBHOST1:7777/soa-infra](http://WEBHOST1:7777/soa-infra)
  - [WEBHOST1:7777/soa/composer](http://WEBHOST1:7777/soa/composer)
  - [WEBHOST1:7777/integration/worklistapp](http://WEBHOST1:7777/integration/worklistapp)
  - [WEBHOST1:7777/sdpMessaging/userprefs-ui](http://WEBHOST1:7777/sdpMessaging/userprefs-ui)
  - [WEBHOST1:7777/b2bconsole](http://WEBHOST1:7777/b2bconsole)
3. Start WLS\_SOA1 from Oracle WebLogic Server Administration Console.
4. Stop WLS\_SOA2.
5. Access the URLs in Step 2 above again and verify the appropriate functionality:

To ensure that routing and failover from the HTTP Server to the OSB\_Cluster is working correctly:

1. While WLS\_OSB2 is running, stop WLS\_OSB1 from Oracle WebLogic Server Administration Console.
2. Access the following URL and verify the appropriate functionality:
  - [WEBHOST1:7777/sbinspection.wsil](http://WEBHOST1:7777/sbinspection.wsil)
3. Start WLS\_OSB1 from Oracle WebLogic Server Administration Console.
4. Stop WLS\_OSB2.
5. Access the URL in Step 2 above again and verify the appropriate functionality:



### 5.14.17 Setting the Front End HTTP Host and Port

You must set the front end HTTP host and port for Oracle WebLogic Server cluster:

1. In the WebLogic Server Administration Console, in the Change Center section, click **Lock & Edit**.
2. In the left pane, select **Environment > Clusters**.
3. Select the SOA\_Cluster.
4. Select **HTTP**.
5. Set the values for the following:
  - **Frontend Host:** soa.mycompany.com  
**Frontend HTTP Port:** 80
  - **Frontend HTTPS Port:** leave it blank

---

**Note:** Make sure this address is correct and available (the load balancing router is up). An incorrect value, for example, **http://** in the address, or trailing / in the hostname, may prevent the SOA system from being accessible even when using the virtual IPs to access it.

---

6. Click **Save**.
7. To activate the changes, click **Activate Changes** in the Change Center section of the Administration Console.
8. Repeat steps 1 through 7 for the OSB\_Cluster.
9. If you have started the server before, notice this change requires a restart of the managed servers in the cluster.

---

**Note:** The SOA system calculates the callback URL as follows:

- If a request to SOA originates from an external or internal service is originating a request to SOA, then SOA uses the callback URL specified by the client.
- If a request to an external or internal service originates from SOA, the callbackURL cannot be populated in the SOA request dynamically because SOA is the originator. Instead, callbackServerURL is used if it is specified as a binding property for the specific reference. (You can set this when modeling the composite or at runtime using the MBeans accessed through Oracle Enterprise Manager console.) This allows different service calls to have different callback URLs. That is, a callback URL from an external service is different than one to an internal service.

However, if the callbackServer URL is not specified as a binding property for that reference, then the system uses the callback URL as specified in soa-infra-config. If the callback URL is not specified in soa-infra-config, then the system uses the callback URL as the front end host specified in WLS. If the front end host is not specified in WLS, the system uses the callback URL as the local host name as provided by WLS MBean APIs.

---

For SOA high availability installations frontended by Oracle HTTP Server, monitoring should be done on the Oracle HTTP Server ports of the real backend servers. This is the case when a deployment is using all the components deployed to the SOA Managed Server. A simple HTTP monitor that pings the HTTP/HTTPS port and expects a pre-determined response in turn should suffice. If only a specific SOA component is being used (such as B2B), then a monitor that does a deeper level check all the way to the managed server can be considered to validate the health of the component in use. Please check with your load balancer vendor on setting up the HTTP monitors with your load balancer.

If you do not set the front end HTTP host and port, you get the following message when trying to retrieve a document definition XSD from Oracle B2B:

```
An error occurred while loading the document definitions.
java.lang.IllegalArgumentException: Cluster address must be set when
clustering is enabled.
```

### 5.14.18 Configuring a Default Persistent Store for Transaction Recovery

Each server has a transaction log, which stores information about committed transactions coordinated by the server that may not have been completed. WebLogic Server uses the transaction log when recovering from system crashes or network failures. To take advantage of the migration capability of the Transaction Recovery Service for servers in a cluster, you must store the transaction log in a location that is available to a server and its backup servers (the nodes the WebLogic Server can failover to), preferably on a dual-ported SCSI disk, or a Storage Area Network (SAN). To configure the default persistent store:

1. In the **Domain Structure** tree, expand **Environment** and select **Servers**.
2. Select the server you want to modify.
3. Select the **Configuration, Services** tab.
4. Under **Default Store**, in the **Directory** field, enter the path to the folder where you want the default persistent store to store its data files.

To enable migration of the Transaction Recovery Service, you must specify a location on a persistent storage solution that is available to other servers in the cluster. Therefore, the directory that you enter must be accessible by both WLS\_SOA1 and WLS\_SOA2 for the SOA\_Cluster and WLS\_OSB1 and WLS\_OSB2 for the OSB\_Cluster. This directory must exist before the server is restarted.

### 5.14.19 Deploying Applications

You can deploy SOA composite applications from Oracle Enterprise Manager Fusion Middleware Control Console with the Deploy SOA Composite wizard. Use the Deploy SOA Composite wizard to deploy any of the following:

- A new SOA composite application for the first time
- A new revision (for example, 2.0) alongside an older revision (for example, 1.0) without impacting the latter. The revision deployed last becomes the new default revision of that composite (unless you specify otherwise at a later step during deployment).
- A bundle (ZIP file) containing multiple SOA composite application revisions (for example, revisions 2.0, 3.0, and 4.0) of a SOA composite application that already has a different revision that is currently deployed (for example, 1.0). This option enables you to deploy revisions 1.0, 2.0, 3.0, and 4.0 at the same time. The bundle

can also contain revisions of different composites. There is no restriction that all revisions must be of the same composite application.

Deployment extracts and activates the composite application in SOA Service Infrastructure. Once an application is deployed, you can perform administration tasks, such as creating instances, configuring properties, monitoring performance, managing instance life cycles, and managing policies and faults.

---

**Note:** If you want to redeploy an *existing* revision of an application, do *not* use this wizard. Instead, use the Redeploy SOA Composite wizard.

---

To deploy applications:

1. Access the Deploy SOA Composite wizard through one of the following options:

| From SOA Infrastructure Menu...               | From the SOA Folder in the Navigator...                                            | From the SOA Infrastructure Home Page...                                                                   | From the SOA Composite Menu...                                  |
|-----------------------------------------------|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| 1. Select <b>SOA Deployment &gt; Deploy</b> . | 1. Right-click <b>soa-infra</b> .<br>2. Select <b>SOA Deployment &gt; Deploy</b> . | 1. Click the <b>Deployed Composites</b> tab.<br>2. Above the <b>Composite</b> table, click <b>Deploy</b> . | 1. Select <b>SOA Deployment &gt; Deploy Another Composite</b> . |

The Select Archive page appears.

2. In the **Archive or Exploded Directory** section, specify the archive of the SOA composite application to deploy. The archive contains the project files of the composite to be deployed (for example, **HelloWorld\_rev1.0.jar** for a single archive or **OrderBooking\_rev1.0.zip** for multiple archives).
3. In the **Configuration Plan** section, optionally specify the configuration plan to include with the archive. The configuration plan enables you to define the URL and property values to use in different environments. During process deployment, the configuration plan is used to search the SOA project for values that must be replaced to adapt the project to the next target environment.

4. Click **Next**.
5. Select the WebLogic Server or cluster to which to deploy the SOA composite application archive. You can deploy to multiple servers and clusters.
6. Click **Next**.
7. Review your selections.
8. Select whether or not to deploy the SOA composite application as the default revision. The default revision instantiates when a new request comes in.
9. Click **Deploy**.  
Processing messages are displayed.
10. When deployment has completed, click **Close**.

**See Also:** *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite* for instructions on creating configuration plans and deploying applications from Oracle JDeveloper

### 5.14.20 Configuring Server Migration for the WLS\_SOA Servers

The high availability architecture for a SOA system uses server migration to protect some singleton services against failures. For more information on whole server migration, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

The WLS\_SOA1 managed server is configured to restart on SOAHOST2 in case of failure, and the WLS\_SOA2 managed server is configured to restart on SOAHOST1 in case of failure. For this configuration the WLS\_SOA1 and WLS\_SOA2 servers listen on specific floating IPs that are failed over by WLS Server Migration. To configure server migration for the WLS\_SOA<sub>n</sub> managed servers, follow these steps:

- Step 1: [Setting Up a User and Tablespace for the Server Migration Leasing Table](#)
- Step 2: [Creating a Multi Data Source Using the Oracle WebLogic Administration Console](#)
- Step 3: [Editing Node Manager's Properties File](#)
- Step 4: [Setting Environment and Superuser Privileges for the wlsifconfig.sh script](#)
- Step 5: [Configuring Server Migration Targets](#)
- Step 6: [Testing Server Migration](#)

#### Step 1 Setting Environment and Superuser Privileges for the wlsifconfig.sh script

1. Grant sudo privilege to the WebLogic user ('oracle') with No Password restriction, and grant execute privilege on /sbin/ifconfig and /sbin/arping binaries.

---

**Note:** On Windows, the script is named wlsifconfig.cmd and it can be run by users with the administrator privilege.

---

Make sure the script is executable by the WebLogic user ('oracle'). The following is an example of an entry inside /etc/sudoers granting sudo execution privilege for oracle and also over the ifconfig and arping:

```
Defaults:oracle !requiretty
oracle ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
```

## Step 2 Testing Server Migration

To verify that Server Migration is working properly, follow these steps:

From SOAHOST1:

1. Force stop the WLS\_SOA1 managed server.

To do this, run this command:

```
SOAHOST1> kill -9 <pid>
```

*pid* specifies the process ID of the managed server. You can identify the *pid* in the node by running this command:

```
SOAHOST1> ps -ef | grep WLS_SOA1
```

---

**Note:** For Windows, the Managed Server can be terminated by using the `taskkill` command. For example:

```
taskkill /f /pid <pid>
```

Where *<pid>* is the process Id of the Managed Server.

To determine the process Id of Managed server run the following command and identify the *pid* of the WLS\_SOA1 Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

2. Watch the Node Manager console: you should see a message indicating that WLS\_SOA1's floating IP has been disabled.
3. Wait for the Node Manager to try a second restart of WLS\_SOA1. Node Manager waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.


From SOAHOST2:

1. Watch the local Node Manager console. After 30 seconds since the last try to restart WLS\_SOA1 on Node 1, Node Manager on Node 2 should prompt that the floating IP for WLS\_SOA1 is being brought up and that the server is being restarted in this node.
2. Access the soa-Service Infrastructure console in the same IP.

Migration can also be verified in Oracle WebLogic Server Administration Console:

1. Log into Oracle WebLogic Server Administration Console.
2. Click on **Domain** on the left pane.
3. Click on the **Monitoring** tab and then on the **Migration** sub-tab.

The Migration Status table provides information on the status of the migration.

**Figure 5–53 Migration Status Screen in Oracle WebLogic Server Administration Console**

**Note:** On Windows, when you manually shut down multiple servers at the same time on the same machine and, on another machine, attempt to start one of the servers that you shut down, the IP bind may not work. This happens because the original machine still has claim to the IP address, even though `netsh` has reported that the IP address has been removed.

To resolve this, you must check the network configuration either by using the `ipconfig` utility or Windows Network Configuration. Either of these may show that one of the virtual/floating IP addresses is still configured even though the servers have been shut down. You can then use Windows Network Configuration to remove the IP address using the following procedure:

1. From Windows Control Panel, select **Network Connections**.
2. Select the appropriate network interface, right-click, and select **Properties**.
3. Select **Internet Protocol (TCP/IP)** and click the **Properties** button.
4. Select **Advanced**.
5. Select the appropriate IP address and click the **Remove** button.

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

## 5.14.21 Configuring Server Migration for the WLS\_OSB Servers

The high availability architecture for an OSB system uses server migration to protect some singleton services against failures. For more information on whole server migration, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

The WLS\_OSB1 managed server is configured to be restarted on OSBHOST2 in case of failure, and the WLS\_OSB2 managed server is configured to be restarted on OSBHOST1 in case of failure. For this configuration the WLS\_OSB1 and WLS\_OSB2 servers listen on specific floating IPs that are failed over by WLS Server Migration. To configure server migration for the WLS\_OSBn managed servers, follow these steps:

---

**Note:** If a tablespace, schema multi data source, and data sources have already been set up for SOA (as described in [Section 5.14.20, "Configuring Server Migration for the WLS\\_SOA Servers"](#) above, then the OSB Cluster can reuse the existing leasing schema and data source and Step 1 and Step 2 below are not required. Instead, simply target the data sources and multi data sources created for SOA to the OSB Cluster.

---

### Step 1 Set Up User and Tablespace for the Server Migration Leasing Table

1. Create a tablespace called `leasing`.

Example: Log on to SQL\*Plus as the `sysdba` user and run the following command:

```
SQL> create tablespace leasing
 logging datafile 'DB_HOMEs/oradata/orcl/leasing.dbf'
 size 32m autoextend on next 32m maxsize 2048m extent management local;
```

2. Create a user named `leasing` and assign to it the `leasing` tablespace.

```
SQL> create user leasing identified by welcome1;
```

```
SQL> grant create table to leasing;
```

```
SQL> grant create session to leasing;
```

```
SQL> alter user leasing default tablespace leasing;
```

```
SQL> alter user leasing quota unlimited on LEASING;
```

3. Create the leasing table using the `leasing.ddl` script.
  - a. Copy the `leasing.ddl` file, located in the `WL_HOME/server/db/oracle/817` or `WL_HOME/server/db/oracle/920` directories, to your database node.
  - b. Connect to the database as the `leasing` user.
  - c. Run the `leasing.ddl` script in SQL\*Plus.

```
SQL> @copy_location/leasing.ddl;
```

### Step 2 Create a Multi Data Source from Oracle WebLogic Server Administration Console

The second step is to create a multi data source for the leasing table from the Oracle WebLogic Server Administration Console:

You create a data source to each of the Oracle RAC database instances during the process of setting up the multi data source, both for these data sources and the global leasing multi data source. When you create a data source:

- Ensure that this is a non-xa data source.
- The names of the multi data sources are in the following format: *<MultiDS>-rac0*, *<MultiDS>-rac1*
- Use Oracle's Driver (Thin) Version 9.0.1, 9.2.0, 10, 11.
- Data sources do not require support for global transactions. Therefore, do *not* use any type of distributed transaction emulation/participation algorithm for the data source (do not choose the **Supports Global Transactions** option, or the **Logging Last Resource, Emulate Two-Phase Commit**, or **One-Phase Commit** options of the **Supports Global Transactions** option), and specify a service name for your database.
- Target these data sources to the OSB cluster.
- Make sure the datasources' connection pool initial capacity is set to 0. To do this, select **Services**, **JDBC**, and then **Datasources**. In the Datasources screen, click the **Datasource Name**, then click the **Connection Pool** tab, and enter 0 in the **Initial capacity** field.

To create a multi data source:

1. From the Domain Structure window in the Oracle WebLogic Server Administration Console, expand the **Services** node, then click **Data Sources**. The Summary of JDBC Data Sources page appears.
2. In the Change Center, click **Lock & Edit**.
3. Click **New**, then click **Multi Data Sources**. The Create a New JDBC Multi Data Source page appears.
4. Enter **leasing** as the **Name**.
5. Enter **jdbc/leasing** as the JNDI name.
6. Select **Failover as algorithm (default)**.
7. Click **Next**.
8. Select **non-XA driver (the default)**.
9. Click **Next**.
10. Click **Create New Data Source**.
11. Enter *leasing-rac0* as name. Enter *jdbc/leasing-rac0* as JNDI name. Enter *oracle* as the database type. For the driver type, enter *Oracle Driver (Thin) for RAC server-Instance connection Version 10,11*.
12. Click **Next**.
13. Deselect **Supports Global Transactions**.
14. Click **Next**.
15. Enter the service name, database name (this is actually the RAC Node instance name, for example: racdb1,racdb2), host port, and password for your leasing schema
16. Click **Next**.
17. Click **Test Configuration** and verify the connection works.



18. Target the data source to the OSB cluster.
19. Select the data source and add it to the right screen.
20. Click **Create a New Data Source** and repeat the steps for the second instance of your Oracle RAC database.
21. Add the second data source to your multi data source.
22. Save and Activate the changes.

### Step 3 Edit the Node Manager's Properties file

For information on Node Manager and whole server migration, see [Section 3.9, "Whole Server Migration."](#)

The `nodemanager.properties` file is located in the `WL_HOME/common/nodemanager` directory. For server migration to work properly, you must add the properties listed in this section:

- `Interface=eth0`

---

**Note:** Do not specify the sub interface, such as `eth0:1` or `eth0:2`. This interface is to be used without the `:0`, or `:1`. The Node Manager's scripts traverse the different `:X` enabled IPs to determine which to add or remove. For example, the valid values in Linux environments are `eth0`, `eth1`, or, `eth2`, `eth3`, `ethn`, depending on the number of interfaces configured.

---

This property specifies the interface name for the floating IP (`eth0`, for example, on Linux).

---

**Note:** For Windows, the Interface should be set to the Network Interface Name. For example: `Interface="Local Area Connection"`.

---

Be sure that the interface provided is the public interface for this node. On multi home nodes, this interface should be the one on which the floating IP can be enabled.

- `NetMask=255.255.255.0`

This property specifies the net mask for the interface for the floating IP. The netmask provided (255.255.255.0) is just an example. The actual value depends on your network.

- `UseMACBroadcast=true`

This property specifies whether or not to use a node's MAC address when sending ARP packets, that is, whether or not to use the `-b` flag in the `arping` command.

After starting Node Manager, verify in Node Manager's output (shell where Node Manager is started) that these properties are being used, or problems may arise during migration. You should see something like this in the Node Manager's output:

```
StateCheckInterval=500
```

```
Interface=eth0 (Linux) or Interface="Local Area Connection" (Windows)
```

```
NetMask=255.255.255.0
```

```
UseMACBroadcast=true
```

---

**Note:** The steps in this section are not required if the server properties (start properties) have been properly set and the Node Manager can start the servers remotely.

---

1. Set the following property in the nodemanager.properties file.
  - StartScriptEnabledSet this property to true.
2. Start the Node Manager on OSBHOST1 and OSBHOST2 by running the startNodeManager.sh script located in the `WL_HOME/server/bin` directory.
3. Validate the changes to the nodemanager.properties file by checking the nodemanager.log file.

#### Step 4 Set Environment and Superuser Privileges for the wlsifconfig.sh script

1. Grant sudo privilege to the WebLogic user ('oracle') with No Password restriction, and grant execute privilege on `/sbin/ifconfig` and `/sbin/arping` binaries.

---

**Note:** On Windows, the script is named `wlsifconfig.cmd` and it can be run by users with the administrator privilege.

---

Make sure the script is executable by the WebLogic user ('oracle'). The following is an example of an entry inside `/etc/sudoers` granting sudo execution privilege for `oracle` and also over the `ifconfig` and `arping`:

```
Defaults:oracle !requiretty
oracle ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
```

#### Step 5 Configure Server Migration Targets

Configuring Cluster Migration sets the `DataSourceForAutomaticMigration` property to true. Follow the steps in this section to configure cluster migration in a migration in a cluster:

1. Log into Oracle WebLogic Server Administration Console
2. In the left pane, expand **Environment** and select **Clusters**.
3. Select the cluster for which you want to configure migration (**OSB\_Cluster**).
4. Click **Migration**.
5. In the Change Center, click **Lock & Edit**.
6. In the **Available** field, select the machine to which to allow migration and click the right arrow. In this case, select **OSBHOST1** and **OSBHOST2**.
7. Select the data source to be used for automatic migration. In this case select the leasing data source.
8. Click **Save**.
9. Set the Candidate Machines for Server Migration. This needs to be done for all the managed servers.

To do this, follow these steps:

- a. In the left pane of Oracle WebLogic Server Administration Console, expand **Environment** and select **Servers**.
- b. Select the server for which you want to configure migration.
- c. Click the **Migration** tab.
- d. In the **Available** field, located in the **Migration Configuration** section, select the machines to which to allow migration and click the right arrow. For **WLS\_OSB1**, select **OSBHOST2**. For **WLS\_OSB2**, select **OSBHOST1**.
- e. Select the checkbox for **Automatic Server Migration Enabled**. This enables the Node Manager to start a failed server on the target node automatically.
- f. Click **Save** and activate the changes.
- g. These changes require restart of the servers and Node Managers to be effective.

### Step 6 Test Server Migration

To verify that Server Migration is working properly, follow these steps:

From OSBHOST1:

1. Force stop the WLS\_OSB1 managed server.

To do this, run this command:

```
OSBHOST1> kill -9 <pid>
```

*pid* specifies the process ID of the managed server. You can identify the *pid* in the node by running this command:

```
OSBHOST1> ps -ef | grep WLS_OSB1
```

---

**Note:** For Windows, the Managed Server can be terminated by using the `taskkill` command. For example:

```
taskkill /f /pid <pid>
```

Where *<pid>* is the process Id of the Managed Server.

To determine the process Id of the Managed Server run the following command and identify the *pid* of the WLS\_OSB1 Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

2. Watch the Node Manager console: you should see a message indicating that WLS\_OSB1's floating IP has been disabled.
3. Wait for the Node Manager to try a second restart of WLS\_OSB1. Node Manager waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

From OSBHOST2:

1. Watch the local Node Manager console. After 30 seconds since the last try to restart WLS\_OSB1 on Node 1, Node Manager on Node 2 should prompt that the

floating IP for WLS\_OSB1 is being brought up and that the server is being restarted in this node.


2. Access the sbinspection.wsil urlconsole in the same IP.

Migration can also be verified in Oracle WebLogic Server Administration Console:

1. Log into Oracle WebLogic Server Administration Console.
2. Click on **Domain** on the left pane.
3. Click on the **Monitoring** tab and then on the **Migration** sub-tab.

The Migration Status table provides information on the status of the migration.

**Figure 5–54 Migration Status Screen in Oracle WebLogic Server Administration Console**



---

**Note:** On Windows, when you manually shut down multiple servers at the same time, on the same machine, and then, on another machine, attempt to start one of the servers that was shut down, the IP bind may not work. This happens because the original machine still has claim to the IP address, even though `netsh` has reported that the IP address has been removed.

To resolve this, you must check the network configuration either by using the `ipconfig` utility or Windows Network Configuration. Either of these may show that one of the virtual/floating IP addresses is still configured even though the servers have been shut down. You can then use Windows Network Configuration to remove the IP address using the following procedure:

1. From Windows Control Panel, select **Network Connections**.
  2. Select the appropriate network interface, right-click, and select **Properties**.
  3. Select **Internet Protocol (TCP/IP)** and click the **Properties** button.
  4. Select **Advanced**.
  5. Select the appropriate IP address and click the **Remove** button.
- 

---

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

---

## 5.14.22 Scaling the Topology

You can scale out or scale up this high availability topology. When you "scale up" the topology, you add new managed servers to nodes that are already running one or more managed servers. When you "scale out" the topology, you add new managed servers to new nodes.

### 5.14.22.1 Scaling Up the Topology (Adding Managed Servers to Existing Nodes)

In this case, you already have a node that runs a managed server configured with SOA and Oracle Service Bus components. The node contains a Middleware home, an Oracle HOME (SOA and Oracle Service Bus) and a domain directory for existing managed servers.

You can use the existing installations (the Middleware home, and domain directories) for creating new WLS\_OSB servers. There is no need to install SOA or Oracle Service Bus binaries in a new location, or to run pack and unpack.

For scaling up the SOA Servers, follow the steps in [Section 5.13.21.1, "Scaling Up the Topology \(Adding Managed Servers to Existing Nodes\)."](#)

Follow these steps for scaling up the topology for the OSB server:

1. Using the Administration Console, clone WLS\_OSBn into a new managed server. The source managed server to clone should be one that already exists on the node where you want to run the new managed server.

To clone a managed server:

- a. Select **Environment** -> **Servers** from the Administration Console.
- b. Select the managed server that you want to clone (for example, WLS\_OSB1).
- c. Select **Clone**.

Name the new managed server WLS\_OSB $n$ , where  $n$  is a number to identify the new managed server.

The rest of the steps assume that you are adding a new server to OSBHOST1, which is already running WLS\_OSB1.

2. For the listen address, assign the virtual host name to use for this new managed server. If you are planning to use server migration as recommended for this server, this virtual host name will enable it to move to another node. The virtual host name should be different from those in use by other managed servers (may be in the same or different domain) that are running in the nodes used by the OSB/SOA domain.

Perform these steps to set the managed server listen address:

- a. Log into the Oracle WebLogic Server Administration Console.
- b. In the Change Center, click **Lock & Edit**.
- c. Expand the **Environment** node in the Domain Structure window.
- d. Click **Servers**. The Summary of Servers page is displayed.
- e. Select the managed server whose listen address you want to update in the **Names** column of the table. The Setting page for that managed server is displayed.
- f. Set the Listen Address to OSBHOST1VHN1.
- g. Click **Save**.
- h. Save and activate the changes.

The changes will not take effect until the managed server is restarted.

3. Update the cluster address to include the new server:
  - a. Select **Environment** -> **Cluster** from the Administration Console.
  - b. Click on the OSB\_Cluster server. The Settings Screen for the OSB\_Cluster appears.
  - c. In the Change Center, click **Lock & Edit**.
  - d. Add the new server's address and port to the **Cluster Address** field. For example:

OSBHOST1VHN1:8011, OSBHOST2VHN1:8011, OSBHOST1VHN2:8011

4. If your Oracle Service Bus configuration includes one or more business services that use JMS request/response functionality, then you must also perform the following procedure using the Oracle Service Bus Console after adding the new managed server to the cluster:
  - a. In the Change Center, click **Create** to create a session.
  - b. Using the Project Explorer, locate and select a business service that uses JMS request/response. Business services of this type display Messaging Service as their Service Type.
  - c. At the bottom of the View Details page, click **Edit**.

- d. If there is a cluster address in the endpoint URI, add the new server to the cluster address.
- e. On the Edit a Business Service - Summary page, click **Save**.
- f. Repeat the previous steps for each remaining business service that uses JMS request/response.
- g. In the Change Center, click **Activate**.
- h. Restart the managed server.
- i. Restart the Administration Server.

The business services are now configured for operation in the extended domain.

---

**Note:** For business services that use a JMS MessageID correlation scheme, you must edit the connection factory settings to add an entry to the table mapping managed servers to queues. For information on how to configure queues and topic destinations, see "JMS Server Targeting" in *Oracle Fusion Middleware Configuring and Managing JMS for Oracle WebLogic Server*.

---

5. If your Oracle Service Bus configuration includes one or more proxy services that use JMS endpoints with cluster addresses, then you must also perform the following procedure using the Oracle Service Bus Console after adding the new managed server to the cluster:
  - a. In the Change Center, click **Create** to create a session.
  - b. Using the Project Explorer, locate and select a proxy service that uses JMS endpoints with cluster addresses.
  - c. At the bottom of the View Details page, click **Edit**.
  - d. If there is a cluster address in the endpoint URI, add the new server to the cluster address.
  - e. On the Edit a Proxy Service - Summary page, click **Save**.
  - f. Repeat the previous steps for each remaining proxy service that uses JMS endpoints with cluster addresses.
  - g. In the Change Center, click **Activate**.
  - h. Restart the managed server.

The proxy services are now configured for operation in the extended domain.

6. Update the Oracle Service Bus result cache Coherence configuration for the new server:
  - a. Log into Oracle WebLogic Server Administration Console. In the Change Center, click **Lock & Edit**.
  - b. In the Domain Structure window, expand the **Environment** node.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Click the name of the server (represented as a hyperlink) in the **Name** column of the table. The settings page for the selected server appears.
  - e. Click the **Server Start** tab.

- f. Enter the following for WLS\_OSBn (on a single line, without a carriage returns):

```
-DOSB.coherence.localhost=osbhost1vhnn -DOSB.coherence.localport=7890
-DOSB.coherence.wka1= osbhost1vhn1 -DOSB.coherence.wka1.port=7890
-DOSB.coherence.wka2= osbhost2vhn1 -DOSB.coherence.wka1.port=7890
```

---

**Note:** The above configuration assumes that servers WLS\_OSB1 and WLS\_OSB2 are running when WLS\_OSBn is started. This will enable WLS\_OSBn to join the coherence cluster started by either WLS\_OSB1 or WLS\_OSB2 using the WKA addresses specified. Also, make sure WLS\_OSB1 and WLS\_OSB2 are started before WLS\_OSBn is started when all three servers are restarted. This will ensure WLS\_OSBn join the cluster started by one of WLS\_OSB1 or WLS\_OSB2. If you want to have a configuration where the order in which the servers are started does not matter, you must add the host and port for WLS\_OSBn as WKA for WLS\_OSB1 and WLS\_OSB2, and also add WLS\_OSBn as WKA for WLS\_OSBn.

---

- g. Save and activate the changes (this change requires the Oracle Service Bus servers to be restarted to be effective).
7. Create JMS Servers and persistent stores for OSB reporting/internal destinations on the new managed server.
- a. Use the WebLogic Server Administration Console to create a new persistent store for the new WseeJMSServer and name it, for example, OSB\_rep\_JMSFileStore\_N. Specify the path for the store. This should be a directory on shared storage, as recommended in [Section 5.14.1.3, "Shared Storage Prerequisites"](#):

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

*ORACLE\_BASE/admin/DOMAIN\_NAME/cluster\_name/jms/OSB\_rep\_JMSFileStore\_N*

- b. Create a new JMS Server for OSB, for example, OSB\_rep\_JMSServer\_N. Use the OSB\_rep\_JMSFileStore\_N for this JMSServer. Target the OSB\_rep\_JMSServer\_N Server to the recently created Managed Server (WLS\_OSBn).
- c. Update the SubDeployment targets for the jmsresources OSB JMS Module to include the recently created OSB JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **jmsResources** (represented as a hyperlink in the Names column of the table). The Settings page for jmsResources appears. Click the SubDeployments tab. The subdeployment module for jmsresources appears.

---

**Note:** This subdeployment module name for destinations is a random name in the form of wlsbjMSServerXXXXXX resulting from the Configuration Wizard JMS configuration for the first two servers (WLS\_OSB1 and WLS\_OSB2).

---



Click on the **wlsbJMSServerXXXXXX** subdeployment and update the targets to include the new **OSB\_rep\_JMSServer\_n** server.

8. Create JMS Servers, persistent stores and destinations for OSB JAX-RPC on the new managed server.

---

**Note:** WebLogic Advanced Web Services for JAX-RPC Extension uses regular (non-distributed) destinations to ensure that a locally processed request on a service gets enqueued only to a local member.

---

- a. Use the WebLogic Server Administration Console to create a new persistent store for the new **WseeJMSServer** and name it, for example, **Wsee\_rpc\_JMSFileStore\_N**. Specify the path for the store. This should be a directory on shared storage, as recommended in [Section 5.14.1.3, "Shared Storage Prerequisites"](#):

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

*ORACLE\_BASE/admin/DOMAIN\_NAME/cluster\_name/jms/Wsee\_rpc\_JMSFileStore\_N*

- b. Create a new JMS Server for OSB JAX-RPC, for example, **OSB\_rpc\_JMSServer\_N**. Use the **Wsee\_rpc\_JMSFileStore\_N** for this JMSServer. Target the **OSB\_rpc\_JMSServer\_N** Server to the recently created Managed Server (**WLS\_OSBn**).
- c. Update the **WseeJMSModule** OSB JMS Module with destinations and the recently created OSB JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **WseeJmsModule** (represented as a hyperlink in the Names column of the table). The Settings page for **WseeJmsModule** appears. Follow steps d through v to complete this step.
- d. In the Change Center, click **Lock and Edit**. Click **New**.
- e. Select **Queue**.
- f. Click **Next**.
- g. Enter **DefaultCallbackQueue-WseeJmsServer\_auto\_n** as name for the queue.
- h. Enter **weblogic.wsee.DefaultCallbackQueue-WseeJmsServer\_auto\_n** as the JNDI name.
- i. Click **Next**.
- j. Click **Create a New Subdeployment**.
- k. Accept the default name.
- l. Click **OK**.
- m. Select **OSB\_rpc\_JMSServer\_n** as the target.
- n. Click **Finish**.
- o. Activate the changes.
- p. Update the local JNDI name for the destination by following steps q through u.

- q. In the Change Center, click **Lock & Edit**.
  - r. In the Settings for the WseeJmsModule page, click on the `DefaultCallbackQueue-WseeJmsServer_auto_n` destination.
  - s. In the general Configuration tab, click on **Advanced**.
  - t. Update the local JNDI name to `weblogic.wsee.DefaultCallbackQueue`.
  - u. Activate the changes.
  - v. Repeat steps d through n for the `DefaultQueue-WseeJmsServer_auto_n` queue, using `weblogic.wsee.DefaultQueue-WseeJmsServer_auto_n` as the JNDI name and `weblogic.wsee.DefaultQueue` as the local JNDI name.
9. Create a new SAF agent and target it to the newly added managed server, as follows:

In the Oracle WebLogic Server Administration Console, expand **Services > Messaging > Store-and-Forward Agents**, and add a new SAF agent `ReliableWseeSAFAgent_auto_N`.

Select persistent store `Wsee_rpc_JMSFileStore_N` (persistent store created for OSB JAX-RPC). Target the SAF Agent to the new managed server and activate changes.
10. Configure TX persistent store for the new server. This should be a location visible from other nodes as indicated in the recommendations about shared storage.

From the Administration Console, select **Server\_name > Services** tab. Under **Default Store**, in **Directory**, enter the path to the folder where you want the default persistent store to store its data files.
11. Disable host name verification for the new managed server. Before starting and verifying the WLS\_OSBn managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in OSBHOSTn. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required (the hostname verification settings is propagated to the cloned server).

To disable host name verification:

  - a. In the Oracle Enterprise Manager Console, select Oracle WebLogic Server Administration Console.
  - b. Expand the **Environment** node in the Domain Structure window.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Select WLS\_OSBn in the Names column of the table. The Settings page for the server appears.
  - e. Click the SSL tab.
  - f. Click **Advanced**.
  - g. Set Hostname Verification to **None**.
  - h. Click **Save**.
12. If it is not already started, start the Node Manager on the node. To start the Node Manager, use the installation in shared storage from the existing nodes as follows:

```
OSBHOSTN> WL_HOME/server/bin/startNodeManager
```

13. Start and test the new managed server from the Administration Console.
  - a. Shut down the existing managed servers in the cluster.
  - b. Ensure that the newly created managed server, WLS\_OSBn, is up.
  - c. Access the application on the newly created managed server (<http://vip:port/sbinspection.wsil>).
14. Configure Server Migration for the new managed server.

---

**Note:** Since this is a scale-up operation, the node should already contain a Node Manager and environment configured for server migration. The floating IP for the new OSB managed server should also be already present.

---

Configure server migration following these steps:

- a. Log into the Administration Console.
  - b. In the left pane, expand **Environment** and select **Servers**.
  - c. Select the name of the new managed server for which you want to configure migration.
  - d. Click the Migration tab.
  - e. In the **Available** field, in the "Migration Configuration" section, select the machines to which to allow migration and click the right arrow. Select the same migration targets as for the servers that already exist on the node.

For example, for new managed servers on OSBHOST1, which is already running WLS\_OSB1, select OSBHOST2. For new managed servers on OSBHOST2, which is already running WLS\_OSB2, select OSBHOST1.

Make sure the appropriate resources are available to run the managed servers concurrently during migration.
  - f. Select the "Automatic Server Migration Enabled" option. This enables the Node Manager to start a failed server on the target node automatically.
  - g. Click **Save**.
  - h. Restart the Administration Server, managed servers, and Node Manager.
15. Test server migration for this new server. Follow these steps from the node where you added the new server:
- a. Stop the WLS\_OSBn managed server.

To do this, run "kill -9 <pid>" on the PID of the managed server. You can identify the PID of the node using "ps -ef | grep WLS\_OSBn".

---

**Note:** For Windows, the Managed Server can be terminated by using the taskkill command. For example:

```
taskkill /f /pid <pid>
```

Where <pid> is the process Id of the Managed Server.

To determine the process Id of the Managed Server run the following command and identify the pid of the WLS\_OSBn Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

- b. Watch the Node Manager Console: you should see a message indicating that WLS\_OSBn's floating IP has been disabled.
- c. Wait for the Node Manager to try a second restart of WLS\_OSBn. Node Manager waits for a fence period of 30 seconds before trying this restart.
- d. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

---

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

---

#### 5.14.22.2 Scaling Out the Topology (Adding Managed Servers to New Nodes)

When you scale out the topology, you add new managed servers configured with OSB to new nodes.

For scaling out the SOA servers, see the steps in [Section 5.13.21.2, "Scaling Out the Topology \(Adding Managed Servers to New Nodes\)."](#)

Follow the steps in this section for scaling out the topology for the OSB server.

Before performing the steps in this section, check that you meet these requirements:

- There must be existing nodes running managed servers configured with OSB within the topology.
- The new node optionally can access the existing home directories for WebLogic Server and OSB installation. (Use the existing installations in shared storage for creating a new WLS\_OSB managed server. You do not need to install WebLogic Server or OSB binaries in every new location in this case, but you do need to run pack and unpack to bootstrap the domain configuration in the new node, unless you are scaling the OSB server to machines containing other servers of the same domain (the SOA servers).)

---

**Note:** If there is no existing installation in shared storage, installing WebLogic Server, SOA, and OSB in the new nodes is required as described in [Section 5.14, "Configuring High Availability for Oracle Service Bus, with SOA Service Infrastructure and Component Service Engines."](#)

---

---

**Note:** When multiple servers in different nodes share an ORACLE\_HOME or WL\_HOME, Oracle recommends keeping the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. To update the oraInventory in a node and attach an installation in a shared storage to it, use `ORACLE_HOME/oui/bin/attachHome.sh`. To update the Middleware home list to add or remove a WL\_HOME, edit the `user_home/boa/beahomelist` file. See the following steps.

---

To scale out the topology:

1. On the new node, mount the existing Middleware home. It should include the OSB and SOA (if homes are shared) installation and the domain directory. Ensure that the new node has access to this directory, just like other nodes in the domain.
2. To attach ORACLE\_HOME in shared storage to the local Oracle Inventory, run this command:

```
OSBHOSTn>cd ORACLE_BASE/product/fmw/soa/
OSBHOSTn>./attachHome.sh -jreLoc ORACLE_BASE/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the `MW_HOME/boa/beahomelist` file and add `ORACLE_BASE/product/fmw` to it.

3. Log in to the Oracle WebLogic Administration Console.
4. Create a new machine for the new node that will be used, and add the machine to the domain.
5. Update the machine's Node Manager's address to map the IP of the node that is being used for scale out.
6. Use the WebLogic Server Administration Console to clone WLS\_OSB1 into a new managed server. Name it WLS\_OSBn, where n is a number and assign it to the new machine.

---

**Note:** These steps assume that you are adding a new server to node *n*, where no managed server was running previously.

---

7. For the listen address, assign the virtual host name to use for this new managed server. If you are planning to use server migration as recommended for this server, this virtual host name will enable it to move to another node. The virtual host name should be different from those in use by other managed servers (may be in the same or different domain) that are running in the nodes used by the OSB/SOA domain.

Perform these steps to set the managed server listen address:

- a. Log into the Oracle WebLogic Server Administration Console.
- b. In the Change Center, click **Lock & Edit**.
- c. Expand the **Environment** node in the Domain Structure window.
- d. Click **Servers**. The Summary of Servers page is displayed.

- e. Select the managed server whose listen address you want to update in the **Names** column of the table. The Setting page for that managed server is displayed.
- f. Set the Listen Address to OSBHOSTnVHN1.
- g. Click **Save**.
- h. Save and activate the changes.

The changes will not take effect until the managed server is restarted.

- 8. Update the cluster address to include the new server:
  - a. Select **Environment -> Cluster** from the Administration Console.
  - b. Click on the OSB\_Cluster server. The Settings Screen for the OSB\_Cluster appears.
  - c. In the Change Center, click **Lock & Edit**.
  - d. Add the new server's address and port to the **Cluster Address** field. For example:

```
OSBHOST1VHN1:8011, OSBHOST2VHN1:8011, OSBHOSTNVHN1:8011
```

- 9. Create JMS servers and persistent stores for OSB reporting/internal destinations on the new managed server.
  - a. Use the WebLogic Server Administration Console to create a new persistent store for the new WseeJMSServer and name it, for example, OSB\_rep\_JMSFileStore\_N. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.14.1.3, "Shared Storage Prerequisites."](#)

---

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

---

```
ORACLE_BASE/admin/domain_name/cluster_name/jms/OSB_rep_JMSFileStore _N
```

- b. Create a new JMS Server for OSB, for example, OSB\_rep\_JMSServer\_N. Use the OSB\_rep\_JMSFileStore\_N for this JMSServer. Target the OSB\_rep\_JMSServer\_N Server to the recently created managed server (WLS\_OSBn).
- c. Update the SubDeployment targets for the jmsresources OSB JMS Module to include the recently created OSB JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **jmsresources** (represented as a hyperlink in the **Names** column of the table). The Settings page for jmsResources appears. Open the **SubDeployments** tab. The subdeployment module for jmsresources appears.

---

---

**Note:** This subdeployment module name is a random name in the form of wlsbJMSServerXXXXXX resulting from the Configuration Wizard JMS configuration for the first two servers (WLS\_OSB1 and WLS\_OSB2).

---

---

Click the **wlsbjMSServerXXXXXX** subdeployment and update the targets to include the new **OSB\_rep\_JMSServer\_N** server.

10. Create JMS Servers, persistent stores and destinations for OSB JAX-RPC on the new managed server.

---

**Note:** WebLogic Advanced Web Services for JAX-RPC Extension uses regular (non-distributed) destinations to ensure that a locally processed request on a service gets enqueued only to a local member.

---

- a. Use the WebLogic Server Administration Console to create a new persistent store for the new **WseeJMSServer** and name it, for example, **Wsee\_rpc\_JMSFileStore\_N**. Specify the path for the store. This should be a directory on shared storage as recommended in [Section 5.14.1.3, "Shared Storage Prerequisites"](#):

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

*ORACLE\_BASE/admin/DOMAIN\_NAME/cluster\_name/jms/Wsee\_rpc\_JMSFileStore\_N*

- b. Create a new JMS Server for OSB JAX-RPC, for example, **OSB\_rpc\_JMSServer\_N**. Use the **Wsee\_rpc\_JMSFileStore\_N** for this JMSServer. Target the **OSB\_rpc\_JMSServer\_N** Server to the recently created Managed Server (**WLS\_OSBn**).
- c. Update the **WseeJMSModule** OSB JMS Module with destinations and the recently created OSB JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **WseeJmsModule** (represented as a hyperlink in the Names column of the table). The Settings page for **WseeJmsModule** appears. Follow steps d through v to complete this step.
- d. In the Change Center, click **Lock and Edit**. Click **New**.
- e. Select **Queue**.
- f. Click **Next**.
- g. Enter **DefaultCallbackQueue-WseeJmsServer\_auto\_n** as name for the queue.
- h. Enter **weblogic.wsee.DefaultCallbackQueue-WseeJmsServer\_auto\_n** as the JNDI name.
- i. Click **Next**.
- j. Click **Create a New Subdeployment**.
- k. Accept the default name.
- l. Click **OK**.
- m. Select **OSB\_rpc\_JMSServer\_n** as the target.
- n. Click **Finish**.
- o. Activate the changes.
- p. Update the local JNDI name for the destination by following steps q through u.

- q. In the Change Center, click **Lock & Edit**.
  - r. In the Settings for the WseeJmsModule page, click on the `DefaultCallbackQueue-WseeJmsServer_auto_n` destination.
  - s. In the general Configuration tab, click on **Advanced**.
  - t. Update the local JNDI name to `weblogic.wsee.DefaultCallbackQueue`.
  - u. Activate the changes.
  - v. Repeat steps d through n for the `DefaultQueue-WseeJmsServer_auto_n` queue, using `weblogic.wsee.DefaultQueue-WseeJmsServer_auto_n` as the JNDI name and `weblogic.wsee.DefaultQueue` as the local JNDI name.
11. Create a new SAF agent and target it to the newly added managed server, as follows:
- In the Oracle WebLogic Server Administration Console, expand **Services > Messaging > Store-and-Forward Agents**, and add a new SAF agent `ReliableWseeSAFAgent_auto_N`.
- Select persistent store `Wsee_rpc_JMSFileStore_N` (persistent store created for OSB JAX-RPC). Target the SAF Agent to the new managed server and activate changes.
12. If your Oracle Service Bus configuration includes one or more business services that use JMS request/response functionality, then you must also perform the following procedure using the Oracle Service Bus Console after adding the new managed server to the cluster:
- a. In the Change Center, click **Create** to create a session.
  - b. Using the Project Explorer, locate and select a business service that uses JMS request/response. Business services of this type display Messaging Service as their Service Type.
  - c. At the bottom of the View Details page, click **Edit**.
  - d. If there is a cluster address in the endpoint URI, add the new server to the cluster address.
  - e. On the Edit a Business Service - Summary page, click **Save**.
  - f. Repeat the previous steps for each remaining business service that uses JMS request/response.
  - g. In the Change Center, click **Activate**.
  - h. Restart the managed server.
  - i. Restart the Administration Server.
- The business services are now configured for operation in the extended domain.

---

**Note:** For business services that use a JMS MessageID correlation scheme, you must edit the connection factory settings to add an entry to the table mapping managed servers to queues. For information on how to configure queues and topic destinations, see "JMS Server Targeting" in *Oracle Fusion Middleware Configuring and Managing JMS for Oracle WebLogic Server*.

---



13. If your Oracle Service Bus configuration includes one or more proxy services that use JMS endpoints with cluster addresses, then you must also perform the following procedure using the Oracle Service Bus Console after adding the new managed server to the cluster:
  - a. In the Change Center, click **Create** to create a session.
  - b. Using the Project Explorer, locate and select a proxy service that uses JMS endpoints with cluster addresses.
  - c. At the bottom of the View Details page, click **Edit**.
  - d. If there is a cluster address in the endpoint URI, add the new server to the cluster address.
  - e. On the Edit a Proxy Service - Summary page, click **Save**.
  - f. Repeat the previous steps for each remaining proxy service that uses JMS endpoints with cluster addresses.
  - g. In the Change Center, click **Activate**.
  - h. Restart the managed server.

The proxy services are now configured for operation in the extended domain.

14. Update the Oracle Service Bus result cache Coherence configuration for the new server:
  - a. Log into Oracle WebLogic Server Administration Console. In the Change Center, click **Lock & Edit**.
  - b. In the Domain Structure window, expand the **Environment** node.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Click the name of the server (represented as a hyperlink) in the **Name** column of the table. The settings page for the selected server appears.
  - e. Click the **Server Start** tab.
  - f. Enter the following for WLS\_OSBn (on a single line, without a carriage returns):

```
-DOSB.coherence.localhost=osbhostnvhn1 -DOSB.coherence.localport=7890
-DOSB.coherence.wka1= osbhost1vhn1 -DOSB.coherence.wka1.port=7890
-DOSB.coherence.wka2= osbhost2vhn1 -DOSB.coherence.wka1.port=7890
```

---

**Note:** The preceding configuration assumes that servers WLS\_OSB1 and WLS\_OSB2 are running when WLS\_OSBn starts. This will enable WLS\_OSBn to join the coherence cluster started by either WLS\_OSB1 or WLS\_OSB2 using the WKA addresses specified. Also, make sure WLS\_OSB1 and WLS\_OSB2 are started before WLS\_OSBn is started when all three servers are restarted. This will ensure WLS\_OSBn join the cluster started by one of WLS\_OSB1 or WLS\_OSB2. If you want to have a configuration where the order in which the servers are started does not matter, you must add the host and port for WLS\_OSBn as WKA for WLS\_OSB1 and WLS\_OSB2, and also add WLS\_OSBn as WKA for WLS\_OSBn.

---

- g. Save and activate the changes (this change requires the Oracle Service Bus servers to be restarted to be effective).

15. Run the pack command on SOAHOST1 to create a template pack as follows:

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin

SOAHOST1> ./pack.sh -managed=true -domain=MW_HOME/user_
projects/domains/soadomain/
-template=soadomaintemplateScale.jar -template_name=soa_domain_templateScale
```

Run the following command on SOAHOST1 to copy the template file created to OSBHOSTn:

```
SOAHOST1> scp soadomaintemplateScale.jar oracle@OSBHOSTN:/ ORACLE_
BASE/product/fmw/soa/common/bin
```

Run the unpack command on OSBHOSTN to unpack the template in the managed server domain directory as follows:

```
OSBHOSTN> cd ORACLE_BASE/product/fmw/soa/common/bin

OSBHOSTN> ./unpack.sh -domain=ORACLE_BASE/product/fmw/user_
projects/domains/soadomain/ -template=soadomaintemplateScale.jar
```

16. Configure TX persistent store for the new server. This should be a location visible from other nodes as indicated in the recommendations about shared storage.

From the Administration Console, select **Server\_name** > **Services** tab. Under **Default Store**, in **Directory**, enter the path to the folder where you want the default persistent store to store its data files.

17. Disable host name verification for the new managed server. Before starting and verifying the WLS\_OSBN managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in OSBHOSTn. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required (the hostname verification setting is propagated to the cloned server).

To disable host name verification:

- a. In the Oracle Enterprise Manager Console, select **Oracle WebLogic Server Administration Console**.
  - b. Expand the **Environment** node in the Domain Structure window.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Select **WLS\_OSBN** in the **Names** column of the table.  
The Settings page for server appears.
  - e. Click the **SSL** tab.
  - f. Click **Advanced**.
  - g. Set Hostname Verification to **None**.
  - h. Click **Save**.
18. Start the Node Manager on the new node. To start the Node Manager, use the installation in shared storage from the existing nodes, and start Node Manager by passing the host name of the new node as a parameter as follows:

```
OSBHOSTN> WL_HOME/server/bin/startNodeManager <new_node_ip>
```

19. Start and test the new managed server from the Oracle WebLogic Server Administration Console:
  - a. Shut down all the existing managed servers in the cluster.
  - b. Ensure that the newly created managed server, WLS\_OSBn, is running.
  - c. Access the application on the newly created managed server (<http://vip:port/sbinspection.wsil>). The application should be functional.
20. Configure server migration for the new managed server.

---

**Note:** Since this new node is using an existing shared storage installation, the node is already using a Node Manager and an environment configured for server migration that includes netmask, interface, wlsifconfig script superuser privileges. The floating IP for the new OSB Managed Server is already present in the new node.

---

Configure server migration following these steps:

- a. Log into the Administration Console.
- b. In the left pane, expand **Environment** and select **Servers**.
- c. Select the server (represented as hyperlink) for which you want to configure migration from the **Names** column of the table. The Setting page for that server appears.
- d. Click the **Migration** tab.
- e. In the **Available** field, in the Migration Configuration section, select the machines to which to allow migration and click the right arrow.

For example, for new managed servers on OSBHOST1, which is already running WLS\_OSB1, select OSBHOST2. For new managed servers on OSBHOST2, which is already running WLS\_OSB2, select OSBHOST1.

---

**Note:** It is recommended to specify the least-loaded machine as the migration target for the new server. The required capacity planning must be completed so that this node has enough available resources to sustain an additional managed server.

---

- f. Select the **Automatic Server Migration Enabled** option. This enables the Node Manager to start a failed server on the target node automatically.
  - g. Click **Save**.
  - h. Restart the Administration Server, managed servers, and Node Manager.
21. Test server migration for this new server. Follow these steps from the node where you added the new server:
    - a. Abruptly stop the WLS\_OSBn managed server.

To do this, run "kill -9 <pid>" on the PID of the managed server. You can identify the PID of the node using "ps -ef | grep WLS\_OSBn".

---

**Note:** For Windows, terminate the Managed Server using the taskkill command. For example:

```
taskkill /f /pid <pid>
```

Where <pid> is the process Id of the Managed Server.

To determine the process Id of the Managed Server, run the following command and identify the pid of the WLS\_OSBN Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

- b. Watch the Node Manager Console for a message indicating that WLS\_OSBN's floating IP has been disabled.
- c. Wait for the Node Manager to try a second restart of WLS\_OSBN. Node Manager waits for a fence period of 30 seconds before trying this restart.
- d. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

---

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

---

## 5.15 Configuring High Availability for Oracle BAM

This section describes how to set up a 2 node Highly Available configuration for BAM. This includes the BAM Server and the BAM Web Applications.

This section includes these topics:

- [Section 5.15.1, "Preparing the Environment: Prerequisite Steps Before Setting up a High Availability Configuration for Oracle BAM"](#)
- [Section 5.15.2, "Installing Oracle HTTP Server on WEBHOST1"](#)
- [Section 5.15.3, "Installing Oracle Fusion Middleware Home"](#)
- [Section 5.15.4, "Enabling VIP0 and VIP1 on BAMHOST1"](#)
- [Section 5.15.5, "Running Oracle Fusion Middleware Configuration Wizard on BAMHOST1 to Create the WebLogic Server Oracle BAM Domain"](#)
- [Section 5.15.6, "Creating boot.properties for the Administration Server and for WLS\\_BAM1 on BAMHOST1"](#)
- [Section 5.15.7, "Starting the Administration Server on BAMHOST1"](#)
- [Section 5.15.8, "Disabling Host Name Verification for the Servers"](#)
- [Section 5.15.9, "Configuring a JMS Persistence Store for BAM UMS"](#)
- [Section 5.15.10, "Configuring a Default Persistence Store for Transaction Recovery"](#)
- [Section 5.15.11, "Untargeting the BAM Server System from BAMHOST2"](#)
- [Section 5.15.12, "Propagating the Domain Configuration from BAMHOST1 with pack/unpack Utilities"](#)

- [Section 5.15.13, "Starting Node Manager on BAMHOST1 and BAMHOST2"](#)
- [Section 5.15.14, "Starting the Oracle BAM System"](#)
- [Section 5.15.15, "Configuring Oracle RAC Failover for the WLS\\_BAM Servers"](#)
- [Section 5.15.16, "Configuring the BAM Web Applications to Use the BAM Server in BAMHOST1"](#)
- [Section 5.15.17, "Configuring Oracle HTTP Servers for the Administration Server and the WLS\\_BAMn Managed Servers"](#)
- [Section 5.15.18, "Validating Access through Oracle HTTP Server"](#)
- [Section 5.15.19, "Configuring Server Migration for the WLS\\_BAM Servers"](#)

The architecture targeted for the configuration steps is as follows:

---

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

---

---

**Note:** The architectures and deployment procedures defined in this guide enable simple clustered deployments. The procedures described in these chapters can be used as a building block to enable this and other similar high availability topologies for these Fusion Middleware components. It is also expected that production deployments will use other required procedures, such as associating security policies with a centralized LDAP server. For complete details of secured, multi-tiered architectures, and deployment procedures, please refer to the Enterprise Deployment Guide for the component you are configuring.

---

---

[Figure 5–55](#) represents the example architecture that configuration steps in this section create.


**Figure 5–55 Oracle BAM High Availability Architecture**

Figure 5–55 illustrates a two-node cluster running Oracle BAM Web Applications. One of them is also running Oracle BAM Server. The Oracle WebLogic Servers are frontended by Oracle HTTP Servers, which load balance incoming requests to BAM Web Applications. A load balancer front ends the system and distributes requests from clients to the two HTTP Servers. An Oracle RAC database stores metadata and BAM schemas.

## 5.15.1 Preparing the Environment: Prerequisite Steps Before Setting up a High Availability Configuration for Oracle BAM

The following sections provide prerequisite steps before setting up an Oracle BAM high availability configuration:

### 5.15.1.1 Database Prerequisites

Oracle BAM requires the presence of a supported database.

To check if your database is certified or to see all certified databases, refer to the "Certified Databases" section in the Certification Document:

[http://www.oracle.com/technology/software/products/ias/files/fusion\\_certification.html](http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html)

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

### 5.15.1.2 VIP and IPs Prerequisites

As shown in Table 5–19, you configure the Administration and the BAM Server in BAMHOST1 to listen on two different virtual IPs, as shown in the architecture diagram. This requires the provisioning of the corresponding VIP in the box and related host names in the DNS system in your network. Ensure that the different VIPs are available and are reachable from BAMHOST1 and BAMHOST2 before running the installation.

**Table 5–19 Virtual Hosts**

| Virtual IP | VIP Maps to... | Description                                                                                                                                                                                                                                                          |
|------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VIP0       | BAMHOST1VHN0   | BAMHOST1VHN0 is the virtual host name that is the listen address for the Administration Server and fails over with manual failover of the Administration Server. It is enabled on the node where the Administration Server process is running (BAMHOST1 by default). |
| VIP1       | BAMHOST1VHN1   | BAMHOST1VHN1 is the virtual host name that maps to the listen address for WLS_BAM1 and fails over with server migration of this managed server. It is enabled on the node where WLS_BAM1 process is running (BAMHOST1 by default).                                   |

### 5.15.1.3 Installing and Configuring the Database Repository

This section describes how to install and configure the database repository.

#### Oracle Clusterware

- For 10g Release 2 (10.2), see the *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

#### Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

- When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

### Oracle Real Application Clusters

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

### Database Initialization Parameters

Ensure that the following initialization parameter is set to the required value. It is checked by Repository Creation Assistant.

**Table 5–20 Required Initialization Parameters**

| Parameter | Required Value | Parameter Class |
|-----------|----------------|-----------------|
| PROCESSES | 300 or greater | Static          |

To check the value of the initialization parameter using SQL\*Plus, you can use the SHOW PARAMETER command.

As the SYS user, issue the SHOW PARAMETER command as follows:

```
SQL> SHOW PARAMETER processes
```

Set the initialization parameter using the following command:

```
SQL> ALTER SYSTEM SET processes=300 SCOPE=SPFILE
```

Restart the database.

---

**Note:** The method that you use to change a parameter's value depends on whether the parameter is static or dynamic, and on whether your database uses a parameter file or a server parameter file. See the *Oracle Database Administrator's Guide* for details on parameter files, server parameter files, and how to change parameter values.

---

### 5.15.1.4 Using Oracle Fusion Middleware Repository Creation Utility to Load Oracle Fusion Middleware Schemas

Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and BAM schemas into a Real Application Cluster database before you install Oracle Fusion Middleware BAM components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. See Oracle Fusion Middleware Installation Concepts and Repository Creation Utility User's Guide for more information about installing RCU.

#### 5.15.1.4.1 Running RCU

Run RCU to install the required metadata for Oracle Fusion Middleware 11g.

1. Start RCU using the following command:

```
RCU_HOME/bin/rcu &
```

2. In the Welcome screen, click **Next**.



3. In the Create Repository screen, select **Create** to load component schemas into a database, click **Next**.
  4. In the Database Connection Details screen, enter connection information for your database:
    - Database Type: Select **Oracle Database**.
    - Host Name: Enter the name of the node that is running the database. For an Oracle RAC database, specify the VIP name, or one of the node names as the host name: **BAMDBHOST1VIRTUAL**.
    - Port: The port number for the database: **1521**
    - Service Name: Enter the service name of the database:  
**bamha.mycompany.com**
    - Username: **SYS**
    - Password: Enter the password for the SYS user.
    - Role: **SYSDBA**
  5. Click **Next**.
  6. If you receive the following warning message:
 

The database you are connecting is with non-UTF8 charset, if you are going to use this DB for multilingual support, you may have data loss. If you are not using for multilingual support you can continue, otherwise we strongly recommend using UTF-8 database.
  7. Click **Ignore** or **Stop**.
  8. In the Select Components screen, select **Create a New Prefix**, and enter a prefix to use for the database schemas, for example, **BAMHA**

Write down the schema names so they are available in later procedures.

    - Select the following schemas:
      - AS Common Schemas:  
Metadata Services
      - SOA Infrastructure:  
Business Activity Monitoring  
User Messaging
  9. Click **Next**.
  10. In the Schema Passwords screen, enter passwords for the main and additional (auxiliary) schema users, and click **Next**.
  11. In the Map Tablespaces screen, choose the tablespaces for the selected components, and click **Next**.
  12. In the Summary screen, click **Create**.
  13. In the Completion Summary screen, click **Close**.
- See *Oracle Fusion Middleware Installation Concepts and Repository Creation Utility User's Guide* for more information about using RCU.

### 5.15.1.5 Configuring Virtual Server Names and Ports for the Load Balancer

This section describes the Load Balancer prerequisites for deploying an Oracle BAM high availability environment.

#### Load Balancers

Oracle BAM uses a hardware load balancer when deployed in a high availability configuration with two Oracle HTTP Servers as web tier components. The hardware load balancer should have the following features:

- Ability to load-balance traffic to a pool of real servers through a virtual host name: Clients access services using the virtual host name (instead of using actual host names). The load balancer can then load balance requests to the servers in the pool.
- Port translation configuration: The load balancer should have the ability to perform port translation, where it allows incoming requests received on one port to be routed to a server process running on a different port. For example, a request received on port 80 can be routed to port 7777.
- Monitoring of ports (HTTP/HTTPS)
- Virtual servers and port configuration: Ability to configure virtual server names and ports on your external load balancer, and the virtual server names and ports must meet the following requirements:
  - The load balancer should allow configuration of multiple virtual servers. For each virtual server, the load balancer should allow configuration of traffic management on more than one port. For example, the load balancer typically needs to be configured with a virtual server and ports for HTTP and HTTPS traffic.
  - The virtual server names must be associated with IP addresses and be part of your DNS. Clients must be able to access the load balancer through the virtual server names.
- Ability to detect node failures and immediately stop routing traffic to the failed node.
- Resource monitoring/port monitoring/ process failure detection: The load balancer must be able to detect service and node failures (through notification or some other means) and to stop directing non-Oracle Net traffic to the failed node. If your load balancer has the ability to automatically detect failures, you should use it.
- Fault-tolerant mode: It is highly recommended that you configure the load balancer to be in fault-tolerant mode.
- Other: It is highly recommended that you configure the load balancer virtual server to return immediately to the calling client when the back-end services to which it forwards traffic are unavailable. This is preferred over the client disconnecting on its own after a timeout based on the TCP/IP settings on the client machine.
- Sticky routing capability: Ability to maintain sticky connections to components based on cookies or URL.
- SSL acceleration: SSL acceleration is a method of offloading the processor-intensive public key encryption algorithms involved in SSL transactions to a hardware accelerator.

This feature is recommended but not required.

Table 5–21 shows some example virtual server names to use for the external load balancer in the Oracle BAM high availability environment:

**Table 5–21 Virtual Server Names and Ports for the BAM External Load Balancer**

| Component                                           | Virtual Server Name |
|-----------------------------------------------------|---------------------|
| Oracle BAM                                          | bam.mycompany.com   |
| WebLogic Server Administration Console              | admin.mycompany.com |
| Oracle Enterprise Manager Fusion Middleware Control | admin.mycompany.com |

### Virtual Server Names

This section describes the virtual server names that should be set up for the high availability deployments described in this section:

#### **bam.mycompany.com**

bam.mycompany.com is a virtual server name that acts as the access point for all HTTP traffic to the runtime BAM Web components. Traffic to the both SSL and non-SSL ports is configured and typically non-SSL is redirected to SSL. Clients access this service using the address bam.mycompany.com:443. This virtual server is defined on the load balancer.

#### **admin.mycompany.com**

This virtual server acts as the access point for all internal HTTP traffic that gets directed to the administration services. The incoming traffic from clients is non-SSL enabled. Thus, the clients access this service using the address admin.mycompany.com:80 and in turn forward these to ports 7777 on WEBHOST1 and WEBHOST2. The services accessed on this virtual host include the WebLogic Administration Server Console and Oracle Enterprise Manager.

In addition, ensure that the virtual server names are associated with IP addresses and are part of your Domain Name System (DNS). The computers on which Oracle Fusion Middleware is running must be able to resolve these virtual server names.

## 5.15.2 Installing Oracle HTTP Server on WEBHOST1

To install Oracle HTTP Server on WEBHOST1:

1. Verify that the servers meet the following requirements:
  - The system, patch, kernel, and other requirements meet the requirements specified in the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite*.
  - Port 7777 is not used by any service on the nodes. You can verify this by running the following command:

Unix:

```
netstat -an | grep LISTEN | grep ":7777"
```

Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7777"
```

If the ports are in use, make them available.

2. On UNIX platforms, if the `/etc/oraInst.loc` or `/var/opt/oracle/oraInst.loc` file exists, check that its contents are correct.

Specifically, check that the inventory directory is correct, and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, skip this step.

3. Start Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:  
For UNIX, run this command: `runInstaller`.  
For Windows, double-click **setup.exe**.
4. In the Welcome screen, click **Next**.
5. In the Select Installation Type screen, select **Install and Configure**, and click **Next**.
6. In the Prerequisite Checks screen, ensure that all the prerequisites are met, and click **Next**.
7. In the Specify Installation Location screen, set the location to:  
`/u01/app/oracle/product/11.1.1/ohs_1`
8. Click **Next**.
9. In the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Do not select **Associate Selected Components with WebLogic Domain**.
  - Click **Next**.
10. In the Specify Component Details screen, enter the following values:
  - Instance Home Location: `/u01/app/oracle/product/11.1.1/ohs_1/instances/ohs_instance1`
  - Instance Name: `ohs_instance1`
  - OHS Component Name: `ohs1`
11. Click **Next**.
12. In the Configure Ports screen, do the following:
  - Select **Specify Ports** using Configuration File and copy the `staticports.ini` template file from your installation disk (the file is located in the `/Disk1/stage/Response` directory) to your user's home. Then use the **Browse** button to select this file.
  - Click **View/Edit File** to open the `staticports.ini` file in an editor.
  - Change the Oracle HTTP Server port in that file to **7777**.
  - Save the file.
13. Click **Next**.
14. In the Configuration Summary screen, ensure that the selections are correct, and click **Install**.
15. In the Installation Progress screen:  
For UNIX systems, a dialog box appears prompting you to run the `oracleRoot.sh` script. Open a command window and run the script, following the prompts.  
Click **Next**.

16. In the Configuration screen, several configuration assistants are launched in succession. When the configuration assistants are finished, the Configuration Completed screen appears.
17. In the Configuration Completed screen, click **Finish** to exit.
18. Repeat the steps for WEBHOST2 and configure your load balancing router with a pool containing both the WEBHOST1 and WEBHOST2 addresses.

#### 5.15.2.1 Validating Oracle HTTP Server

To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URL a browser:

**HTTP://WEBHOST1:7777/**

If Oracle HTTP Server is set up correctly, the **Hello World** page appears in the browser.

### 5.15.3 Installing Oracle Fusion Middleware Home

This section describes how to install Oracle Fusion Middleware on all nodes in the application tier that run Oracle WebLogic Server and Oracle Fusion Middleware BAM.

Install the following Oracle Fusion Middleware components:

- Oracle WebLogic Server (see [Section 5.15.3.1, "Installing Oracle WebLogic Server."](#))
- Oracle SOA Suite (see [Section 5.15.3.2, "Installing Oracle BAM Using the Oracle SOA Suite Installer"](#))

#### 5.15.3.1 Installing Oracle WebLogic Server

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

To install Oracle WebLogic Server on all nodes in the application tier:

1. Start Oracle WebLogic Server Installer.
2. In the Welcome screen, click **Next**.
3. In the Choose Middleware Home Directory screen:
  - Select **Create a New Middleware Home**
  - For the **Middleware Home Directory** field, enter **MW\_HOME**
  - Click **Next**.
4. In the Register for Security Updates screen, enter your contact information for security update notifications, and click **Next**.
5. In the Choose Install Type screen, select **Custom**, and click **Next**.
6. In the Choose Products and Components screen, click **Next**.
7. In the JDK Selection screen, select only **Oracle JRockit 160\_20\_D1.0.1-2124 SDK**, and click **Next**.
8. In the Choose Product Installation Directories screen, accept the following directory:

WL\_HOME

Click **Next**.

9. In the Installation Summary screen, click **Next**.
10. In the Installation Complete screen, deselect **Run QuickStart**, and click **Done**.

### 5.15.3.2 Installing Oracle BAM Using the Oracle SOA Suite Installer

Oracle BAM is installed as part of the Oracle SOA Suite. The 11g Oracle SOA Suite Installer is used to install the required Oracle BAM binaries on all the nodes in the application tier:

1. On Linux platforms, if the `/etc/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, you can skip this step.

2. Start the Oracle FMW 11g SOA Suite Installer:

On UNIX:

```
BAMHOST1> runInstaller
```

On Windows:

```
BAMHOST1> setup.exe
```

When the Oracle FMW 11g SOA Suite Installer prompts you for a **JRE/JDK location** enter Oracle SDK location created in Oracle WebLogic Server installation, for example, `MW_HOME/jrockit_160_<version>`.

3. In the Welcome screen, click **Next**.
4. In the Prerequisite Checks screen, verify that the checks complete successfully, and click **Next**.
5. In the Specify Installation Location screen:
  - For Middleware Home, enter **MW\_HOME**
  - For Oracle Home Directory, enter **soa**Click **Next**.
6. In the Installation Summary screen, click **Install**.
7. In the Installation Complete screen, click **Finish**.

### 5.15.4 Enabling VIP0 and VIP1 on BAMHOST1

For information about configuring virtual IPs for the Administration Server and configuring the Administration Server for high availability, see [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

### 5.15.5 Running Oracle Fusion Middleware Configuration Wizard on BAMHOST1 to Create the WebLogic Server Oracle BAM Domain

Run Oracle Fusion Middleware Configuration Wizard from the SOA home directory to create a domain containing the Administration Server and Oracle BAM components. Ensure that the database where you installed the repository is running. For Oracle RAC databases, all the instances must be running.

1. Change the directory to the location of Oracle Fusion Middleware Configuration Wizard, located in the SOA/BAM home directory:

```
BAMHOST1> cd ORACLE_HOME/common/bin
```

2. Start Oracle Fusion Middleware Configuration Wizard

For UNIX:

```
BAMHOST1> ./config.sh
```

For Windows:

```
BAMHOST1> config.cmd
```

3. In the Welcome screen, select **Create a New WebLogic Domain**, and click **Next**.

4. In the Select Domain Source screen, select **Generate a domain configured automatically to support the following products**, and select the following products (**Basic WebLogic Server domain - 10.3.6.0** is selected by default):

- Oracle Business Activity Monitoring
- Oracle WSM-PM (selected by default)
- Enterprise Manager
- JRF (selected by default)

Click **Next**.

5. In the Specify Domain Name and Location screen, make the following entries:

- Domain Name: **bamdomain**
- Domain Location: accept the default
- Application Location: accept the default

Click **Next**.

6. In the Configure Administrator Username and Password screen, enter the username and password to be used for the domain's administrator. Click **Next**.

7. In the Configure Server Start Mode and JDK screen, make the following selections:

- WebLogic Domain Startup Mode: select **Production Mode**
- JDK Selection: select **Oracle SDK1.6.0\_05**.

Click **Next**.

8. In the Configure JDBC Component Schema screen:

Figure 5–56 Configure JDBC Component Schema

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Vendor:  DBMS/Service:

Driver:  Host Name:

Schema Owner:  Port:

Schema Password:

RAC configuration for component schemas:

☐ Convert to GridLink ☐ Convert to RAC multi data source ☒ Don't convert

|                                     | Component Schema       | DBMS/Service | Host Name        | Port | Schema Owner   | Schema Password |
|-------------------------------------|------------------------|--------------|------------------|------|----------------|-----------------|
| <input checked="" type="checkbox"/> | SOA Infrastructure     | orcl         | dbhostexample.co | 1521 | SOAHA SOAINFRA | *****           |
| <input checked="" type="checkbox"/> | User Messaging Service | orcl         | dbhostexample.co | 1521 | SOAHA ORASDPM  | *****           |
| <input checked="" type="checkbox"/> | OWSM MDS Schema        | orcl         | dbhostexample.co | 1521 | SOAHA MDS      | *****           |
| <input checked="" type="checkbox"/> | SOA MDS Schema         | orcl         | dbhostexample.co | 1521 | SOAHA MDS      | *****           |

Exit Help Previous Next

- a. Ensure that the following data source appears on the screen. The user names shown in Table 5–22, assume that **bamha** was used as the prefix for schema creation from RCU.

Table 5–22 Configuring Values for Data Sources

| Data Source        | User Name     |
|--------------------|---------------|
| BAMDataSource      | bamha_orabam  |
| OraSDPMDDataSource | bamha_orasdpm |
| mds-owsm           | bamha_mds     |

- b. For the RAC configuration for component schemas, select **Convert to GridLink** or **Convert to RAC multi data source**. For single database configuration, select **Don't convert**. If you select **Don't Convert**, continue to Step 10.
- c. Click **Next**.
9. In the Configure JDBC Component Schema screen, enter values for the following fields, specifying connect information for the Oracle RAC database that was seeded with RCU:

For multi data sources, enter values for these fields:



**Figure 5–57 Configure RAC Multi Data Source Component Schema Screen**

**Fusion Middleware Configuration Wizard**

### Configure RAC Multi Data Source Component Schema

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for RAC Service-Instance

Service Name: soaha.mycompany.com

Username: Varies among component schemas

Password: \*\*\*\*\*

| Host Name     | Instance Name | Port |
|---------------|---------------|------|
| soadbhost1vir | soahadb1      | 1521 |
| soadbhost2vir | soahadb2      | 1521 |

Add Delete

|                                     | RAC Component Schema   | Service Name        | Schema Owner   | Schema Password |
|-------------------------------------|------------------------|---------------------|----------------|-----------------|
| <input checked="" type="checkbox"/> | SOA Infrastructure     | soaha.mycompany.com | SOAHA SOAINFRA | *****           |
| <input checked="" type="checkbox"/> | User Messaging Service | soaha.mycompany.com | SOAHA ORASDPM  | *****           |
| <input checked="" type="checkbox"/> | OWSM MDS Schema        | soaha.mycompany.com | SOAHA MDS      | *****           |
| <input checked="" type="checkbox"/> | SOA MDS Schema         | soaha.mycompany.com | SOAHA MDS      | *****           |

Exit Help Previous Next

For GridLink data sources, enter values for these fields:

**Figure 5–58 Configure GridLink RAC Component Schema**

**Fusion Middleware Configuration Wizard**

### Configure GridLink RAC Component Schema

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for GridLink Connections; Vers

Service Name: mycompany.com

Username: Varies among component schemas

Password: \*\*\*\*\*

Enable FAN: ☒

Enable SSL: ☐

Wallet File: Enter a value Browse

Wallet Password: Enter a value

| Service Listener | Port | Protocol |
|------------------|------|----------|
| soadbhost12-scn  | 1521 | TCP      |

| ONS Host        | Port |
|-----------------|------|
| soadbhost12-scn | 6200 |

Add Delete

|                                     | RAC Component Schema   | Service Name  | Schema Owner  | Schema Password |
|-------------------------------------|------------------------|---------------|---------------|-----------------|
| <input checked="" type="checkbox"/> | BAM Schema             | mycompany.com | BAMHA ORABAM  | *****           |
| <input checked="" type="checkbox"/> | User Messaging Service | mycompany.com | BAMHA ORASDPM | *****           |
| <input checked="" type="checkbox"/> | OWSM MDS Schema        | mycompany.com | BAMHA MDS     | *****           |

Exit Help Previous Next

With all schemas selected, enter values for the following fields, specifying the connect information for the Oracle RAC database that was seeded with RCU.

- a. Driver: *For RAC:* Select **Oracle driver (Thin) for RAC Service-Instance connections, Versions:10, 11**. *For GridLink:* **Oracle Driver (Thin) for GridLink Connections, Versions:10 and later**
- b. Service Name: Enter the service name of the database. **For a GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example,  
`<mydbservice>.us.oracle.com`

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example,  
`us.oracle.com`

---

- c. Username: Enter the complete user name (including the prefix) for the schemas.
- d. Password: enter the password to use to access the schemas.
- e. **For a GridLink data source**, enter the SCAN address in the Listener Address field and SCAN port in the Port field. Enter the ONS host and port information in the ONS Host and Port fields, respectively.

---

**Note:** Oracle recommends that you use SCAN addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

---

**For a multi data source**, click **Add**. Enter details for Oracle RAC instances.

- f. Update each schema: select one data source at a time and add the appropriate details.  
  
Ensure that the information is entered for all schemas: **BAM Schema Only**, **User Messaging Service Schema Only**, **OWSM MDS Schema Only**.
- g. Click **Next**.
10. In the Test JDBC Data Sources screen, the connections are tested automatically. The Status column shows the results. Ensure that all connections were successful. If not, click **Previous** to correct your entries.  
  
Click **Next** when all the connections are successful.
11. In the Select Optional Configuration screen, select **Administration Server**, **Managed Servers**, **Cluster and Machines** and **Deployment and services**, and click **Next**.
12. In the Configure the Administration Server screen, enter the following values:
  - Name: **AdminServer**

- Listen Address: **VIP0**
- Listen Port: **7001**
- SSL listen port: N/A
- SSL enabled: leave unchecked

Click **Next**.

13. In the Configure Managed Servers screen, add the following managed servers:

**Table 5–23 Configuring Managed Servers**

| Name     | Listen Address                                               | Listen Port | SSL Listen Port | SSL Enabled |
|----------|--------------------------------------------------------------|-------------|-----------------|-------------|
| WLS_BAM1 | BAMHOST1VHN1                                                 | 8001        | n/a             | No          |
| WLS_BAM2 | BAMHOST2 (BAM Server WLS_BAM2 does not use server migration) | 8001        | n/a             | No          |

---

**Note:** Although the standard recommendation is to run custom applications and other systems in a separate WebLogic Managed Server, the creation of custom WLS managed servers described in [Figure 5–55](#) is not addressed here.

---

Click **Next**.

14. In the Configure Clusters screen, add the following cluster:

- Name: **BAM\_Cluster**
- Cluster Messaging Mode: **unicast**
- Multicast Address: N/A
- Multicast Port: N/A
- Cluster Address: Leave empty

Click **Next**.

15. In the Assign Servers to Clusters screen, assign the following servers to BAM\_Cluster:

- WLS\_BAM1
- WLS\_BAM2

Click **Next**.

16. In the Configure Machines screen:

- Delete the **LocalMachine** that appears by default.
- Click the **Unix Machine** tab, and add the following machines:

**Table 5–24 Configuring Machines**

| Name     | Node Manager Listen Address |
|----------|-----------------------------|
| BAMHOST1 | BAMHOST1                    |
| BAMHOST2 | BAMHOST2                    |

Leave all other fields to their default values, and click **Next**.

17. In the Assign Servers to Machines screen, assign servers to machines as follows:

- BAMHOST1: **AdminServer, WLS\_BAM1**
- BAMHOST2: **WLS\_BAM2**

Click **Next**.

18. In the Target Deployments to Cluster or Services screen, ensure the following targets:

- User Messaging Deployments should be targeted only to BAM\_Cluster. (The usermessaging-xmpp, usermessaging-smpp, and usermessaging-voicexml applications are optional.)
- wsm-pm should be targeted only to BAM\_Cluster.
- The DMS Application should be targeted to BAM\_Cluster and Administration Server.
- The oracle.rules.\*, oracle.sdp.\* and oracle.bam.\* deployments should be targeted only to BAM\_Cluster.
- The oracle.wsm.seedpolicies library should be targeted only to the BAM\_Cluster.

19. In the Target Services to Cluster or Servers screen, ensure the following targets:

- WSM Startup Class should be targeted only to BAM\_Cluster.
- mds-wsm, mds-wsm-rac0, and mds-wsm-rac1 should be targeted to both BAM\_Cluster and AdminServer.
- OrasDPMDatasource, OrasDPMDatasource-rac0, and OrasDPMDatasource-rac1 should be targeted to the BAM\_Cluster.
- OWSM Startup Class is only targeted to BAM\_Cluster \*DMS Startup is targeted both to BAM\_Cluster and AdminServer\*mds-wsm, mds-wsm-rac0, and mds-wsm-rac1 should be targeted to both BAM\_Cluster and AdminServer
- mds-soa, mds-soa-rac0, and mds-soa-rac1 should be targeted to both BAM\_Cluster and AdminServer.

20. In the Configuration Summary screen, click **Create**.

21. In the Creating Domain screen, click **Done**.

### 5.15.6 Creating boot.properties for the Administration Server and for WLS\_BAM1 on BAMHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a boot.properties file for the Administration Server on BAMHOST1.

For the Administration Server:

1. Create the following directory:

```
mkdir ORACLE_BASE/product/fmw/user_
projects/domains/bamdomain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` in the directory created in the previous step, and enter the following lines in the file:

```
username=adminuser
password=password
```

---

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible in order for the entries to be encrypted.

---

---

### 5.15.7 Starting the Administration Server on BAMHOST1

To start the Administration Server on BAMHOST1 run the following commands:

For Linux:

```
BAMHOST1> cd MW_HOME/user_projects/domains/bamdomain/bin
BAMHOST1> ./startWebLogic.sh
```

For Windows:

```
startWebLogic.cmd
```

To verify that the Administration Server is properly configured:

1. In a browser, go to `http://VIP0:7001/console`.
2. Log in as the administrator.
3. Verify that the WLS\_BAM1 and WLS\_BAM2 managed servers are listed.
4. Verify that the BAM\_Cluster cluster is listed.
5. Verify that you can access Enterprise Manager at `http://VIP0:7001/em`.

### 5.15.8 Disabling Host Name Verification for the Servers

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and validating the topology, and enable them again once the high availability topology configuration is complete.

To disable host name verification:

1. In Oracle WebLogic Server Administration Console, select **Administration Server**, **SSL**, and then **Advanced**.
2. Set **Hostname Verification** to **None**.
3. In Oracle WebLogic Server Administration Console, select **WLS\_BAM1**, **SSL**, and then **Advanced**.
4. Set **Hostname Verification** to **None**.
5. In Oracle WebLogic Server Administration Console, select **WLS\_BAM2**, **SSL**, and then **Advanced**.
6. Set **Hostname Verification** to **None**.
7. Restart the Administration Server for the changes to take affect.

8. Save and activate the changes.

### 5.15.9 Configuring a JMS Persistence Store for BAM UMS

Configure the location for all of the persistence stores as a directory that is visible from both BAMHOST1 and BAMHOST2. This is required in order to enable the resume of transactions when a server is migrated to a different node. By using a shared location for the persistent stores from both nodes, it is guaranteed that no messages are lost should a failover take place.

---

**Note:** See [Considerations for Using File Stores on NFS](#) for important information about JMS messages and transaction logs, as well as releasing locks on file stores.

---

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Services** node and then click the **Persistence Stores** node.  
The Summary of Persistence Stores page appears.
3. Select the a persistent store (represented as a hyperlink) from the Name column of the table.  
The Settings page for the persistence store appears.
4. In the Configuration tab, enter the location on a persistent storage solution (such as NAS or SAN) that is available to other servers in the cluster in the Directory field. Specifying this location enables pending JMS messages to be sent.
5. Click **Save**.
6. Repeat these steps for all persistent stores.
7. Save and Activate the changes.

### 5.15.10 Configuring a Default Persistence Store for Transaction Recovery

Each server has a transaction log that stores information about committed transactions that are coordinated by the server that may not have been completed. The WebLogic Server uses this transaction log for recovery from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to a server and its backup servers.

---

**Note:** Preferably, this location should be a dual-ported SCSI disk or on a Storage Area Network (SAN).

---

To set the location for the default persistence store:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node and then click the **Servers** node.  
The Summary of Servers page appears.
3. Click the name of the server (represented as a hyperlink) in Name column of the table.

The settings page for the selected server appears and defaults to the Configuration tab.

4. Click the **Services** tab.
5. In the Default Store section of the page, enter the path to the folder where the default persistent stores will store its data files.
6. Save and Activate the changes.

---

**Note:** To enable migration of the Transaction Recovery Service, specify a location on a persistent storage solution that is available to other servers in the cluster. This directory must also exist before you restart the server.

---

### 5.15.11 Untargeting the BAM Server System from BAMHOST2

Because the BAM server component in BAM is a singleton, you must untarget it from one of the WLS\_BAM servers before you configure it for server migration. Otherwise the system would use two active BAM Servers which could cause different data inconsistencies. This way, BAM Web applications run in both BAMHOST1 and BAMHOST2, but BAM Server is initially active only in BAMHOST1.

In this step, you remove the BAM server runtime from WLS\_BAM2. To untarget the BAM server artifacts from WLS\_BAM2:

1. Log into the Oracle WebLogic Administration Console at `http://BAMHOST1VHN0:7001/console`.
2. In the Domain Structure window, choose **Environment** and then **Servers**.  
The Summary of Servers page appears.
3. Select **WLS\_BAM2** in Name column of the table.  
The Settings page for WLS\_BAM2 appears.
4. Click the **Deployments** tab.
5. Select the **oracle-bam** application from the Name column of the table.  
The Settings page for the oracle-bam application appears.
6. Click the **Targets** tab.
7. In the Change Center, click **Lock & Edit**.
8. Change the targets for the modules as described in [Table 5-25](#).  
Save and Activate the changes

---

**Note:** You must target all of these components as described in [Table 5-25](#), as incorrect targeting can prevent the BAM system from starting.

---

**Table 5–25** *oracle-bam Component Target Types*

| Component                         | Type                   | Target      |
|-----------------------------------|------------------------|-------------|
| oracle-bam(11.1.1)                | Enterprise Application | BAM_Cluster |
| /oracle/bam                       | WEBAPP                 | WLS_BAM1    |
| oracle-bam-adc-ejb.jar            | EJB                    | WLS_BAM1    |
| oracle-bam-ems-ejb.jar            | EJB                    | WLS_BAM1    |
| oracle-bam-eventengine-ejb.jar    | EJB                    | WLS_BAM1    |
| oracle-bam-reportcache-ejb.jar    | EJB                    | WLS_BAM1    |
| oracle-bam-statuslistener-ejb.jar | EJB                    | WLS_BAM1    |
| OracleBAM                         | WEBAPP                 | BAM_Cluster |
| OracleBAMWS                       | WEBAPP                 | BAM_Cluster |
| sdpMessagingClient-ejb.jar        | EJB                    | WLS_BAM1    |

### 5.15.12 Propagating the Domain Configuration from BAMHOST1 with pack/unpack Utilities

Follow these steps to propagate the domain configuration to BAMHOST1 using Pack/Unpack utilities:

1. Run the following pack command on BAMHOST1 to create a template pack:

```
BAMHOST1> cd ORACLE_COMMON_HOME/common/bin
BAMHOST1> ./pack.sh -managed=true
domain=ORACLE_BASE/product/fmw/user_projects/domains/bamdomain/
-template=bamdomaintemplate.jar
-template_name=bam_domain_template
```

2. Run the unpack command on BAMHOST2 to unpack the propagated template:

```
BAMHOST2> cd ORACLE_COMMON_HOME/common/bin
BAMHOST2> ./unpack.sh
-domain=MW_HOME/user_projects/domains/bamdomain/
-template=bamdomaintemplate.jar
```

### 5.15.13 Starting Node Manager on BAMHOST1 and BAMHOST2

To start Node Manager on BAMHOST1 and BAMHOST2:

1. Run the `setNMProps.sh` script, located in the `ORACLE_COMMON_HOME/common/bin` directory.

This script sets the `StartScriptEnabled` property to **true** before starting Node Manager, and allows starting servers from the Administration Console (the environment variables required for the BAM Servers are set in the default start script in the domain directory):

```
BAMHOST1> cd ORACLE_COMMON_HOME/common/bin
BAMHOST1> ./setNMProps.sh
```

---

**Note:** You must use the `StartScriptEnabled` property to avoid class loading failures and other problems.

---



2. Start Node Manager using the following command:

```
BAMHOST1> cd WL_HOME/server/bin
BAMHOST1> ./startNodeManager.sh
```

3. Repeat step1 and 2 for Node Manager in BAMHOST2.

### 5.15.14 Starting the Oracle BAM System

To start the WLS\_BAM1 managed server on BAMHOST1:

1. Start the WLS\_BAM1 managed servers:
  - a. Log into the WebLogic Server Administration Console using the following URL:

```
http://bamhost1vhn0:7001/console
```

- b. In the Domain Structure window, expand the **Environment** node and then select **Servers**.

The Summary of Servers page appears.

- c. Click the **Control** tab.
  - d. Select **WLS\_BAM1** from the Servers column of the table.
  - e. Click **Start**.
2. Access <http://bamhost1vhn1:9001/OracleBAM> to verify status of WLS\_BAM1.

3. Start the WLS\_BAM2 managed servers:

- a. Log into the WebLogic Server Administration Console at <http://bamhost1vhn0:7001/console>.
  - b. In the Domain Structure window, expand the **Environment** node and then select **Servers**.

The Summary of Servers page appears.

- c. Click the **Control** tab.
  - d. Select **WLS\_BAM2** from the Servers column of the table.
  - e. Click **Start**.
4. Access [HTTP://BAMHOST2:9001/OracleBAM](http://BAMHOST2:9001/OracleBAM) to verify status of WLS\_BAM2.

If the managed servers fails to start with the following message:

```
Listener refused the connection with the following error:
ORA-12519, TNS:no appropriate service handler found
The Connection descriptor used by the client was <db_connect_string>
```

Verify that the PROCESSES initialization parameter for the database is set to a high enough value.

### 5.15.15 Configuring Oracle RAC Failover for the WLS\_BAM Servers

Oracle BAM allows customizing the behavior of a BAM server when Oracle RAC is used as the repository for the BAM schemas and a failure occurs in the database. The properties that allow this customization can be adjusted depending on the application and based on the desired expected behavior for each BAM system. The properties are

configured in the Fusion Middleware Control Console System MBean Browser, or in the corresponding Oracle BAM-specific XML configuration file.

If you want to completely disable Oracle BAM's failover to the database in an Oracle RAC configuration set `UseDBFailover` to `false` in the Fusion Middleware Control Console System MBean Browser for your BAM server. The default value of this property is `true`, therefore, failover is provided. You can also increase or decrease the number of retries in the access to the database when there is a database instance failure (for BAM to retry the in flight transactions). To adjust the number of retries change the `MaxDBNodeFailoverRetries` in the Fusion Middleware Control Console System MBean Browser. The default value for `MaxDBNodeFailoverRetries` is 5 times. See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite* for more details on the properties that can be configured for Oracle BAM.

### 5.15.16 Configuring the BAM Web Applications to Use the BAM Server in BAMHOST1

To configure the OracleBamWeb (WLS\_BAM1) and OracleBamWeb (WLS\_BAM2) applications to use the BAM server in BAMHOST1:

1. Access Oracle Enterprise Manager Fusion Middleware Control at the following URL  
  
`http://bamhost1vhn0:7001/em`
2. Expand **BAM** in the navigation tree.
3. Right-click **OracleBamWeb(WLS\_BAM1)**.
  - Select **BAM Web Properties** from the context menu.  
The BAM Web Properties page appears.
  - Define the following properties:
    - Enter **WEBHOST:7777** for the application URL.
    - Enter **BAMHOST1VHN1** for the server name.
  - Modify the listening port for the BAM server using the Mbean browser. To do this, follow these steps in the order shown:
    - Log onto the Oracle Enterprise Manager Fusion Middleware Control.
    - Expand the domain name in the left navigation tree.
    - Expand the **BAM** item in the left navigation tree.
    - On the BAM drop down menu at the top right, select **Mbean Browser**.
    - Navigate to the **oracle.bam.web -> Server -> Application -> Config -> BAMWebConfig** on the right.
    - In the **ServerPort** field, replace the "DEFAULT" value with 9001.
4. Select **OracleBamWeb(WLS\_BAM2)** from the navigation tree and repeat these steps.

### 5.15.17 Configuring Oracle HTTP Servers for the Administration Server and the WLS\_BAMn Managed Servers

To enable Oracle HTTP Server to route to the BAM\_Cluster, which contains the WLS\_BAMn managed servers, you must set the `WebLogicCluster` parameter to the list of nodes in the cluster:

1. On WEBHOST1 and WEBHOST2, add the following lines to the `ORACLE_BASE/admin/<instance_name>/config/OHS/<component_name>/mod_wl_ohs.conf` file:

```
WSM-PM
<Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicCluster bamhost1vhn1:9001,bamhost2:9001
</Location>

UMS
<Location /sdpmessaging/ >
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

UMS WS
<Location /ucs/messaging/webservice >
 SetHandler weblogic-handler
 WebLogicCluster SOAHOST1VHN1:8001,SOAHOST2VHN1:8001
</Location>

BAM Web app
<Location /OracleBAM>
 SetHandler weblogic-handler
 WebLogicCluster bamhost1vhn1:9001,bamhost2:9001
</Location>

BAM Web Services
<Location /OracleBAMWS>
 SetHandler weblogic-handler
 WebLogicCluster bamhost1vhn1:9001,bamhost2:9001
</Location>
```

Make sure the `httpd.conf` file located in the same directory as the `mod_wl_ohs.conf` file contains the following lines pointing to your load balancing router addresses:

```
NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName https://bam.mycompany.com:443
 ServerAdmin you@your.address
 RewriteEngine On
 RewriteOptions inherit
</VirtualHost>

NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName admin.mycompany.com:80
 ServerAdmin you@your.address
 RewriteEngine On
 RewriteOptions inherit
</VirtualHost>
```

---

**Note:** Values in this document such as `bam.mycompany.com:443`, `7777`, `admin.mycompany:80`, and `you@youraddress` are examples only. Enter values based on the actual environment.

---

## 2. Restart Oracle HTTP Server on WEBHOST1 and WEBHOST2.

```
WEBHOST1> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs1
```

```
WEBHOST2> ORACLE_BASE/admin/<instance_name>/bin/opmnctl restartproc
ias-component=ohs2
```

### 5.15.18 Validating Access through Oracle HTTP Server

Verify that the BAM servers' status is reported as "Running" in the Admin Console. If the server is shown as "Starting" or "Resuming," wait for the server status to change to "Started." If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors. Verify that you can access these URLs, where 'WEBHOSTN' specifies the name of each Oracle HTTP Server host (for example, WEBHOST1, WEBHOST2):

- <http://WEBHOST1:7777/wsm-pm>
- <http://WEBHOST2:7777/wsm-pm>
- <http://WEBHOST1:7777/sdpMessaging/userprefs-ui>
- <http://WEBHOST2:7777/sdpMessaging/userprefs-ui>
- <http://WEBHOST1:7777/OracleBAM>
- <http://WEBHOST2:7777/OracleBAM>

Verify these URLs also using the load balancing router address:

- <http://bam.mycompany.com:80/wsm-pm>
- <http://bam.mycompany.com:80/sdpMessaging/userprefs-ui>
- <http://bam.mycompany.com:80/OracleBAM>

Verify the following URLs to ensure that routing and failover from the HTTP Server to the BAM cluster is working correctly:

1. While WLS\_BAM2 is running, stop WLS\_BAM1 from Oracle WebLogic Server Administration Console.
2. Access the following URLs and verify the appropriate functionality:
  - <http://WEBHOST1:7777/wsm-pm>
  - <http://WEBHOST1:7777/sdpMessaging/userprefs-ui>
  - <http://WEBHOST1:7777/OracleBAM>
3. Start WLS\_BAM1 from Oracle WebLogic Server Administration Console.
4. Stop WLS\_BAM2.
5. Access the URLs from step 2 above again and verify the appropriate functionality.

### 5.15.19 Configuring Server Migration for the WLS\_BAM Servers

The high availability architecture for Oracle BAM uses server migration to protect singleton services against failures. The WLS\_BAM1 managed server is configured so that it can be restarted on BAMHOST2 in case of failure. For this configuration, WLS\_BAM1 listens on a specific floating IP address that is failed over by Oracle WebLogic Server Migration. To configure server migration for the WLS\_BAM1 managed servers, follow these steps:

---

**Note:** If server migration was configured previously for SOA, the BAM system can use the same data sources and database schemas. In that case, steps 1 through 5 may not be required, but you must also target the corresponding server-migration/leasing datasources to the BAM Cluster.

---

### 5.15.19.1 Setting Up the User and Tablespace for the Server Migration Leasing Table

To create the user and tablespace:

1. Create a tablespace called **leasing**. For example, log on to SQL\*Plus as the sysdba user and run the following command:

Example: Log on to SQL\*Plus as the sysdba user and run the following command:

```
SQL> create tablespace leasing
 logging datafile 'DB_HOME/oradata/orcl/leasing.dbf'
 size 32m autoextend on next 32m maxsize 2048m extent management local;
```

2. Create a user named **leasing** and assign to it the leasing tablespace.

```
SQL> create user leasing identified by welcome1;
```

```
SQL> grant create table to leasing;
```

```
SQL> grant create session to leasing;
```

```
SQL> alter user leasing default tablespace leasing;
```

```
SQL> alter user leasing quota unlimited on LEASING;
```

3. Create the leasing table using the `leasing.ddl` script.
  - a. Copy the `leasing.ddl` file located in the `WL_HOME/server/db/oracle` directory to your database node.
  - b. Connect to the database as the leasing user.
  - c. Run the `leasing.ddl` script in SQL\*Plus.

```
SQL> @copy_location/leasing.ddl;
```

### 5.15.19.2 Creating a GridLink or Multi Data Source from the WebLogic Server Administration Console

You create a data source to each of the Oracle RAC database instances during the process of setting up the multi data source, both for these data sources and the global leasing multi data source. When you create a data source:

- Ensure that this is a non-xa data source.
- Use Oracle's Driver (Thin) Version 9.0.1, 9.2.0, 10, 11.
- Do *not* use any type of distributed transaction emulation/participation algorithm for the data source (do not choose the **Supports Global Transactions** option, or the **Logging Last Resource, Emulate Two-Phase Commit**, or **One-Phase Commit** options of the **Supports Global Transactions** option), and specify a service name for your database. Data sources do not require support for global transactions.
- Target these data sources to the cluster(s).

- Make sure the datasources' connection pool initial capacity is set to **0**. To do this, select **Services**, **JDBC**, and then **Datasources**. In the Datasources screen, click the **Datasource Name**, then click the **Connection Pool** tab, and enter **0** in the **Initial capacity** field.
- Ensure that an ONS daemon is running on your database servers at all times. You can start the ONS daemon on a database server by running the `onsctl` command:  
`start`

### Creating a GridLink Data Source

To create a GridLink Data Source:

1. Log in to the Oracle WebLogic Server Administration Console.
2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the **Domain Structure** tree, expand **Services**, then select **Data Sources**.
4. On the **Summary of Data Sources** page, click **New** and select **GridLink Data Source**.
5. Enter the following information:
  - a. Enter a logical name for the datasource in the Name field. For example, **gridlink**
  - b. Enter a name for JNDI. For example, **jdbc/gridlink**.
  - c. Click **Next**.
6. In the **Transaction Options** page, de-select **Supports Global Transactions**, and click **Next**.
7. Select **Enter individual listener information** and click **Next**.
8. Enter the following connection properties:
  - d. **Service Name:** Enter the service name of the database. You must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example, `<mydbservice>.us.oracle.com`

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example, `us.oracle.com`

---

- **Host Name** - The DNS name or IP address of the server that hosts the database. For an Oracle GridLink service-instance connection, this must be the same for each data source in a given multi data source.
- **Port** - The port on which the database server listens for connections requests.
- **Database User Name:** The database user name. For example, `myDataBase`.
- **Password:** The password. For example, `myPassword1`.
- Confirm Password and click **Next**.

**Tip:** For more information, see the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

The console automatically generates the complete JDBC URL. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=left)(PORT=1234))(ADDRES
S=(PROTOCOL=TCP)(HOST=right)(PORT=1234))(ADDRESS=(PROTOCOL
=TCP)(HOST=center)(PORT=1234)))(CONNECT_DATA=(SERVICE_
NAME=myService)))
```

9. On the **Test GridLink Database Connection** page, review the connection parameters and click **Test All Listeners**.

Oracle WebLogic attempts to create a connection from the Administration Server to the database. Results from the connection test appear at the top of the page. If the test is unsuccessful, correct any configuration errors and retry the test.

Click **Next**.

10. In the **ONS Client Configuration** page, do the following:

- Select **Fan Enabled** to subscribe to and process Oracle FAN events.
- In ONS host and port, enter a comma-separated list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.
- Click **Next**.

11. On the **Test ONS client configuration** page, review the connection parameters and click **Test All ONS Nodes**.

Click **Next**.

12. In the **Select Targets** page, select **Dept1\_Cluster1** as the target and **All Servers in the cluster**.

13. Click **Finish**. Click **Activate Changes**.

### Creating a Multi Data Source

To create a multi data source, complete these steps:

1. From the Domain Structure window in the Oracle WebLogic Server Administration Console, expand the **Services** node, then click **Data Sources**. The Summary of JDBC Data Sources page appears.
2. In the Change Center, click **Lock & Edit**.
3. Click **New**, then click **Multi Data Sources**. The Create a New JDBC Multi Data Source page appears.
4. Enter **leasing** as the Name
5. Enter **jdbc/leasing** as the JNDI name.
6. Select **Failover** as algorithm (default).
7. Click **Next**.
8. Select **BAM\_Cluster** as the target.
9. Click **Next**.
10. Select **non-XA driver** (the default).
11. Click **Next**.
12. Click **Create New Data Source**.

13. Enter *leasing-rac0* as name. Enter *jdbc/leasing-rac0* as JNDI name. Enter *oracle* as the database type. For the driver type, enter *Oracle Driver (Thin) for RAC server-Instance connection Version 10,11*.

---

**Note:** When creating the multi data sources for the leasing table, enter names in the format of *<MultiDS>-rac0*, *<MultiDS>-rac1*, and so on.

---

14. Click **Next**.
15. Deselect **Supports Global Transactions**.
16. Click **Next**.
17. Enter the service name, database name (the RAC Node instance name, for example: *racdb1,racdb2*), host port, and password for your leasing schema.
18. Click **Next**.
19. Click **Test Configuration** and verify the connection works.
20. Target the data source to the BAM cluster.
21. Select the data source and add it to the right screen.
22. Click **Create a New Data Source** second instance of your Oracle RAC database, target it to *BAM\_Cluster*, repeating the steps for the second instance of your Oracle RAC database.
23. Add the second data source to your multi data source.
24. Click **Save**.
25. Click **Activate Changes**.

### 5.15.19.3 Edit the Node Manager's Properties File

The `nodemanager.properties` file is located in the `WL_HOME/server/bin` directory.

```
Interface=eth0
NetMask=255.255.255.0
UseMACBroadcast=true
```

- `Interface=eth0`

This property specifies the interface name for the floating IP (*eth0*, for example, on Linux).

---

**Note:** Do not specify the sub interface, such as *eth0:1* or *eth0:2*. This interface is to be used without the `:0` or `:1`. The Node Manager's scripts traverse the different `:X` enabled IPs to determine which to add or remove. For example, the valid values in Linux environments are *eth0*, *eth1*, or, *eth2*, *eth3*, *ethn*, depending on the number of interfaces configured.

---

---

**Note:** For Windows, the Interface should be set to the Network Interface Name. For example: `Interface="Local Area Connection"`.

---



- NetMask

This property specifies the net mask for the interface for the floating IP.

- UseMACBroadcast

This property specifies whether or not to use a node's MAC address when sending ARP packets, that is, whether or not to use the -b flag in the arping command.

Perform this configuration in the two nodes BAMHOST1 and BAMHOST2. Verify in the output of Node Manager (the shell where the Node Manager is started) that these properties are in use. Otherwise, problems may occur during migration. The output should be similar to the following:

```
...
StateCheckInterval=500
Interface=eth0 (Linux) or Interface="Local Area Connection" (Windows)
NetMask=255.255.255.0
...
```

#### 5.15.19.4 Set Environment and Superuser Privileges for the wlsifconfig.sh Script

To set the environment and superuser privileges for the wlsifconfig.sh script:

---

**Note:** On Windows, the script is named `wlsifconfig.cmd` and it can be run by users with the administrator privilege.

---

1. Ensure that your PATH environment variable includes the files listed in [Table 5–26](#).

**Table 5–26** Files Required to be in the PATH Environment Variable

| File                | Located in this directory                                    |
|---------------------|--------------------------------------------------------------|
| wlsifconfig.sh      | MW_HOME/user_projects/domains/bamdomain/bin/server_migration |
| wlscontrol.sh       | WL_HOME/common/bin                                           |
| nodemanager.domains | WL_HOME//common                                              |

2. Grant sudo configuration for the wlsifconfig.sh script.

- Configure sudo to work without a password prompt.
- For security reasons, sudo should be restricted to the subset of commands required to run the wlsifconfig.sh script. For example, to set the environment and superuser privileges for the wlsifconfig.sh script, complete these steps:
  - Grant sudo privilege to the WebLogic user ('oracle') with no password restriction, and grant execute privilege on the /sbin/ifconfig and /sbin/arping binaries.
  - Make sure the script is executable by the WebLogic user ('oracle'). The following example shows an entry inside /etc/sudoers granting sudo execution privilege for oracle and also over ifconfig and arping:

```
Defaults:oracle !requiretty
oracle ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
```

---

**Note:** Ask the system administrator for the sudo and system rights as appropriate to this step.

---

### 5.15.19.5 Configure Server Migration Targets

Configuring Cluster Migration sets the `DataSourceForAutomaticMigration` property to **true**. To configure cluster migration in a cluster:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand **Environment** and select **Clusters**.
3. Click the cluster for which you want to configure migration (BAM\_Cluster) in the Name column of the table.
4. Click the **Migration** tab. In the Change Center, click **Lock & Edit**.
5. In the Available field, select the machine to which to allow migration and click the right arrow. In this case, select **BAMHOST1** and **BAMHOST2**.
6. Select the data source to be used for automatic migration. In this case select the **leasing** data source.
7. Click **Save**.
8. Set the Candidate Machines for **Server Migration**.

You must perform this task for WLS\_BAM1:

- a. In Domain Structure window of the Oracle WebLogic Server Administration Console, expand **Environment** and select **Servers**.
- b. Select the server which you want to configure migration for. Click the **Migration** tab
- c. In the Available field, located in the Migration Configuration section, select the machines to which to allow migration and click the right arrow. Select **BAMHOST2** for **WLS\_BAM1**.
- d. Select **Automatic Server Migration Enabled**.  

This enables the Node Manager to start a failed server on the target node automatically.
- e. Click **Save** and Activate the changes.
- f. Restart the Administration Server and the WLS\_BAM1 server.

### 5.15.19.6 Test Server Migration

To verify that Server Migration is working properly, follow these steps:

From BAMHOST1:

1. Force stop the WLS\_BAM1 managed server using the following command:

```
BAMHOST1> kill -9 <pid>
```

`pid` specifies the process ID of the managed server. You can identify the `pid` in the node by running the following command:

```
BAMHOST1> ps -ef | grep WLS_BAM1
```

---

**Note:** For Windows, terminate the Managed Server using the taskkill command. For example:

```
taskkill /f /pid <pid>
```

Where <pid> is the process Id of the Managed Server.

To determine the process Id of the Managed Server, run the following command and identify the pid of the WLS\_BAM1 Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

2. In the Node Manager console, a message indicates that WLS\_BAM1's floating IP has been disabled.
3. Wait for the Node Manager to try a second restart of WLS\_BAM1.  
Node Manager waits for a fence period of thirty seconds before trying this restart.
4. After Node Manager restarts the server, stop it again.  
Node Manager should now log a message indicating that the server will not be restarted again locally.

From BAMHOST2:

1. Watch the local Node Manager console.  
After thirty seconds since the last try to restart WLS\_BAM1 on Node 1, Node Manager on BAMHOST2 should prompt that the floating IP for WLS\_BAM1 is being brought up and that the server is being restarted in this node.
2. Access the Oracle BAM console using BAMHOST1VHN1/OracleBAM and WEBHOST1:7777/OracleBAM.

#### Verification From the Administration Console

To verify Migration in the Administration Console:

1. Log into the Administration Console. Click on **Domain** on the left console.
2. Click the **Monitoring** tab and then on the **Migration** tab.

The Migration Status table provides information on the migration status.

---

**Note:** To fail a server back to its original node/machine after it migrates, use the Oracle WebLogic Administration Console to stop and restart the managed server. The appropriate Node Manager starts the managed server on the machine that it was originally assigned to.

---

### 5.15.20 Configuring Clients Connecting to the BAM System

When accessing a BAM server using the BAM Adapter (RMI), the virtual hostname of the BAM server (BAMHOST1VNH1) should be used for the connection. SOAP requests come through HTTP, therefore it is required to use the load balancing router addresses when using the BAM Adapter in this situation.



---

# Configuring High Availability for Oracle ADF and WebCenter Portal Applications

This chapter describes high availability concepts and configuration procedures for Oracle Application Development Framework (Oracle ADF) and Oracle WebCenter Portal.

This chapter includes the following sections:

- [Section 6.1, "Oracle ADF and High Availability Concepts"](#)
- [Section 6.2, "Configuring an Oracle ADF High Availability Deployment"](#)
- [Section 6.3, "WebCenter Portal and High Availability Concepts"](#)
- [Section 6.4, "Configuring High Availability for WebCenter Portal"](#)
- [Section 6.5, "Configuring High Availability for WebCenter Portal Applications"](#)

## 6.1 Oracle ADF and High Availability Concepts


Oracle ADF is an end-to-end application framework that builds on Java Platform, Enterprise Edition (Java EE) standards and open-source technologies to simplify and accelerate implementing service-oriented applications. Oracle ADF is suitable for enterprise developers who want to create applications that search, display, create, modify, and validate data using web, wireless, desktop, or web services interfaces. Used in tandem, Oracle JDeveloper 11g and Oracle ADF provide an environment that covers the full development lifecycle from design to deployment, with drag-and-drop data binding, visual UI design, and team development features built in.

### 6.1.1 Understanding Oracle ADF

In line with community best practices, applications built using the Fusion web technology stack achieve a clean separation of business logic, page navigation, and user interface by adhering to a model-view-controller (MVC) architecture supported by the following layers.

As shown in [Figure 6–1](#), in an MVC architecture:

- The model layer represents the data values related to the current page
- The view layer contains the UI pages used to view or modify that data
- The controller layer processes user input and determines page navigation
- The business service layer handles data access and encapsulates business logic

**Figure 6–1 Overview of Oracle ADF Architecture**


#### 6.1.1.1 Oracle ADF Components

The core module in the framework is Oracle ADF Model. The Oracle ADF Model layer enables a unified approach to bind any user interface to any business service, without the need to write code.

The other modules that make up a Fusion web application technology stack are:

- Oracle ADF Business Components, which simplifies building business services.
- Oracle ADF Faces, which offers a rich library of AJAX-enabled UI components for Web applications built with JavaServer Faces (JSF).
- Oracle ADF Controller, which integrates JSF with Oracle ADF Model. The ADF Controller extends the standard JSF controller by providing additional functionality, such as reusable task flows that pass control not only between JSF pages, but also between other activities, for instance method calls or other task flows.

Figure 6–2 illustrates where each Oracle ADF module fits in the Fusion web application architecture.

**Figure 6–2 Simple Oracle ADF Architecture**

**6.1.1.1.1 ADF Business Components** ADF Business Components are prebuilt application objects that accelerate the job of delivering and maintaining high-performance, richly functional, database-centric services. When building service-oriented Java EE applications, developers implement the core business logic as one or more business services. These backend services provide clients with a way to query, insert, update, and delete business data as required while enforcing appropriate business rules. ADF Business Components provides a ready-to-use implementation of Java EE design patterns and best practices.

Oracle ADF Business Components provides the following key components to simplify building database-centric business services:

- **Entity object**

An entity object represents a row in a database table and simplifies modifying its data by handling all data manipulation language (DML) operations. It can encapsulate business logic to ensure that business rules are consistently enforced. Developers can associate an entity object with others to reflect relationships in the underlying database schema to create a layer of business domain objects to reuse in multiple applications.

- **View object**

A view object represents a SQL query and simplifies working with its results. Developers use the SQL language to join, project, filter, sort, and aggregate data into the shape required by the end-user task being represented in the user interface. This includes the ability to link a view object with other entity objects to create master-detail hierarchies of any complexity. When end users modify data in the user interface, view objects collaborate with entity objects to consistently validate and save the changes.

- **Application module**

An application module is the transactional component that UI clients use to work with application data. It defines an updateable data model and top-level procedures and functions (called service methods) related to a logical unit of work related to an end-user task.

For more information about Oracle ADF Business Components, see the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

**6.1.1.1.2 ADF Model Layer** In the model layer, Oracle ADF Model implements data controls that abstract the implementation technology of a business service. Standard metadata interfaces describe the service's operations and data collections, including information about the properties, methods, and types involved. In Oracle JDeveloper, developers see that information as icons that they can drag and drop onto a page. When a developer drags the representation of a service onto a page, Oracle JDeveloper automatically creates the bindings from the page to the services. At runtime, the ADF Model layer reads the information describing the application's data controls and data bindings from appropriate XML files and implements the two-way connection between the user interface and the application's business service.

Oracle ADF provides out-of-the-box data control implementations for the most common business service technologies. Using Oracle JDeveloper and Oracle ADF together provides a declarative, drag-and-drop data binding experience for building user interfaces. Along with support for ADF Business Components application modules, ADF Model also supports the following service technologies:

- Enterprise JavaBeans (EJB) session beans and JPA entities
- JavaBeans
- Web services
- XML
- CSV files

For more information about Oracle ADF Model, see the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

**6.1.1.1.3 ADF Controller** In the controller layer, where handling page flow of the web applications is a key concern, ADF Controller provides an enhanced navigation and state management model on top of JSF. JDeveloper supports declarative creation of task flows that can manage application control between different types of activities, such as pages, methods on managed beans, declarative case statements, or calls to other task flows.

For more information about Oracle ADF Controller, see the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

**6.1.1.1.4 ADF Faces Rich Client** ADF Faces rich client (ADF Faces) is a set of standard JSF components that include built-in AJAX functionality. AJAX is a combination of asynchronous JavaScript, dynamic HTML (DHTML), XML, and `XmlHttpRequest` communication channel. This combination allows requests to be made to the server without fully rerendering the page. While AJAX allows rich client-like applications to use standard internet technologies, JSF provides server-side control, which reduces the dependency on an abundance of JavaScript often found in typical AJAX applications.

ADF Faces provides over 100 rich components, including hierarchical data tables, tree menus, in-page dialogs, accordions, dividers, and sortable tables. ADF Faces also provides ADF Data Visualization components, which are Flash- and SVG-enabled components capable of rendering dynamic charts, graphs, gauges, and other graphics that provide a real-time view of underlying data. Each component also supports customization and skinning, along with internationalization and accessibility.

To achieve these front-end capabilities, ADF Faces components use a rendering kit that handles displaying the component and also provides the JavaScript objects needed for



the rich functionality. This built-in support enables developers to build rich applications without needing extensive knowledge of the individual technologies on the front or back end.


For more information about ADF Faces, including the architecture and detailed information about each of the components, see the *Oracle Fusion Middleware Web User Interface Developer's Guide for Oracle Application Development Framework*.

### 6.1.1.2 Oracle ADF Single Node Architecture

You can install the Oracle ADF runtime to the Oracle WebLogic Server using either the Oracle JDeveloper Installer or the Oracle Fusion Middleware Application Developer Installer. The Application Developer Installer also lets you optionally install Fusion Middleware Control to provide web-based administration support for all Managed Servers in the domain. The Oracle JDeveloper installer does not install Fusion Middleware Control. See the deployment chapter of the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

When you use the Application Developer Installer to install the Oracle ADF runtime, it creates an Oracle Application Developer home directory (by default, `Oracle_APPDEV1`) under the Middleware home. After you use the domain configuration wizard to create an Application Developer domain (`base_domain`) based on the JRF domain template, you can configure the server topology. In a typical set up, the domain has an Administration server containing the WLS Administration Console and Fusion Middleware Control. Typically, the Oracle ADF runtime libraries (part of the Java Required Files) get deployed to the Managed Servers, in addition to the user-facing custom Fusion web applications. To provide customization and personalization features, you can install an optional MDS repository and configure it separately. [Figure 6–3](#) shows a basic single-node Oracle ADF architecture.

**Figure 6–3 Basic Single-Node Oracle ADF Architecture**



For more information about domains and servers, see the *Oracle Fusion Middleware Administrator's Guide*.

### 6.1.1.3 Oracle ADF External Dependencies

If the Fusion web application involves customization using Oracle Metadata Services (MDS), you should register your MDS repository with the Oracle WebLogic Server domain before you deploy the application. For information about registering MDS, see the *Oracle Fusion Middleware Administrator's Guide*.

Then, when you deploy the application, JDeveloper prompts you to choose the target metadata repository or shared metadata repository. You will be able to choose from the

list of metadata repositories registered with the Oracle WebLogic Administration Server.

To ensure that you receive the metadata repository prompt, the application's `adf-config.xml` file must define a `cust-config` element in the `mds-config` section. This element specifies an ordered and named list of customization classes. A customization class is the interface that MDS uses to define which customization applies to the base definition metadata. In JDeveloper, you can use the overview editor for the `adf-config.xml` file to define a `cust-config` element.

For information about configuring the `adf-config.xml` file for MDS, see the chapter on customizing with MDS in the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

For more information about the MDS architecture and metadata repositories (database and file-based) and archives (EAR, MAR), refer to the *Oracle Fusion Middleware Administrator's Guide*.

#### 6.1.1.4 Oracle ADF Log File

The operations performed by the Fusion web application are logged directly to the WebLogic Managed Server where the application is running:

```
DOMAIN_HOME/servers/server_name/logs/server_name-diagnostic.log
```

The log files for the different WebLogic Managed Servers are also available from the Oracle WebLogic Server Administration Console. To verify the logs, access the Oracle WebLogic Server Administration Console `http://<admin_server_host>:<port>/console` and click on **Diagnostics-Log Files**.

This log's granularity and logging properties can be changed using Oracle Enterprise Manager Fusion Middleware Control (Fusion Middleware Control). Fusion Middleware Control is a Web browser-based, graphical user interface that you can use to monitor and administer a farm. To receive high availability warning diagnostic messages for Oracle ADF, the level should be set to `FINE`, as described in [Section 6.1.4.3, "Troubleshooting Oracle ADF Replication and Failover Issues."](#)

For more information about the level of diagnostics you can specify for Fusion web applications, see the chapter on testing and debugging in the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

For details about using Fusion Middleware Control to change the log settings of WebLogic Managed Servers and Oracle ADF, see the *Oracle Fusion Middleware Administrator's Guide*.

## 6.1.2 Oracle ADF High Availability Considerations

Fusion web applications built on the Oracle ADF technology stack are Java EE applications (and J2EE applications). You should observe the same best practices for Fusion web applications as you would any other Java EE application in a high availability environment.

### 6.1.2.1 Oracle ADF Scope and Session State

At runtime, ADF objects such as the binding container and managed beans are instantiated. Each of these objects has a defined life span set by its scope attribute.

There are six types of scopes in a Fusion web application:

- Application scope: The object is available for the duration of the application.
- Session scope: The object is available for the duration of the session.

- Page flow scope: The object is available for the duration of a bounded task flow.
- Request scope: The object is available from the time an HTTP request is made until a response is sent back to the client.
- Backing bean scope: Used for managed beans for page fragments and declarative components only, the object is available from the time an HTTP request is made until a response is sent back to the client.
- View scope: The object is available until the view ID for the current view activity changes. This scope can be used to hold values for a given page. However, unlike request scope, which can be used to store a value needed from one page to the next, anything stored in view scope will be lost once the view ID changes.

When the Fusion web application runs in a clustered environment, a portion of the application's state is serialized and copied to another server or a data store at the end of each request so that the state is available to other servers in the cluster.

When you design an application to run in a clustered environment, you must:

- Ensure that all managed beans with a life span longer than one request are serializable (that is, they implement the `java.io.Serializable` interface). Specifically, beans stored in session scope, page flow scope, and view scope must be serializable.
- Verify that Oracle ADF is aware of changes to managed beans stored in ADF scopes (view scope and page flow scope) and enable tracking changes to ADF memory scopes.

When you modify a value within a managed bean in either view scope or page flow scope, the application must notify Oracle ADF so that it can ensure the bean's new value is replicated.

In [Example 6-1](#), an attribute of an object in view scope is modified.

#### **Example 6-1 Code that Modifies an Object in viewScope**

```
Map<String, Object> viewScope =
 AdfFacesContext.getCurrentInstance().getViewScope();
MyObject obj = (MyObject)viewScope.get("myObjectName");
Obj.setFoo("newValue");
```

Without additional code, Oracle ADF will be unaware of this change and will not know that a new value must replicate within the cluster. To inform Oracle ADF of the change and need for replication, use the `markScopeDirty()` method shown in [Example 6-2](#). The `markScopeDirty()` method accepts only `viewScope` and `pageFlowScope` as parameters.

#### **Example 6-2 Additional Code to Notify Oracle ADF of Changes to an Object**

```
controllerContext ctx = ControllerContext.getInstance();
ctx.markScopeDirty(viewScope);
```

This code is needed for any request that modifies an existing object in one of the ADF scopes. If the scope itself is modified by the scope's `put()`, `remove()`, or `clear()` methods, it is not necessary to notify Oracle ADF.

To enable ADF Controller to track changes to ADF memory scopes and replicate the page flow scope and view scope within the server cluster, you can enable the `<adf-scope-ha-support>` parameter in the `adf-config.xml` file, as described in [Section 6.1.3.3, "Configuring adf-config.xml."](#)

For more information about ADF object scopes, see the chapter on Fusion page lifecycle in the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

### 6.1.2.2 Oracle ADF Failover and Expected Behavior

An Oracle WebLogic cluster provides application high availability. If one member of the cluster is unavailable, any other available member of the cluster is able to handle the request.

#### Session Failover Requirements

For seamless failover of a Fusion web application, the application must meet the following conditions:

- The application is in a cluster and at least one member of the application cluster is available to serve the request.
- For stateful applications, state replication is configured correctly as described in [Section 6.1.3, "Configuring Oracle ADF for High Availability."](#)
- If you are using Oracle HTTP Server, the server is configured with the WebLogicCluster directive to balance among all available application instances.
- If you are using a hardware load balancer, the load balancer is:
  - Routing traffic to all available instances
  - Configured correctly with a health monitor to mark unavailable instances
  - Configured to support persistence of session state

#### Expected Behavior for Application Failover

If the environment has been configured correctly, application users do not notice when an application instance in a cluster becomes unavailable. The sequence of events in an application failover is, for example, as follows:

1. A user makes a request and is routed by a hardware load balancer to instance A of the application.
2. Instance A of the application becomes unavailable because of node failure, process failure, or network failure.
3. The hardware load balancer marks instance A as unavailable.
4. The user makes a subsequent request. The request is routed to instance B.
5. Instance B is configured as a replication partner of Instance A and has the user's session state.
6. The application resumes using the session state on Instance B and the user continues working without interruption.

### 6.1.2.3 Oracle ADF Active Data Services

The Fusion technology stack includes the Active Data Service (ADS), which allows you to bind ADF Faces components to an active data source using the ADF Model layer. In JDeveloper, you configure individual components in your JSF pages to display active data. When you configure components to use active data, data is pushed to the client whenever a change event is raised by the data source. The data is pushed from the server to the client and displayed by the browser.

To support failover for pages configured to display active data, it is necessary to enable failover for the ADF Business Components application module and to enable ADF Controller to track changes to the ADF memory scopes. With these ADF settings configured, when failover occurs, the ADF server will detect the failover and request all pages configured to display active data to refresh themselves, after that ADS restarts and data is pushed to the client.

For details about how to enable failover for the Fusion web application, see [Section 6.1.3, "Configuring Oracle ADF for High Availability."](#)

For more information about using Oracle ADF Faces components with an active data service, see the chapter on ADS in the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

#### 6.1.2.4 Configuring the ADF Application Module for Oracle RAC

When configuring the ADF application module to access a highly available database system, such as redundant databases or Oracle Real Application Clusters (Oracle RAC) as the backend, the data source must be container-defined. In this scenario, it is required to use a GridLink data source or multi data source. However, from the standpoint of the application module configuration, the naming convention for the multi data source or GridLink data source is the same as it is for a non-multi data source or GridLink data source. This naming convention ensures that the correct data source is used at runtime. For details about configuring GridLink data sources for high availability applications, see [Section 4.1.4, "Configuring GridLink Data Sources with Oracle RAC."](#) For multi data sources, see [Section 4.1.5, "Configuring Multi Data Sources with Oracle RAC."](#)

### 6.1.3 Configuring Oracle ADF for High Availability

To support automatic replication and failover for web applications within a clustered environment, Oracle WebLogic Server supports mechanisms for replicating HTTP session state across clusters. You can configure Oracle ADF to ensure the Fusion web application's state can be restored from any server in the cluster.

#### 6.1.3.1 Configuring Application Modules

An application module is the transactional component that UI clients use to work with application data. It defines an updateable data model and top-level procedures and functions (called service methods) related to a logical unit of work related to an end-user task. An application module supports passivating, or storing, its transaction state as a snapshot in the database. It also supports the reverse operation of activating the transaction state from one of these saved snapshots.

For more information on the management of application module state, see "Introduction to Fusion Web Application State Management" in *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

To enable support for ADF Business Components failover, set the `jbo.dofailover` parameter to `true` so that the application module state is saved on release. This allows Oracle ADF to restore the application module state from a snapshot saved from a previous checkin. By contrast, when the failover feature is disabled, which it is by default, then application module state is saved only when the application is reused by a subsequent user session and only when the application module pool cannot find an unused application module.

You can set this parameter in your application module configuration on the **Pooling and Scalability** tab of the Edit Business Components Configuration dialog.

To configure application modules for high availability:

1. Launch JDeveloper and open the application.
2. In the Application Navigator, expand the project that contains the data model and locate the application module.
3. Right-click the application module and select **Configurations**.
4. Click **Edit**.
5. Click the **Pooling and Scalability** tab.
6. Select the **Failover Transaction State Upon Managed Release** checkbox.
7. Click **OK** to close the **Edit Business Components Configuration** dialog.
8. Click **OK** to close the **Manage Configurations** dialog.

For more information about Oracle ADF state management facilities and how to use them, see the chapter on application state management in the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

### 6.1.3.2 Configuring weblogic.xml

To enable support for replicating HTTP session state, you must assign a value to the `persistent-store-type` element in the Oracle WebLogic Server `weblogic.xml` file. The value `replicated_if_clustered` ensures that the in-effect persistent store type will be replicated so that sessions on the clustered environment are stored in accordance with the value set for the cluster of servers to which this server belongs.

To configure the `weblogic.xml` file for high availability:

1. Launch JDeveloper and open the application.
2. In the Application Navigator, expand the project that contains the web application and expand the **WEB-INF** folder.
3. Double-click the **weblogic.xml** file, and click the **Source** tab to edit the file.
4. In the file, add the `persistent-store-type` definition to the `session-descriptor` element:

```
<weblogic-web-app>
 <session-descriptor>
 <persistent-store-type>
 replicated_if_clustered
 </persistent-store-type>
 </session-descriptor>
</weblogic-web-app>
```

### 6.1.3.3 Configuring adf-config.xml

When you design an application to run in a clustered environment, you must ensure that Oracle ADF is aware of changes to managed beans stored in ADF scopes (view scope and page flow scope).

When a value within a managed bean in either view scope or page flow scope is modified, the application must notify Oracle ADF so that it can ensure the bean's new value is replicated.

To enable ADF Controller to track changes to ADF memory scopes and replicate the page flow scope and view scope within the server cluster, you must set the ADF Controller parameter `<adf-scope-ha-support>` in the application's `adf-config.xml` file to `true`. For example, when set to `true` for an application and

that application adds or removes a bean from a page flow scope during a request, the change will automatically replicated within a cluster.

The `adf-config.xml` file is the central configuration file for all ADF components. The file contains sections to configure the runtime behavior for ADF components, including, ADF Controller.

---

**Note:** If your application uses MDS and will use an Oracle database that supports failover, Oracle recommends enabling MDS retry on failover. To do this, add the following `retry-connection` entry to the MDS configuration section of `adf-config.xml`.

---

```
<persistence-config>
 <metadata-namespaces>...
 <metadata-store-usages>...
 <external-change-detection enabled="false"
polling-interval-secs="30"/>
 <read-only-mode enabled="true"/>
 <retry-connection enabled="true"/>
</persistence-config>
```

---

To configure the `adf-config.xml` file for high availability:

1. Launch JDeveloper and open the application.
2. In the Application Navigator, expand the **Application Resources**.
3. Select **Descriptors**, and then select the **ADF META-INF** node.
4. Double-click the **adf-config.xml** file, and click the **Source** tab to edit the file.
5. Add the following to the file:

```
<adf-controller-config xmlns="http://xmlns.oracle.com/adf/controller/config">
 <adf-scope-ha-support>true</adf-scope-ha-support>
</adf-controller-config>
```

For more information about using the `adf-config.xml` file to configure ADF, see the chapter on creating complex task flows in the *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

#### 6.1.3.4 Configuring `org.apache.myfaces.trinidad.CHECK_FILE_MODIFICATION`

The `org.apache.myfaces.trinidad.CHECK_FILE_MODIFICATION` parameter must not be set to `true` when running in a high availability environment. Setting this context parameter to `true` can lead to errors after failover occurs.

## 6.1.4 Troubleshooting Oracle ADF High Availability

This section describes procedures for troubleshooting possible issues with Oracle ADF.

### 6.1.4.1 Troubleshooting Oracle ADF Development Issues

When you develop the Fusion web application in Oracle JDeveloper, the integrated development environment provides support for detecting potential High Availability issues. The warnings that JDeveloper provides are generated by the audit framework and will be triggered to display in the JDeveloper source editors. The warnings the editors display are based on the audit rules for High Availability applications.

The High Availability audit rules that JDeveloper enables by default are:

- **ADF Controller Configuration - High Availability for ADF Scopes is not Enabled** simply warns the developer that the `adf-scope-ha-support` flag in the `adf-config.xml` file is set is not set to `true`. This audit rule fires only when the `<adf-controller-config>` element is present the ADF application-level configuration file (`adf-config.xml`).
- **ADF Page Flows - Bean in Scope Map is Modified** warns the developer when the some code calls a setter method on a bean to indicate that the code did not subsequently call the `ControllerContext.markScopeDirty()` method. This audit rule fire only when the `adf-scope-ha-support` flag in the `adf-config.xml` file is set to `true`.
- **ADF Page Flows - EL Bean is Modified** warns the developer when some code evaluates an EL expression that mutates a bean to indicate that the code did not subsequently call the `ControllerContext.markScopeDirty()` method. This audit rule fire only when the `adf-scope-ha-support` flag in the `adf-config.xml` file is set to `true`.
- **ADF Page Flows - Managed Bean Class Not Serializable** warns the developer that a managed bean has a non-serializable class defined in `viewScope`, `pageFlowScope`, or `sessionScope`. This audit rule fire only when the `adf-scope-ha-support` flag in the `adf-config.xml` file is set to `true`.

You can modify the High Availability audit rule settings using the Preference dialog in JDeveloper. From the JDeveloper toolbar, choose **Tools - Preferences**, under **Audit - Profiles** expand **ADF Controller Configuration** or **ADF Pages Flows** and make the desired audit rule selections.

You can also trigger the audit by choosing **Build - Audit *project.jpr*** from the JDeveloper toolbar.

#### 6.1.4.2 Troubleshooting Oracle ADF Deployment Issues

Fusion web applications are deployed when the managed server is first started. Use the Oracle WebLogic Server Administration Console first to check that all application deployments were successful:

Click **Deployments** in the left hand pane. The right hand pane shows the application deployments and their status. The state of all applications, assuming all the servers are running, should be **ACTIVE**.

If an application deployment has failed, the server logs may provide some indication of why the application was not deployed successfully. The server logs are located in the `DOMAIN_HOME/servers/server_name/logs` directory. Common issues include:

- Unavailability of external resources, such as database resources. Examine the error, fix it, and attempt to redeploy the application.
- The appropriate applications or libraries are not targeted correctly to the right managed server or Cluster.

#### 6.1.4.3 Troubleshooting Oracle ADF Replication and Failover Issues

State Replication is most prominent in failover scenarios. A user working on one server may discover that, upon failover:

- Windows may close or the state might reset.
- Screens may require a reset.
- The application may redirect to the logon screen.



To diagnose and troubleshoot state replication issues.

1. Confirm that this is not a known replication issue.

See [Section 6.1.2.2, "Oracle ADF Failover and Expected Behavior"](#) for possible expected behaviors. Before proceeding to further diagnose the issue, first confirm that the failover behavior is not an expected behavior.

2. Check load balancer settings.

For replication and failover to function correctly, the load balancer must be configured with the appropriate persistence settings. For more details on configuring Hardware Load Balancers for Oracle WebLogic Server, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

3. Check the cluster status.

Replication occurs within the context of a cluster. For failover to be successful, there must be at least one other healthy member of the cluster available. You can check cluster status in one of two ways:

- Check the cluster status using the Oracle WebLogic Server Administration Console - In the Left-hand pane, click on **Servers**. Verify the state of all servers in the cluster.
- Check the cluster status using weblogic.Admin utility The weblogic.Admin command can be used to query the state of all servers in a specific cluster. For example:

```
$ java weblogic.Admin -url Adminhost:7001 -username <username> -password
<password> CLUSTERSTATE -clustername Spaces_Cluster
```

This example returns:

```
There are 2 server(s) in cluster: Spaces_Cluster
The alive servers and their respective states are listed below:
Application Server---RUNNING
Managed Server---RUNNING
```

4. Check cluster communications.

Although Cluster members may all be running, there may be communication issues which prevent them from communicating replication information to each other. There are two types of cluster communication configurations. Troubleshooting depends on the cluster type:

- Checking Unicast cluster communications - For Unicast clusters, managed servers must be able to access each other's hosts and each other's default listening port.  
  
Ensure that all individual managed servers have their Listen Address set correctly. You can find this setting by selecting **Configuration, General** for each managed server.
- Checking Multicast cluster communications - For multicast clusters, servers must be able to intercept the same multicast traffic. Ensure that multicast is configured correctly by running the WebLogic utility `utils.MulticastTest` on each machine. For example:

```
$ java utils.MulticastTest -H
```

5. Confirm Oracle WebLogic Server application configuration.

Oracle WebLogic Server is not configured by default for failover. In-memory replication takes place only with the proper setting in the `weblogic.xml` file:

```
<session-descriptor>
<persistent-store-type>replicated_if_clustered</persistent-store-type>
</session-descriptor>
```

A `persistent-store-type` of `replicated` is also acceptable. This setting can be made in JDeveloper, as described in [Section 6.1.3.2, "Configuring weblogic.xml."](#)

**6. Confirm Oracle ADF Business Components configuration.**

Oracle ADF is not configured by default for failover. Failover is supported only with the proper setting in the ADF Business Components configuration file (`bc4j.xcfg`):

```
<AppModuleConfig ...
 <AM-Pooling jbo.dofailover="true"/>
</AppModuleConfig>
```

This setting is made in JDeveloper through the Edit Business Components Configuration dialog, as described in [Section 6.1.3.1, "Configuring Application Modules."](#)

**7. Confirm Oracle WebLogic Server connection pool parameter.**

Set an appropriate value for the `weblogic-application.xml` deployment descriptor parameter `inactive-connection-timeout-seconds` on the element `<connection-check-params> pool-params`.

When enabling application module state passivation, a failure can occur when Oracle WebLogic Server is configured to forcibly release connections back into the pool. The failure creates an exception "Connection has already been closed" that gets saved to the server log. The user interface does not show this exception.

Set `inactive-connection-timeout-seconds` to several minutes. In most cases, this setting avoids forcing the inactive connection timeout and passivation failure. Adjust the setting as needed for your environment.

**8. Confirm Oracle ADF Controller configuration.**

Oracle ADF is not configured by default to replicate changes to ADF objects in ADF memory scopes. ADF object replication is supported only with the proper setting in the ADF application-level configuration file (`adf-config.xml`):

```
<adfc:adf-controller-config>
 <adfc:adf-scope-ha-support>true</adfc:adf-scope-ha-support>
</adfc:adf-controller-config>
```

This setting is made in JDeveloper through the source editor. See [Section 6.1.3.1, "Configuring Application Modules."](#)

**9. Check default logger messages.**

By default the ADF log shows high-level messages (INFO level). The default logging often reports problems with serialization and replication without the need to enable more detailed log messages. For more information about the log, see [Section 6.1.1.4, "Oracle ADF Log File."](#)

**10. Enable log messages for ADF high availability applications.**

Configure the ADF logger to output runtime messages for high availability. By default the ADF log display high-level messages (INFO level). You enable high

availability diagnostics for ADF Controller by setting the logging level in Fusion Middleware Control to FINE.

When enabled, the logger outputs a warning if the `adfc:adf-scope-ha-support` setting in the `adf-config.xml` file is not set. For more information about the ADF logger, see [Section 6.1.1.4, "Oracle ADF Log File."](#)

## 11. Enable debug.

Check the server logs for any unusual messages on managed server startup. In particular, if the managed server is unable to locate other members of the cluster. The server logs are located in the `DOMAIN_HOME/servers/SERVER_NAME/logs` directory.

For further debugging, enable the flags `DebugCluster`, `DebugClusterAnnouncements`, `DebugFailOver`, `DebugReplication`, and `DebugReplicationDetails`. Each flag can be enabled with the `weblogic.Admin` utility:

```
$ java weblogic.Admin -url Adminhost:7001 -username <username> -password
<password> SET -type ServerDebug -property DebugCluster true
```

## 12. Enable component state serialization checking.

Enable server checking to ensure no unserializable state content on session attributes is detected. This check is disabled by default to reduce runtime overhead. Serialization checking is supported by the Java server system property `org.apache.myfaces.trinidad.CHECK_STATE_SERIALIZATION`.

[Table 6–1](#) shows the options that can be used with the property. Use commas to delimit the options.

**Table 6–1** *CHECK\_STATE\_SERIALIZATION Options*

| Option      | Description                                                                                                                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tree        | Checks whether the entire component tree is serializable. This is the fastest component check. Most testing should be performed with this flag enabled.                                                                                             |
| component   | Checks each component individually for serializability. This option is much slower than "tree." It is typically turned on only after testing with "tree" has reported an error. This option is used to narrow down the problematic component.       |
| property    | Checks each component attribute individually for serializability. This is slower than "component" and is used to narrow the down the specific problematic component attribute after a failure has been detected in the "tree" or "component" modes. |
| session     | Checks that all attributes in the JSF Session Map that are marked as Serializable are in fact serializable.                                                                                                                                         |
| application | Checks that all attributes in the JSF Application Map that are marked as Serializable are in fact serializable.                                                                                                                                     |
| beans       | Checks that any serializable object in the appropriate map has been marked as dirty if the serializable content of the object has changed during the request.                                                                                       |
| all         | Checks everything.                                                                                                                                                                                                                                  |

For high availability testing, start off by validating that the Session and JSF state is serializable by launching the application server with the system property:

```
-Dorg.apache.myfaces.trinidad.CHECK_STATE_SERIALIZATION=session,tree
```

Add the `beans` option to check that any serializable object in the appropriate map has been marked as dirty if the serialized content of the object has changed during the request:

```
-Dorg.apache.myfaces.trinidad.CHECK_STATE_SERIALIZATION=session,tree,beans
```

If a JSF state serialization failure is detected, relaunch the application server with the system property to enable component and property flags and rerun the test:

```
-Dorg.apache.myfaces.trinidad.CHECK_STATE_SERIALIZATION=all
```

Since these are Java system properties, they need to be specified when the application server is started.

## 6.2 Configuring an Oracle ADF High Availability Deployment

This section describes how to configure an example Oracle ADF high availability deployment.

---

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

### 6.2.1 Terminology for Directories and Directory Environment Variables

This list describes the directories and variables used in this section:

- **ORACLE\_BASE:** This environment variable and related directory path refers to the base directory under which Oracle products are installed.
- **MW\_HOME:** This environment variable and related directory path refers to the location where Oracle Fusion Middleware resides.
- **WL\_HOME:** This environment variable and related directory path contains installed files necessary to host a WebLogic Server.
- **ORACLE\_HOME:** This environment variable and related directory path refers to the location where Oracle SOA Suite is installed.
- **ORACLE\_COMMON\_HOME:** This environment variable and related directory path refers to the Oracle home that contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF).
- **DOMAIN directory:** This directory path refers to the location where the Oracle WebLogic domain information (configuration artifacts) is stored.
- **ORACLE\_INSTANCE:** An Oracle instance contains one or more system components, such as Oracle Web Cache, Oracle HTTP Server, or Oracle Internet Directory. An Oracle instance directory contains updateable files, such as configuration files, log files, and temporary files.

The values used and recommended for consistency for this directories are:

- **ORACLE\_BASE:** /u01/app/oracle

- MW\_HOME (application tier): `ORACLE_BASE/product/fmw`
- WL\_HOME: `MW_HOME/wlserver_10.3`
- ORACLE\_HOME: `MW_HOME/adf`

## 6.2.2 Using RCU to Load Fusion Middleware Schemas in the Database

This step is required only if your ADF application needs to use any of schemas that are part of Oracle Fusion Middleware. Typically, this is done if the ADF application uses MDS repository, in which case you must install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store into a Real Application Clusters (Oracle RAC) database before you install Oracle Fusion Middleware. Oracle Fusion Middleware provides a tool, the Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database.

Use the latest version of RCU to install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

To check if your database is certified or to see all certified databases, refer to the "Certified Databases" section in the Certification Document:

[http://www.oracle.com/technology/software/products/ias/files/fusion\\_certification.html](http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html)

### 6.2.2.1 Running RCU

To install the required metadata for Oracle Fusion Middleware 11g:

1. Start RCU using the following command:

```
RCU_HOME/bin/rcu &
```

2. In the Welcome screen, click **Next**.
3. In the Create Repository screen, select **Create** to load component schemas into a database, and click **Next**.
4. In the Database Connection Details screen, enter connection information for your database:

- **Database Type:** Select `Oracle Database`
- **Host Name:** Enter the name of the node that is running the database. For an Oracle RAC database, specify the VIP name or one of the node names as the host name: `ADFDATABASE1VIRTUAL`.
- **Port:** The port number for the database: `1521`
- **Service Name:** Enter the service name of the database:  
`adfha.mycompany.com`
- **Username:** `SYS`
- **Password:** Enter the password of the SYS user.
- **Role:** `SYSDBA`

Click **Next**.

5. If you see the following message, click **Ignore** or **Stop**:

The database you are connecting is with non-UTF8 charset, if you are going to use this database for multilingual support, you may have data loss. If you are not using for multilingual support you can continue, otherwise, Oracle strongly recommends using a UTF-8 database.

6. In the Select Components screen, select **Create a New Prefix** and enter a prefix to use for the database schemas, for example, ADFHA.

Write down the schema names so they are available in later procedures.

Select the following schemas:

- **AS Common Schemas**
- **Metadata Services**

Click **Next**.

7. In the Schema Passwords screen, enter passwords for the main and additional (auxiliary) schema users and click **Next**.
8. In the Map Tablespaces screen, choose tablespaces for the selected components and click **Next**.
9. In the Summary screen, click **Create**.
10. In the Completion Summary screen, click **Close**.

See the *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about installing RCU.

## 6.2.3 Installing Oracle HTTP Server on WEBHOST1

To install Oracle HTTP Server on WEBHOST1:

1. Verify that the servers meet the following requirements:
  - The system, patch, kernel, and other requirements meet the requirements specified in the *Oracle Fusion Middleware Repository Creation Utility User's Guide*.
  - This example uses port 7777. If you choose port 7777, ensure that it is not used by any service on the nodes by running the following command:

UNIX:

```
netstat -an | grep LISTEN | grep ":7777"
```

Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7777"
```

2. If port 7777 is in use, choose another port, or make it available.
3. On UNIX platforms, if the `/etc/oraInst.loc` or `/var/opt/oracle/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, skip this step.

4. Start Oracle Universal Installer for Oracle Fusion Middleware 11g Webtier Utilities CD installation as follows:

For UNIX, run this command: `./runInstaller`

For Windows, double-click `setup.exe`.

5. In the Specify Inventory Directory screen, enter the location for the inventory and the user group, and click **OK**.
6. Execute the `root` privileged actions as indicated in the dialog, and click **OK**.
7. In the Welcome screen, click **Next**.
8. In the Select Installation Type screen, select **Install and Configure**, and click **Next**.
9. In the Prerequisite Checks screen, ensure that all the prerequisites are met and click **Next**.
10. In the Specify Installation Location screen, set the location to:

```
/u01/app/oracle/product/11.1.1/ohs_1
```

Click **Next**.

11. In the Configure Components screen:

- Select **Oracle HTTP Server**.
- Do not select **Associate Selected Components with WebLogic Domain**.

Click **Next**.

12. In the Specify Component Details screen, enter the following values:

- **Instance Home Location:**

```
/u01/app/oracle/product/11.1.1/ohs_1/instances/ohs_instance1
```

- **Instance Name:** `ohs_instance1`

- **OHS Component Name:** `ohs1`

Click **Next**.

13. In the Specify Webtier Port Details screen:

- Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
- Enter Oracle HTTP Server port. For example, enter `7777`.

Click **Next**.

---

**Note:** For more information about setting ports, refer to the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

---

14. In the Configuration Summary screen, ensure that the selections are correct, and click **Install**.

15. In the Installation Progress screen:

For UNIX systems, a dialog box appears prompting you to run the `oracleRoot.sh` script. Open a command window and run the script, following the prompts.

Click **Next**.

16. In the Configuration screen, several configuration assistants are started in succession. When the configuration assistants are finished, the Configuration Completed screen appears.

17. In the Configuration Completed screen, click **Finish** to exit.

### 6.2.3.1 Validating Oracle HTTP Server

To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by entering the following URL in a web browser:

`WebHost1:7777/`

If Oracle HTTP Server is set up correctly, the Hello World page appears in the browser.

## 6.2.4 Installing the Oracle Fusion Middleware Home

Use the information in these sections to install Oracle Fusion Middleware components:

- [Section 6.2.4.1, "Installing Oracle WebLogic Server"](#)
- [Section 6.2.4.2, "Installing Oracle Fusion Middleware for Oracle ADF Applications"](#)

### 6.2.4.1 Installing Oracle WebLogic Server

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

To install Oracle WebLogic Server on all nodes in the application tier:

1. On UNIX platforms, if the `/etc/oraInst.loc` or `/var/opt/oracle/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory.

If the `/etc/oraInst.loc` file does not exist, skip this step.

2. Start the Oracle WebLogic Server installer.
3. In the Welcome screen, click **Next**.
4. In the Choose Middleware Home Directory screen:
  - Select **Create a New Middleware Home**.
  - For the **Middleware Home Directory** field, enter:

`ORACLE_BASE/product/fmw`

Click **Next**.

5. In the Register for Security Updates screen, enter your contact information for security update notifications, and click **Next**.
6. In the Choose Install Type screen, select **Custom**, and click **Next**.
7. In the JDK Selection screen, select only **Oracle JRockit 1.6.0\_<version> SDK**. Click **Next**.
8. In the Choose Product Installation Directories screen, accept the following directory:

`ORACLE_BASE/product/fmw/wlserver_10.3`

Click **Next**.

9. In the Application Server screen, select **WebLogic**. Click **Next**.



10. In the Installation Summary screen, click **Next**.
11. In the Installation Complete screen, deselect **Run QuickStart**, and click **Done**.

#### 6.2.4.2 Installing Oracle Fusion Middleware for Oracle ADF Applications

To install Oracle Fusion Middleware for Oracle ADF, use the Application Developer Install and perform the following on all the nodes in the application tier:

1. Start the Oracle Fusion Middleware for Oracle Fusion Middleware 11g Application Developer installer:

On UNIX (Linux used in this example):  
 APPHOST1> runInstaller

On Windows:  
 APPHOST1> setup.exe

When Oracle Fusion Middleware 11g Application Developer installer prompts you for a JRE/JDK location enter the Oracle SDK location created in the Oracle WebLogic Server installation in [Section 6.2.4.1, "Installing Oracle WebLogic Server,"](#) for example:

`ORACLE_BASE/product/fmw/jrockit_160_<version>`

2. In the Welcome screen, click **Next**.
3. In the Prerequisite Check screen, verify that the checks complete successfully, and click **Next**.
4. In the Specify Installation Location screen:
  - For **Middleware Home**, enter: `ORACLE_BASE/product/fmw`
  - For **Oracle Home Directory**, enter the directory you want to use, for example: `adf`
 Click **Next**.
5. In the Application Server screen, select **WebLogic**.
6. In the Installation Summary screen, click **Install**.
7. In the Installation Complete screen, click **Finish**.

---

**Note:** Before you run the Configuration Wizard by following the instructions in [Section 6.2.6, "Running the Configuration Wizard on APPHOST1 to Create the WebLogic Server ADF Domain,"](#) ensure that you applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home, so that you have the latest version of Oracle Fusion Middleware.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

---

### 6.2.5 Administration Server High Availability

For information about configuring Oracle WebLogic Server Administration Server, see [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability."](#)

## 6.2.6 Running the Configuration Wizard on APPHOST1 to Create the WebLogic Server ADF Domain

Run the Oracle Fusion Middleware Configuration Wizard from the `adf` directory in the Middleware home to create a domain containing the Administration Server and Oracle components.

1. Start Oracle Fusion Middleware Configuration Wizard from the `MW_HOME/common/bin` directory using the following command:  

```
APPHOST1> ./config.sh
```
2. In the Welcome screen, select **Create a New WebLogic Domain** and click **Next**.
3. In the Select Domain Source screen, select **Generate a domain configured automatically to support the following products**, and select the following products:
  - **Oracle Enterprise Manager - 11.1.1.0**
  - **Oracle JRF - 11.1.1.0**Click **Next**.
4. Enter the **Domain Name**, **Domain Location**, and **Application Location** and click **Next**.
5. In the Configure Administrator Username and Password screen, enter the username and password to be used for the domain's administrator, and click **Next**.
6. In the Configure Server Start Mode and JDK screen, make the following selections:
  - **WebLogic Domain Startup Mode**: select **Production Mode**
  - **JDK Selection**: select **Oracle JRockit 1.6.0\_<version>**Click **Next**.
7. In the Select Optional Configuration screen, select the following:
  - **Deployments and Services**
  - **Managed Servers, Clusters and Machines**
  - **Administration Server**Click **Next**.
8. In the Customize Server and Cluster Configuration screen, select **Yes**, and click **Next**.
9. In the Configure the Administration Server screen, enter the following values:
  - **Name**: AdminServer
  - **Listen Address**: APPHOST1
  - **Listen Port**: 7001
  - **SSL listen port**: NA
  - **SSL enabled**: Leave uncheckedClick **Next**.
10. In the Configure Managed Servers screen, add the following managed servers:

| Managed Server Name | Listen Address       | Listen Port | SSL Listen Port | SSL Enabled |
|---------------------|----------------------|-------------|-----------------|-------------|
| WLS_ADF1            | Hostname of APPHOST1 | 8889        | NA              | unchecked   |
| WLS_ADF2            | Hostname of APPHOST2 | 8889        | NA              | unchecked   |

Click **Next**.

11. In the Configure Clusters screen, add the following cluster:

- **Name:** ADF\_CLUSTER
- **Cluster Messaging Mode:** unicast
- **Cluster Address Enabled:** Leave blank

Click **Next**.

12. In the Assign Servers to Clusters screen, assign the following servers to the Cluster:

- ADF\_CLUSTER:
  - WLS\_ADF1
  - WLS\_ADF2

Click **Next**.

13. In the Configure Machines screen:

- Delete the LocalMachine that appears by default.
- Click the **Unix Machine** tab, and add the following machines:

| Name     | Node Manager Listen Address |
|----------|-----------------------------|
| APPHOST1 | Hostname of APPHOST1        |
| APPHOST2 | Hostname of APPHOST2        |

Click **Next**.

14. In the Assign Servers to Machines screen, assign servers to machines as follows:

- **APPHOST1:** AdminServer, WLS\_ADF1
- **APPHOST2:** WLS\_ADF2

15. In the Target Deployment to Clusters or Servers screen, click **Next**. In the Target Services to Clusters or Servers screen, click **Next**.

16. In the Configuration Summary screen, click **Create**.

17. In the Creating Domain screen, click **Done**.

### 6.2.6.1 Creating boot.properties for the Administration Server and Managed Servers on APPHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a `boot.properties` file for the Administration Server and for the managed servers on APPHOST1.

For the Administration Server, follow these steps:

1. Create the following directory:

```
APPHOST1> mkdir -p MW_HOME/wls/user_
projects/domains/adfdomain/servers/AdminServer/security
```

2. Use a text editor to create a file named `boot.properties` in the directory created in the previous step, and enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

For example:

```
username=weblogic
password=weblogic
```

---

---

**Note:** When you start the Administration Server or Managed Server, username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible in order for the entries to be encrypted.

---

---

For the WLS\_ADF Managed Servers, complete the following:

Copy the file you created for the Administration Server to all servers.

## 6.2.7 Starting the System in APPHOST1

This section describes procedures for starting the system in APPHOST1.

### 6.2.7.1 Starting the Administration Server on APPHOST1

To start the Administration Server on APPHOST1 run the following commands:

```
APPHOST1> cd ORACLE_BASE/product/fmw/user_projects/domains/adfdomain/bin
APPHOST1> ./startWebLogic.sh
```

### 6.2.7.2 Validating the Administration Server

To verify that the Administration Server is properly configured, follow these steps:

1. In a Web browser, go to `http://VIP1:7001/console`.
2. Log in as the administrator.
3. Verify that the WLS\_ADF1 and WLS\_ADF2 managed servers are listed.
4. Verify that the ADF\_Cluster cluster is listed.
5. Verify that you can access Enterprise Manager at `http://VIP1:7001/em`.

### 6.2.7.3 Disabling Host Name Verification for the Administration Server and Managed Servers for APPHOST1 and APPHOST2

You must perform this procedure if you have not set up SSL communication between the Administration Server and the Node Manager. If SSL is not set up, you receive an error message unless you disable host name verification.

You can re-enable host name verification when you have set up SSL communication between the Administration Server and the Node Manager.

To disable host name verification on APPHOST1:

1. In Oracle WebLogic Server Administration Console, select **Administration Server**, **SSL** and then **Advanced**.
2. In the Change Center click **Lock & Edit**.
3. When prompted, save the changes and activate them.
4. Set **Hostname Verification** to **None**.
5. Select **WLS\_ADF1**, **SSL** and then **Advanced**.
6. Set **Hostname Verification** to **None**.
7. Restart the Administration Server and the WLS\_ADF1 Managed Server.

To disable host name verification on APPHOST2:

1. In Oracle WebLogic Server Administration Console, select **WLS\_ADF2**, **SSL**, and then **Advanced**.
2. Set **Hostname Verification** to **None**.
3. Restart the Administration Server and the WLS\_ADF2 Managed Server.

#### 6.2.7.4 Starting Node Manager on APPHOST1

To start Node Manager on APPHOST1:

1. Run the `setNMProps.sh` script located in the `ORACLE_COMMON_HOME/common/bin` directory to set the `StartScriptEnabled` property to `true` before starting Node Manager:

```
OAHOST1> cd ORACLE_COMMON_HOME/common/bin
APPHOST1> ./setNMProps.sh
```

---

**Note:** You must use the `StartScriptEnabled=true` property to avoid class loading failures and other problems.

---

2. Start Node Manager:

```
APPHOST1> cd WL_HOME/server/bin
APPHOST1> ./startNodeManager.sh
```

## 6.2.8 Installing Oracle WebLogic Server and Oracle ADF on APPHOST2

Repeat the procedures for installing WebLogic Server and Oracle ADF for APPHOST2, start with [Section 6.2.3, "Installing Oracle HTTP Server on WEBHOST1."](#) The directory paths for binary files and domains used when installing new nodes must be exactly the same as those used for the first node. If these paths and domains are not exactly the same as those used for the first node, failover does not occur.

## 6.2.9 Propagating the Domain Configuration to APPHOST2 with pack/unpack Utilities

To propagate the domain configuration to APPHOST2 using pack/unpack utilities:

1. Run the following pack command on APPHOST1 to create a template pack:

```
APPHOST1> cd ORACLE_COMMON_HOME/common/bin
APPHOST1> ./pack.sh -managed=true -domain=ORACLE_BASE/product/fmw/user_
projects/domains/adfdomain/
-template=adfdomaintemplate.jar
-template_name=adf_domain_template
```

2. Copy the template file created in the previous step from APPHOST1 to APPHOST2. For example, on a Unix platform:

```
APPHOST1> scp adfdomaintemplate.jar
user@node2:ORACLE_COMMON_HOME/common/bin
```

3. Run the unpack command on APPHOST2 to unpack the propagated template:

```
APPHOST2> cd ORACLE_COMMON_HOME/common/bin
APPHOST2> ./unpack.sh
-domain=ORACLE_BASE/product/fmw/user_projects/domains/adfdomain/
-template=adfdomaintemplate.jar
```

### 6.2.9.1 Creating boot.properties for the Administration Server and Managed Servers on APPHOST2

To create a `boot.properties` file for the Administration Server and for the Managed Servers on APPHOST2:

1. Create the following directories:

```
APPHOST1> mkdir -p MW_HOME/wls/user_projects/domains/adfdomain/servers/WLS_ADF2
APPHOST2> mkdir -p MW_HOME/wls/user_projects/domains/adfdomain/servers/WLS_
ADF2/security
```

2. Use a text editor to create a file named `boot.properties` in the directory created in the previous step, and enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

For example:

```
username=weblogic
password=weblogic
```

---

---

**Note:** When you start the Administration Server or Managed Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible in order for the entries to be encrypted.

---

---

### 6.2.9.2 Starting Node Manager on APPHOST2

To start the Node Manager on APPHOST2, repeat the steps from [Section 6.2.7.4, "Starting Node Manager on APPHOST1"](#) on APPHOST2.

### 6.2.9.3 Configuring the ADF Application for Replication

Use the procedures in this section to configure your application for replication.

### Clustering Requirement

The application must be deployed to an Oracle WebLogic Cluster. This automatically establishes a replication channel for the multiple instances of the application.

---

**Note:** In a Unicast cluster, the default replication channel is configured using the Listen address of each managed server. Therefore, the Listen address should be configured to be a specific IP address or host name, instead of being configured to listen on Any.

---

### Oracle ADF Replication

It is essential that Oracle ADF is configured properly. The following tag should be present in the `adf-config.xml` file, one of the Application Resources, for a stateful application:

```
<adfc:adf-controller-config><adfc:adf-scope-ha-support>true</adfc:adf-scope-ha-sup
port></adfc:adf-controller-config>
```

Applications must also have replication enabled. Oracle WebLogic Server allows several types of persistent stores for replication. For more information on persistent stores, see the *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server* manual.

The ADF application can be enabled by for this by default with the following setting in the `weblogic.xml` file:

```
<session-descriptor>
<persistent-store-type>replicated_if_clustered</persistent-store-type>
</session-descriptor>
```

The `replicated_if_clustered` setting disables replication for standalone application environments, and uses in-memory replication within a cluster environment.

Ensure that any custom application is configured for in-memory replication.

#### 6.2.9.4 Deploying the ADF Application

After the cluster is set up, you can deploy the ADF application. Be aware of the following:

- Deploy the ADF application to an EAR file using the Administration Server Console.
- Ensure that the deployment is to a cluster.
- If your application uses MDS, register the MDS with the domain using the instructions in the "Registering a Database-Based MDS Repository" section in the *Oracle Fusion Middleware Administrator's Guide* manual.

#### 6.2.9.5 Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers

Enable Oracle HTTP Server to route to the Administration Server that contains WebCenter Portal managed servers by setting the `WebLogicCluster` parameter to the list of nodes in the cluster. Follow these steps:

1. Add the following lines to the `OHS_HOME/instances/ohs_instance1/config/OHS/ohs1/mod_wl_ohs.conf` file:
 

```
Spaces
```

```
<Location /applicationMountpoint>
 WebLogicCluster apphost1.com:8888,apphost2.com:8889
 SetHandler weblogic-handler
</Location>
```

## 2. Restart Oracle HTTP Server on WEBHOST1:

```
WEBHOST1> OHS_HOME/instances/ohs_instance1/bin/opmnctl restartproc
ias-component=OHS_COMPONENT1
```

### 6.2.9.6 Validating Access through Oracle HTTP Server

Verify the URLs to ensure that appropriate routing and failover is working from the HTTP Server to WebCenter Portal cluster. Follow these steps:

1. Start WC\_Spaces1, WC\_Spaces2, WC\_Collaboration1, WC\_Collaboration2, WC\_Uutilities1, WC\_Uutilities2, WC\_Portlet1, and WC\_Portlet2 from the WebLogic Server Administration Console as follows:

- a. Access the Administration Console at the following URL:

```
http://APPHOST1/console
```

- b. Click on one of the Managed Servers, for example, WC\_Spaces1.

- c. Select the **Control** tab.

- d. Select **Start** to start the Managed Server.

- e. Repeat the previous steps for each Managed Server.

- f. Verify direct access to the Managed Servers using the following URLs:

```
apphost1:8888/applicationMountpoint
```

```
apphost2:8888/applicationMountpoint
```

```
apphost1:8889/applicationMountpoint
```

```
apphost2:8889/applicationMountpoint
```

2. While WLS\_ADF2 is running, stop WLS\_ADF1 from Oracle WebLogic Server Administration Console.

3. Access the following URL and verify the appropriate functionality:

```
WebHost1:7777/applicationMountpoint
```

4. Start WLS\_ADF1 from the WebLogic Server Administration Console.

5. Stop WLS\_ADF2.

6. Access the following URL and verify the appropriate functionality:

```
WebHost2:7777/applicationMountpoint
```

### 6.2.10 Scaling the Topology

You can scale out and scale up an Oracle ADF topology. When you scale up the topology, you add new managed servers to nodes that are already running one or more managed servers. When you scale out the topology, you add new managed servers to new nodes.



### 6.2.10.1 Scaling Up the Topology (Adding Managed Servers to Existing Nodes)

In this case, you already have a node that runs a managed server configured with Oracle ADF. The node contains a Middleware home in shared storage.

You can use the existing installations (Middleware home and domain directories) for creating new servers. There is no need to install Oracle Fusion Middleware binaries in a new location or to run pack and unpack.

To scale up the topology:

1. Using the Oracle WebLogic Server Administration Console, clone WLS\_ADF1 into a new managed server. The source managed server to clone should be one that already exists on the node where you want to run the new managed server.

To clone a managed server:

- a. Select **Environment > Servers** from the Administration Console.
- b. Select the managed server that you want to clone.
- c. Select **Clone**.

Name the new managed server WLS\_ADFn, where n is a number to identify the new managed server.

2. For the listen address, assign the host name or IP to use for this new managed server.
3. Ensure that the port number for this managed server is available on this node.
4. Reconfigure the Oracle HTTP Server module with the new member in the cluster. See [Section 6.2.9.5, "Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers."](#)

### 6.2.10.2 Scaling Out the Topology (Adding Managed Servers to New Nodes)

In scaling out your topology, you add new managed servers configured with Oracle ADF applications to new nodes.

Before performing the steps in this section, check that you meet these requirements:

- In your topology, there are existing nodes running managed servers configured with ADF applications.
- ADF managed servers are clustered and the new managed server will also be part of that cluster.
- The new node can access the existing home directories for WebLogic Server. You use the existing installations in shared storage for creating a new managed server. There is no need to install WebLogic Server binaries in a new location, although you must run pack and unpack to create a managed server domain.

To scale out the topology:

1. On the new node, mount the existing Middleware home, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the `MW_HOME/boa/beahomelist` file and add `ORACLE_BASE/product/fmw` to it.
3. Log into the Oracle WebLogic Administration Console.
4. Create a new machine for the new node that will be used, and add the machine to the domain.

5. Update the machine's Node Manager's address to map the IP of the node that is being used for scale out.
6. Use the Oracle WebLogic Server Administration Console to clone WLS\_ADF1 into a new managed server. Name it WLS\_ADFn, where n is a number and assign it to the new machine.
7. For the listen address, assign the host name or IP to use for the new managed server.

To set the managed server listen address:

- a. Log into the Oracle WebLogic Server Administration Console.
- b. In the Change Center, click **Lock & Edit**.
- c. Expand the **Environment** node in the Domain Structure window.
- d. Click **Servers**. The Summary of Servers page is displayed.
- e. Select the managed server whose listen address you want to update in the **Names** column of the table. The Setting page for that managed server is displayed.
- f. Set the Listen Address to APPHOSTn where APPHOSTn is the DNS name of your new machine.
- g. Click **Save**.
- h. Save and activate the changes.

The changes do not take effect until the managed server restarts.

8. Run the pack command on APPHOST1 to create a template pack as follows:

```
APPHOST1> cd ORACLE_COMMON_HOME/common/bin
APPHOST1> ./pack.sh -managed=true -domain=MW_HOME/user_projects/domains/
adfdomain/ -template=adfdomaintemplateScale.jar
-template_name=adf_domain_templateScale
```

Run the following command on APPHOST1 to copy the template file created to APPHOSTn:

```
APPHOST1> scp adfdomaintemplateScale.jar oracle@APPHOSTn:/
ORACLE_COMMON_HOME/common/bin
```

Run the unpack command on APPHOSTn to unpack the template in the managed server domain directory as follows:

```
APPHOSTn> cd ORACLE_COMMON_HOME/common/bin
APPHOSTn> ./unpack.sh
-domain=ORACLE_BASE/product/fmw/user_projects/domains/adfdomain/
-template=adfdomaintemplateScale.jar
```

9. Start the Node Manager on the new node. To start the Node Manager, use the installation in shared storage from the existing nodes, and start Node Manager by passing the host name of the new node as a parameter as follows:

```
APPHOSTn> WL_HOME/server/bin/startNodeManager <new_node_ip>
```

10. Start and test the new managed server from the Oracle WebLogic Server Administration Console:
  - a. Shut down all the existing managed servers in the cluster.
  - b. Ensure that the newly created managed server is running.

- c. Access the application on the newly created managed server. The application should be functional.

## 6.3 WebCenter Portal and High Availability Concepts

The information in this section guides you through the issues and considerations necessary for designing a WebCenter Portal high availability cluster.

### 6.3.1 Understanding WebCenter Portal

WebCenter Portal combines the standards-based, declarative development of Java Server Faces (JSF), the flexibility and power of portals, and a set of integrated Web 2.0 services.

#### 6.3.1.1 WebCenter Portal Components

WebCenter Portal includes the following components.

- **Oracle WebCenter Portal: Spaces** offers a single, integrated, Web-based environment for social networking, communication, collaboration, and personal productivity through a robust set of services and applications.  
  
The Spaces application is built using JSF, Oracle ADF, Oracle WebCenter Portal: Framework, Oracle WebCenter Portal Services, and Composer. The Spaces application provides:
  - A browser-based platform for creating enterprise portals, multiple sites and communities.
  - A Home space for each user, providing a private work area for storing personal content, keeping notes, viewing and responding to business process assignments, emailing, and so on.
  - Threaded discussions, blogs, wikis, worklists, announcements, RSS, recent activities, search, and more.
- **Oracle WebCenter Portlets Framework** supports deployment and execution of both standards-based portlets (JSR 168, WSRP 1.0 and 2.0), and traditional Oracle PDK-Java based portlets. WebCenter Portal provides several out-of-the-box producers, such as OmniPortlet, Web Clipping, and WSRP Tools.
- **Oracle WebCenter Portal:Framework** provides the following capabilities:
  - Run-time application customization (you can make in-place changes to WebCenter Portal applications using Composer without re-deploying the application)
  - Support for JSR-168 and JSR-286 standards-based WSRP portlets, and PDK-Java portlets
  - Content integration through JCR (JSR170), to content repositories such as Oracle WebCenter Content Server, Oracle Portal, and file systems
  - Oracle JSF Portlet Bridge, which lets you expose JSF pages and Oracle ADF task flows as standards-based portlets
- **Oracle WebCenter Portal's Discussions server** provides the ability to integrate discussion forums and announcements into your applications.
- **Oracle WebCenter Portal's Analytics Collector** provides users with the ability to view reports on the various user activities within WebCenter Portal applications, for example:

- Logins
  - Page views
  - Portlet views
  - Document views
  - Search metrics
  - Response times
- **Oracle WebCenter Portal's Activity Graph** service provides users with the ability to analyze the statistics collected by Oracle WebCenter Portal's Analytics. The output of an Activity Graph analysis is the collected scores for objects and users, which are used to give recommendations. The scores are stored in the Activities database.
  - **Oracle WebCenter Portal's Personalization Server** is a lightweight service that client applications access by means of RESTful web service. The Personalization server enables you to deliver application content to targeted users based on selected criteria.


WebCenter Portal applications can also integrate with WebCenter Portal: Services. See the table Services Data Repositories in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*. This guide also describes how to configure WebCenter Portal services.

### 6.3.1.2 WebCenter Portal Single-node Architecture

WebCenter Portal installation creates a WebCenter directory under the Middleware home directory. The installation creates a WebCenter domain (`wc_domain`), which contains the Administration Server and four WebLogic Managed Servers: `WC_Spaces1` (which hosts the Spaces application), `WC_Portlets` (which hosts several out-of-the-box portlet producers), `WC_Collaboration1` (which host the Discussions server), `WC_Uilities1` (which hosts Analytics, Activity Graph, and the Personalization server), and any other WebCenter Portal services that you choose to integrate).

WebCenter Portal applications that you build using WebCenter Portal: Framework, are referred to as "Framework applications". When you create additional custom managed servers, the appropriate libraries provision them so that they can draw on the same external resources as WebCenter Portal: Spaces. For more information about managed servers, see "Understanding Oracle Fusion Middleware Concepts" in the *Oracle Fusion Middleware Administrator's Guide*. [Figure 6-4](#) shows a basic single-node WebCenter Portal architecture.

**Figure 6-4 Basic Single-Node WebCenter Portal Architecture**



### **6.3.1.3 WebCenter Portal: State and Configuration Persistence**

See WebCenter Portal State and Configuration Persistence in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.


### **6.3.1.4 WebCenter Portal Log File Locations**

See WebCenter Portal Log File Locations in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

## **6.3.2 WebCenter High Availability Architecture and Failover Considerations**

An Oracle WebLogic cluster provides high availability for applications. When one member of the cluster is unavailable, another member of the cluster handles the request.

You can deploy each of the managed servers in a WebCenter Portal deployment as a cluster, with different cluster members on the same node or on different nodes. In [Figure 6–5](#), either cluster member can serve all requests sent to the cluster equally.

**Figure 6–5 WebCenter Portal Two-Node High Availability Architecture**

The following sections contain more information on the runtime and configuration aspects of a WebCenter Portal cluster.

### 6.3.2.1 WebCenter Portal Applications

During WebCenter Portal installation, the managed servers are provisioned with system libraries and ADF libraries. [Table 6–2](#) lists the managed servers and the applications which run on them.

**Table 6–2 WebCenter Portal’s Managed Servers and Applications**

| Managed Server   | Installed Application(s)                                                                      |
|------------------|-----------------------------------------------------------------------------------------------|
| WC_Spaces        | Spaces<br>Spaces Online Help                                                                  |
| WC_Portlets      | OmniPortlet and Web Clipping<br>WSRP Tools<br>Pagelet Producer<br>WebCenter Services Producer |
| WC_Collaboration | Discussions Server                                                                            |
| WC_Utilities     | Analytics Collector<br>Activity Graph Engines<br>Personalization Services                     |

### 6.3.2.2 Oracle WebCenter Startup Order

When you start a managed server, applications and libraries start in the following order:

1. Oracle System libraries, known as the JRF libraries.
2. Oracle ADF libraries.
3. Instrumentation applications such as Oracle DMS.
4. WebCenter Portal applications.

The startup order is also the order of dependency. If a dependent component does not deploy successfully, a later component may not function correctly. WebCenter Portal application startup does not depend on the availability of external services such as a discussions server, or other back-end servers.

### 6.3.2.3 Deploying WebCenter Portal Application on a Cluster

For a WebCenter Portal cluster deployment such as the one shown in [Figure 6–5](#), follow these rules for the targeting of applications, libraries, and system resources:

- Target applications and libraries to the cluster target. For example, target the Spaces application to the Spaces cluster.
- Target JDBC resources to the cluster target.

WebCenter Portal applications and the WebCenter Portal Discussions server deploy as Oracle WebLogic *stage* applications. During the initial deployment of each application, the deployment files are received from the Administration Server and deployed locally.

### Cluster Communications

By default, each WebCenter Portal cluster is configured as unicast. To configure your WebCenter Portal cluster for multicast, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

### 6.3.2.4 WebCenter Portal Analytics Collector

Analytics in WebCenter Portal consists of individual Analytics Collectors that receive requests from analytics clients such as the Spaces application.

This section describes how an Analytics Collector cluster works.

In a clustered environment, producers (clients of the collector) and collectors are configured with a cluster-specific channel name. Each collector periodically broadcasts a heartbeat with its location to the cluster-specific channel. The producer listens to the channel for these collector heartbeats and, when it hears one, adds the collector to its list of known collectors. When the producer needs to send an event, it uses a round robin algorithm to select a collector from its list and sends the event to that collector. If a collector stops (intentionally or due to failure), it stops broadcasting a heartbeat. When the producer stops hearing the heartbeat it removes the collector from its list and stops sending events to that collector. If the producer does not hear any collector heartbeats, it does not send any events.

WebCenter Portal Analytics uses UDP and multicast on a configured set of ports to communicate with its clients. For a single node setup, the client is configured with a WLST command that has the server host/port location and simply transmits all events to this location via UDP. For a multiple node setup, the server is configured to broadcast (via UDP multicast) the location(s) of the various servers running on the cluster; the client is configured with the same WLST command, so it receives the server locations and keeps the list of available servers (which persists in memory.) If a client does not receive a periodic heartbeat from a server, it removes that server from its list of known servers and stops sending events to it.

Oracle recommends that you configure Analytics Collectors as unicast, forming a 1-1 relationship with the WebCenter Portal application client.

### 6.3.2.5 WebCenter Portal State Replication

WebCenter Portal relies on Oracle ADF, which has several stateful components. Oracle WebCenter itself is also a stateful application. Therefore, you must configure state replication in cluster scenarios.

For more information on how state replication works in Oracle WebLogic Server, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

Oracle WebLogic Server supports two types of state replication:

- **In-memory replication** - Using in-memory replication, Oracle WebLogic Server copies a session state from one server instance to another. The primary server creates a primary session state on the server to which the client first connects, and a secondary replica on another WebLogic Server instance in the cluster. The replica is kept up to date so that it may be used if the server that hosts the servlet fails.
- **JDBC-based persistence** - In JDBC-based persistence, Oracle WebLogic Server maintains the HTTP session state of a servlet or JSP using file-based or database-based persistence.

To configure state replication successfully, you must configure both the environment and the application properly. For information on state replication behavior under failover conditions see [Section 6.3.2.8, "Expected Behavior for Application Failover."](#)

For more information on diagnosing state replication issues see [Section 6.3.2.5, "WebCenter Portal State Replication."](#)

### 6.3.2.6 Understanding the Distributed Java Object Cache

WebCenter Portal applications use a distributed Java Object Cache for greater performance. Configure this cache across all WebCenter Portal clusters, with one distributed cache per cluster.

[Table 6–3](#) lists examples of the object types that WebCenter Portal places in the Java Object Cache.



**Table 6–3 WebCenter Portal Object Types for Java Object Cache**

| <b>WebCenter Portal Component</b>      | <b>Object Cached</b>                                                                                                        |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Discussions service                    | Topics and forums                                                                                                           |
| Announcements service                  | Announcements                                                                                                               |
| Instant Messaging and Presence service | Presence subscription lists<br>User's presence/subscription status                                                          |
| Worklist service                       | Called Worklist items, such that cached data is used until refresh is called.                                               |
| Content Integration (Oracle Portal)    | JCR: type information and metadata obtained from the repository.                                                            |
| Service Framework                      | User profile. Queried usernames.                                                                                            |
| Recent Activity service                | Recent activity results per user.                                                                                           |
| Spaces application                     | Global list of spaces and space templates in the application.<br>List of spaces and space templates that a user can access. |
| Page service                           | List of pages in a scope.                                                                                                   |
| WSRP server                            | Preference store values for WSRP producers.                                                                                 |
| Documents service                      | Provisioning and configuration checks for the Documents service configured for the Spaces application                       |
| Profile management                     | Lightweight user profile objects                                                                                            |
| Navigation service                     | List of active Navigation Model objects                                                                                     |
| Portlet Consumer                       | Portlet markup                                                                                                              |

### **Collaboration**

Collaboration services cache objects in the Java Object Cache on a per user session basis. These cached user sessions are destroyed when the HTTP session is destroyed.

### **Worklist**

Worklist caches the called items so that unless you click the refresh button or trigger the fifteen-minute refresh poll, the same items are displayed as the user changes the sort and group by settings. Group and sort order settings also cache the display, updating when a change occurs. This is read only data which, if not present, is fetched and stored on the cache.

### **Spaces Application**

The list of all templates and public spaces are maintained in the Java Object Cache. All users share the list, in addition to the list of spaces and templates that a particular user can access. These are cached per user. A template created on one JVM does not show up if the template cache has been initialized in another JVM unless JOC distributes the data or an administrator in the second JVM does an explicit refresh where the cache is rebuilt.

### **Documents Service**

The Documents service uses Java Object Cache to cache provisioning and configuration checks for document services, as applies to Spaces application configuration. For provisioning, cached objects are flagged as distributed. They are

replicated by a correctly configured Java Object Cache in a high availability environment, the configuration cached state is kept locally. All cached objects are flagged to expire after one minute. Caching reduces the number of times Oracle WebCenter Content calls are made to check the state of the Oracle WebCenter Content Server, as the Spaces application repeatedly checks provisioning and configuration to control service rendering in the UI.

### **Portlet Consumer**

The portlet consumer caches portlet mark-up in the Java Object Cache according to the cache headers in the portlet response. The cache can be expires based or validation based.

Profile Management caches the lightweight user profile objects in Java Object Cache. If particular profile data is not found, it will queried from the backend and the cache would be populated. The number of objects stored have an upper bound of 1000 and are in the Java Object Cache for one hour.

Navigation caches the Navigation Model objects in the Java Object Cache on a per user session basis. These cached objects are destroyed when the HTTP session is destroyed.

### **Recent Activity**

The list of recent activity results are cached for each user to prevent a requery of results each time the recent activity task flow or RSS feed is viewed. The cache is automatically refreshed every fifteen minutes, or can be manually refreshed using the refresh icon in the recent activity task flow.

For Java Object Cache configuration procedures, see [Section 6.4.13, "Configuring the Java Object Cache."](#)

### **6.3.2.7 WebCenter Portal Protection from Failover and Expected Behavior**

An Oracle WebLogic cluster provides application high availability. If one member of the cluster is unavailable, any other available member of the cluster is able to handle the request.

### **Session Failover Requirements**

For seamless failover of a WebCenter Portal application, the application must meet the following conditions:

- Must be in a cluster and at least one member of the application cluster is available to serve the request.
- If stateful, state replication is configured correctly as described in [Section 6.4.15, "Configuring WebCenter Portal for Replication."](#)
- If you using Oracle HTTP Server, the server configuration is configured with the WebLogicCluster directive to balance among all available application instances.
- If you are using a hardware load balancer, the load balancer is:
  - Routing traffic to all available instances
  - Configured correctly with a health monitor to mark unavailable instances
  - Configured to support persistence of session state

### 6.3.2.8 Expected Behavior for Application Failover

If the environment is configured correctly, application users do not notice when an application instance in a cluster becomes unavailable. An example of the sequence of events during an application failover follows:

1. A user makes a request and is routed by a hardware load balancer to instance A of the application.
2. Instance A becomes unavailable because of node, process, or network failure.
3. The hardware load balancer marks Instance A as unavailable.
4. The user makes a subsequent request that is routed to instance B.
5. Instance B is configured as a replication partner of Instance A and has the user's session state.
6. The application resumes using the session state on Instance B and the user continues working without interruption.

### Exceptions to Expected Behavior

For WebCenter Portal, known exceptions are as follows:

- **Oracle ADF Pop-ups** - Open pop-ups close on failover, affecting the following components which otherwise have no exceptions:
  - Composer (Property Inspector, EL Builder, Task Flow parameter picker, Resource String Editor)
  - Save (confirmation dialog box)
  - Lists
  - Links (link deletion pop-up)

When failover occurs, you must repeat the action that leads to the pop-up to make it reappear. The specific ways in which this appears in Spaces and suggested remedies are listed in [Table 6–4](#).

**Table 6–4 WebCenter Portal Troubleshooting Scenarios**

| Action Before Failover                                                                                        | After Failover                                                                                                                                                                                                     | Suggested Remedy                                                                                 |
|---------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Go to any page and click <b>Create Page</b> .                                                                 | Type in a name, select a theme, and click <b>OK</b> . When you select the theme, the page creation pop-up closes.                                                                                                  | Repeat the operation.                                                                            |
| Launch Manage Pages.                                                                                          | Perform any operation within the pop-up, except for closing the pop-up, for example, click <b>Page Actions</b> . When you perform any operation, the Manage Pages pop-up is closed.                                | Repeat the operation.                                                                            |
| Launch Manage Pages, click <b>Page Actions</b> against a page, and then the <b>Delete</b> option in the menu. | Click <b>OK</b> on the confirmation pop-up. Clicking <b>OK</b> closes the confirmation pop-up and the Manage Pages pop-up. The result of the deletion (which may have gone through) is not visible among the tabs. | Relaunch the Manage Pages pop-up to see if the page is deleted. If not, try deleting once again. |

**Table 6–4 (Cont.) WebCenter Portal Troubleshooting Scenarios**

| Action Before Failover                                                                                                                                                          | After Failover                                                                                                                                                                                                                                         | Suggested Remedy                                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Launch <b>Personalize Applications</b> pop-up.<br>Perform any operation that sends a request, other than clicking <b>OK</b> , for example, expand the <b>Applications</b> node. | Perform any operation that sends a request, other than clicking <b>OK</b> , for example, expand the <b>Applications</b> node. When you perform any operation that sends a request to the server, the <b>Personalize Applications</b> pop-up is closed. | Repeat the operation                                                                                             |
| Launch Preferences.                                                                                                                                                             | Switch between the <b>Preferences</b> tabs (General, Accounts, Messaging, Search). When you switch between the <b>Preference</b> tabs, the <b>Preferences</b> pop-up is closed.                                                                        | Repeat the operation                                                                                             |
| Launch Manage Favorites.<br>Stop the server, perform any operation other than closing the pop-up. For example, expand a folder, click <b>Edit favorite information</b> .        | Perform any operation other than closing the pop-up, for example, expand a folder, and click <b>Edit favorite information</b> . When you perform any operation, the <b>Manage Favorites</b> pop-up is closed.                                          | Relaunch the <b>Manage Favorites</b> pop-up to see if the operation was successful. If not, retry the operation. |
| Choose to edit applications and create a new folder                                                                                                                             | An MDS exception opens.                                                                                                                                                                                                                                | Retry the operation                                                                                              |
| Work in Composer Customization Manager.                                                                                                                                         | Oracle Composer loses the state associated with the Customization Manager                                                                                                                                                                              | Relaunch Customization Manager.                                                                                  |
| Select a component while in Oracle Composer Source View                                                                                                                         | Oracle Composer loses track of the component.                                                                                                                                                                                                          | Repeat the operation                                                                                             |

- **Component Specific Issues** - Other issues which are specific to different components in WebCenter Portal are listed in [Table 6–5](#).

**Table 6–5 WebCenter Portal Exceptions to Expected Behavior**

| WebCenter Portal Component | Exceptions to Expected Behavior                                                                                                                                                                                                                                                                                                           |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Pagelet Producer           | Unsaved data or pending changes are not preserved on failover. If failover occurs, administrators must reestablish their administrative session and users may need to reestablish their session if the proxy is required to have a state. If SSO is configured, credentials are automatically provided, and the session is reestablished. |
| Space Events               | When changing certain fields (start or end date/hour/minute), the event form closes when failover occurs.                                                                                                                                                                                                                                 |
| Portlets                   | Failures are fully transparent.                                                                                                                                                                                                                                                                                                           |
| Lists                      | Failures are fully transparent.                                                                                                                                                                                                                                                                                                           |
| Links                      | Failures are fully transparent.                                                                                                                                                                                                                                                                                                           |

**Table 6–5 (Cont.) WebCenter Portal Exceptions to Expected Behavior**

| WebCenter Portal Component | Exceptions to Expected Behavior                                                                                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Search                     | In the midst of a long-running query. The results that come back are not guaranteed (note there is no "write" operation here), the user can just reissue the query.                                                                                                                                                 |
| Tagging                    | Failures are fully transparent.                                                                                                                                                                                                                                                                                     |
| Recent Activities          | The open/close state of the tree nodes is not replicated. After after failover, the tree of results closes all of the nodes. The nodes need to be reopened.                                                                                                                                                         |
| Worklist                   | Failures are fully transparent.                                                                                                                                                                                                                                                                                     |
| Document Manager           | When a user uploads a document and a document with the same name already exists, the user is taken to the Confirm new version screen. While on that screen, the file is stored in a temporary local location. If failover occurs at that moment, the uploaded file is lost and you must restart the upload process. |

### 6.3.2.9 Monitoring Logging of Application Deployments

Use Oracle WebLogic Server Administration Console to check the status of the application deployments. You can also use Oracle WebLogic Server infrastructure and Enterprise Manager Fusion Middleware Control for starting stopping, and monitoring WebCenter Portal processes.

### 6.3.2.10 WebCenter Portal Cluster-wide Configuration Changes

For the Spaces application and portlet producers, all configuration data is stored in the MDS repository and portlet producers. Additional cluster deployments automatically read the latest configuration upon application startup.

For the discussions server, the configuration information should be moved over from an existing cluster server. The pack/unpack utility of Oracle WebLogic Server does this automatically. Oracle recommends this procedure.

### 6.3.2.11 Maintaining Configuration in a Clustered Environment

For the Spaces application and portlet producers, all configuration data is stored in the MDS repository and portlet producers. Any changes made to the configuration of one server in the cluster are immediately visible to all other members.

Changes to discussions server are not frequent, however, when they do occur, the changes must be reapplied to other members of the cluster. You can do this by connecting directly to each discussions server, instead of using the load balancer, and making the necessary administration changes.

## 6.4 Configuring High Availability for WebCenter Portal

[Figure 6–5](#) illustrates a two-node WebCenter Portal cluster running on two Oracle WebLogic Servers in one WebLogic Server domain. Oracle WebLogic Servers are front-ended by Oracle HTTP Server which load balances incoming requests to them. An Oracle RAC database stores metadata and schemas. VIPs are used for the Administration Server (for manual failover).

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

**Note:** All command examples are for k or bash shell. No C shell examples are provided.

---

Steps to configure WebCenter Portal for a high availability configuration are:

- [Section 6.4.1, "Preparing the Environment: Prerequisite Steps Before Setting up a WebCenter Portal High Availability Configuration"](#)
- [Section 6.4.2, "Installing Oracle HTTP Server on WEBHOST1"](#)
- [Section 6.4.3, "Installing Oracle Fusion Middleware Home"](#)
- [Section 6.4.4, "Enabling the Administration Server VIP"](#)
- [Section 6.4.5, "Running Oracle Fusion Middleware Configuration Wizard on APPHOST1 to Create the WebLogic Server WebCenter Domain"](#)
- [Section 6.4.6, "Creating boot.properties for the Administration Server and for Managed Servers on APPHOST1"](#)
- [Section 6.4.7, "Starting the System in APPHOST1"](#)
- [Section 6.4.8, "Install WebLogic Server and WebCenter Portal on APPHOST2"](#)
- [Section 6.4.9, "Propagating the Domain Configuration to APPHOST2 with pack/unpack Utilities"](#)
- [Section 6.4.10, "Starting Node Manager on APPHOST2"](#)
- [Section 6.4.11, "Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers"](#)
- [Section 6.4.12, "Configuring Manual Failover of the Administration Server to APPHOST2"](#)
- [Section 6.4.13, "Configuring the Java Object Cache"](#)
- [Section 6.4.14, "Configuring the Distributed Notifications for the MDS Repository"](#)
- [Section 6.4.15, "Configuring WebCenter Portal for Replication"](#)
- [Section 6.4.16, "Configuring Analytics"](#)
- [Section 6.4.17, "Configuring Activity Graph"](#)
- [Section 6.4.18, "Configuring Clustering for Discussions Server"](#)
- [Section 6.4.19, "Scaling the Topology"](#)
- [Section 6.4.20, "Troubleshooting WebCenter Portal High Availability"](#)
- [Section 6.4.21, "Converting Discussions from Multicast to Unicast"](#)

## 6.4.1 Preparing the Environment: Prerequisite Steps Before Setting up a WebCenter Portal High Availability Configuration

The following sections provide prerequisite steps before setting up a WebCenter Portal high availability configuration.

For information about platform-specific commands, see the *Oracle Fusion Middleware Installation Guide for Oracle WebCenter*.

#### 6.4.1.1 Database Prerequisites

WebCenter Portal requires the presence of a supported database and schemas.

To check if your database is certified or to see all certified databases, refer to the "Certified Databases" section in the Certification Document:

[http://www.oracle.com/technology/software/products/ias/files/fusion\\_certification.html](http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html)

To determine the database version, execute the following query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

#### 6.4.1.2 VIP and IP Prerequisites

To configure a virtual IP for the Administration Server, see [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

#### 6.4.1.3 Installing and Configuring the Database Repository

This section describes how to install and configure the database repository.

##### Oracle Clusterware

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

##### Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.
- When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

##### Oracle Real Application Clusters

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

You must install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database before you install the Oracle Fusion Middleware components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. You install RCU in its own, separate Middleware home.

Use the latest version of RCU to install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

### Database Initialization Parameters

Ensure that the following initialization parameter is set to the required value. It is checked by Repository Creation Assistant.

**Table 6–6 Required Initialization Parameters**

| Parameter | Required Value | Parameter Class |
|-----------|----------------|-----------------|
| PROCESSES | 300 or greater | Static          |

To check the value of the initialization parameter using SQL\*Plus, you can use the SHOW PARAMETER command.

As the SYS user, issue the SHOW PARAMETER command as follows:

```
SQL> SHOW PARAMETER processes
```

Set the initialization parameter using the following command:

```
SQL> ALTER SYSTEM SET processes=300 SCOPE=SPFILE
```

Restart the database.

---

**Note:** The method that you use to change a parameter's value depends on whether the parameter is static or dynamic and on whether your database uses a parameter file or a server parameter file. See the *Oracle Database Administrator's Guide* for details on parameter files, server parameter files, and how to change parameter values.

---

#### 6.4.1.4 Installing and Configuring an LDAP Provider

For production environments, it is a mandatory requirement for WebCenter Portal high availability topologies to have an external LDAP policy store. For more information on the supported policy stores and instructions to configure LDAP, see the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

#### 6.4.1.5 Terminology for Directories and Directory Environment Variables

The follow list describes the directories and variables used in this chapter:

- **ORACLE\_BASE:** This environment variable and related directory path refers to the base directory under which Oracle products are installed.
- **MW\_HOME:** This environment variable and related directory path refers to the location where Fusion Middleware (FMW) resides.
- **WL\_HOME:** This environment variable and related directory path contains installed files necessary to host a WebLogic Server.
- **ORACLE\_HOME:** This environment variable and related directory path refers to the location where WebCenter Portal is installed.
- **ORACLE\_COMMON\_HOME:** This environment variable and related directory path refers to the Oracle home that contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF).



- **DOMAIN\_HOME:** This directory path refers to the location where the Oracle WebLogic Domain information (configuration artifacts) is stored.

Recommended values for consistency for these directories are:

**ORACLE\_BASE:**

`/u01/app/oracle`

**MW\_HOME (application tier):**

`ORACLE_BASE/product/fmw`

**WL\_HOME:**

`MW_HOME/wlserver_10.3`

**ORACLE\_HOME:**

`MW_HOME/WC1`

**ORACLE\_COMMON\_HOME:**

`MW_HOME/oracle_common`

#### 6.4.1.6 Using Oracle Fusion Middleware Repository Creation Utility to Load the Fusion Middleware Schemas in the Database

Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and WebCenter Portal schemas into a Real Application Cluster database before you install WebCenter Portal components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database.

Use the latest version of RCU to install the 11g (11.1.1) Oracle Fusion Middleware Repository into a Real Application Clusters database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

##### 6.4.1.6.1 Running RCU

Run RCU to install the required metadata for Oracle Fusion Middleware 11g.

1. Start RCU using the following command:

`RCU_HOME/bin/rcu`

2. In the Welcome screen, click **Next**.
3. In the Create Repository screen, select **Create** to load component schemas into a database and click **Next**.
4. In the Database Connection Details screen, enter connection information for your database:
  - Database Type: Select **Oracle Database**.
  - Host Name: Enter the name of the node that is running the database. For an Oracle RAC database, specify the VIP name, or one of the node names as the host name: **WCDBHOST1VIRTUAL**.
  - Port: The port number for the database: **1521**
  - Service Name: Enter the service name of the database: **wcha.mycompany.com**

- Username: **SYS**
  - Password: Enter the password for the SYS user.
  - Role: **SYSDBA**
5. Click **Next**.
  6. If you receive the following warning message, click **Ignore** or **Stop**:

The database you are connecting is with non-UTF8 charset, if you are going to use this DB for multilingual support, you may have data loss. If you are not using for multilingual support you can continue, otherwise we strongly recommend using UTF-8 database.
  7. In the Select Components screen, select **Create a New Prefix**, and enter a prefix to use for the database schemas, for example, **WCHA**

Write down the schema names so they are available in later procedures.

    - Select the following schemas:
      - AS Common schemas:  
**Metadata Services**
      - WebCenter Infrastructure:  
**WebCenter Portal**
  8. Click **Next**.
  9. In the Schema Passwords screen, enter passwords for the main and additional (auxiliary) schema users. Click **Next**.
  10. In the Map Tablespaces screen, choose the tablespaces for the selected components. Click **Next**.
  11. In the Summary screen, click **Create**.
  12. In the Completion Summary screen, click **Close**.

See *Oracle Fusion Middleware Installation Concepts and Repository Creation Utility User's Guide* for more information about using RCU.

## 6.4.2 Installing Oracle HTTP Server on WEBHOST1

To install Oracle HTTP Server on WEBHOST1:

1. Verify that the servers meet the following requirements:
  - The system, patch, kernel, and other requirements meet the those specified in the *Oracle Fusion Middleware Repository Creation Utility User's Guide*.
  - This example uses port 7777. If you choose port 7777, ensure that it is not used by any service on the nodes. Verify this by running the following command:

Unix:

```
netstat -an | grep LISTEN | grep ":7777"
```

Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7777"
```

If port 7777 is in use, choose another port, or make it available.

2. On UNIX platforms, if the `/etc/oraInst.loc` or `/var/opt/oracle/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory.  
If the `/etc/oraInst.loc` file does not exist, skip this step.
3. Start Oracle Universal Installer for Oracle Fusion Middleware 11g Webtier Utilities CD installation as follows:  
For UNIX, run this command: `./runInstaller`.  
For Windows, double-click **setup.exe**.
4. In the Specify Inventory Directory screen, enter the location for the inventory and the user group, and click **OK**.
5. Execute the root privileged actions as indicated in the dialog, and click **OK**.
6. In the Welcome screen, click **Next**.
7. In the Select Installation Type screen, select **Install and Configure**. Click **Next**.
8. In the Prerequisite Checks screen, ensure that all prerequisites are met. Click **Next**.
9. In the Specify Installation Location screen, set the location to:  
`/u01/app/oracle/product/11.1.1/ohs_1`
10. Click **Next**.
11. In the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Do not select **Associate Selected Components with WebLogic Domain**.
  - Click **Next**.
12. In the Specify Component Details screen, enter the following values:
  - Instance Home Location: **app/oracle/product/11.1.1/ohs\_1/instances/ohs\_instance1**
  - Instance Name: **ohs\_instance1**
  - OHS Component Name: **ohs1**
13. Click **Next**.
14. In the **Specify Webtier Port Details** screen, do the following:
  - Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File** and then use the **Browse** function to select the file.
  - Enter Oracle HTTP Server port. For example, enter **7777**.
 Click **Next**.

---

**Note:** For more information on setting ports, refer to *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite*.

---

15. Click **Next**.
16. In the Configuration Summary screen, ensure that the selections are correct, and click **Install**.

17. In the Installation Progress screen:

For UNIX systems, a dialog box appears prompting you to run the `oracleRoot.sh` script. Open a command window and run the script, following the prompts.

Click **Next**.

18. In the Configuration screen, several configuration assistants are started in succession. When the configuration assistants are finished, the Configuration Completed screen appears.

19. In the Configuration Completed screen, click **Finish** to exit.

#### 6.4.2.1 Validating Oracle HTTP Server

To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URL a browser:

**WebHost1:7777/**

If Oracle HTTP Server is set up correctly, the **Hello World** page appears in the browser.

### 6.4.3 Installing Oracle Fusion Middleware Home

Use the following procedures to install Oracle Fusion Middleware components:

- Oracle WebLogic Server (see "[Section 6.4.3.1, \"Installing Oracle WebLogic Server\"](#)")
- WebCenter Portal (see [Section 6.4.3.2, \"Installing Oracle Fusion Middleware for WebCenter Portal\"](#))

#### 6.4.3.1 Installing Oracle WebLogic Server

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

To install Oracle WebLogic Server on all nodes in the application tier:

1. On UNIX platforms, if the `/etc/oraInst.loc` or `/var/opt/oracle/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for it.  
  
If the `/etc/oraInst.loc` file does not exist, skip this step.
2. Start Oracle WebLogic Server Installer.
3. In the Welcome screen, click **Next**.
4. In the Choose Middleware Home Directory screen:
  - Select **Create a New Middleware Home**
  - For the **Middleware Home Directory** field, enter `MW_HOME`
  - Click **Next**.
5. In the Register for Security Updates screen, enter your contact information for security update notifications, and click **Next**.
6. In the Choose Install Type screen, select **Custom**, and click **Next**.
7. In the JDK Selection screen, select only **Oracle JRockit 1.6.0\_<version> SDK** and click **Next**.

8. In the Choose Product Installation Directories screen, accept the following directory:  
`WL_HOME`  
 Click **Next**.
9. In the Installation Summary screen, click **Next**.
10. In the Installation Complete screen, deselect **Run QuickStart**, and click **Done**.

### 6.4.3.2 Installing Oracle Fusion Middleware for WebCenter Portal

To install Oracle Fusion Middleware for WebCenter Portal on all the nodes in the application tier:

1. Start the installer for Oracle Fusion Middleware for WebCenter Portal.

On UNIX, (Linux used in this example):

```
APPHOST1> runInstaller
```

On Windows:

```
APPHOST1> setup.exe
```

When Oracle Fusion Middleware 11g WebCenter Installer prompts you for a **JRE/JDK location** enter Oracle SDK location created in the Oracle WebLogic Server installation, for example, `MW_HOME/jrockit_160_<version>`.

2. In the Welcome screen, click **Next**.
3. In the Prerequisite Check screen, verify that the checks complete successfully, and click **Next**.
4. In the Specify Installation Location screen:
  - For Middleware Home, enter **MW\_HOME**
  - For Oracle Home Directory, enter the directory you want to use, for example, **wc**
 Click **Next**.
5. In the Installation Summary screen, click **Install**.
6. In the Installation Complete screen, click **Finish**.

---

**Note:** Before you run the Configuration Wizard by following the instructions in [Section 6.4.5, "Running Oracle Fusion Middleware Configuration Wizard on APPHOST1 to Create the WebLogic Server WebCenter Domain,"](#) ensure that you have applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home, so that you have the latest version of Oracle Fusion Middleware.

---

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

---

## 6.4.4 Enabling the Administration Server VIP

For information about configuring virtual IPs for the Administration Server, see [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability."](#)

## 6.4.5 Running Oracle Fusion Middleware Configuration Wizard on APPHOST1 to Create the WebLogic Server WebCenter Domain

Run Oracle Fusion Middleware Configuration Wizard from the WebCenter directory in the Middleware home to create a domain containing the Administration Server and WebCenter Portal components. Verify that the database where you installed the repository is running. For Oracle RAC databases, Oracle recommends having all the instances running.

1. Start Oracle Fusion Middleware Configuration Wizard from the `ORACLE_HOME/common/bin` directory using the following command:  

```
APPHOST1> ./config.sh
```
2. In the Welcome screen, select **Create a New WebLogic Domain** and click **Next**.
3. In the Select Domain Source screen, select **Generate a domain configured automatically to support the following products**. Select the following products:

When you select WebCenter Spaces, Oracle WSM Policy Manager is selected automatically.

- Oracle WebCenter Spaces - 11.1.1.0
- Oracle WebCenter Services Portlets - 11.1.1.0
- Oracle WebCenter Pagelet Producer - 11.1.1.0
- Oracle Enterprise manager - 11.1.1.0
- Oracle Portlet Producers - 11.1.1.0
- Oracle WebCenter Discussion Server - 11.1.1.0
- Oracle WebCenter ActivityGraph Engines - 11.1.1.0
- Oracle WebCenter Personalization - 11.1.1.0
- Oracle WebCenter Analytics Collector - 11.1.1.0
- Oracle WSM Policy Manager
- Oracle JRF - 11.1.1.1

Click **Next**.

4. Enter the **Domain Name**, **Domain Location**, and **Application Location** and click **Next**.
5. In the Configure Administrator Username and Password screen, enter the username and password to use for the domain's administrator. Click **Next**.
6. In the Configure Server Start Mode and JDK screen, select the following:
  - WebLogic Domain Startup Mode: select **Production Mode**
  - JDK Selection: select **Oracle JRockit 1.6.0\_<version>**.

Click **Next**.

7. In the Configure JDBC Data Sources screen:

- a. Select all component schemas in the table at the bottom of the screen.
- b. For the RAC configuration for component schemas, select **Convert to GridLink** or **Convert to RAC multi data source**. For single database configuration, select **Don't convert**.
- c. Ensure that the following data source appear on the screen. You provide a custom prefix for these schemas. [Table 6–7](#) lists the data sources, the schemas used and the managed servers to which they are assigned.

For information about GridLink data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC"](#); for multi data source, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

**Table 6–7 Oracle WebCenter Portal Data Sources**

| Data Source                | User Name        |
|----------------------------|------------------|
| WebCenterDS                | wcha_webcenter   |
| ActivitiesDS               | wcha_activities  |
| DiscussionsDS              | wcha_discussions |
| PersonalizationDS          | wcha_webcenter   |
| PortletDS                  | wcha_portlet     |
| Portlet-ServicesProducerDS | wcha_portlet     |
| WC-ServicesProducerDS      | wcha_webcenter   |
| WebCenterMDS               | wcha_mds         |
| PersonalizationMDS         | wcha_mds         |
| mds-PageletProducerDS      | wcha_mds         |
| mds-ServicesProducerDS     | wcha_mds         |

- d. Click **Next**.
8. In the Configure JDBC Component Schema screen, enter values for the following fields, specifying connect information for the Oracle RAC database that was seeded with RCU.

**For multi data sources**, enter values for these fields:

**Figure 6–6 Configure RAC Multi Data Source Component Schema Screen**

**Fusion Middleware Configuration Wizard**

**Configure RAC Multi Data Source Component Schema**

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for GridLink Connections; V

Service Name: wcdb.my.company.com

Username: Varies among component schemas

Password: \*\*\*\*\*

| Host Name    | Instance Name | Port |
|--------------|---------------|------|
| wcdbhost-vp1 | wcdb1         | 1521 |
| wcdbhost-vp2 | wcdb2         | 1521 |

Add Delete

| RAC Component Schema                                                  | Service Name        | Schema Owner    | Schema Password |
|-----------------------------------------------------------------------|---------------------|-----------------|-----------------|
| <input checked="" type="checkbox"/> ActivitiesDS Schema               | wcdb.my.company.com | DEV_ACTIVITIES  | *****           |
| <input checked="" type="checkbox"/> DiscussionDS Schema               | wcdb.my.company.com | DEV_DISCUSSIONS | *****           |
| <input checked="" type="checkbox"/> PersonalizationDS Schema          | wcdb.my.company.com | DEV_WEBCENTER   | *****           |
| <input checked="" type="checkbox"/> Portlet-ServicesProducerDS Schema | wcdb.my.company.com | DEV_PORTLET     | *****           |
| <input checked="" type="checkbox"/> WC-ServicesProducerDS Schema      | wcdb.my.company.com | DEV_WEBCENTER   | *****           |
| <input checked="" type="checkbox"/> WebCenterDS Schema                | wcdb.my.company.com | DEV_WEBCENTER   | *****           |
| <input checked="" type="checkbox"/> mds-PageletProducerDS Schema      | wcdb.my.company.com | DEV_MDS         | *****           |
| <input checked="" type="checkbox"/> mds-ServicesProducerDS Schema     | wcdb.my.company.com | DEV_MDS         | *****           |
| <input checked="" type="checkbox"/> WebCenter MDS Schema              | wcdb.my.company.com | DEV_MDS         | *****           |

Exit Help Previous Next

For GridLink data sources, enter values for these fields:

**Figure 6–7 Configuring GridLink RAC Component Schema**

**Fusion Middleware Configuration Wizard**

**Configure GridLink RAC Component Schema**

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for GridLink Connections; Versior

Service Name: wcdb.my.company.com

Username: Varies among component schemas

Password: \*\*\*\*\*

Enable FAN: ☒

Enable SSL: ☐

Wallet File: Enter a value Browse

Wallet Password: Enter a value

| Service Listener         | Port | Protocol |
|--------------------------|------|----------|
| wcdbhost-r-mycompany.com | 1521 | TCP      |

| ONS Host                 | Port |
|--------------------------|------|
| wcdbhost-r-mycompany.com | 6200 |

Add Delete

| RAC Component Schema                                                  | Service Name        | Schema Owner    | Schema Password |
|-----------------------------------------------------------------------|---------------------|-----------------|-----------------|
| <input checked="" type="checkbox"/> ActivitiesDS Schema               | wcdb.my.company.com | DEV_ACTIVITIES  | *****           |
| <input checked="" type="checkbox"/> DiscussionDS Schema               | wcdb.my.company.com | DEV_DISCUSSIONS | *****           |
| <input checked="" type="checkbox"/> PersonalizationDS Schema          | wcdb.my.company.com | DEV_WEBCENTER   | *****           |
| <input checked="" type="checkbox"/> Portlet-ServicesProducerDS Schema | wcdb.my.company.com | DEV_PORTLET     | *****           |
| <input checked="" type="checkbox"/> PortletDS Schema                  | wcdb.my.company.com | DEV_PORTLET     | *****           |
| <input checked="" type="checkbox"/> WC-ServicesProducerDS Schema      | wcdb.my.company.com | DEV_WEBCENTER   | *****           |
| <input checked="" type="checkbox"/> WebCenterDS Schema                | wcdb.my.company.com | DEV_WEBCENTER   | *****           |

Exit Help Previous Next

- a. Choose a driver: *For RAC:* Select **Oracle driver (Thin) for RAC Service-Instance connections, Versions:10, 11.** *For GridLink:* **Oracle Driver (Thin) for GridLink Connections, Versions:10 and later**



- b. Enter the **Service name** of the database. For a GridLink data source, you must enter the RAC service name in lower-case letters followed by the domain name us.oracle.com. For example, <mydbservice>.us.oracle.com

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example, us.oracle.com

---

- c. Enter the **Prefix\_Username** for WebCenter Portal schemas.
- d. Enter the **Password** for the schemas.
- e. **For a GridLink data source**, enter the SCAN address in the Listener Address field and SCAN port in the Port field. Enter the ONS host and port information in the ONS Host and Port fields, respectively.

---

**Note:** Oracle recommends that you use SCAN addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

---

**For a multi data source**, click **Add**. Enter details for Oracle RAC instances.

- f. Update each multi data source schema by selecting one data source at a time and adding the appropriate details.

Click **Next**.

9. In the Test JDBC Data Sources screen, the connections are tested automatically. (For GridLink data sources, ONS connections are also tested.) The **Status** column shows the results. Ensure that all connections were successful. If not, click **Previous** to return to the previous screen and correct your entries.

Click **Next** when all the connections are successful.

10. In the Select Optional Configuration screen, select the following:

- **Administration Server**
- **Managed Servers, Clusters and Machines**

11. In the Configure the Administration Server screen, enter the following values:

- Name: **AdminServer**
- Listen Address: Enter the VIP address used in [Section 6.4.4](#).
- Listen Port: **7001**
- SSL listen port: N/A
- SSL enabled: leave unchecked

Click **Next**.

12. In the Configure Managed Servers screen, add the following managed servers:

**Table 6–8 Configuring Managed Servers**

| Name              | Listen Address       | Listen Port | SSL Listen Port | SSL Enabled |
|-------------------|----------------------|-------------|-----------------|-------------|
| WC_Portlet1       | Hostname of APPHOST1 | 8889        | n/a             | unchecked   |
| WC_Portlet2       | Hostname of APPHOST2 | 8889        | n/a             | unchecked   |
| WC_Spaces1        | Hostname of APPHOST1 | 8888        | n/a             | unchecked   |
| WC_Spaces2        | Hostname of APPHOST2 | 8888        | n/a             | unchecked   |
| WC_Collaboration1 | Hostname of APPHOST1 | 8890        | n/a             | unchecked   |
| WC_Collaboration2 | Hostname of APPHOST2 | 8890        | n/a             | unchecked   |
| WC_Uilities1      | Hostname of APPHOST1 | 8891        | n/a             | unchecked   |
| WC_Uilities2      | Hostname of APPHOST2 | 8891        | n/a             | unchecked   |

Click **Next**.

13. In the Configure Clusters screen, add the following clusters

Portlet\_Cluster

- Name: **Portlet\_Cluster**
- Cluster Messaging Mode: **unicast**
- Cluster Address Enabled: **leave blank**

Spaces\_Cluster

- Name: **Spaces\_Cluster**
- Cluster Messaging Mode: **unicast**

---

**Note:** By default, each WebCenter Portal cluster is configured as unicast. To configure your WebCenter Portal cluster for multicast, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

---

- Cluster Address Enabled: **leave blank**

Collaboration\_Cluster

- Name: **Collaboration\_Cluster**
- Cluster Messaging Mode: **unicast**
- Cluster Address Enabled: **leave blank**

Utilities\_Cluster

- Name: **Utilities\_Cluster**
- Cluster Messaging Mode: **unicast**
- Cluster Address Enabled: **leave blank**

Click **Next**.

14. In the Assign Servers to Clusters screen, assign the following servers to clusters:

- **Spaces\_Cluster**

- WC\_Spaces1
- WC\_Spaces2
- Portlet\_Cluster
  - WC\_Portlet1
  - WC\_Portlet2
- Collaboration\_Cluster
  - WC\_Collaboration1
  - WC\_Collaboration2
- Utilities\_Cluster
  - WC\_Uilities1
  - WC\_Uilities2

Click **Next**.

15. In the Configure Machines screen:


- Delete the **LocalMachine** that appears by default.
- Click the **Unix Machine** tab, and add the following machines:

**Table 6–9 Configuring Machines**

| Name     | Node Manager Listen Address |
|----------|-----------------------------|
| APPHOST1 | Hostname of APPHOST1        |
| APPHOST2 | Hostname of APPHOST2        |

Click **Next**.

16. In the Assign Servers to Machines screen, assign servers to machines as follows:

**Figure 6–8 Assign Servers to Machines Screen**

- APPHOST1: WC\_Spaces1, WC\_Portlet1, WC\_Collaboration1, WC\_Uilities1
- APPHOST2: WC\_Spaces2, WC\_Portlet2, WC\_Collaboration2, WC\_Uilities2

Click **Next**.

17. In the Configuration Summary screen, click **Create**.

18. In the Creating Domain screen, click **Done**.

## 6.4.6 Creating boot.properties for the Administration Server and for Managed Servers on APPHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a `boot.properties` file for the Administration Server and for the managed servers on APPHOST1.

For the Administration Server:

1. Create the following directory:

```
APPHOST1> mkdir -p MW_HOME/user_projects/domains/wc_
domain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` in the directory created in the previous step, and enter the following lines in the file:

```
username=adminuser
password=password
```

---

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible for the entries to be encrypted.

---

---

For the WC\_Spaces1 managed server:

1. Create the following directories:

```
APPHOST1> mkdir ORACLE_BASE/product/fmw/user_projects/domains/wc_
domain/servers/WC_Spaces1/security
```

2. Use a text editor to create a file called `boot.properties` in the security directory created in the previous step, and enter the following lines in the file:

```
username=adminuser
password=password
```

3. Repeat Steps 2 and 3 for the WC\_Portlet1 and WC\_Collaboration1 Managed Servers on APPHOST1.

---

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible in order for the entries to be encrypted.

---

---

## 6.4.7 Starting the System in APPHOST1

This section describes procedures for starting the system in APPHOST1

### 6.4.7.1 Starting the Administration Server on APPHOST1

To Start the Administration Server on APPHOST1 run the following commands:

```
APPHOST1> cd ORACLE_BASE/product/fmw/user_projects/domains/wc_domain/bin
```

```
APPHOST1> ./startWebLogic.sh
```

### 6.4.7.2 Validating the Administration Server

To verify that the Administration Server is properly configured:

1. In a browser, go to `http://VIP1:7001/console`.
2. Log in as the administrator.
3. Verify that all the managed servers (WC\_Spaces1, WC\_Spaces2, and so on) are listed.
4. Verify that all Clusters are listed.
5. Verify that you can access Enterprise Manager at `http://VIP1:7001/em`.
6. Log in to EM Console with the username and password you used in step 2.

#### 6.4.7.3 Disabling Host Name Verification for the Administration Server and the Managed Servers for APPHOST1 and APPHOST2

This step is required if you have not set up SSL communication between the Administration Server and the Node Manager. If SSL is not set up, you receive an error message unless you disable host name verification.

You can re-enable host name verification when you have set up SSL communication between the Administration Server and the Node Manager.

To disable host name verification:

1. In Oracle WebLogic Server Administration Console, select **Servers**, and then **AdminServer**.
2. Select **SSL**, and then **Advanced**.
3. In the Change Center, click **Lock & Edit**.
4. When prompted, save the changes and activate them.
5. Set **Hostname Verification** to **None**.
6. Select **WC\_Spaces1**, **SSL**, and then **Advanced**.
7. Set **Hostname Verification** to **None**.
8. Repeat Steps 6 and 7 for all the Managed Servers.
9. Restart the AdminServers and all the Managed Servers.

#### 6.4.7.4 Starting Node Manager on APPHOST1

Perform these steps to start Node Manager on APPHOST1:

1. Run the `setNMProps.sh` script, which is located in the `ORACLE_COMMON_HOME/common/bin` directory, to set the `StartScriptEnabled` property to `true` before starting Node Manager:

```
APPHOST1> cd ORACLE_COMMON_HOME/common/bin
APPHOST1> ./setNMProps.sh
```

---

---

**Note:** You must use the `StartScriptEnabled` property to avoid class loading failures and other problems.

---

---

2. Start Node Manager:

```
APPHOST1> cd WL_HOME/server/bin
APPHOST1> ./startNodeManager.sh
```

### 6.4.8 Install WebLogic Server and WebCenter Portal on APPHOST2

Repeat the procedures for installing WebLogic Server and WebCenter Portal for APPHOST2, start with [Section 6.4.3.1, "Installing Oracle WebLogic Server"](#). The directory paths for binary files and domains used when installing new nodes must be exactly the same as those used for first node. If these paths and domains are not exactly the same as those used for the first node, failover is does not occur.

## 6.4.9 Propagating the Domain Configuration to APPHOST2 with pack/unpack Utilities

Follow these steps to propagate the domain configuration to APPHOST2 using Pack/Unpack utilities:

1. Run the following pack command on APPHOST1 to create a template pack:

```
APPHOST1> cd WL_HOME/common/bin
APPHOST1> ./pack.sh -managed=true -domain=ORACLE_BASE/product/fmw/user_
projects/domains/wc_domain/
-template=wc_domaintemplate.jar
-template_name=wc_domain_template
```

2. Run the following command on APPHOST1 to copy the template file created in the previous step to APPHOST2 using, in this example, scp:

```
APPHOST1> scp wc_domaintemplate.jar
APPHOST2:WL_HOME/common/bin
```

3. Run the unpack command on APPHOST2 to unpack the propagated template:

```
APPHOST2> cd WL_HOME/common/bin
APPHOST2> ./unpack.sh
-domain=ORACLE_BASE/product/fmw/user_projects/domains/wc_domain/
-template=wc_domaintemplate.jar
```

## 6.4.10 Starting Node Manager on APPHOST2

To start the Node Manager on APPHOST2, repeat the steps from [Section 6.4.7.4, "Starting Node Manager on APPHOST1."](#) on APPHOST2.

## 6.4.11 Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers

Enable Oracle HTTP Server to route to the Administration Server that contains WebCenter Portal managed servers by setting the WebLogicCluster parameter to the list of nodes in the cluster.

1. Add the following lines to the OHS\_HOME/instances/ohs\_instance1/config/OHS/ohs1/mod\_wl\_ohs.conf file:

```
Spaces
<Location /webcenter>
 WebLogicCluster apphost1.com:8888,apphost2.com:8888
 SetHandler weblogic-handler
</Location>

<Location /webcenterhelp>
 WebLogicCluster apphost1.com:8888,apphost2.com:8888
 SetHandler weblogic-handler
</Location>

<Location /rss>
 WebLogicCluster apphost1.com:8888,apphost2.com:8888
 SetHandler weblogic-handler
</Location>

<Location /rest>
 WebLogicCluster apphost1.com:8888,apphost2.com:8888
 SetHandler weblogic-handler
```

```
</Location>

Pagelet Producer
<Location /pagelets>
 WebLogicCluster apphost1.com:8889,apphost2.com:8889
 SetHandler weblogic-handler
</Location>

Activity Graph
The WebLogicHost below should be set to the Host on which ActivityGraph is
running
<Location /activitygraph-engines>
 WebLogicCluster apphost1.com:8891
 SetHandler weblogic-handler
</Location>

Portlet
<Location /portalTools>
 WebLogicCluster apphost1.com:8889,apphost2.com:8889
 SetHandler weblogic-handler
</Location>

<Location /wsrp-tools>
 WebLogicCluster apphost1.com:8889,apphost2.com:8889
 SetHandler weblogic-handler
</Location>

SES Search
<Location /rsscrawl>
 SetHandler weblogic-handler
 WeblogicHost ses.example.com
 WeblogicPort 7777
</Location>
<Location /sesUserAuth>
 SetHandler weblogic-handler
 WeblogicHost ses.example.com
 WeblogicPort 7777
</Location>

Personalization
<Location /wcps>
 SetHandler weblogic-handle
 WeblogicHost webcenter.example.com
 WeblogicPort 8889
</Location>

Discussions
<Location /owc_discussions>
 WebLogicCluster apphost1.com:8890,apphost2.com:8890
 SetHandler weblogic-handler
</Location>

#AdminServer and EM
<Location /console>
 SetHandler weblogic-handler
 WebLogicHost VIP1
 WeblogicPort 7001
</Location>

<Location /consolehelp>
```



```

 SetHandler weblogic-handler
 WebLogicHost VIP1
 WeblogicPort 7001
 </Location>

 <Location /em>
 SetHandler weblogic-handler
 WebLogicHost VIP1
 WeblogicPort 7001
 </Location>

```

## 2. Restart Oracle HTTP Server on WEBHOST1:

```

WEBHOST1> OHS_HOME/instances/ohs_instance1/bin/opmnctl restartproc
ias-component=OHS_COMPONENT1

```

### 6.4.11.1 Configuring a Virtual Host for Oracle Pagelet Producer and Sharepoint

The WebCenter suite includes two applications that use / as the web context root:

- Sharepoint Root - deployed on WC\_Spaces
- Pagelet Producer Root - deployed on WC\_Portlet

**6.4.11.1.1 Virtual Hosts Requirement** To route the Sharepoint Root application without virtual hosts by using Oracle HTTP Server (OHS), add the following entry:

```

<Location />
 SetHandler weblogic-handler
 WebLogicHost webcenter.example.com
 WebLogicPort 8889
</Location>

```

This entry grabs all context roots that are not defined explicitly. Sharepoint requires mapping, which isn't possible to do in a single OHS. For this reason, you need a *virtual host configuration*, a single system that runs more than one web site such as `www.company1.com` and `www.company2.com`. Virtual hosts are *IP-based* (a unique IP address for each web site) or *name-based* (multiple names run on each IP address). It is not obvious to end users that the web sites run on the same physical server. For more information on virtual hosts, see Apache HTTP Server documentation.

### Oracle HTTP Server Configuration

The following entry adds two name-based hosts: one for Sharepoint Root and one for Pagelet Producer Root.

1. Find OHS `httpd.conf` in `$WEBTIER_INSTANCE_HOME/config/OHS/<ohs>`.
2. Find sample `VirtualHost` configuration and add the following entry:

```

NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName webhost.example.com
</VirtualHost>

<VirtualHost *:7777>
 ServerName webtier-spaces.example.com
 <Location />
 SetHandler weblogic-handler
 WebLogicCluster apphost1:8888,apphost2:8888
 </Location>
 <Location /webcenter>

```

```
 Deny from all
 </Location>
 <Location /webcenterhelp>
 Deny from all
 </Location>
 <Location /rest>
 Deny from all
 </Location>
</VirtualHost>
```

3. Ensure that the new virtual hosts `webtier-pageletproducer` and `webtier-spaces` have been added to DNS.

**6.4.11.1.2 Additional Configuration** In a Virtual Host setup, you must configure additional properties to use applications routed via virtual host.

### Sharepoint

For single sign-on setups, including integration with Oracle Access Manager 10g or Oracle Access Manager 11g, see the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal*.

### 6.4.11.2 Validating Access through Oracle HTTP Server

Verify the URLs to ensure that appropriate routing and failover is working from the HTTP Server to the WebCenter Portal cluster.

1. Start `WC_Spaces1`, `WC_Spaces2`, `WC_Portlet1` and `WC_Portlet2` from the WebLogic Server Administration Console as follows:
  - a. Access the Administration Console at the following URL  
`http://APHHOST1/console`
  - b. Click **Servers**.
  - c. Open the **Control** tab.
  - d. Select `WC_Spaces1`, `WC_Spaces2`, `WC_Portlet1` and `WC_Portlet2`.
  - e. Click **Start**.
  - f. Verify direct access to the managed servers using the following URLs:  
`apphost1:8888/webcenter`  
`apphost2:8888/webcenter`  
`apphost1:8889/portalTools`  
`apphost2:8889/portalTools`
2. While `WC_Spaces2` and `WC_Portlet2` are running, stop `WC_Spaces1` and `WC_Portlet1` from Oracle WebLogic Server Administration Console.
3. Access the following URLs and verify the appropriate functionality:
  - `WebHost1:7777/webcenter`
  - `WebHost1:7777/portalTools`
4. Start `WC_Spaces1` and `WC_Portlet1` from the WebLogic Server Administration Console.
5. Stop `WC_Spaces2` and `WC_Portlet2`.
6. Access the following URLs and verify the appropriate functionality:

- WebHost1:7777/webcenter
- WebHost1:7777/portalTools

## 6.4.12 Configuring Manual Failover of the Administration Server to APPHOST2

For information about configuring the Administration Server for high availability, see [Section 12.4, "Transforming the Administration Server in an Existing Domain for Cold Failover Cluster."](#)

## 6.4.13 Configuring the Java Object Cache

The Java Object Cache (JOC) should be configured among all the servers running WebCenter Spaces. This local cache is provided to increase the performance of the Spaces application.

For instructions on configuring the Java Object Cache among all the servers running the Spaces application, see [Section F.1, "Configuring the Java Object Cache."](#)

## 6.4.14 Configuring the Distributed Notifications for the MDS Repository

In high availability environments, Oracle recommends that you configure distributed notifications for the MDS repository.

For instructions on configuring distributed notifications for the MDS repository, see [Appendix G, "Configuring Distributed Notifications for MDS."](#)

## 6.4.15 Configuring WebCenter Portal for Replication

Use the procedures in this section to configure WebCenter Portal for replication.

### Clustering Requirement

The application must be deployed to an Oracle WebLogic Cluster. This automatically establishes a replication channel for the multiple instances of the application.

---

**Note:** In a Unicast cluster, the default replication channel is configured using the Listen address of each managed server. Therefore, the Listen address should be configured to be a specific IP address or host name, instead of being configured to listen on Any.

---

### Oracle ADF Replication

WebCenter Portal relies on Oracle ADF components, therefore, is it essential that you configure Oracle ADF properly. The following tag should be present in the `adf-config.xml` file, one of the Application Resources, for a stateful application:

```
<adfc:adf-controller-config>
<adfc:adf-scope-ha-support>true</adfc:adf-scope-ha-support>
</adfc:adf-controller-config>
```

In WebCenter Portal applications, this is enabled by default.

### Application Replication

Applications must also have replication enabled. Oracle WebLogic Server allows several types of persistent stores for replication. For more information on persistent stores, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

WebCenter Portal applications are enabled by default with the following setting in the `weblogic.xml` file:

```
<session-descriptor>
<persistent-store-type>replicated_if_clustered</persistent-store-type>
</session-descriptor>
```

The `replicated_if_clustered` setting disables replication for stand-alone application environments, and uses in-memory replication within a cluster environment.

Ensure that all WebCenter Portal applications are configured for in-memory replication.

## 6.4.16 Configuring Analytics

Analytics Collectors are configured out-of-the-box so all you must do is configure a connection between the Spaces application and the Analytics Collector.

1. Open the WLST shell:

```
ORACLE_HOME/common/bin/wlst.sh
```

2. Connect to WLS Server:

```
connect('weblogic_admin_username', 'weblogic_admin_pwd', 'APPHOST1:8888')
```

Note that you are connecting to the host and port of the Spaces Server.

3. Create the Analytics Collector connection and make it the default connection:

```
createAnalyticsCollectorConnection(appName='webcenter',
connectionName='HAConn1', isUnicast=1,
collectorHost='localhost', collectorPort=31314, isEnabled=1, timeout=30,
default=1)
```

4. List the changes made:

```
listDefaultAnalyticsCollectorConnection(appName='webcenter')
```

## 6.4.17 Configuring Activity Graph

The Activity Graph Engines application should run as a singleton. In a cluster environment, all instances of the Activity Graph Engines application should be disabled except for one.

To disable the Activity Graph Engines application:

1. Log on to the Oracle WebLogic Administration Console.
2. Shut down WC\_Spaces, WC\_Portlets, and WC\_Uutilities. Select **Deployments**.
3. In the Change Center, click **Lock & Edit**.
4. Alter the targets for each of the following deployments:
  - `activitygraph-engines` (11.1.1.6.0)
  - `oracle.webcenter.activitygraph.engine1lib` (11.1.1,11.1.1)
  - `oracle.webcenter.activitygraph.lib` (11.1.1,11.1.1)
    - a. Select the deployment.
    - b. Select the Targets tab. Click on **Change Targets**.

- c. Ensure that the deployment is only targeted to **Part of the Cluster/one of the Managed Servers**. Click **OK** to save the changes.
- d. Click **Activate All Changes**.

Since Activity Graph runs on one node only, it is unavailable if this node is lost or the Managed Server is not available. In these cases, Oracle recommends that you deploy Activity Graph on an active Managed Server. You can automate this process by configuring Service Migration. For an example, see [Section 5.13.20, "Configuring Server Migration for the WLS\\_SOA Servers."](#)

## 6.4.18 Configuring Clustering for Discussions Server

This procedure requires a unicast cluster.

Before you perform this procedure, see [Section 6.4.21, "Converting Discussions from Multicast to Unicast"](#):

- To convert a multicast cluster to a unicast cluster (if necessary).
- To verify that an existing unicast cluster has the correct settings.

Ensure that all members of the Discussions server cluster can communicate with each other using the Discussions server Administration Console:

1. Log into each member of the cluster at:

`http://<host>:<port>/owc_discussions/admin`

2. Go to **Cache Settings**.

3. At the bottom of the page, in the **Cache Features** section, ensure that **Clustering** is set to **Enabled**.

The top of the page should now list all members of the cluster.

4. Again, towards the end of the page, under the **Cache Tools** section, do **Cluster wide cache reset and the Cache warm up Task**. Repeat the Cache warm up task on all members of the cluster.

## 6.4.19 Scaling the Topology

You can scale out and scale up a WebCenter Portal topology. When you scale up the topology, you add new managed servers to nodes that are already running one or more managed servers. When you scale out the topology, you add new managed servers to new nodes.

### 6.4.19.1 Scaling Up the Topology (Adding Managed Servers to Existing Nodes)

In this case, you already have a node that runs a managed server configured with WebCenter portal components. The node contains a Middleware home and a WebCenter directory in shared storage.

You can use the existing installations (Middleware home and domain directories) for creating new WebCenter Portal managed servers. There is no need to install WebCenter Portal binaries in a new location, or to run pack and unpack. Running multiple managed servers on one node is only supported for the WC\_Spaces servers and the WC\_Portlet servers.

Follow these steps for scaling up the topology:

1. Using the Administration Console, clone WC\_Spaces1 or WC\_Portlet1 into a new managed server. The source managed server to clone should be one that already exists on the node where you want to run the new managed server.

To clone a managed server:

- a. Select **Environment** -> **Servers** from the Administration Console.
- b. Select the managed server that you want to clone (for example, WC\_Spaces1 or WC\_Portlet1).
- c. Select **Clone**.

Name the new managed server *SERVER\_NAME*<sub>n</sub>, where n is a number to identify the new managed server.

2. For the listen address, assign the host name or IP to use for this new managed server.

Ensure the port number for this managed server is available on this node.

3. Add the new managed server to the Java Object Cache Cluster (see [Section 6.4.13, "Configuring the Java Object Cache"](#)).
4. Reconfigure the Oracle HTTP Server module with the new member in the cluster. See [Section 6.4.11, "Configuring Oracle HTTP Server for the Administration Server and WebCenter Portal Managed Servers."](#)

#### 6.4.19.2 Scaling Out the Topology (Adding Managed Servers to New Nodes)

In scaling out your topology, you add new managed servers configured with WebCenter Portal applications to new nodes.

Before performing the steps in this section, check that you meet these requirements:

- In your topology, there are existing nodes running managed servers configured with WebCenter Portal applications.
- The new node can access the existing home directories for WebLogic Server and WebCenter Portal. You use the existing installations in shared storage for creating a new managed server. There is no need to install WebLogic Server or WebCenter Portal binaries in a new location, although you must run pack and unpack to create a managed server domain.

Follow these steps for scaling out the topology:

1. On the new node, mount the existing Middleware home, which should include the WebCenter Portal installation and the domain directory, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To attach *ORACLE\_HOME* in shared storage to the local Oracle Inventory, execute the following commands:

```
WCHOSTn> cd ORACLE_BASE/product/fmw/wc/
WCHOSTn> ./attachHome.sh -jreLoc ORACLE_BASE/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the *MW\_HOME/boa/beahomelist* file and add *ORACLE\_BASE/product/fmw* to it.

3. Log into the Oracle WebLogic Administration Console.
4. Create a new machine for the new node that will be used, and add the machine to the domain.

5. Update the machine's Node Manager's address to map the IP of the node that is being used for scale out.
6. Use the Oracle WebLogic Server Administration Console to clone either WC\_Spaces1 or WC\_Portlet1 or WC\_Collaboration1 or WC\_Uutilities1 into a new managed server. Name it WC\_XXXn, where n is a number and assign it to the new machine.
7. For the listen address, assign the host name or IP to use for the new managed server. Perform these steps to set the managed server listen address:
  - a. Log into the Oracle WebLogic Server Administration Console.
  - b. In the Change Center, click **Lock & Edit**.
  - c. Expand the **Environment** node in the Domain Structure window.
  - d. Click **Servers**. The Summary of Servers page is displayed.
  - e. Select the managed server whose listen address you want to update in the **Names** column of the table. The Setting page for that managed server is displayed.
  - f. Set the Listen Address to WCHOSTn where WCHOSTn is the DNS name of your new machine.
  - g. Click **Save**.
  - h. Save and activate the changes.

The changes will not take effect until the managed server is restarted.

8. Run the pack command on WCHOST1 to create a template pack as follows:

```
WCHOST1> cd ORACLE_COMMON_HOME/common/bin
WCHOST1> ./pack.sh -managed=true -domain=MW_HOME/user_projects/
domains/wc_domain/ -template=webcenterdomaintemplateScale.jar
-template_name=webcenter_domain_templateScale
```

Run the following command on WCHOST1 to copy the template file created to WCHOSTn:

```
WCHOST1> scp wc_domaintemplateScale.jar oracle@WCHOSTn:/ ORACLE_BASE/product/
fmw/wc/common/bin
```

Run the unpack command on WCHOSTn to unpack the template in the managed server domain directory as follows:

```
WCHOSTn> cd ORACLE_BASE/product/fmw/wc/common/bin
WCHOSTn> ./unpack.sh -domain=ORACLE_BASE/product/fmw/user_projects/domains/
wc_domain/ -template=wc_domaintemplateScale.jar
```

9. Start the Node Manager on the new node. To start the Node Manager, use the installation in shared storage from the existing nodes, and start Node Manager by passing the host name of the new node as a parameter as follows:

```
WCHOSTn> WL_HOME/server/bin/startNodeManager <new_node_ip>
```

10. If this is a new Collaboration managed server, ensure that you have followed the steps in [Section 6.4.18, "Configuring Clustering for Discussions Server"](#) to configure clustering for the new Discussions server. Ensure also that the steps in [Section 6.4.21, "Converting Discussions from Multicast to Unicast"](#) are performed, using the hostname of the new host for the coherence.localhost parameter.

11. If this is a new Utilities managed server, ensure that Activity Graph is disabled by following the steps in [Section 6.4.17, "Configuring Activity Graph."](#) Ensure also that the steps for configuring a new Analytics Collector in [Section 6.4.16, "Configuring Analytics"](#) have been followed.
12. Start and test the new managed server from the Oracle WebLogic Server Administration Console:
  - a. Shut down all the existing managed servers in the cluster.
  - b. Ensure that the newly created managed server is running.
  - c. Access the application on the newly created managed server. The application should be functional.

## 6.4.20 Troubleshooting WebCenter Portal High Availability

This section describes procedures for troubleshooting WebCenter Portal.

### 6.4.20.1 Troubleshooting WebCenter Portal Deployment Issues

WebCenter Portal applications deploy when the managed server first starts; use Oracle WebLogic Server Administration Console first to check that all application deployments are successful:

Click **Deployments** in the left hand pane. The right hand pane shows the application deployments and their status. The state of all applications, assuming all the servers are running, should be ACTIVE.

If an application deployment fails, the server logs may provide some indication of why the application was not deployed successfully. The server logs are located in the `DOMAIN_HOME/servers/SERVER_NAME/logs` directory. Common issues include:

- Unavailability of external resources, such as database resources. Examine the error, fix it, and attempt to redeploy the application.
- The appropriate applications or libraries are not targeted correctly to the right managed server or cluster.

### 6.4.20.2 Troubleshooting WebCenter Portal Replication and Failover Issues

State replication is most prominent in failover scenarios. A user working on one server may discover that, upon failover:

- Windows may be closed or state might be reset.
- Screens may require a reset.
- The application may be redirecting to the logon screen.

To troubleshoot and diagnose state replication issues.

1. Confirm that this is not a known replication issue.

See [Section 6.3.2.8, "Expected Behavior for Application Failover"](#) for possible expected behaviors. Before proceeding to further diagnose the issue, first confirm that the failover behavior is not an expected behavior.

2. Check load balancer settings.

For replication and failover to function correctly, the load balancer must be configured with the appropriate persistence settings. For more details on configuring Hardware Load Balancers for Oracle WebLogic Server, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.



### 3. Check the cluster status.

Replication occurs within the context of a cluster. For failover to be successful, there must be at least one other healthy member of the cluster available. You can check cluster status in one of two ways:

- Check the cluster status using Oracle WebLogic Server Administration Console - In the Left-hand pane, click on **Servers**. Verify the state of all servers in the cluster.
- Check the cluster status using weblogic.Admin utility The weblogic.Admin command can be used to query the state of all servers in a specific cluster. For example:

```
$ java weblogic.Admin -url Adminhost:7001 -username <username> -password
<password> CLUSTERSTATE -clustername Spaces_Cluster
```

This example returns:

```
There are 2 server(s) in cluster: Spaces_Cluster
The alive servers and their respective states are listed below:
WC_Spaces1---RUNNING
WC_Spaces2---RUNNING
```

### 4. Check cluster communications.

Although cluster members may all be running, there may be communication issues which prevent them from communicating replication information to each other. There are two types of cluster communication configurations.

Troubleshooting depends on the cluster type:

- Checking Unicast cluster communications - For Unicast clusters, managed servers must be able to access each other's hosts and each other's default listening port.  
  
Ensure that all individual managed servers have their Listen Address set correctly. You can find this setting by selecting **Configuration, General** for each managed server.
- Checking Multicast cluster communications - For multicast clusters, servers must be able to intercept the same multicast traffic. Ensure that multicast is configured correctly by running the WebLogic utility `utils.MulticastTest` on each machine. For example:

```
$ java utils.MulticastTest -H
```

### 5. Confirm application configuration.

WebCenter Portal applications are configured correctly by default. For user applications, in-memory replication take place only with the proper configuration in `weblogic.xml`:

```
<session-descriptor>
<persistent-store-type>replicated_if_clustered</persistent-store-type>
</session-descriptor>
```

A persistent-store-type of replicated is also acceptable.

### 6. Enable Debug.

Check the server logs for any unusual messages on managed server startup. In particular, if the managed server is unable to locate other members of the cluster.

The server logs are located in the *DOMAIN\_HOME/servers/SERVER\_NAME/logs* directory.

For further debugging, enable the flags `DebugCluster`, `DebugClusterAnnouncements`, `DebugFailOver`, `DebugReplication`, and `DebugReplicationDetails`. Each flag can be enabled with the `weblogic.Admin` utility:

```
$ java weblogic.Admin -url Adminhost:7001 -username <username> -password
<password> SET -type ServerDebug -property DebugCluster true
```

### 6.4.20.3 Troubleshooting Lost Changes to Policies

Policies are refreshed at an interval defined in `jps-config.xml` by the variable `oracle.security.jps.ldap.policystore.refresh.interval`. By default, this interval is ten minutes. If a server is lost, requests are routed to another server in the cluster where the original policies are cached. Recent policy changes that occurred since the last refresh may appear to be lost. For example, roles created immediately before the failover may be seen as unavailable.

The policies are still in the back-end policy store. Refresh the cache to restore policy information. Re-try after a few minutes, or log in and log out of the application to force a policy cache refresh.

### 6.4.20.4 Troubleshooting JOC Configuration

If attempts to start a Managed Server that uses the Java Object Cache, such as OWSM or the Spaces application Managed Servers, fail, and the following errors appear in the logs:

```
J2EE JOC-058 distributed cache initialization failure
J2EE JOC-043 base exception:
J2EE JOC-803 unexpected EOF during read.
```

Another process is using the same port that JOC is attempting to obtain. To solve this problem, either stop that process, or reconfigure JOC for this cluster to use another port in the recommended port range.

## 6.4.21 Converting Discussions from Multicast to Unicast

To convert Discussions from multicast to unicast:

### Step 1: Add the startup parameters

To add the relevant startup parameters:

1. In the Oracle WebLogic Server Administration Console, select **Servers**, **WC\_Collaboration1**, **Configuration**, and then **Server Start**.
2. In the **Arguments** box, add the following:  

```
-Dtangosol.coherence.wka1=WCPHOST1
-Dtangosol.coherence.wka2=WCPHOST2
-Dtangosol.coherence.localhost=WCPHOST1
-Dtangosol.coherence.wka1.port=8089
-Dtangosol.coherence.wka2.port=8089
-Dtangosol.coherence.localport=8089
```

Where **Host1** is where **WC\_Collaboration1** is running.

To use a port other than 8088 for WebCenter Coherence communications, specify a different port number using the `-Dtangosol.coherence.wka1.port` and `-Dtangosol.coherence.wka2.port` arguments in the example above.

3. Repeat steps 1 and 2 for WC\_Collaboration2, setting the localhost parameter to Host2 instead.
4. Restart the WC\_Collaboration servers.

### Step 2: Validate the changes

To validate the changes:

1. Log in to the Discussions Forum Administration Console at `http://<host>:<port>/owc_discussions/admin`.
2. Select **Cache Settings** in the left pane.
3. At the bottom of the screen, ensure that **Clustering** is set to **Enabled**.
4. Repeat steps 1 through 3 for all members of the cluster.

As servers join the cluster they appear at the top of the screen.

## 6.5 Configuring High Availability for WebCenter Portal Applications

WebCenter Portal Applications are deployed on servers and clusters that are provisioned with the necessary libraries. This provisioning occurs by applying a domain template to an existing WebCenter Portal domain.

The extension can be done only once per domain. This means that any WebCenter Portal application servers must be either created at one time or as clones of existing WebCenter Portal application servers.

### 6.5.1 Configuring a Cluster for WebCenter Portal Applications

To configure a cluster to deploy WebCenter Portal applications:

1. Start Oracle Fusion Middleware Configuration Wizard from the `ORACLE_HOME/common/bin` directory using the following command:  

```
APPHOST1> ./config.sh
```
2. In the Welcome screen, select **Extend an Existing WebLogic Domain**, and click **Next**.
3. In the next screen, select your Weblogic Domain Directory.
4. In the next screen, choose **Extend my Domain using an existing Extension template**.
5. Choose either of these two templates, depending on whether you are creating your own Framework applications or Portlet Producer applications:

- For Framework applications:

```
$WC_ORACLE_HOME/common/templates/applications/oracle.wc_custom_portal_
template_11.1.1.jar
```

- For Portlet Producer applications:

```
$WC_ORACLE_HOME/common/templates/applications/oracle.wc_custom_services_
producer_template_11.1.1.jar
```

6. Configure the new data sources (Custom Portal Activities, Custom Portal WebCenter, Custom Portal MDS) as Oracle RAC by selecting all the data sources and then checking the RAC datasources checkbox. The applications use existing schemas.
7. Enter the Database connection information and proceed to check the status of the data sources.
8. In the Options Configuration Screen, select **Managed Servers, Clusters and Machines**.
9. In the Managed Servers screen:
  - If in step 5. you selected the custom\_portal template, rename the managed server WC\_CustomPortal to WC\_CustomPortal\_Cluster1. If you selected the custom\_services\_producer template, rename the managed server WC\_CustomServicesProducer to WC\_CustomServicesProducer\_Cluster1
  - Click **Add** to create a new managed server and name it WC\_CustomPortal\_Cluster2. Set the Listen Port to be the same as that of WC\_CustomPortal\_Cluster1.
10. In the Configure Clusters screen, add the cluster WC\_CustomServicesProducer\_Cluster1.
11. In the Assign Servers to Clusters screen, add the two newly created servers to the Cluster.
12. In the Assign Server to Machines screen, place each of the two servers on separate machines.
13. Click **Extend** and then **Done** to Extend the Domain.

When you complete these steps, you can start the managed server using Oracle WebLogic Server Administration Console and then deploy WebCenter Portal applications. Deploy these WebCenter Portal applications to the cluster target, not the managed server target.

## 6.5.2 Adding More WebCenter Portal Applications Servers

New servers can only be added to a domain that has already been extended by cloning an existing server:

1. Access the Oracle WebLogic Server Administration Console at <http://APPHOST1:7001/console> and log in.
2. In the Domain Structure window, choose **Environment > Servers** to show all the existing managed servers on the Summary of Servers page.
3. Select one of the custom managed servers. Click **Clone**.
4. Enter a server name (for example, WC\_CustomPortal3). Set the Listen Address and Port where this new server resides.

The new server automatically becomes a member of the cluster (for example, WC\_CustomPortalCluster).

## 6.5.3 Configuring Distributed Notifications for the MDS Repository

For high availability WebCenter Portal applications, Oracle recommends that you configure distributed notifications for the MDS repository. For instructions on

configuring distributed notifications for the MDS repository, see [Appendix G](#), "Configuring Distributed Notifications for MDS."



---

# High Availability for Oracle Data Integrator

Oracle Data Integrator provides a complete set of components for designing, deploying and managing data integration processes. Data integration processes move and transform data from source data servers to target data servers, using an Extract-Load-Transform approach that eliminates the need for a transformation engine by delegating all transformations to the source and target data servers.

This chapter provides a description of Oracle Data Integrator components from a high availability perspective. The sections in this chapter outline the single instance concepts that are important for designing high availability deployment.

This chapter includes the following topics:

- [Section 7.1, "Introduction to Oracle Data Integrator"](#)
- [Section 7.2, "Oracle Data Integrator Single Instance Characteristics"](#)
- [Section 7.3, "Oracle Data Integrator High Availability and Failover Considerations"](#)
- [Section 7.4, "Configuring High Availability for Oracle Data Integrator"](#)

## 7.1 Introduction to Oracle Data Integrator


As shown in [Figure 7-1](#), Oracle Data Integrator provides a comprehensive set of features for developing, deploying and monitoring data integration processes. Oracle Data Integrator provides a unified run-time architecture based on Java components.

The run-time architecture consists of:

- **Repository:** The repository stores the data integration design-time objects, the run-time objects (data integration scenarios to execute) and the sessions corresponding to execution instances of these design-time and run-time objects.
- **Run-time agents:** These are the standalone agent or Java EE agents deployed in an application server, which execute the sessions. The agent connects the repository as well as the source and target data servers when processing the data integration sessions. The agent also provides a web service for starting and monitoring scenario executions from third party SOA-enabled applications.
- **Oracle Data Integrator Console:** This web application allows users to browse, monitor and manage the artifacts stored in the repository. It also performs run-time operations, such as starting a scenario on run-time agents.

## 7.2 Oracle Data Integrator Single Instance Characteristics

[Figure 7-1](#) shows the Oracle Data Integrator single instance architecture.

**Figure 7-1 Oracle Data Integrator Single Instance Architecture**

Oracle Data Integrator run-time agents manage integration processes. Oracle Data Integrator agents are components that run the integration jobs deployed in a production configuration as scenarios stored in a repository.

Oracle Data Integrator agents process each scenario execution instance as a session. Each session exists in the agent as a separate thread of the agent Java process.

Agents store very basic information about the session they run. Most of the session data is stored in the repository. When a scenario is executed on an agent, the agent creates a session in the repository that corresponds to this scenario's instance. The agent reads each task of this session from the repository, processes it, and writes the result - the return code, message and tasks metrics such as the duration or number of rows processed - into the repository.

The repository consists of two database schemas, one containing the master repository, and one containing the work repository. The master repository contains all topology and security related information (such as the source data server definition, target data server definition, and user credential). The work repository contains development and run-time data (such as sessions and scenarios). The master repository also contains the connection information to the work repository. To connect to a work repository, an agent first connects to the master repository, checks the Oracle Data Integrator user's credentials, reads the work repository connection information, and then connects to the work repository. A typical topology includes one master repository and possibly several work repositories (for example, for test and production).

Sessions can be initiated on the agent:

- From another Oracle Data Integrator component (such as the agent or Oracle Data Integrator Studio) over HTTP.
- Via the agent's web service interface.
- From an external scheduler or from a command line.
- From a Java program using the Agent Invocation SDK

The agent is always attached to a master repository. It connects to this master repository at startup and is able to start sessions on any of the work repositories attached to this master. It also acts as a scheduler. On startup, the agent reads from the different work repositories the schedules defined for the agent, and stores this scheduling information. The agent is able to initiate sessions from this in-memory schedule on the appropriate work repositories.

Agents can interact with one another through remote scenario startup (over HTTP) or via the load balancing feature. Load balancing allows defining hierarchies of



parent/child agents. In this hierarchy, parent agents can delegate the processing of their sessions to their child agents.

The agent is a Java program that is provided as a Java EE agent and as a standalone agent. The Java EE agent is a web application that can be deployed in a Java EE application server, along with other web applications within the same JVM. This agent can use this server's data sources to connect the source, target and repository databases.

The standalone agent is provided as a standalone Java process started from a command line interface. This standalone agent is similar to the Java EE agent, but is embedded in a lightweight container. The main difference is that unlike the Java EE agent, the standalone agent can connect the source and target data servers using only direct JDBC connection.

## 7.2.1 Oracle Data Integrator Sessions Lifecycle and Recovery

When an execution request arrives to a run-time agent, the agent connects the master repository to check the user credentials and then the work repository to create the session and all its tasks, and marks them as "waiting." Then it creates the connections to all the data servers that will be used during this session.

When execution starts, the agent reads the first task in the work repository to be executed, and marks both the session and this task as "running." This task can start an operation on the data servers or on the operating system. When the task is complete, the agent writes into the work repository the execution result for this task, moves it to a finished state ("Done", "Warning" or "Error") and proceeds to the next task in the session. Note that errors cases can be handled in the ODI packages, and an error does not necessarily halt a session. When the session completes (either because of an unmanaged error, or by reaching a final step), the agent moves the session to a finished state ("Done", "Warning" or "Error") and releases all the connections. At this point, the session is finished.

### 7.2.1.1 Sessions Interruption

Sessions can be interrupted when:

- A user requests the agent to stop the session.
- An agent is stopped by the administrator. All sessions for this agent are stopped, depending on the agent stop mode selected.
- A critical event occurs on the agent or the repository.

Any session that is stopped due to user or administrator action is moved to an error state and marked as "Stopped."

In the case of an agent or repository crash, a session that cannot be stopped properly still appears in a running state in the repository. These sessions are called stale sessions, because they are marked as running, but are no longer handled by any agent. Stale sessions are automatically moved to an error state when an agent restarts and detects that these sessions are incorrectly marked in the repository as being executed by this agent.

### 7.2.1.2 Recovering Sessions

Oracle Data Integrator uses JDBC transactions when interacting with source and target data servers, and any open transaction state is not persisted when a session finishes in error state. The appropriate restart point is the task that started the unfinished transaction(s). If such a restart point is not identifiable, it is recommended that you

start a fresh session by executing the scenario instead of restarting existing sessions that are in error state.

By default, a session restarts from the last task that failed to execute (typically a task in error or in waiting state). A session may need to be restarted in order to proceed with existing staging tables and avoid re-running long loading phases. In that case the user should take into consideration transaction management, which is KM specific. A general guideline is: If a crash occurs during a loading task, you can restart from the loading task that failed. If a crash occurs during an integration phase, restart from the first integration task, because integration into the target is within a transaction. This guideline applies only to one interface at a time. If several interfaces are chained and only the last one performs the commit, then they should all be restarted because the transaction runs over several interfaces.

To restart from a specific task or step:


1. In Operator Navigator, navigate to this task or step, edit it and switch it to Waiting state.
2. Set all tasks and steps after this one in the Operator tree view to Waiting state.
3. Right-click the session and click **Restart**.

The session restarts from the first task in waiting state.

## 7.2.2 Agent Startup and Shutdown Cycle

Figure 7–2 shows the agent startup cycle.

**Figure 7–2 Oracle Data Integrator Agent Startup Cycle**



When the Oracle Data Integrator agent starts, it first reads its configuration, which includes master repository connection information. Then the agent connects to each of the work repositories attached to this master repository and removes stale sessions. Stale sessions are sessions that are incorrectly indicated in the work repository as running on this given agent. Stale sessions may result from an agent being stopped without being able to stop these sessions gracefully. As the agent restarts, it identifies the stale sessions and moves them to an error state.

From that point, the agent can retrieve and compute the schedules available for it in each work repository. Once this phase is complete, the agent waits for incoming

sessions requests to start their processing, and is also able to start sessions based on its schedules.

### 7.2.3 Oracle Data Integrator External Dependencies

Oracle Data Integrator depends on the Oracle Data Integrator master repository and work repository database schemas.

If advanced features are being used, these other dependencies may exist:

- Other Oracle Data Integrator agents: If the load balancing feature is configured and the agent needs to delegate the execution of sessions to its child agents.
- If External Password Storage is enabled for this agent's master repository, the agent depends on the credential store for retrieving the source and target data servers' passwords to connect these data servers during session execution.
- If External Authentication is enabled for this agent's master repository, the run-time agents as well as Oracle Data Integrator Console depend on the Identity Store service that stores the Oracle Data Integrator user accounts.

These components must be available for the Oracle Data Integrator system to start and run properly.

### 7.2.4 Oracle Data Integrator Startup and Shutdown Process

The Oracle Data Integrator Java EE agent application is started by default whenever any Oracle WebLogic Managed Server to which it has been deployed is started. Normally, you should not need to stop Oracle Data Integrator or any of its components by themselves. Some operations may require the Oracle WebLogic Managed Server where Oracle Data Integrator runs to be rebooted. Only some patching scenarios should require stopping the application.

You can use Oracle WebLogic Server Administration Console to verify the status and to start and stop Oracle WebLogic Server. You can also use the WebLogic Server WLST command line to control the application. Oracle Enterprise Manager Fusion Middleware Control also allows multiple operations and configuration of Oracle Data Integrator as well as monitoring its status.

### 7.2.5 Oracle Data Integrator Configuration Artifacts

This section describes Oracle Data Integrator configuration artifacts.

#### 7.2.5.1 Java EE Agent Configuration

The configuration parameters for the Java EE agent are stored in both the agent application and in the master repository to which the agent is attached.

As shown in [Figure 7-3](#), the agent configuration stored in the master repository can be edited from the Oracle Data Integrator Studio and includes the following key parameters:

- Host
- Port
- Web Application Context

**Figure 7–3 Agent Configuration in the Master Repository Displayed in Oracle Data Integrator Studio**

The screenshot shows the Oracle Data Integrator Studio interface for configuring an agent. The title bar indicates the agent is 'AGENT\_001'. Below the title bar, there are tabs for 'Scheduling information', 'Update Scheduling', 'Test', and 'Generate WLS Template'. The 'Definition' tab is active, showing a sidebar with 'Definition' selected. The main panel displays the agent configuration details:

| Agent                       |                   |
|-----------------------------|-------------------|
| Name:                       | AGENT_001         |
| Host:                       | wls001.oracle.com |
| Port:                       | 8090              |
| Web application context:    | oraclediagent     |
| Protocol:                   | http              |
| Maximum number of sessions: | 1000              |

Using this information and valid credentials, any component connecting the repository is able to request an execution from this agent on the <host> : <port> / <web\_application\_context> URL. Several agents can be started on the same host and port pair by specifying a different **Web application context** value.

In addition to this information, the Master Repository also contains:

- The names of the data sources used by this agent to connect the source, target and repository database hosts.
- The hierarchy of agents for the load balancing feature.

Oracle Data Integrator Studio provides a feature to generate an Oracle WebLogic template for a given agent. This template contains the agent binaries, configuration and data source definitions. Oracle WebLogic administrators can use these templates to deploy pre-configured agents over a cluster.

Other configuration options at the container level, such as creation or modification of data sources are available as WebLogic Server domain configuration, and are synchronized across a cluster of Oracle WebLogic Servers by the Oracle WebLogic Server core infrastructure.

See the chapter on setting up the topology in *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator* for more information on the Java EE agent deployment options.

### 7.2.5.2 Standalone Agent Configuration

The configuration parameters for the standalone agent are stored both in the standalone agent configuration file as well as in the master repository to which this agent is attached. The repository part of the agent configuration is similar to the Java EE agent.

The agent configuration differs as it is stored in a configuration file, stored in its installation folder:

- For UNIX: `ODI_HOME/oracledi/agent/bin/odiparams.sh`
- For Windows: `ODI_HOME/oracledi/agent/bin/odiparams.bat`

This configuration file includes the following parameters that must be edited manually as part of the standalone agent configuration:

- JDBC connection information to connect the master repository
- ODI Supervisor user and encrypted password

The agent started in standalone mode is also configured with the command line used to start it. This command line can include the following parameters:

- Listening Port
- Agent Name
- Trace Level

A component willing to send an execution request to this agent will send it on the following URL: <host>:<port>/oraclediagent. Note that the standalone deployed application name is set to oraclediagent. As a consequence, to start several standalone agents on a single host, different ports must be used.

See the chapter on setting up the topology in *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator* for more information on the agent configuration details.

### 7.2.5.3 Oracle Data Integrator Console Configuration

Oracle Data Integrator Console configuration consists of connection definitions to the master and work repositories that can be browsed using this web application.

The list of connections is stored in the `repositories.xml` file in the following directory:

```
user_projects/domains/domainName/config/oracledi
```

Connections can be added, edited, or deleted from the Oracle Data Integrator Console management pages.

---

**Note:** Oracle Data Integrator Console is used as the entry point for Enterprise Manager to discover Oracle Data Integrator targets in a domain. The discovery process works in the following way: Enterprise Manager identifies Oracle Data Integrator Console. Using the Oracle Data Integrator Console configuration, Enterprise Manager identifies the master and work repositories as well as the run-time agents in the domain.

---

### 7.2.5.4 Oracle Data Integrator Log Locations and Configuration

This section provides information about Oracle Data Integrator log locations and configuration.

**7.2.5.4.1 Oracle Data Integrator Session Logs** Oracle Data Integrator session execution logs are stored in the work repositories against which the sessions are started. This session shows Oracle Data Integrator session details, such as the executed code and the number of processed rows. This log can be displayed from the Oracle Data Integrator Studio's Operator Navigator, in the **Session List** accordion, or from Oracle Data Integrator Console's **Browse** tab, under **Run-Time > Sessions**.

**7.2.5.4.2 Java EE Agent Log Files** The operations performed by the Oracle Data Integrator Java EE agent are logged by Oracle WebLogic Managed Server where the agent application is running. You can find these logs at the following location:

```
DOMAIN_HOME/servers/WLS_ServerName/logs/oracledi/odiagent.log
```

The log files for the different Oracle WebLogic Server Managed Servers are also available from Oracle WebLogic Server Administration Console. To verify the logs,

access Oracle WebLogic Server Administration Console using the following URL:  
`admin_server_host:port/console`. Click **Diagnostics-Log Files**.

It is also important to verify the output of the Oracle WebLogic Managed Server where Oracle Data Integrator is running. This information is stored at the following location:

`DOMAIN_HOME/servers/WLS_ServerName/logs/WLS_ServerName.out`

Additionally, a diagnostic log is produced in the log directory for the managed server. This log's granularity and logging properties can be changed through the following file:

`DOMAIN_HOME/config/fmwconfig/logging/oraclediagent-logging.xml`

**7.2.5.4.3 Standalone Agent Log Files** The operations performed by the Oracle Data Integrator standalone agent are logged by the lightweight container running the standalone agent. By default, logs are traced on the console and in the `ODI_HOME/oracledi/log/` folder.

The logging method and the logging level can be configured by editing the `ODI_HOME/oracledi/agent/bin/ODI-logging-config.xml` file.


**7.2.5.4.4 Oracle Data Integrator Console Log Files** Oracle Data Integrator Console logging operations are logged by Oracle WebLogic Managed Server where the agent application is running, like the Java EE agent log files described in [Section 7.2.5.4.2, "Java EE Agent Log Files."](#)

## 7.3 Oracle Data Integrator High Availability and Failover Considerations

This section describes Oracle Data Integrator high availability and failover considerations.

### 7.3.1 Oracle Data Integrator Clustered Deployment

[Figure 7-4](#) shows a two-node Oracle Data Integrator cluster running on two Oracle WebLogic servers. Oracle WebLogic Servers are front ended by Oracle HTTP Servers, which load balance incoming requests to them.

**Figure 7-4 Oracle Data Integrator High Availability Architecture**

The main characteristics of this configuration are:

- Oracle Data Integrator applications run on two clustered WebLogic Server managed servers. The WebLogic Server cluster synchronizes configuration for common artifacts of WebLogic Server used by Oracle Data Integrator, such as data sources.
- To avoid duplicate schedule processing, only one of these agents behaves like a scheduler. A Coherence cache is used to handle scheduler service uniqueness and migration.

The agent provides failover scheduling capabilities. For example, if a schedule is supposed to start at 9 AM, and the cycle is to run job X every hour for four hours, and the agent fails at 9:55 AM, it should compute where it was in the cycle and continue. However, if a single job is scheduled to start at 9 AM, and the agent fails at 8:59 AM, and then recovers at 9:01 AM, then it will not run the job that was scheduled at 9 AM.

- Requests to the Oracle Data Integrator agent in a cluster must be routed via a load balancer or via an HTTP proxy server. The address of this fronting server is used by clients to connect transparently to any of the Oracle Data Integrator servers in

the cluster. This address must be specified in the agent definition in the master repository. The scheduler singleton also routes all scheduled sessions startup requests to this address so that they are load balanced over the cluster.

- Oracle Data Integrator's master and work repositories database is configured with Oracle Real Application Clusters (Oracle RAC) to protect from database failures. Oracle Data Integrator components perform the appropriate reconnection and operations retries if database instance failure occurs.

## 7.3.2 Standalone Agent High Availability with OPMN

The standalone agent is a standalone Java process started from a command line interface. This agent is typically deployed locally on the source or target machines for optimal integration flow performances. You can use OPMN to start, stop and protect the standalone agent in this situation.

To add a standalone agent to OPMN:

1. Edit the `agentcreate.properties` file contained in the `ODI_HOME/oracledi/agent/bin/` directory to match your agent and OPMN configuration.
2. Run the script that adds this agent to the OPMN configuration.
  - For UNIX: `ODI_HOME/oracledi/agent/bin/opmn_addagent.sh`
  - For Windows: `ODI_HOME/oracledi/agent/bin/opmn_addagent.bat`

You can determine the status of Oracle HTTP Server using `opmnctl`:

```
opmnctl status
```

To start all agent components in an Oracle instance using `opmnctl`:

```
opmnctl startproc process-type=odiagent
```

To start a specific agent component, such as `odiagent1`, using `opmnctl`:

```
opmnctl startproc ias-component=odiagent1
```

To stop all agent components in an Oracle instance using `opmnctl`:

```
opmnctl stopproc process-type=odiagent
```

To stop a specific agent component, such as `odiagent1`, using `opmnctl`:

```
opmnctl stopproc ias-component=odiagent1
```

## 7.3.3 Oracle Data Integrator Protection from Failure and Expected Behavior

This section describes how an Oracle Data Integrator high availability cluster deployment and OPMN protects components from failure. This section also describes expected behavior in the event of component failure.

### 7.3.3.1 WebLogic Server or Standalone Agent Crash

If a WebLogic Server crashes, Node Manager attempts to restart it locally. If repeated restarts fail, the WebLogic Server infrastructure attempts to perform a server migration of the server to the other node in the cluster. While the failover takes place, the other WebLogic instance becomes the scheduler and is able to read, compute, and execute the schedule for all work repositories. A Coherence cache is used to handle the scheduler lifecycle. Locking guarantees the uniqueness of the scheduler, and event



notification provides scheduler migration. Note that when an agent restarts and computes its schedule, it takes into account schedules in progress (those in the middle of an execution cycle). These are automatically continued in their execution cycle beyond the server startup time. New sessions will be triggered as if the scheduler was never stopped.

Stale sessions are moved to an error state and are treated as such when restarted. This session recovery/restart is described in [Section 7.2.1.1, "Sessions Interruption"](#) and [Section 7.2.1.2, "Recovering Sessions."](#)

If a standalone agent crashes, OPMN restarts it locally. Existing sessions on the failing agent running become stale sessions, which are removed when the agent is restarted.

Oracle Data Integrator agents may be down due to failure in accessing resources, or other issues unrelated to whether the managed server is running. Therefore, Oracle recommends that administrators monitor the managed server logs for cluster errors caused by the application. For information about log file locations, see [Section 7.2.5.4, "Oracle Data Integrator Log Locations and Configuration."](#)

The Oracle Data Integrator Console does not support HTTP session failover. The user must log into the Oracle Data Integrator Console again after a failure.

### 7.3.3.2 Repository Database Failure

The Oracle Data Integrator repositories are protected against failures in the database by using multi data sources. These multi data sources are typically configured during the initial set up of the system (Oracle Fusion Middleware Configuration Wizard allows you to define these multi-pools directly at installation time) and guarantee that when an Oracle RAC database instance that hosts a repository fails, the connections are re-established with available database instances. The multi data source allows you to configure connections to multiple instances in an Oracle RAC database.

The Java EE agent uses WebLogic multi data sources that are configured during initial setup. The standalone agent uses the Oracle RAC JDBC connection string specified in `odiparams.sh`.

For additional information about multi data source configuration with Oracle RAC, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

Oracle Data Integrator implements a retry logic that allows in-flight sessions to proceed if a repository instance becomes unavailable and is restored at a later time. In an Oracle RAC enabled configuration, both in-flight and incoming session execution requests are served as long as an Oracle RAC node is available. This is supported in both the standalone and Java EE agents using the Retry Connection Count *number* and Connection Retry Delay *time* parameters. Users can configure these parameters when generating the WebLogic Server template for the Java EE agent and by editing `odiparams.sh` or `odiparams.bat` for the standalone agent.

If Oracle Data Integrator Studio loses its connection to an Oracle RAC database, you will lose any Oracle Data Integrator Studio work performed since the last save operation. As a general practice, save your work on a regular basis when you use Oracle Data Integrator Studio.

### 7.3.3.3 Scheduler Node Failure

In an Oracle Data Integrator agent cluster, when the agent node that is the scheduler node crashes, another node in the WebLogic Server cluster takes over as the scheduler node. The new scheduler node reinitializes all the schedules from that point and continues executing the scheduled scenarios from that point forward.

However, an issue arises in this situation if a scheduled scenario with a repeatable execution cycle was running on the first scheduler node when that node crashed. When the new scheduler node takes over, the scheduler scenario that was running on the first scheduler node will not continue its iterations on the new scheduler node from the point at which the first scheduler node crashed.

For example, if the scheduled scenario is configured to repeat the execution ten times after an interval of two minutes, and the first scheduler node crashes in the middle of the third execution, the new scheduler node should continue the execution of the scenario for the next eight executions. However, the new scheduler node does not continue the remaining executions of the scenario.

## 7.4 Configuring High Availability for Oracle Data Integrator

This section describes the installation and configuration steps for setting up an Oracle Data Integrator high availability configuration.

### 7.4.1 Running RCU to Create the Master and Work Repositories

Use the latest version of the Repository Creation Utility (RCU) to create the necessary schemas for Oracle Data Integrator in an Oracle RAC database repository.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

Start RCU by running this command:

```
RCU_HOME/bin/rcu
```

Then follow these steps:

1. On the Create Repository page, choose **Create and Load new schemas**.
2. On the Database Connection Details page, enter Database details. For Oracle RAC, only one mode is necessary.
3. On the Select Components page, select **Oracle Data Integrator**. This causes **Master and Work Repository** to be automatically selected.
4. On the Schema Passwords page, specify a password for the schema.
5. On the Custom Variables page, specify at least the Supervisor and Work Repository passwords. All other fields can be left as default.
6. On the Map Tablespaces screen, click **OK**.
7. On the Summary screen, click **OK** and complete the repository installation.

### 7.4.2 Installation and Configuration of the First Oracle Data Integrator Host

This section describes the installation and configuration steps for the first Oracle Data Integrator host (for example, APPHOST1).

#### 7.4.2.1 Installing the Oracle WebLogic Server on APPHOST1

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

Start the Oracle WebLogic Server installer.

Then follow these steps in the installer to install Oracle WebLogic Server on the computer:

1. On the Welcome screen, click **Next**.
2. On the Choose Middleware Home Directory screen, select **Create a New Middleware Home**. Then choose a directory on your computer into which the Oracle WebLogic software is to be installed.  
Click **Next**.
3. On the Register for Security Updates screen, enter your "My Oracle Support" User Name and Password so that you can be notified of security updates
4. On the Choose Install Type screen, the installation program prompts you to indicate whether you wish to perform a complete or a custom installation.  
Choose **Typical**.  
Click **Next**.
5. On the Choose Product Installation Directories screen, accept the default locations for WebLogic Server and Coherence.  
Click **Next**.
6. On the Installation Summary screen, the window contains a list of the components you selected for installation, along with the approximate amount of disk space to be used by the selected components once installation is complete.  
Click **Next**.
7. Click **Done** to exit the installer.

#### 7.4.2.2 Install Oracle Data Integrator on APPHOST1

Start the Oracle Data Integrator installer:

```
./runInstaller
```

When the installer prompts you for a JDK location, enter the full path to the SDK under the Middleware home.

Then perform these installation steps:

1. If the Specify Inventory Directory screen appears, enter an Oracle inventory location and OS Group name, and click **OK**.
2. On the Select Installation Type screen:
  - Select **J2EE Installation**.
  - Optionally, select **Developer** installation if Oracle Data Integrator Studio is required.
3. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
4. On the Specify Installation Location screen, select the Middleware Home from the list.  
Click **Next**.
5. On the Repository Configuration screen, choose **Skip Repository Configuration** and click **Next**.
6. On the Installation Summary screen, click **Install** to start the installation.

7. On the Installation Complete screen, click **Finish**.

### 7.4.2.3 Create the High Availability Domain

To start the Configuration Wizard, run `config.sh` in the `ODI_HOME/common/bin` directory. Then follow these steps:

1. On the Welcome screen, choose **Create a new WebLogic Domain**.
2. On the Select Domain Source screen, choose these components:
  - Oracle Data Integrator - SDK Web Services - 11.1.1.0
  - Oracle Data Integrator - Console - 11.1.1.0
  - Oracle Data Integrator - Agent - 11.1.1.0

Oracle Enterprise Manager and the Oracle Enterprise Manager Plug-in for ODI should also optionally be selected.
3. On the Specify Domain Name and Location screen, enter the domain name and click **Next**.
4. On the Configure Administrator Username and Password screen, enter a preferred username and password for the domain.
5. On the Configure Server Start Mode and JDK screen, choose **Production Mode**. Click **Next**.
6. In the Configure JDBC Component Schema screen:

- a. Select all component schemas that appear in the table at the bottom: **ODI Master Schema** and **ODI Work Schema**.
- b. For the RAC configuration for component schemas, select **Convert to GridLink** or **Convert to RAC multi data source**. For single database configuration, select **Don't convert**.
- c. Ensure that the following data source appears on the screen. The user names in [Table 7-1](#) assume that you used **DEV** as the prefix for schema creation from RCU.

For information about GridLink data source configuration with Oracle RAC, see [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC."](#)

For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

**Table 7-1 Configuring Values for Data Sources**

| Component Schema  | Schema Owner |
|-------------------|--------------|
| ODI Master Schema | DEV_ODI_REPO |
| ODI Work Schema   | DEV_ODI_REPO |

- d. Click **Next**.
7. In the Configure JDBC Component Schemas screen, enter values for the following fields, specifying connect information for the Oracle RAC database that was seeded with RCU:

**For multi data sources**, enter values for these fields:

**Figure 7-5 Configure RAC Multi Data Source Component Schema Screen**

**Fusion Middleware Configuration Wizard**

**Configure RAC Multi Data Source Component Schema**

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for RAC Service-Instance co

Service Name: Enter a value

Username: DEV\_ODI\_REPO

Password: \*\*

| Host Name | Instance Name | Port |
|-----------|---------------|------|
|           |               |      |

Add Delete

| RAC Component Schema                                  | Service Name | Schema Owner | Schema Password |
|-------------------------------------------------------|--------------|--------------|-----------------|
| <input checked="" type="checkbox"/> ODI Master Schema |              | DEV_ODI_REPO | **              |
| <input checked="" type="checkbox"/> ODI Work Schema   |              | DEV_ODI_REPO | **              |

Exit Help Previous Next

For GridLink data sources, enter values for these fields:

**Figure 7-6 Configure GridLink RAC Component Schema**

**Fusion Middleware Configuration Wizard**

**Configure GridLink RAC Component Schema**

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for RAC Service-Instance connect

Service Name: Enter a value

Username: DEV\_ODI\_REPO

Password: \*\*

Enable FAN: ☒

Enable SSL: ☐

Wallet File: Enter a value

Wallet Password: Enter a value

| Service Listener | Port | Protocol |
|------------------|------|----------|
|                  |      | TCP      |

| ONS Host | Port |
|----------|------|
|          |      |

Add Delete

| RAC Component Schema                                  | Service Name | Schema Owner | Schema Password |
|-------------------------------------------------------|--------------|--------------|-----------------|
| <input checked="" type="checkbox"/> ODI Master Schema |              | DEV_ODI_REPO | **              |
| <input checked="" type="checkbox"/> ODI Work Schema   |              | DEV_ODI_REPO | **              |

Exit Help Previous Next

- a. Driver: For RAC: Select **Oracle driver (Thin) for RAC Service-Instance connections, Versions:10, 11.** For GridLink: **Oracle Driver (Thin) for GridLink Connections, Versions:10 and later.**

- b. Service Name: Enter the service name of the database. **For a GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example,  
`<mydbservice>.us.oracle.com`

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example,  
`us.oracle.com`

---

- c. Username prefix: Enter the prefix for the schemas. The user names in [Table 7–1](#) assume that **DEV** was used as prefix for schema creation from RCU.
- d. Password and Confirm Password: Enter the password to access the schemas.
- e. **For a GridLink data source**, enter the SCAN address in the Listener Address field and SCAN port in the Port field. Enter the ONS host and port information in the ONS Host and Port fields, respectively.

---

**Note:** Oracle recommends that you use SCAN addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

---

**For a multi data source**, click **Add**. Enter details for Oracle RAC instances.

- f. Update each schema: select one data source at a time and add the appropriate details.  
  
Ensure that you enter information for all schemas: **ODI Master** and **ODI Work**.
- g. Click **Next**.
- 8. In the Test JDBC Data Sources screen, the connections are tested automatically. The **Status** column shows the results. Ensure that all connections were successful. If not, click **Previous** to correct your entries.  
  
Click **Next** when all the connections are successful.
- 9. On the Select Optional Configuration screen, check the following:
  - **Administration Server**
  - **Managed Server, Clusters and Machines**
  - **Deployments and Services**Click **Next**.
- 10. On the Configure Administration Server screen, set the Listen Address and Listen Port for the Administration Server.
- 11. On the Configure Managed Servers page:
  - Click to add a new server: `odi_server2`

- Set the Listen Address for each server to the hostname on which each server will run.
  - Ensure that both servers have the same Listen Port.
12. On the Configure Clusters page, create a cluster named `odi_cluster` and accept the default of unicast.
  13. On the Assign Servers to Cluster page, add both servers to the cluster.
  14. On the Configure Machines page, create two machines.
  15. On the Assign Machines page:
    - Add Admin Server and `odi_server1` to Machine 1.
    - Add `odi_server2` to Machine 2.
  16. On the Configuration Summary page, click **Create** to create the domain.

#### 7.4.2.4 Start the Administration Server

Start the WebLogic Administrator Server on APPHOST1 using this command:

```
DOMAIN_HOME/bin/startWebLogic.sh
```

#### 7.4.2.5 Configure the Credential Store

You can configure the credential store using WLST or Enterprise Manager.

##### 7.4.2.5.1 Configuring Credentials Using WLST

Run the following script:

```
MW_HOME/oracle_common/common/bin/wlst.sh
```

and then issue these commands:

```
connect('weblogic','welcome1','t3://localhost:7001')
createCred(map="oracle.odi.credmap", key="SUPERVISOR", user="SUPERVISOR",
password="SUNOPSIS", desc="ODI SUPERVISOR Credential")
exit()
```

The user above is the Supervisor user that was created when the repository was created using RCU. The password to specify is the password for that Supervisor user.

##### 7.4.2.5.2 Configuring Credentials using Enterprise Manager

To configure credentials using Enterprise Manager, follow these steps:

1. Log into Oracle Enterprise Manager and navigate to **Domain > Security > Credentials** to display the Credentials page.
2. Click **Create Map** to display the Create Map dialog.
3. In the Create Map dialog, enter the name of the map for the credential being created: `oracle.odi.credmap`
4. Click **OK** to return to the Credentials page. The new credential map name is displayed with a map icon in the table.
5. Click **Create Key** to display the Create Key dialog.
6. In this dialog, select a map of **oracle.odi.credmap**, enter a key in the text box Key of 'SUPERVISOR', select a type from the menu Type (the appearance of the dialog changes according to the type selected), enter the required data:

```
user="SUPERVISOR", password="your_password"
```

7. Click **OK** when finished to return to the Credentials page. The new key is shown under the map icon corresponding to the map you selected.

#### 7.4.2.6 Configure the Default Agent

Before you start the Oracle Data Integrator managed servers, create the agent definition for Oracle Data Integrator Agent. Follow these steps:

1. Start the Oracle Data Integrator Studio at `ODI_HOME/oracledi/client/odi.sh`
2. After Oracle Data Integrator Studio comes up, click **Connect to Repository** on the left-hand pane.
3. A Login window appears. Click on the pencil icon to edit the Master Repository logon information.
4. On the 'Repository Connection Information' screen, enter the following:  
ODI Connection:
  - **Login Name:** Master Repository
  - **User:** SUPERVISOR
  - **Password:** (The password used when creating the credential store)DB Connection:
  - **User:** (ODI Schema user such as DEV\_ODI\_REPO)
  - **Password:** (ODI Schema password)
  - **Drive Name:** oracle.jdbc.driver.OracleDriver
  - **URL:** (connection URL in the form jdbc:oracle:thin:@host:port:sid)
5. Click **Test** to test the connection.
6. Click **OK** to return to the logon screen.
7. Ensure the User/Password are correct and then click **OK** to connect to the Master Repository.
8. Once connected, open up the Physical Architecture section on the left-hand pane.
9. Select **Agents** and right-click to create a New Agent.
10. Create a new Agent with the following properties:
  - **Name:** OracleDI Agent
  - **Host:** (The host where odi\_server1 will be running)
  - **Port:** (The port that the odi\_server1 managed server is running on)
  - **Web Application Context:** oraclediagent
11. Open up the DataSources tab. On the left hand pane, expand the Repositories section and drag the **Work Repository** from there into the Datasources area of the Agent.
12. Edit the Work Repository datasource to add the JNDI name used by the datasource in the WebLogic Server installer. This should be `jdbc/odiWorkRepository`.
13. Save and exit Oracle Data Integrator Studio.



### 7.4.2.7 Configure Coherence for the Cluster

You must configure coherence to enable communication among the cluster members. To configure Coherence:

1. In the Administration Console, select **Environment > Servers** from the left hand tab.
2. Select `odi_server1`.
3. Click the Server Startup tab.
4. In the Arguments box, enter the following (all on a single line):

```
-Doracle.odi.coherence.wka1=machine1 -Doracle.odi.coherence.wka1.port=9088
-Doracle.odi.coherence.wka2=machine2 -Doracle.odi.coherence.wka2.port=9088
-Dtangosol.coherence.localport=9088
```

where `machine1` and `machine2` are the hostnames of the two machines in the cluster.

---

**Note:** Use 9088 as the Coherence port if it is unused on the machine. Otherwise, choose another port to configure as the Coherence port.

---

5. Click **Save**.
6. Repeat the steps above for `odi_server2`.

### 7.4.2.8 Configure Node Manager and Start `odi_server1`

Configure and start Node Manager using these commands:

```
$ MW_HOME/oracle_common/common/bin/setNMProps.sh
$ MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
```

Start `odi_server1` as follows:

1. In the Administration Console, choose **Environment > Servers > Control** tab.
2. Check `odi_server1` and click **Start**.

### 7.4.2.9 Verify the Oracle Data Integrator Agent is Running

Verify that the Oracle Data Integrator agent is running by entering the following URL in a web browser:

```
http://APPHOST1:PORT/oraclediagent
```

## 7.4.3 Installation and Configuration of the Second Oracle Data Integrator Host

This section describes the installation and configuration steps for the second Oracle Data Integrator host (for example, `APPHOST2`).

### 7.4.3.1 Installing the Oracle WebLogic Server on `APPHOST2`

Follow the instructions in [Section 7.4.2.1, "Installing the Oracle WebLogic Server on `APPHOST1`."](#)

Use the same directory paths as you used on `APPHOST1`.

### 7.4.3.2 Pack and Unpack the Domain from APPHOST1 to APPHOST2

Use the following commands to pack the domain that was created on the first host (APPHOST1):

```
$ cd MW_HOME/oracle_common/common/bin
$./pack.sh -domain=<full path to the domain> -template=mytemplate.jar
-template_name=<descriptive template name> -managed=true
```

Copy the jar file you created on the first host to the second host (APPHOST2) and unpack it:

```
$ cd MW_HOME/oracle_common/common/bin
$./unpack.sh -domain=<full path to the domain> -template=mytemplate.jar
```

### 7.4.3.3 Configure Node Manager and Start odi\_server2

Configure and start Node Manager on APPHOST2 using these commands:

```
$ MW_HOME/oracle_common/common/bin/setNMProps.sh
$ MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
```

Start odi\_server2 as follows:

1. In the Administration Console, choose **Environment > Servers > Control** tab.
2. Check odi\_server2 and click **Start**.

### 7.4.3.4 Verify the Oracle Data Integrator Agent is Running

Verify that the Oracle Data Integrator agent is running by entering the following URL in a web browser:

```
http://APPHOST2:PORT/oraclediagent
```

## 7.4.4 Installing Oracle HTTP Server

Install Oracle HTTP Server on another host (for example, WEBHOST1).

Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.

Start the installer for Oracle Web Tier components.

```
HOST> ./runInstaller
```

When the installer prompts you for a JRE/JDK location, enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, /u01/app/product/fmw/jrockit\_160\_<version>.

Then perform these installation steps:

1. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
2. On the Specify Installation Location screen, enter the following values:

**MW\_HOME:** Enter the value of the MW\_HOME, for example:

```
/u01/app/product/fmw
```

Select the previously installed Middleware Home from the drop-down list. For the Oracle HTTP Server Oracle Home (OHS\_ORACLE\_HOME) directory, enter the directory name `WEB`.

Click **Next**.

3. On the Summary screen, click **Install**.

When prompted, on Linux and UNIX installations, run the script `oracleRoot.sh` as the root user.

4. On the Installation Complete screen, click **Finish**.

Repeat these steps to install the Oracle HTTP Server on another web tier host (for example, `WEBHOST2`).

#### 7.4.4.1 Upgrading the Oracle HTTP Server Oracle Home

This section provides the steps to upgrade your Oracle HTTP Server software. For more information, see Applying the Latest Oracle Fusion Middleware Patch Set in the *Oracle Fusion Middleware Patching Guide*.

To upgrade the OHS\_ORACLE\_HOME:

1. Start the Web Tier Upgrade Installer by running `./runinstaller`.
2. On the Welcome screen, click **Next**.
3. On the Prerequisite Checks screen, click **Next**.
4. On the Specify Install Location screen, provide the path to the Oracle Middleware Home and the name of the Oracle HTTP Server Oracle Home directory.
5. On the Installation Summary screen, validate your selections, and then click **Install**.
6. The Installation Progress screen shows the progress of the install.

Once the installation is done, the `oracleRoot.sh` confirmation dialog box appears. This dialog box advises you that a configuration script needs to be run as root before the installation can proceed. Leaving the confirmation dialog box open, open another shell window, log in as root, and run this script file:  
`/u01/app/oracle/product/fmw/id/oracleRoot.sh`. After the script completes, click **OK** on the Confirmation Dialog box.

7. On the Installation Complete screen, click **Finish** to exit.

These steps should be performed on each of the web tier hosts on which Oracle HTTP Server is installed (for example, `WEBHOST1` and `WEBHOST2`).

#### 7.4.4.2 Configuring Oracle HTTP Server

After the installation add the following lines to the `OHS_HOME/instances/ohs_instance1/config/OHS/ohs11/mod_wl_ohs.conf` file on the host or hosts on which you installed Oracle HTTP Server:

```
ODI Agent
<Location /oraclediagent>
 WebLogicCluster host1:8002,host2:8002
 SetHandler weblogic-handler
</Location>

#ODI Explorere
<Location /odirepex>
 WebLogicCluster host1:8002,host2:8002
```

```
 SetHandler weblogic-handler
 </Location>

 #ODI Webservices
 <Location /oraclediskws>
 WebLogicCluster host1:8002,host2:8002
 SetHandler weblogic-handler
 </Location>
```

These steps should be performed on each of the web tier hosts on which Oracle HTTP Server is installed (for example, WEBHOST1 and WEBHOST2).

#### 7.4.4.3 Configuring the Load Balancer

The load balancer should be configured to round-robin requests across the two HTTP Servers on WEBHOST1 and WEBHOST2.

Also, follow these load balancer configuration recommendations:

- Persistence should not be enabled on the load balancer. This is provided by the HTTP Servers.
- Monitors should be configured to monitor the HTTP Server ports so that the load balancer can provide failover.

#### 7.4.4.4 Verify the Oracle Data Integrator Agent is Running

Verify that the Oracle Data Integrator agent is running by entering the following URL in a web browser:

```
http://LBRHOST:PORT/oraclediagent
```

#### 7.4.4.5 Reconfigure Agents

Agent definitions should point to the load balancer address instead of the individual server addresses.

This should be kept in mind when creating new agents.

For the default agent, connect to Oracle Data Integrator Studio again:

1. Start Oracle Data Integrator Studio at `ODI_HOME/oracledi/client/odi.sh`.
2. After Oracle Data Integrator Studio comes up, click on **Connect to Repository** on the left-hand pane.
3. When the Login window appears, click **OK** to logon.
4. When you are connected, open the Physical Architecture section on the left-hand pane.
5. Select **Agents** and then select the OracleDI Agent.
6. Edit the following properties:
  - Host: The load balancer virtual server address.
  - Port: The load balancer virtual address listening port.
7. Click **Test** to test the agent connection.
8. Save and exit Oracle Data Integrator Studio.

---

# Configuring High Availability for Identity Management Components

The Oracle Identity Management products enable you to configure and manage the identities of users, devices, and services across diverse servers, to delegate administration of these identities, and to provide end users with self-service privileges. You can use these products to configure single sign-on across applications and to process users' credentials to ensure that only users with valid credentials can log into and access online resources.

It is critical to configure high availability for the Oracle Identity Management products because other enterprise-level applications depend on them. If the Oracle Identity Management products fail, applications depending on them also fail.


This chapter describes configuring the Identity Management products for high availability in an active-active configuration.

This chapter includes the following topics:

- [Section 8.1, "Identity Management Product Components and High Availability Concepts"](#)
- [Section 8.2, "Prerequisites for Oracle Identity Management High Availability Configuration"](#)
- [Section 8.3, "Oracle Internet Directory High Availability"](#)
- [Section 8.4, "Oracle Virtual Directory High Availability"](#)
- [Section 8.5, "Oracle Directory Integration Platform High Availability"](#)
- [Section 8.6, "Oracle Directory Services Manager High Availability"](#)
- [Section 8.7, "Collocated Architecture High Availability"](#)
- [Section 8.8, "Oracle Access Manager High Availability"](#)
- [Section 8.9, "Oracle Identity Manager High Availability"](#)
- [Section 8.10, "Authorization Policy Manager High Availability"](#)
- [Section 8.11, "Oracle Identity Navigator High Availability"](#)
- [Section 8.12, "Oracle Adaptive Access Manager High Availability"](#)
- [Section 8.13, "Oracle Identity Federation High Availability"](#)
- [Section 8.14, "Starting and Stopping Oracle Identity Management Components"](#)

## 8.1 Identity Management Product Components and High Availability Concepts

[Figure 8–1](#) shows a sample Oracle Fusion Middleware 11g Oracle Identity Management high availability architecture. This architecture includes a web tier, application tier, and directory tier.

**Figure 8–1 Oracle Fusion Middleware 11g Oracle Identity Management High Availability Architecture**

In [Figure 8–1](#), the web tier includes the **WEBHOST1** and **WEBHOST2** computers.

An Oracle HTTP Server instance is installed on WEBHOST1 and an Oracle HTTP Server instance is installed on WEBHOST2. A load balancing router routes requests to the Oracle HTTP Server instances on WEBHOST1 and WEBHOST2.

The application tier includes the IDMHOST1 and IDMHOST2 computers.

On IDMHOST1, the following installations have been performed:

- An Oracle Directory Services Manager instance and an Oracle Directory Integration Platform instance have been installed in the WLS\_ODS1 Managed Server. The Oracle RAC database has been configured in a JDBC multi data source to protect the instances from Oracle RAC node failure.
- An Oracle Access Manager Access Server instance has been installed in the WLS\_OAM1 Managed Server.
- A WebLogic Administration Server has been installed. Under normal operations, this is the active Administration Server. The Administration Server is a singleton application. The Oracle Access Manager Console has also been installed as a singleton application.

On IDMHOST2, the following installations have been performed:

- An Oracle Directory Services Manager instance and an Oracle Directory Integration Platform instance have been installed in the WLS\_ODS2 Managed Server. The Oracle RAC database has been configured in a JDBC multi data source to protect the instances from Oracle RAC node failure.

The instances in the WLS\_ODS2 Managed Server on IDMHOST2 and the instances in the WLS\_ODS1 Managed Server on IDMHOST1 are configured as the Cluster\_ODS cluster.

- An Oracle Access Manager Access Server instance has been installed in the WLS\_OAM2 Managed Server.

The Access Server instance in the WLS\_OAM2 Managed Server on IDMHOST2 and the Access Server instance in the WLS\_OAM1 Managed Server on IDMHOST1 are configured as the Cluster\_OAM cluster.

- A WebLogic Administration Server has been installed. Under normal operations, this is the passive Administration Server. You will make this Administration Server active if the Administration Server on IDMHOST1 becomes unavailable. See [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability"](#) for more information about active-passive configurations. The Oracle Access Manager Console has also been installed as a singleton application, and it is passive until the Administration Server on IDMHOST2 becomes active.

The directory tier includes the OIDHOST1 and OIDHOST2 computers.

On OIDHOST1, an Oracle Internet Directory instance and an Oracle Virtual Directory instance have been installed. Transparent Application Failover (TAF) is used to connect the Oracle Internet Directory instance with the Oracle RAC database that serves as the security metadata repository. The database is enabled for server-side TAF and HA Events Notification.

On OIDHOST2, an Oracle Internet Directory instance and an Oracle Virtual Directory instance have been installed. Transparent Application Failover (TAF) is used to connect the Oracle Internet Directory instance with the Oracle RAC database that serves as the security metadata repository. The database is enabled for server-side TAF and HA Events Notification.

The Oracle Internet Directory instances on OIDHOST1 and OIDHOST2 have been configured as a cluster.



An Oracle Real Applications Cluster (Oracle RAC) database serves as the security metadata repository.

### 8.1.1 About the 11g Oracle Identity Management Products

[Table 8–1](#) summarizes the Oracle Identity Management products that you can install using the suite-level installation program for 11g. See the introductory chapter of the *Oracle Fusion Middleware Quick Installation Guide for Oracle Identity Management* for details:

**Table 8–1 The 11g Identity Management Components and Product Suites**

| Product                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Product Suite                                                    |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Oracle Internet Directory</b>             | Oracle Internet Directory is an LDAP Version 3-enabled service that enables fast retrieval and centralized management of information about dispersed users, network configuration, and other resources.                                                                                                                                                                                                                                                                                                                                     | Oracle Identity Management Platform and Directory Services Suite |
| <b>Oracle Virtual Directory</b>              | Oracle Virtual Directory is an LDAP Version 3-enabled service that provides an abstracted view of one or more enterprise data sources.<br><br>Oracle Virtual Directory consolidates multiple data sources into a single directory view, enabling you to integrate LDAP-aware applications with diverse directory server data stores.                                                                                                                                                                                                        | Oracle Identity Management Platform and Directory Services Suite |
| <b>Oracle Directory Integration Platform</b> | The Oracle Directory Integration Platform is a J2EE application that enables you to synchronize data between various directories and Oracle Internet Directory. Oracle Directory Integration Platform includes services and interfaces that allow you to deploy synchronization solutions with other enterprise repositories. It can also be used to provide Oracle Internet Directory interoperability with third party metadirectory solutions.                                                                                           | Oracle Identity Management Platform and Directory Services Suite |
| <b>Oracle Directory Services Manager</b>     | Oracle Directory Services Manager is a unified graphical user interface (GUI) for Oracle Virtual Directory and Oracle Internet Directory. Oracle Directory Services Manager simplifies the administration and configuration of Oracle Virtual Directory and Oracle Internet Directory by allowing you to use web-based forms and templates.<br><br>Oracle Directory Services Manager is available from either the Oracle Enterprise Manager Fusion Middleware Control or from its own URL.                                                  | Oracle Identity Management Platform and Directory Services Suite |
| <b>Oracle Access Manager</b>                 | Oracle Access Manager 11g is the successor product to both Oracle Access Manager 10g (access only) and Oracle Single Sign-On 10g. Oracle Access Manager 11g provides a single authoritative source for all authentication and authorization services. The core service provided is the checking of valid session tokens, the requesting of credentials if the session token is invalid or missing, and the issuing of session tokens, intercepting resource requests and evaluating access control policies to control access to resources. | Oracle Identity and Access Management Suite                      |

**Table 8–1 (Cont.) The 11g Identity Management Components and Product Suites**

| <b>Product</b>                        | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>Product Suite</b>                                             |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>Oracle Identity Manager</b>        | Oracle Identity Manager is a user provisioning and administration solution that automates the process of adding, updating, and deleting user accounts from applications and directories. It also improves regulatory compliance by providing granular reports that attest to who has access to what. Oracle Identity Manager is available as a stand-alone product or as part of Oracle Identity and Access Management Suite.                         | Oracle Identity and Access Management Suite                      |
| <b>Authorization Policy Manager</b>   | Authorization Policy Manager is a graphical interface tool for administering application policies. The intended users of Authorization Policy Manager are security administrators. With this tool, an administrator can view and manage policies across enterprise applications. Administrators can be specified to manage all applications running in the domain or just a subset of them.                                                           | Oracle Identity and Access Management Suite                      |
| <b>Oracle Identity Navigator</b>      | Oracle Identity Navigator is an administrative portal designed to act as a launch pad for Oracle Identity Management products. It does not replace the individual product consoles. Rather, it allows you to access the Oracle Identity Management consoles from one site.                                                                                                                                                                            | Oracle Identity and Access Management Suite                      |
| <b>Oracle Adaptive Access Manager</b> | Oracle Adaptive Access Manager (OAAM) is Oracle Identity Management's solution for web access real-time fraud detection and multifactor online authentication security for the enterprise. OAAM enables real-time blocking of fraudulent access requests, delivers advanced alerting mechanisms, and protects businesses and their customers from attacks such as phishing, Trojans, viruses, fraudulent transactions, and Man-in-the-Middle attacks. | Oracle Identity and Access Management Suite                      |
| <b>Oracle Identity Federation</b>     | Oracle Identity Federation enables companies to provide services and share identities across their respective security domains, while providing protection from unauthorized access.                                                                                                                                                                                                                                                                  | Oracle Identity Management Platform and Directory Services Suite |

## 8.2 Prerequisites for Oracle Identity Management High Availability Configuration

This section describes the prerequisite steps that you must complete before setting up an Oracle Identity Management high availability configuration.

### 8.2.1 Oracle Home Requirement

The Oracle home for the Identity Management components must be the same across all the nodes. For example, if you choose `/u01/app/oracle/product/fmw/idm` as the Oracle home on Node1, then you must choose `/u01/app/oracle/product/fmw/idm` as the Oracle home on all subsequent nodes.

## 8.2.2 Database Prerequisites

Several Oracle Identity Management components require the presence of a supported database and schemas.

To check if your database is certified or to see all certified databases, refer to the "Certified Databases" section in the Certification Document:

[http://www.oracle.com/technology/software/products/ias/files/fusion\\_certification.html](http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html)

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

## 8.2.3 Installing and Configuring the Database Repository

The database used to store the metadata repository should be highly available in its own right. For maximum availability, Oracle recommends the use of an Oracle Real Application Clusters (Oracle RAC) database.

Ideally, the database will use Oracle Automatic Storage Management for the storage of data; however, this is not necessary.

If you use Oracle ASM, then Oracle ASM should be installed into its own Oracle Home and have two disk groups:

- One for the Database files.
- One for the Flash Recovery Area.

If you are using Oracle ASM, the best practice is to also use Oracle Managed Files.

This section provides links to instructions for installing and configuring a database repository.

### Oracle Clusterware

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

### Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

### Oracle Real Application Clusters

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

Many of the Oracle Fusion Middleware components require the existence of schemas in a database prior to installation. Oracle Fusion Middleware provides a tool, the Repository Creation Utility (RCU), to create the component schemas in an existing database. For high availability environments, these schemas must be created and loaded into a Real Application Clusters (Oracle RAC) database.

See the next section for information on using RCU to load the Oracle Identity Management schemas into an Oracle RAC database repository.

## 8.2.4 Obtaining the Repository Creation Utility Software

To obtain the latest version of RCU, go to the Oracle Fusion Middleware 11gR1 Software Download page on Oracle Technology Network:

[http://www.oracle.com/technology/software/products/middleware/htdocs/fmw\\_11\\_download.html](http://www.oracle.com/technology/software/products/middleware/htdocs/fmw_11_download.html)

Look for Repository Creation Utility in the "Required Additional Software" table. After downloading the .zip file, extract the contents to a directory of your choice; this directory is referred to as the *RCU\_HOME* directory.

---

---

**Note:** On Windows operating systems, ensure that you do not unzip the RCU .zip file to a directory name containing spaces.

---

---

Before you install any of the Oracle Identity Management components described in this chapter, run RCU to create the schemas used by the component into an Oracle RAC database. These schemas are required for the high availability Oracle Identity Management configurations described in this chapter.

For additional information about RCU, see *Oracle Fusion Middleware Installation Planning Guide* and *Oracle Fusion Middleware Repository Creation Utility User's Guide*.

### 8.2.4.1 Executing the Repository Creation Utility

Use the Repository Creation Utility (RCU) that is version compatible with the product you are installing.

For additional information about running RCU, see *Oracle Fusion Middleware Installation Planning Guide* and *Oracle Fusion Middleware Repository Creation Utility User's Guide*.

The schemas you create depend on the Identity Management products you wish to install and configure, for example:

- If the database is for Oracle Identity Manager, select **Identity Management - Oracle Identity Manager**.

---

---

**Note:** the SOA and the MDS Schemas are automatically selected.

---

---

- If the database is for Oracle Internet Directory, select **Identity Management - (Oracle Internet Directory - ODS)**.
- If the database is for Oracle Access Manager, select **Identity Management - Oracle Access Manager**.
- If the database is for Oracle Identity Federation, select **Identity Management - Oracle Identity Federation**.

- If the database is for Oracle Adaptive Access Manager, select **Identity Management - Oracle Adaptive Access Manager**.

## 8.2.5 Configuring the Database for Oracle Fusion Middleware 11g Metadata

Create the Oracle Real Application Clusters database to store Oracle Fusion Middleware 11g metadata with the following characteristics:

- It should be in archive log mode to facilitate backup and recovery.
- Optionally, flashback should be enabled.
- It should be created with the ALT32UTF8 character set.

The value of the static PROCESSES initialization parameter must be 500 or greater for Oracle Internet Directory. This value is checked by the Repository Creation Utility.

To check the value, you can use the SHOW PARAMETER command in SQL\*Plus:

```
prompt> sqlplus "sys/password as sysdba"
SQL> SHOW PARAMETER processes
```

One common method of changing the parameter value is to use a command similar to the following and then stop and restart the database to make the parameter take effect:

```
prompt> sqlplus "sys/password as sysdba"
SQL> ALTER SYSTEM SET PROCESSES=500 SCOPE=SPFILE;
```

The method that you use to change a parameter's value depends on whether the parameter is static or dynamic, and on whether your database uses a parameter file or a server parameter file. See the *Oracle Database Administrator's Guide* for details on parameter files, server parameter files, and how to change parameter values.

### 8.2.5.1 Database Examples in This Chapter

[Table 8–2](#) shows the values used for database configuration examples in this chapter.

**Table 8–2 Databases Used in Identity Management Configuration Examples**

| Component                             | Database Service Name | Database Instance Name |
|---------------------------------------|-----------------------|------------------------|
| Oracle Internet Directory             | oid.mycompany.com     | oiddb1, oiddb2         |
| Oracle Virtual Directory              | N/A                   | N/A                    |
| Oracle Directory Integration Platform | oid.mycompany.com     | oiddb1, oiddb2         |
| Oracle Directory Services Manager     | N/A                   | N/A                    |
| Oracle Access Manager                 | oam.mycompany.com     | oamdb1, oamdb2         |
| Oracle Identity Manager               | oim.mycompany.com     | oimdb1, oimdb2         |
| Authorization Policy Manager          | apm.mycompany.com     | apmdb1, apmdb2         |
| Oracle Identity Navigator             | N/A                   | N/A                    |
| Oracle Adaptive Access Manager        | oaam.mycompany.com    | oaamdb1, oaamdb2       |
| Oracle Identity Federation            | oif.mycompany.com     | oifdb1, oifdb2         |

### 8.2.5.2 Database Services

Oracle recommends using the Oracle Enterprise Manager Cluster Managed Services Page to create database services that client applications will use to connect to the database. For complete instructions on creating database services, see the chapter on Workload Management in the *Oracle Real Application Clusters Administration and Deployment Guide*.

You can also use SQL\*Plus to configure your Oracle RAC database to automate failover for Oracle Internet Directory using the following instructions. Note that each of the following commands only has to be run on one node in the cluster:

1. Use the CREATE\_SERVICE subprogram to both create the database service and enable high availability notification and configure server-side Transparent Application Failover (TAF) settings:

```
prompt> sqlplus "sys/password as sysdba"

SQL> EXECUTE DBMS_SERVICE.CREATE_SERVICE
(SERVICE_NAME => 'idm.mycompany.com',
NETWORK_NAME => 'idm.mycompany.com',
AQ_HA_NOTIFICATIONS => TRUE,
FAILOVER_METHOD => DBMS_SERVICE.FAILOVER_METHOD_BASIC,
FAILOVER_TYPE => DBMS_SERVICE.FAILOVER_TYPE_SELECT,
FAILOVER_RETRIES => 5, FAILOVER_DELAY => 5);
```

The EXECUTE DBMS\_SERVICE command above must be entered on a single line to execute properly.

2. Add the service to the database and assign it to the instances using srvctl:

```
prompt> srvctl add service -d idmdb -s idm -r idmdb1,idmdb2
```

3. Start the service using srvctl:

```
prompt> srvctl start service -d idmdb -s idm
```

---

---

**Note:** For more information about the SRVCTL command, see the *Oracle Real Application Clusters Administration and Deployment Guide*.

---

---

If you already have a service in the database, ensure that it is enabled for high availability notifications and configured with the proper server-side Transparent Application Failover (TAF) settings. Use the DBMS\_SERVICE package to modify the service to enable high availability notification to be sent through Advanced Queuing (AQ) by setting the AQ\_HA\_NOTIFICATIONS attribute to TRUE and configure server-side Transparent Application Failover (TAF) settings, as shown below:

```
prompt> sqlplus "sys/password as sysdba"

SQL> EXECUTE DBMS_SERVICE.MODIFY_SERVICE
(SERVICE_NAME => 'idm.mycompany.com',
AQ_HA_NOTIFICATIONS => TRUE,
FAILOVER_METHOD => DBMS_SERVICE.FAILOVER_METHOD_BASIC,
FAILOVER_TYPE => DBMS_SERVICE.FAILOVER_TYPE_SELECT,
FAILOVER_RETRIES => 5, FAILOVER_DELAY => 5);
```

The EXECUTE DBMS\_SERVICE command above must be entered on a single line to execute properly.

---

**Note:** For more information about the DBMS\_SERVICE package, see the *Oracle Database PL/SQL Packages and Types Reference*.

---

When using a 11.2 database, please follow the steps in the "Creating and Deleting Database Services with SRVCTL" section of the *Oracle Database Administrator's Guide* for 11g Release 2 (11.2).

### 8.2.5.3 Verifying Transparent Application Failover (TAF)

This section describes how to validate the Transparent Application Failover (TAF) configuration settings made earlier.

After the Oracle Internet Directory process has been started, you can query the FAILOVER\_TYPE, FAILOVER\_METHOD, and FAILED\_OVER columns in the V\$SESSION\_VIEW to obtain information about connected clients and their TAF status.

For example, use the following SQL statement to verify that TAF is correctly configured:

```
SELECT MACHINE, FAILOVER_TYPE, FAILOVER_METHOD, FAILED_OVER, COUNT(*)
FROM V$SESSION
GROUP BY MACHINE, FAILOVER_TYPE, FAILOVER_METHOD, FAILED_OVER;
```

The output before failover is similar to this:

| MACHINE  | FAILOVER_TYPE | FAILOVER_M | FAI | COUNT(*) |
|----------|---------------|------------|-----|----------|
| oidhost1 | SELECT        | BASIC      | NO  | 11       |
| oidhost1 | SELECT        | BASIC      | NO  | 1        |

The output after failover is similar to this:

| MACHINE  | FAILOVER_TYPE | FAILOVER_M | FAI | COUNT(*) |
|----------|---------------|------------|-----|----------|
| oidhost2 | SELECT        | BASIC      | NO  | 11       |
| oidhost2 | SELECT        | BASIC      | NO  | 1        |

### 8.2.5.4 Configuring Virtual Server Names and Ports for the Load Balancer

This section describes the network prerequisites for deploying an Oracle Identity Management high availability environment.

**8.2.5.4.1 Load Balancers** All components in the Oracle Identity Management software stack require a hardware load balancer when deployed in a high availability configuration. The hardware load balancer should have the following features:

- Ability to load-balance traffic to a pool of real servers through a virtual host name: Clients access services using the virtual host name (instead of using actual host names). The load balancer can then load balance requests to the servers in the pool.
- Port translation configuration: The load balancer should have the ability to perform port translation, where it allows incoming requests received on one port to be routed to a server process running on a different port. For example, a request received on port 80 can be routed to port 7777.
- Protocol translation: The load balancer should support protocol translation between systems running different protocols. It enables users on one network to

access hosts on another network, despite differences in the native protocol stacks associated with the originating device and the targeted host. For example, incoming requests can be HTTPS, and outgoing requests can be HTTP.

This feature is recommended but not required.

- SSL acceleration: SSL acceleration is a method of offloading the processor-intensive public key encryption algorithms involved in SSL transactions to a hardware accelerator.

This feature is recommended but not required.

- Monitoring of ports (HTTP, HTTPS, LDAP, LDAPS)
- Virtual servers and port configuration

Ability to configure virtual server names and ports on your external load balancer, and the virtual server names and ports must meet the following requirements:

- The load balancer should allow configuration of multiple virtual servers. For each virtual server, the load balancer should allow configuration of traffic management on more than one port. For example, for Oracle Internet Directory clusters, the load balancer needs to be configured with a virtual server and ports for LDAP and LDAPS traffic.
- The virtual server names must be associated with IP addresses and be part of your DNS. Clients must be able to access the load balancer through the virtual server names.
- Ability to detect node failures and immediately stop routing traffic to the failed node
- Resource monitoring / port monitoring / process failure detection  
The load balancer must be able to detect service and node failures (through notification or some other means) and to stop directing non-Oracle Net traffic to the failed node. If your load balancer has the ability to automatically detect failures, you should use it.
- Fault-tolerant mode  
It is highly recommended that you configure the load balancer to be in fault-tolerant mode.
- Other  
Oracle recommends that you configure the load balancer virtual server to return immediately to the calling client when the back-end services that it forwards traffic to are unavailable. This is preferred over the client disconnecting on its own after a timeout based on the TCP/IP settings on the client machine.
- Sticky routing capability

Ability to maintain sticky connections to components based on cookies or URL.

Table 8–3 shows the virtual server names to use for the external load balancer in the Oracle Identity Management high availability environment.

**Table 8–3 Virtual Server Names for the External Load Balancer**

| Component                 | Virtual Server Name |
|---------------------------|---------------------|
| Oracle Internet Directory | oid.mycompany.com   |
| Oracle Virtual Directory  | ovd.mycompany.com   |



**Table 8–3 (Cont.) Virtual Server Names for the External Load Balancer**

| Component                                 | Virtual Server Name |
|-------------------------------------------|---------------------|
| Oracle Identity Federation                | oif.mycompany.com   |
| Oracle Directory Services Manager Console | admin.mycompany.com |
| Oracle Access Manager                     | sso.mycompany.com   |
| Oracle Adaptive Access Manager            | oaam.mycompany.com  |
| Oracle Identity Manager                   | sso.mycompany.com   |

**8.2.5.4.2 Virtual Server Names** This section describes the virtual server names that should be set up for the high availability deployments described in this chapter.

Ensure that the virtual server names are associated with IP addresses and are part of your Domain Name System (DNS). The computers on which Oracle Fusion Middleware is running must be able to resolve these virtual server names

#### **oid.mycompany.com**

This virtual server acts as the access point for all LDAP traffic to the Oracle Internet Directory servers in the directory tier. Traffic to both the SSL and non-SSL ports is configured. The clients access this service using the address oid.mycompany.com:636 for SSL and oid.mycompany.com:389 for non-SSL.

Monitor the heartbeat of the Oracle Internet Directory processes on OIHOST1 and OIHOST2. If an Oracle Internet Directory process stops on OIHOST1 or OIHOST2, or if either host OIHOST1 or OIHOST2 is down, the load balancer must continue to route the LDAP traffic to the surviving computer.

#### **ovd.mycompany.com**

This virtual server acts as the access point for all LDAP traffic to the Oracle Virtual Directory servers in the directory tier. Traffic to both the SSL and non-SSL port is configured. The clients access this service using the address ovd.mycompany.com:7501 for SSL and ovd.mycompany.com:6501 for non-SSL.

Monitor the heartbeat of the Oracle Virtual Directory processes on OVDHOST1 and OVDHOST2. If an Oracle Virtual Directory process stops on OVDHOST1 or OVDHOST2, or if either host OVDHOST1 or OVDHOST2 is down, the load balancer must continue to route the LDAP traffic to the surviving computer.

#### **oif.mycompany.com**

This virtual server acts as the access point for all HTTP traffic to the Oracle Identity Federation servers in the application tier.

#### **oaam.mycompany.com**

This virtual server acts as the access point for all Oracle Adaptive Access Manager traffic that gets directed to the web site.

#### **sso.mycompany.com**

This virtual server acts as the access point for all Oracle Access Manager traffic that gets directed to the web site.

This virtual server acts as the access point for all HTTP traffic that gets directed to the single sign on services.

This virtual host must be configured to preserve the client IP address for a request. In some load balancing routers, this can be achieved by enabling the load balancing

router to insert the original client IP address of a request in an X-Forwarded-For HTTP header.

## 8.3 Oracle Internet Directory High Availability

This section provides an introduction to Oracle Internet Directory and describes how to design and deploy a high availability environment for Oracle Internet Directory.

This section includes the following topics:


- [Section 8.3.1, "Oracle Internet Directory Component Architecture"](#)
- [Section 8.3.2, "Oracle Internet Directory High Availability Concepts"](#)
- [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps"](#)
- [Section 8.3.4, "Validating Oracle Internet Directory High Availability"](#)
- [Section 8.3.5, "Oracle Internet Directory Failover and Expected Behavior"](#)
- [Section 8.3.6, "Troubleshooting Oracle Internet Directory High Availability"](#)
- [Section 8.3.7, "Additional Oracle Internet Directory High Availability Issues"](#)

### 8.3.1 Oracle Internet Directory Component Architecture

Oracle Internet Directory is an LDAP store that can be used by Oracle components such as Directory Integration Platform, Oracle Directory Services Manager, JPS, and also by non-Oracle components. These components connect to Oracle Internet Directory using the LDAP or LDAPS protocols.

The Oracle directory replication server uses LDAP to communicate with an Oracle directory (LDAP) server instance. To communicate with the database, all components use OCI/Oracle Net Services. Oracle Directory Services Manager and the command-line tools communicate with the Oracle directory servers over LDAP.

[Figure 8–2](#) shows Oracle Internet Directory in a non-high availability architecture.

**Figure 8–2 Oracle Internet Directory in a Non-High Availability Architecture**

An Oracle Internet Directory node consists of one or more directory server instances connected to the same directory store. The directory store—that is, the repository of the directory data—is an Oracle database.

Figure 8–2 shows the various directory server elements and their relationships running on a single node.

Oracle Net Services is used for all connections between the Oracle database server and:

- The Oracle directory server non-SSL port (389 for this topology)
- The Oracle directory server SSL-enabled port (636 for this topology)
- The OID Monitor

LDAP is used for connections between the directory server and:

- Oracle Directory Services Manager
- Oracle directory replication server

The Oracle directory server instance and the Oracle directory replication server connect to OID Monitor by way of the operating system.

As shown in Figure 8–2, an Oracle Internet Directory node includes the following major elements:

**Table 8–4 An Oracle internet Directory Node**

| Element                          | Description                                                                                                                                                                                                                                                                                               |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle directory server instance | Also called either an LDAP server instance or a directory server instance, it services directory requests through a single Oracle Internet Directory dispatcher process listening at specific TCP/IP ports. There can be more than one directory server instance on a node, listening on different ports. |

**Table 8–4 (Cont.) An Oracle internet Directory Node**

| Element                                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle directory replication server                   | <p>Also called a replication server, it tracks and sends changes to replication servers in another Oracle Internet Directory system. There can be only one replication server on a node. You can choose whether to configure the replication server. If there are multiple instances of Oracle Internet Directory that use the same database, only one of them can be running replication. This is true even if the Oracle Internet Directory instances are on different nodes.</p> <p>The replication sever process is a process within Oracle Internet Directory. It only runs when replication is configured.</p> <p>For more information about Oracle Internet Directory replication, refer to <a href="#">Chapter 10, "Configuring Identity Management for Maximum High Availability."</a></p>                                                                                                                                                                                                                                                                                                           |
| Oracle Database Server                                | Stores the directory data. Oracle strongly recommends that you dedicate a database for use by the directory. The database can reside on the same node as the directory server instances.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Oracle Process Manager and Notification Server (OPMN) | Manages Oracle Internet Directory as an Oracle Fusion Middleware component. OPMN uses the directives in the OID component snippet in <code>ORACLE_INSTANCE/opmn.xml</code> and invokes OIDMON and OIDCTL as required. The command-line utility is <code>opmnctl</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| OID Monitor (OIDMON)                                  | <p>Initiates, monitors, and terminates the LDAP server and replication server processes. When you invoke process management commands, such as <code>oidctl</code> or <code>opmnctl</code>, or when you use Fusion Middleware Control to start or stop server instances, your commands are interpreted by this process.</p> <p>OIDMON also monitors servers and restarts them if they have stopped running for abnormal reasons.</p> <p>OIDMON starts a default instance of <code>OIDLDAPD</code>. If the default instance of <code>OIDLDAPD</code> is stopped using the <code>OIDCTL</code> command, then <code>OIDMON</code> stops the instance. When <code>OIDMON</code> is restarted by <code>OPMN</code>, <code>OIDMON</code> restarts the default instance.</p> <p>All <code>OID Monitor</code> activity is logged in the file <code>ORACLE_INSTANCE/diagnostics/log/OID/component_id/oidmon-xxxx.log</code>. This file is on the Oracle Internet Directory server file system.</p> <p><code>OID Monitor</code> checks the state of the servers through mechanisms provided by the operating system.</p> |
| OID Control Utility (OIDCTL)                          | Communicates with <code>OID Monitor</code> by placing message data in Oracle Internet Directory server tables. This message data includes configuration parameters required to run each Oracle directory server instance. Normally used from the command line only to stop and start the replication server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### 8.3.1.1 Oracle Internet Directory Component Characteristics

Oracle Internet Directory, which is Oracle's LDAP store, is a C-based component that uses a database as its persistence store. It is a stateless process and stores all of the data and the majority of its configuration information in the back-end database. It uses Oracle Net Services to connect to the database.

**8.3.1.1.1 Runtime Processes** Oracle Internet Directory has the following runtime processes:

- **OIDLDAPD:** This is the main process for Oracle Internet Directory. OIDLDAPD consists of a dispatcher process and a server process. The dispatcher process spawns the OIDLDAPD server processes during startup. Each OIDLDAPD dispatcher process has its own SSL and non-SSL ports for receiving requests. Every OID instance has one dispatcher and one server process by default. The number of server processes spawned for an instance is controlled by the `orclserverprocs` attribute.
- **OIDMON:** OIDMON is responsible for the process control of an Oracle Internet Directory instance. This process starts, stops, and monitors Oracle Internet Directory. During startup OIDMON spawns the OIDLDAPD dispatcher process and the replication server process, if replication is configured for the instance.
- **Replication server process:** This is a process within Oracle Internet Directory that runs only when replication is configured. The replication server process is spawned by OIDMON during startup.
- **OPMN:** The Oracle Process Manager and Notification Server (OPMN) is a daemon process that monitors Oracle Fusion Middleware components, including Oracle Internet Directory. Oracle Enterprise Manager Fusion Middleware Control uses OPMN to stop or start instances of Oracle Internet Directory. If you stop or start Oracle Internet Directory components from the command line, you use `opmnctl`, the command-line interface to OPMN.

OPMN is responsible for the direct start, stop, restart and monitoring of OIDMON. It does not start or stop the server process directly.

**8.3.1.1.2 Process Lifecycle** OPMN is responsible for the direct start, stop, restart and monitoring of the daemon process, OIDMON (`ORACLE_HOME/bin/oidmon`). OIDMON is responsible for the process control of an Oracle Internet Directory instance. In 11g Release 1 (11.1.1), you can have multiple instances of Oracle Internet Directory on the same Oracle instance on the same node. For details, refer to *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

### Process Status Table

Oracle Internet Directory process information is maintained in the `ODS_PROCESS_STATUS` table in the ODS database user schema. OIDMON reads the contents of the table at a specified interval and acts upon the intent conveyed by the contents of that table. The interval is controlled by the value of the sleep command line argument used at OIDMON startup, and the default value is 10 seconds.

### Starting and Stopping Oracle Internet Directory

An Oracle Internet Directory instance can be started and stopped using the Oracle Enterprise Manager Fusion Middleware Control or the command `opmnctl`.

#### Start Process

The start process for Oracle Internet Directory is:

1. Upon receiving the start command, OPMN issues an `oidmon start` command with appropriate arguments, as specified in the `opmn.xml` file.
2. OIDMON then starts all Oracle Internet Directory Server instances whose information in the `ODS_PROCESS_STATUS` table has state value 1 or 4 and `ORACLE_INSTANCE`, `COMPONENT_NAME`, `INSTANCE_NAME` values matching the environment parameters set by OPMN.

#### Stop Process

The stop process for Oracle Internet Directory is:

1. Upon receiving the stop command, OPMN issues an oidmon stop command.
2. For each row in the ODS\_PROCESS\_STATUS table that matches the environment parameters ORACLE\_INSTANCE, COMPONENT\_NAME, and INSTANCE\_NAME, the oidmon stop command kills OIDMON, OIDLDAPD, and OIDREPLD processes and updates the state to 4.

### Monitoring

OPMN does not monitor server processes directly. OPMN monitors OIDMON and OIDMON monitors the server processes. The events are:

- When you start OIDMON through OPMN, OPMN starts OIDMON and ensures that OIDMON is up and running.
- If OIDMON goes down for some reason, OPMN brings it back up.
- OIDMON monitors the status of the Oracle Internet Directory dispatcher process, LDAP server processes, and replication server process and makes this status available to OPMN and Oracle Enterprise Manager Fusion Middleware Control.

**8.3.1.1.3 Request Flow** Once the Oracle Internet Directory process starts up, clients access Oracle Internet Directory using the LDAP or LDAPS protocol. There is no impact on other running instances when an Oracle Internet Directory instance starts up.

Oracle Internet Directory listener/dispatcher starts a configured number of server processes at startup time. The number of server processes is controlled by the `orclserverprocs` attribute in the instance-specific configuration entry. The default value for `orclserverprocs` is 1. Multiple server processes enable Oracle Internet Directory to take advantage of multiple processor systems.

The Oracle Internet Directory dispatcher process sends the LDAP connections to the Oracle Internet Directory server process in a round robin fashion. The maximum number of LDAP connections accepted by each server is 1024 by default. This number can be increased by changing the attribute `orclmaxldapconns` in the instance-specific configuration entry, which has a DN of the form:

```
cn=componentname,cn=osdldapd,cn=subconfigsubentry
```

Database connections from each server process are spawned at server startup time, depending on the value set for the instance configuration parameters `ORCLMAXCC` and `ORCLPLUGINWORKERS`. The number of database connections spawned by each server equals `ORCLMAXCC + ORCLPLUGINWORKERS + 2`. The Oracle Internet Directory server processes communicate with the Oracle database server through Oracle Net Services. An Oracle Net Services listener/dispatcher relays the request to the Oracle database. For more information, refer to *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

**8.3.1.1.4 Configuration Artifacts** The storage location requires a DB connect string. `TNSNAMES.ORA` is stored in `ORACLE_INSTANCE/config`. The wallet is stored in `ORACLE_INSTANCE/OID/admin` (The DB ODS user password is stored in the wallet).

**8.3.1.1.5 External Dependencies** Oracle Internet Directory uses an Oracle database to store configuration information as well as data. It uses the ODS schema to store this information.

The Oracle directory replication server uses LDAP to communicate with an Oracle directory (LDAP) server instance. To communicate with the database, all components

use OCI/Oracle Net Services. Oracle Directory Services Manager and the command-line tools communicate with the Oracle directory servers over LDAP.

**8.3.1.1.6 Oracle Internet Directory Log File** Log files for Oracle Internet Directory are under the following directory:

`ORACLE_INSTANCE/diagnostics/log/OID`

Table 8–5 shows Oracle Internet Directory processes and the log file name and location for the process.

**Table 8–5 Locations of Oracle Internet Directory Process Log Files**

| Process                                 | Log File Location                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Directory server (oidldapd)             | <p><code>ORACLE_INSTANCE/diagnostics/logs/OID/componentName/oidldapd00sPID-XXXX.log</code> where:</p> <p>00 is the instance number (00 by default)</p> <p>s stands for server</p> <p>PID is the server process identifier</p> <p>XXXX is a number from 0000 to <code>orclmaxlogfilesconfigured</code>. Once the <code>orclmaxlogfilesconfigured</code> value is reached, it starts over again from 0000. When it starts over, it truncates the file to 0 bytes.</p> <p><code>ORACLE_INSTANCE/diagnostics/logs/OID/componentName/oidstackInstNumberPID.log</code></p> |
| LDAP dispatcher (oidldapd)              | <p><code>ORACLE_INSTANCE/diagnostics/logs/OID/componentName/oidldapd00-XXXX.log</code> where:</p> <p>00 is the instance number (00 by default)</p> <p>XXXX is a number from 0000 to <code>orclmaxlogfilesconfigured</code></p>                                                                                                                                                                                                                                                                                                                                       |
| OID Monitor (OIDMON)                    | <p><code>ORACLE_INSTANCE/diagnostics/logs/OID/componentName/oidmon-XXXX.log</code> where:</p> <p>XXXX is a number from 0000 to <code>orclmaxlogfilesconfigured</code></p>                                                                                                                                                                                                                                                                                                                                                                                            |
| Directory replication server (oidrepld) | <p><code>ORACLE_INSTANCE/diagnostics/logs/OID/componentName/oidrepld-XXXX.log</code> where:</p> <p>XXXX is a number from 0000 to <code>orclmaxlogfilesconfigured</code></p>                                                                                                                                                                                                                                                                                                                                                                                          |

For more information on using the log files to troubleshoot Oracle Internet Directory issues, see [Section 8.3.6, "Troubleshooting Oracle Internet Directory High Availability"](#).

## 8.3.2 Oracle Internet Directory High Availability Concepts

This section provides conceptual information about using Oracle Internet Directory in a high availability two-node Cluster Configuration. See [Section 8.3.2.3, "Oracle Internet Directory Prerequisites"](#) for prerequisites and [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps"](#) for specific steps for setting up the two-node Cluster Configuration.

### 8.3.2.1 Oracle Internet Directory High Availability Architecture

Figure 8–3 shows the Oracle Internet Directory Cluster Configuration high availability architecture in an active-active configuration.


**Figure 8–3 Oracle Internet Directory Cluster Configuration High Availability Architecture**

Figure 8–3 shows Oracle Internet Directory in the directory tier in a Cluster Configuration high availability architecture. Clustering is set up at installation time. The load balancing router routes LDAP client requests to the two Oracle Internet Directory instances that are clustered on OIDHOST1 and OIDHOST2.

Transparent Application Failover (TAF) is used to connect the Oracle Internet Directory instances with the Oracle RAC database that serves as the security metadata repository. The Oracle RAC database is configured in TNSNAMES.ORA. High availability event notification is used for notification when an Oracle RAC instance becomes unavailable. See [Section 4.1.7.1, "Oracle Internet Directory"](#) for more information about using Oracle Internet Directory with Oracle RAC.

**8.3.2.1.1 Starting and Stopping the Cluster** In the Cluster Configuration, OPMN commands are used to start each Oracle Internet Directory instance. There is no impact to Oracle Internet Directory at startup. A new database connection is spawned when Oracle Internet Directory starts.

When the cluster is stopped using OPMN, Oracle Internet Directory disconnects from the database and the Oracle Internet Directory server stops.

**8.3.2.1.2 Cluster-Wide Configuration Changes** Configuration changes can be done at a cluster level to any instance in the Cluster Configuration. All the nodes in the Cluster Configuration that share the same database read the same configuration information. The OIDMON process polls for configuration changes on each Oracle Internet Directory server and updates the database repository about configuration changes. OIDMON and other Oracle Internet Directory servers pull the changes from the database repository. In this way, any change made at a cluster member level is propagated to every Oracle Internet Directory server in the cluster.

The instance-specific configuration attributes for an Oracle Internet Directory LDAP server configuration are stored in this LDAP entry:

```
cn=<component-name>,cn=configsets,cn=osldlapd,cn=subconfigsubentry
```

Oracle Internet Directory server configuration aspects such as the number of servers, database connections, sizelimit, and timelimit are part of the instance-specific server configuration entry.

The configuration attributes that are common to all Oracle Internet Directory instances in a cluster are stored in the LDAP entry:



```
cn=dsaconfig,cn=configsets,cn=osldlapd,cn=oracle internet directory
```

If you want to retain instance-specific server configuration attributes for each Oracle Internet Directory instance in the cluster, then you should choose a distinct Oracle Internet Directory component name for each Oracle Internet Directory instance at install/configuration time; for example, oid1 on node1 and oid2 on node2. In this case, the configuration entries will be cn=oid1,cn=osldlapd,cn=subconfigsubentry and cn=oid2,cn=osldlapd,cn=subconfigsubentry respectively and they need to be updated separately for each Oracle Internet Directory instance.

On the other hand, if you choose to have a common set of server configuration attributes for both Oracle Internet Directory instances in the cluster, then you should choose the same Oracle Internet Directory component name for both Oracle Internet Directory instances, for example, oid1 on both Oracle Internet Directory node1 and node2. In this case, there will be one common configuration entry cn=oid1,cn=osldlapd,cn=subconfigsubentry.

Oracle Internet Directory LDAP server instances cache certain LDAP metadata artifacts such as Schema, ACLs, and Password Policy. Multiple Oracle Internet Directory LDAP server processes on a given node keep their caches in sync via semantics built around a shared memory segment managed by Oracle Internet Directory on each node. OIDMON keeps these caches in sync across nodes by ensuring that these shared memory segments are in sync across the nodes, which is achieved using the Oracle Internet Directory database.

Oracle Internet Directory also caches metadata and metadata changes trigger notification across the nodes. The ldapmodify utility is used to change metadata. The Oracle Internet Directory server that gets the ldapmodify request for the metadata change notifies other Oracle Internet Directory servers about the change of metadata (including OIDMON). OIDMON is responsible for notifying OIDMON on other nodes about the metadata changes.

### 8.3.2.2 Protection from Failures and Expected Behavior

This section discusses protection from different types of failure in an Oracle Internet Directory Cluster Configuration.

**8.3.2.2.1 Oracle Internet Directory Process Failure** OIDMON monitors Oracle Internet Directory processes. If the Oracle Internet Directory process goes down, OIDMON tries to restart it.

OPMN monitors OIDMON. If OIDMON goes down, OPMN restarts OIDMON.

If an Oracle Internet Directory process cannot be restarted, then the front-ending load balancing router detects failure of Oracle Internet Directory instances in the Cluster Configuration and routes LDAP traffic to surviving instances. In case of failure, the LDAP client retries the transaction. If the instance fails in the middle of a transaction, the transaction is not committed to the database. When the failed instance comes up again, the load balancing router detects this and routes requests to all the instances.

If an Oracle Internet Directory instance in the Cluster Configuration gets hung, the load balancing router detects this and routes requests to surviving instances.

If one of the Oracle Internet Directory instances in the two-node Cluster Configuration fails (or if one of the computers hosting an instance fails), the load balancing router routes clients to the surviving Oracle Internet Directory instance.

**8.3.2.2.2 Expected Client Application Behavior When Failure Occurs** Oracle Internet Directory server failure is usually transparent to Oracle Internet Directory clients as

they continue to get routed through the load balancer. External load balancers are typically configured to perform a health check of Oracle Internet Directory processes. If a request is received before the load balancer detects process unavailability, clients application could receive a error. If the client application performs a retry, the load balancer should route it to a healthy Oracle Internet Directory instance and the request should be successful.

In Oracle Internet Directory active-active configurations, if you are doing ldapadd operations through the LDIF file at the time of failover, your operation would fail even if you are doing this operation through a load balancer host and port. This is because Oracle Internet Directory is down for a fraction of a second. For most applications, this will not be an issue because most applications have the ability to retry the connection a fixed number of times.

**8.3.2.2.3 External Dependency Failure** This section describes the protection available for Oracle Internet Directory from database failure.

By default, the `tnsnames.ora` file configured in Oracle Internet Directory's `ORACLE_INSTANCE` ensures that Oracle Internet Directory's connections to the database are load balanced between the Oracle RAC database instances. For example, if an Oracle Internet Directory instance establishes four database connections, two connections are made to each database instance.

Oracle Internet Directory uses database high availability event notification to detect database node failure and to fail over to a surviving node.

If Transparent Application Failover (TAF) is configured, then upon a database instance failure, Oracle Internet Directory will fail over its database connections to the surviving database instance, which enables the LDAP search operations that were in progress during the failover to be continued.

If both Transparent Application Failover (TAF) and high availability event notification are configured, Transparent Application Failover (TAF) is used for failover and high availability event notifications are used only for logging the events. The high availability event notifications are logged in `OIDLDAPD` log file.

Oracle Internet Directory also has a mechanism to detect stale database connections, which allows Oracle Internet Directory to reconnect to the database.

If none of the database instances are available for a prolonged period, then the Oracle Internet Directory LDAP and REPL processes will automatically be shut down. However, `OIDMON` and `OPMN` will continue to ping for the database instance availability and when the database becomes available, the Oracle Internet Directory processes (LDAP and REPL) are automatically restarted by `OIDMON`.

While all the database instances are down, `OIDMON` will continue to be up and an `opmnctl status` command will show that `OIDLDAPD` instances are down. When a database instance becomes available, `OIDMON` will restart all configured Oracle Internet Directory instances.

All database failover induced activity for Oracle Internet Directory is recorded in the `OIDMON` log file.

### **8.3.2.3 Oracle Internet Directory Prerequisites**

This section describes prerequisites for setting up the Oracle Internet Directory high availability architecture.

**8.3.2.3.1 Synchronizing the Time on Oracle Internet Directory Nodes** Before setting up Oracle Internet Directory in a high availability environment, you must ensure that the time on the individual Oracle Internet Directory nodes is synchronized.

Synchronize the time on all nodes using Greenwich Mean Time so that there is a discrepancy of no more than 250 seconds between them.

If OID Monitor detects a time discrepancy of more than 250 seconds between the two nodes, the OID Monitor on the node that is behind stops all servers on its node. To correct this problem, synchronize the time on the node that is behind in time. The OID Monitor automatically detects the change in the system time and starts the Oracle Internet Directory servers on its node.

If there are more than two nodes, the same behavior is followed. For example, assume that there are three nodes, where the first node is 150 seconds ahead of the second node, and the second node is 150 seconds ahead of the third node. In this case, the third node is 300 seconds behind the first node, so the OID Monitor will not start the servers on the third node until the time is synchronized.

**8.3.2.3.2 Using RCU to Create Oracle Internet Directory Schemas in the Repository** Before you can install the Oracle Internet Directory instances on OIDHOST1 and OIDHOST2, you must use the latest version of the Repository Creation Utility (RCU) to create the collection of schemas used by Oracle Identity Management and Management Services.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

Follow these steps to run RCU and create the Identity Management schemas in an Oracle RAC database repository:

1. Issue this command:

```
prompt> RCU_HOME/bin/rcu &
```

2. On the Welcome screen, click **Next**.
3. On the Create Repository screen, select the **Create** operation to load component schemas into an existing database.  
Click **Next**.
4. On the Database Connection Details screen, enter connection information for the existing database as follows:

**Database Type:** Oracle Database

**Host Name:** Name of the computer on which the database is running. For an Oracle RAC database, specify the VIP name or one node name. Example:  
INFRADBHOST1-VIP or INFRADBHOST2-VIP

**Port:** The port number for the database. Example: 1521

**Service Name:** The service name of the database. Example: oid.mycompany.com

**Username:** SYS

**Password:** The SYS user password

**Role:** SYSDBA

Click **Next**.

5. On the Select Components screen, create a new prefix and select the components to be associated with this deployment:

**Create a New Prefix:** `idm` (Entering a prefix is optional if you select only **Identity Management** (Oracle Internet Directory - ODS) in the **Components** field)

**Components:** Select **Identity Management** (Oracle Internet Directory - ODS). De-select all other schemas.

Click **Next**.

6. On the Schema Passwords screen, enter the passwords for the main and additional (auxiliary) schema users.

Click **Next**.

7. On the Map Tablespaces screen, select the tablespaces for the components.

8. On the Summary screen, click **Create**.

9. On the Completion Summary screen, click **Close**.

**8.3.2.3.3 Load Balancer Virtual Server Names for Oracle Internet Directory** When Oracle Internet Directory is deployed in a high availability configuration, it is recommended to use an external load balancer to front-end the Oracle Internet Directory instances and load balance requests between the various Oracle Internet Directory instances.

Refer to [Section 8.2.5.4, "Configuring Virtual Server Names and Ports for the Load Balancer"](#) for details.

### 8.3.3 Oracle Internet Directory High Availability Configuration Steps

Oracle Internet Directory can be deployed in a highly availability configuration either in a standalone mode or as a part of a WebLogic Server domain.

The standalone mode should be chosen for deployments where Oracle Internet Directory will be managed entirely using command line tools, and where Oracle Enterprise Manager Fusion Middleware Control and Oracle Directory Services Manager are not deemed mandatory. Later, if desired, the standalone Oracle Internet Directory instances can be registered with a remote Oracle WebLogic Server domain using `opmnctl` commands. Oracle Enterprise Manager Fusion Middleware Control and Oracle Directory Services Manager can be used to manage Oracle Internet Directory instances that are configured with an Oracle WebLogic Server domain.

It is recommended that Oracle Internet Directory be set up in a clustered deployment, where the clustered Oracle Internet Directory instances access the same Oracle RAC database repository.

#### 8.3.3.1 Installing Oracle Fusion Middleware Components

This section describes how to install the required binaries for the Oracle WebLogic Server (`WL_HOME`) and Oracle Home for (`ORACLE_HOME`) for Oracle Identity Management.

Oracle strongly recommends that you read the release notes for any additional installation and deployment considerations prior to starting the setup process.

**8.3.3.1.1 Install Oracle WebLogic Server** The first step in the installation procedure is to install Oracle WebLogic Server.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

Start the Oracle WebLogic Server installer.

Then follow these steps in the installer to install Oracle WebLogic Server on the computer:

1. On the Welcome screen, click **Next**.
2. On the Choose Middleware Home Directory screen, select **Create a New Middleware Home**. Then choose a directory on your computer into which the Oracle WebLogic software is to be installed.

For the **Middleware Home Directory**, specify this value:

```
/u01/app/oracleproduct/fmw
```

Click **Next**.

3. On the Register for Security Updates screen, enter your "My Oracle Support" User Name and Password so that you can be notified of security updates
4. On the Choose Install Type screen, the installation program prompts you to indicate whether you wish to perform a complete or a custom installation.

Choose **Custom**.

Click **Next**.

5. On the Choose Products and Components screen, select only **Oracle JRockit SDK** and click **Next**.
6. On the Choose Product Installation Directories screen, accept the directory `/u01/app/oracle/fmw/wlserver_10.3`.

Click **Next**.

7. On the Installation Summary screen, the window contains a list of the components you selected for installation, along with the approximate amount of disk space to be used by the selected components once installation is complete.

Click **Next**.

8. On the Installation Complete screen, deselect the **Run Quickstart** checkbox.

Click **Done**.

**8.3.3.1.2 Installing Oracle Fusion Middleware for Identity Management** The next step is to install the Oracle Fusion Middleware components.

---

**Note:** Because the installation is performed on shared storage, the two `MW_HOME` installations are accessible and used by the remaining servers in the topology.

---

Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.

On Linux platforms, if the `/etc/oraInst.loc` file exists, verify that its contents are correct. Specifically, check that the inventory directory is correct and that you have

write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

Start the installer for Oracle Fusion Middleware components.

On UNIX, issue the command:

```
HOST1> runInstaller
```

On Windows, double-click `setup.exe`.

Before starting the install, ensure that the following environment variables are not set:

- `LD_ASSUME_KERNEL`
- `ORACLE_INSTANCE`

On the Specify Inventory Directory screen, do the following:

- Enter `HOME/oraInventory`, where *HOME* is the home directory of the user performing the installation (this is the recommended location).
- Enter the OS group for the user performing the installation.
- Click **Next**.

For a UNIX install, follow the instructions on screen to execute `createCentralInventory.sh` as root.

Click **OK**.

Proceed as follows:

1. On the Specify Oracle Inventory Directory screen, enter `HOME/oraInventory`, where *HOME* is the home directory of the user performing the installation (This is the recommended location).

Enter the OS group for the user performing the installation.

Click **Next**.

2. On the Welcome screen, click **Next**.
3. On the Select Installation Type screen, select **Install-Do Not Configure**.

Click **Next**.

4. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.

5. On the Specify Installation Location screen, enter the following values:

**MW\_HOME:** Enter the value of the `MW_HOME`, for example:

```
/u01/app/product/fmw
```

Select the previously installed Middleware Home from the drop-down list. For the Oracle Home directory, enter the directory name `IDM`.

Click **Next**.

6. On the Summary screen, click **Install**.

When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the root user.

7. On the Installation Complete screen, click **Finish**.

**8.3.3.1.3 Upgrading Oracle Identity Management** This section provides the steps to upgrade your Oracle Identity Management software. For more information, see *Applying the Latest Oracle Fusion Middleware Patch Set* in the *Oracle Fusion Middleware Patching Guide*.

1. Start the IDM Upgrade Installer by running `./runinstaller`.
2. On the Welcome screen, click **Next**.
3. On the Prerequisite Checks screen, click **Next**.
4. On the Specify Install Location screen, provide the path to the Oracle Middleware Home and the name of the Oracle Home directory.
5. On the Installation Summary screen, validate your selections, and then click **Install**.
6. The Installation Progress screen shows the progress of the install.  
Once the installation is done, the oracleRoot.sh confirmation dialog box appears. This dialog box advises you that a configuration script needs to be run as root before the installation can proceed. Leaving the confirmation dialog box open, open another shell window, log in as root, and run this script file:  
`/u01/app/oracle/product/fmw/id/oracleRoot.sh`. After the script completes, click **OK** on the Confirmation Dialog box.
7. On the Installation Complete screen, click **Finish** to exit.

### 8.3.3.2 Configuring Oracle Internet Directory Without a WebLogic Domain

This section describes the steps to deploy Oracle Internet Directory without a WebLogic Server domain.

**8.3.3.2.1 Configuring Oracle Internet Directory on OIDHOST1** Ensure that the schema database is running and that RCU has been used to seed the ODS database schema, then follow these steps to configure the Oracle Internet Directory instance on OIDHOST1:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OIDHOST1 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Ensure that ports 389 and 636 are not in use by any service on the computer by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":389"
```

```
netstat -an | grep LISTEN | grep ":636"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":389"
```

```
netstat -an | findstr "LISTEN" | findstr ":636"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 389 and 636 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Internet Directory):

```
The Non-SSL port for OID
Oracle Internet Directory Port No = 389
The SSL port for OID
Oracle Internet Directory (SSL) Port No = 636
```

7. Start the Oracle Identity Management 11g Configuration Assistant under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Configure without a Domain** and then click **Next**.
10. On the Specify Installation Location screen, specify the following values:

- **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.

- **Oracle Home Directory:**

`/u01/app/oracle/product/fmw/idm`

- **Oracle Instance Location:**

`/u01/app/oracle/admin/oid_instance1`

- **Oracle Instance Name:**

`oid_instance1`

---

**Note:** Ensure that the Oracle Home Location directory path for `OIDHOST1` is the same as the Oracle Home Location path for `OIDHOST2`. For example, if the Oracle Home Location directory path for `OIDHOST1` is:

`/u01/app/oracle/product/fmw/idm`

then the Oracle Home Location directory path for `OIDHOST2` must be:

`/u01/app/oracle/product/fmw/idm`

---



Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Internet Directory**, deselect all the other components, and click **Next**.
13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

14. On the Specify Schema Database screen, select **Use Existing Schema** and specify the following values:

- **Connect String:**

```
infradbhost1-vip.mycompany.com:1521:oiddb1^infradbhost2-vip.mycompany.com:1521:oiddb2@oid.mycompany.com
```

---

---

**Note:** The Oracle RAC database connect string information needs to be provided in the format  
*host1:port1:instance1^host2:port2:instance2@servicename.*

During this installation, it is not required for all the Oracle RAC instances to be up. If one Oracle RAC instance is up, the installation can proceed.

It is required that the information provided above is complete and accurate. Specifically, the correct host, port, and instance name must be provided for each Oracle RAC instance, and the service name provided must be configured for all the specified Oracle RAC instances.

Any incorrect information entered in the Oracle RAC database connect string has to be corrected manually after the installation.

---

---

- **User Name:** ODS
- **Password:** \*\*\*\*\*

Click **Next**.

15. On the Create Oracle Internet Directory screen, specify the Realm and enter the Administrator (`cn=orcladmin`) password and click **Next**.
16. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens) and click **Install**.

17. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

18. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
19. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

**8.3.3.2.2 Oracle Internet Directory Component Names Assigned by Oracle Identity Management Installer** When you perform an Oracle Internet Directory installation using Oracle Identity Management 11g installer, the default component name that the installer assigns to the Oracle Internet Directory instance is `oid1`. You cannot change this component name.

The instance-specific configuration entry for this Oracle Internet Directory instance is `cn=oid1, cn=osldapd, cn=subconfigsubentry`.

If you perform a second Oracle Internet Directory installation on another computer and that Oracle Internet Directory instance uses the same database as the first instance, the installer detects the previously installed Oracle Internet Directory instance on the other computer using the same Oracle database, so it gives the second Oracle Internet Directory instance a component name of `oid2`.

The instance-specific configuration entry for the second Oracle Internet Directory instance is `cn=oid2, cn=osldapd, cn=subconfigsubentry`. Any change of properties in the entry `cn=oid2, cn=osldapd, cn=subconfigsubentry` will not affect the first instance (`oid1`).

If a third Oracle Internet Directory installation is performed on another computer and that instance uses the same database as the first two instances, the installer gives the third Oracle Internet Directory instance a component name of `oid3`, and so on for additional instances on different hosts that use the same database.

Note that the shared configuration for all Oracle Internet Directory instances is `cn=dsaconfig, cn=configsets, cn=oracle internet directory`. Any change in this entry will affect all the instances of Oracle Internet Directory.

This naming scheme helps alleviate confusion when you view your domain using Oracle Enterprise Manager by giving different component names to your Oracle Internet Directory instances.

**8.3.3.2.3 Configuring Oracle Internet Directory on `OIDHOST2`** Ensure that the Oracle Internet Directory repository is running and then follow these steps to configure the Oracle Internet Directory instance on `OIDHOST2`:

---

**Note:** The instructions in this section can also be used to scale out Oracle Internet Directory in your 11g Oracle Identity Management high availability configuration.

---

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.

2. Ensure that Oracle Identity Management software has been installed and upgraded on OIDHOST2 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. On OIDHOST1, ports 389 and 636 were used for Oracle Internet Directory. The same ports should be used for the Oracle Internet Directory instance on OIDHOST2. Therefore, ensure that ports 389 and 636 are not in use by any service on OIDHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":389"
```

```
netstat -an | grep LISTEN | grep ":636"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":389"
```

```
netstat -an | findstr "LISTEN" | findstr ":636"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 389 and 636 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Internet Directory):

```
The Non-SSL port for OID
Oracle Internet Directory Port No = 389
The SSL port for OID
Oracle Internet Directory (SSL) Port No = 636
```

7. Start the Oracle Identity Management 11g Configuration Assistant under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: `./config.sh`

On Windows, double-click `config.exe`

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Configure without a Domain** and then click **Next**.
10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.
  - **Oracle Home Directory:**

```
/u01/app/oracle/product/fmw/idm
```

- **Oracle Instance Location:**

/u01/app/oracle/admin/oid\_instance2

- **Oracle Instance Name:**

oid\_instance2

---

---

**Note:** Ensure that the Oracle Home Location directory path for OIDHOST2 is the same as the Oracle Home Location directory path for OIDHOST1. For example, if the Oracle Home Location directory path for OIDHOST1 is:

/u01/app/oracle/product/fmw/idm

then the Oracle Home Location directory path for OIDHOST2 must be:

/u01/app/oracle/product/fmw/idm

---

---

Click **Next**.

**11.** On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

**12.** On the Configure Components screen, select **Oracle Internet Directory**, deselect all the other components, and click **Next**.

**13.** On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

**14.** On the Specify Schema Database screen, select **Use Existing Schema** and specify the following values.

- **Connect String:**

infradbhost1-vip.mycompany.com:1521:oiddb1^infradbhost2-vip.mycompany.com:1521:oiddb2@oid.mycompany.com

---

**Note:** The Oracle RAC database connect string information needs to be provided in the format

*host1:port1:instance1^host2:port2:instance2@servicename.*

During this installation, it is not required for all the Oracle RAC instances to be up. If one Oracle RAC instance is up, the installation can proceed.

It is required that the information provided above is complete and accurate. Specifically, the correct host, port, and instance name must be provided for each Oracle RAC instance, and the service name provided must be configured for all the specified Oracle RAC instances.

Any incorrect information entered in the Oracle RAC database connect string has to be corrected manually after the installation.

---

- User Name: ODS
- Password: \*\*\*\*\*

Click **Next**.

15. The ODS Schema in use message appears. The ODS schema chosen is already being used by the existing Oracle Internet Directory instance. Therefore, the new Oracle Internet Directory instance being configured would reuse the same schema.

Choose **Yes** to continue.

16. On the Specify OID Administrator Password screen, specify the OID Administrator password and click **Next**.
17. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.
18. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

19. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
20. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

**8.3.3.2.4 Registering Oracle Internet Directory with a WebLogic Domain** If you want to manage an Oracle Internet Directory component with Oracle Enterprise Manager Fusion Middleware Control, you must register the component and the Oracle Fusion Middleware instance that contains it with an Oracle WebLogic Server domain. You can register an Oracle Fusion Middleware instance with a WebLogic domain during installation or Oracle instance creation, but you are not required to do so. If an Oracle Fusion Middleware instance was not previously registered with a WebLogic domain, you can register it by using `opmnctl registerinstance`.

Before using the `opmnctl registerinstance` command to register an Oracle Internet Directory instance with an Oracle WebLogic Server domain, ensure that the WebLogic Server is already installed.

Then execute the `opmnctl registerinstance` command in this format (note that the `ORACLE_HOME` and `ORACLE_INSTANCE` variables have to be set for each installation before executing this command in the home directory for the Oracle Internet Directory instance):

```
opmnctl registerinstance -adminHost WLSTHostName -adminPort WLSPort
-adminUsername adminUserName
```

For example:

```
opmnctl registerinstance -adminHost idmhost1.mycompany.com -adminPort 7001
-adminUsername weblogic
```

```
Command requires login to weblogic admin server (idmhost1.mycompany.com)
Username: weblogic
Password: *****
```

For additional details on registering Oracle Internet Directory components with a WebLogic domain, see the "Registering an Oracle Fusion Middleware Instance or Component with the WebLogic Server" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

### 8.3.3.3 Configuring Oracle Internet Directory With a WebLogic Domain

This section describes the steps to deploy Oracle Internet Directory in a high availability configuration as part of a WebLogic Server domain.

In this configuration, Oracle Internet Directory and a WebLogic Server domain is configured on the first host, and only Oracle Internet Directory is configured on the second host. The Oracle Internet Directory instance on the second host joins the domain created on the first host.

**8.3.3.3.1 Configuring Oracle Internet Directory on OIDHOST1** Ensure that the schema database is running and that RCU has been used to seed the ODS database schema, then follow these steps to configure the Oracle Internet Directory instance on `OIDHOST1`:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on `OIDHOST1` as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Ensure that ports 389 and 636 are not in use by any service on the computer by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":389"
```

```
netstat -an | grep LISTEN | grep ":636"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":389"
```

```
netstat -an | findstr "LISTEN" | findstr ":636"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 389 and 636 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Internet Directory):

```
The Non-SSL port for OID
Oracle Internet Directory Port No = 389
The SSL port for OID
Oracle Internet Directory (SSL) Port No = 636
```

7. Start the Oracle Identity Management 11g Configuration Assistant under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Create New Domain**. Specify the values for.
  - **User Name:** weblogic
  - **Password:** <password for the weblogic user>
  - **Confirm Password:** <confirm the password for the weblogic user>
  - **Domain Name:** IDMDomain
10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.
  - **Oracle Home Directory:**

```
/u01/app/oracle/product/fmw/idm
```
  - **WebLogic Server Directory:**

```
/u01/app/oracle/product/fmw/wlserver_10.3
```
  - **Oracle Instance Location:**

```
/u01/app/oracle/admin/oid_inst1
```
  - **Oracle Instance Name:**

```
oid_inst1
```

---

---

**Note:** Ensure that the Oracle Home Location directory path for OIDHOST1 is the same as the Oracle Home Location path for OIDHOST2. For example, if the Oracle Home Location directory path for OIDHOST1 is:

`/u01/app/oracle/product/fmw/idm`

then the Oracle Home Location directory path for OIDHOST2 must be:

`/u01/app/oracle/product/fmw/idm`

---

---

Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Internet Directory**, deselect all the other components, and click **Next**.

13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

---

---

**Note:** The default Oracle WebLogic Server clustering mode set by the installer is unicast (not multicast).

---

---

14. On the Specify Schema Database screen, select **Use Existing Schema** and specify the following values:

- **Connect String:**  
`infradbhost1-vip.mycompany.com:1521:oiddb1^infradbhost2-vip.mycompany.com:1521:oiddb2@oid.mycompany.com`



---

**Note:** The Oracle RAC database connect string information needs to be provided in the format

*host1:port1:instance1^host2:port2:instance2@servicename.*

During this installation, it is not required for all the Oracle RAC instances to be up. If one Oracle RAC instance is up, the installation can proceed.

It is required that the information provided above is complete and accurate. Specifically, the correct host, port, and instance name must be provided for each Oracle RAC instance, and the service name provided must be configured for all the specified Oracle RAC instances.

Any incorrect information entered in the Oracle RAC database connect string has to be corrected manually after the installation.

---

- User Name: ODS
- Password: \*\*\*\*\*

Click **Next**.

15. On the Configure OID screen, specify the Realm and enter the Administrator (cn=orcladmin) password and click **Next**.
16. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.
17. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.  
Click **Next**.
18. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
19. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

---

**Note:** For information on the Oracle Internet Directory component names assigned by Oracle Identity Management 11g Installer, refer to [Section 8.3.3.2.2, "Oracle Internet Directory Component Names Assigned by Oracle Identity Management Installer."](#)

---

**8.3.3.3.2 Creating boot.properties for the Administration Server on OIDHOST1** This section describes how to create a `boot.properties` file for the Administration Server on `OIDHOST1`. The `boot.properties` file enables the Administration Server to start without prompting for the administrator username and password.

Follow these steps to create the `boot.properties` file:

1. On `OIDHOST1`, go the following directory:

`MW_HOME/user_projects/domains/domainName/servers/AdminServer/security`

For example:

`cd /u01/app/oracle/product/fmw/user_`

```
projects/domains/IDMDomain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` under the security directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

3. Stop the Administration Server if it is running.  
See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.
4. Start the Administration Server on OIDHOST1 using the `startWebLogic.sh` script located under the `MW_HOME/user_projects/domains/domainName/bin` directory.
5. Validate that the changes were successful by opening a web browser and accessing the following pages:
  - WebLogic Server Administration Console at:  
`http://oidhost1.mycompany.com:7001/console`
  - Oracle Enterprise Manager Fusion Middleware Control at:  
`http://oidhost1.mycompany.com:7001/em`

Log into these consoles using the `weblogic` user credentials.

**8.3.3.3.3 Configuring Oracle Internet Directory on OIDHOST2** Ensure that the Oracle Internet Directory repository is running and then follow these steps to configure the Oracle Internet Directory instance on OIDHOST2:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OIDHOST2 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. On OIDHOST1, ports 389 and 636 were used for Oracle Internet Directory. The same ports should be used for the Oracle Internet Directory instance on OIDHOST2. Therefore, ensure that ports 389 and 636 are not in use by any service on OIDHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":389"
```

```
netstat -an | grep LISTEN | grep ":636"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":389"
```

```
netstat -an | findstr "LISTEN" | findstr ":636"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 389 and 636 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Internet Directory):

```
The Non-SSL port for OID
Oracle Internet Directory Port No = 389
The SSL port for OID
Oracle Internet Directory (SSL) Port No = 636
```

7. Start the Oracle Identity Management 11g Configuration Assistant under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Extend Existing Domain** and specify the following values.
  - **HostName:** `idmhost1.mycompany.com` (This is the host where the WebLogic Administration Server is running.)
  - **Port:** `7001` (This is the WebLogic Administration Server port)
  - **User Name:** `weblogic`
  - **Password:** `<password for the weblogic user>`
10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.
  - **Oracle Home Directory:** `idm`
  - **WebLogic Server Directory:**

```
/u01/app/oracle/product/fmw/wlserver_10.3
```
  - **Oracle Instance Location:**

```
/u01/app/oracle/admin/oid_inst2
```

- **Oracle Instance Name:** oid\_inst2

---

**Note:** Ensure that the Oracle Home Location directory path for OIDHOST1 is the same as the Oracle Home Location path for OIDHOST2. For example, if the Oracle Home Location directory path for OIDHOST1 is:

```
/u01/app/oracle/product/fmw/idm
```

then the Oracle Home Location directory path for OIDHOST2 must be:

```
/u01/app/oracle/product/fmw/idm
```

---

Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Internet Directory** and click **Next**.

13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

14. On the Specify Schema Database screen, select **Use Existing Schema** and specify the following values:

- **Connect String:**

```
infradbhost1-vip.mycompany.com:1521:oiddb1^infradbhost2-vip.mycompany.com:1521:oiddb2@oid.mycompany.com
```

---

**Note:** The Oracle RAC database connect string information needs to be provided in the format

*host1:port1:instance1^host2:port2:instance2@servicename.*

During this installation, it is not required for all the Oracle RAC instances to be up. If one Oracle RAC instance is up, the installation can proceed.

It is required that the information provided above is complete and accurate. Specifically, the correct host, port, and instance name must be provided for each Oracle RAC instance, and the service name provided must be configured for all the specified Oracle RAC instances.

Any incorrect information entered in the Oracle RAC database connect string has to be corrected manually after the installation.

---

- User Name: ODS
- Password: \*\*\*\*\*

Click **Next**.

15. The ODS Schema in use message appears. The ODS schema chosen is already being used by the existing Oracle Internet Directory instance. Therefore, the new Oracle Internet Directory instance being configured would reuse the same schema.

Choose **Yes** to continue.

16. On the Specify OID Admin Password screen, specify the Oracle Internet Directory Administrator password and click **Next**.
17. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.
18. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

19. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
20. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

### 8.3.4 Validating Oracle Internet Directory High Availability

Use the `ldapbind` command-line tool to ensure that you can connect to each Oracle Internet Directory instance and the LDAP Virtual Server. The `ldapbind` tool enables you to determine whether you can authenticate a client to a server.

---

**Note:** See the "Configuring Your Environment" section of *Oracle Fusion Middleware Reference for Oracle Identity Management* for a list of the environment variables you must set before using the `ldapbind` command.

---

For non-SSL:

```
ldapbind -h oidhost1.mycompany.com -p 389 -D "cn=orcladmin" -q
ldapbind -h oidhost2.mycompany.com -p 389 -D "cn=orcladmin" -q
ldapbind -h oid.mycompany.com -p 389 -D "cn=orcladmin" -q
```

---

**Note:** The -q option above prompts the user for a password. LDAP tools have been modified to disable the options -w *password* and -P *password* when the environment variable LDAP\_PASSWORD\_PROMPTONLY is set to TRUE or 1. Use this feature whenever possible.

---

For SSL:

```
ldapbind -h oidhost1.mycompany.com -p 636 -D "cn=orcladmin" -q -U 1
ldapbind -h oidhost2.mycompany.com -p 636 -D "cn=orcladmin" -q -U 1
ldapbind -h oid.mycompany.com -p 636 -D "cn=orcladmin" -q -U 1
```

where -U is an optional argument used to specify the SSL authentication mode. These are the valid values for the SSL authentication mode:

- 1 = No authentication required
- 2 = One way authentication required. With this option, you must also supply a wallet location (-W "file:/home/my\_dir/my\_wallet") and wallet password (-P *wallet\_password*).
- 3 = Two way authentication required. With this option, you must also supply a wallet location (-W "file:/home/my\_dir/my\_wallet") and wallet password (-P *wallet\_password*).

For more information about the ldapbind command, see the ldapbind section in *Oracle Fusion Middleware Reference for Oracle Identity Management*.

For information about setting up SSL for Oracle Internet Directory, see "Configuring Secure Sockets Layer (SSL)" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* manual.

WebLogic Server Administration Console:

<http://oidhost1.mycompany.com:7001/console>

Oracle Enterprise Manager Fusion Middleware Console:

<http://oidhost1.mycompany.com:7001/em>

## 8.3.5 Oracle Internet Directory Failover and Expected Behavior

This section includes steps for performing a failover of Oracle Internet Directory and for performing a failover of Oracle RAC.

### 8.3.5.1 Performing an Oracle Internet Directory Failover

Follow these steps to perform a failover of an Oracle Internet Directory instance and to check the status of Oracle Internet Directory:

1. On OIDHOST1, use the opmnctl command to stop the Oracle Internet Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=oid1
```

2. On OIDHOST2, check the status of Oracle Internet Directory using the load balancing router:

---

**Note:** See the "Configuring Your Environment" section of *Oracle Fusion Middleware Reference for Oracle Identity Management* for a list of the environment variables you must set before using the `ldapbind` command.

---

```
ldapbind -h oid.mycompany.com -p 389 -D "cn=orcladmin" -q
```

---

**Note:** The `-q` option above prompts the user for a password. LDAP tools have been modified to disable the options `-w password` and `-P password` when the environment variable `LDAP_PASSWORD_PROMPTONLY` is set to `TRUE` or `1`. Use this feature whenever possible.

---

3. On OIDHOST1, use the `opmnctl` command to start the Oracle Internet Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl start (if OPMN is not running)
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=oid1
```

4. On OIDHOST2, use the `opmnctl` command to stop the Oracle Internet Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=oid1
```

5. On OIDHOST1, check the status of Oracle Internet Directory using the load balancing router:

```
ldapbind -h oid.mycompany.com -p 389 -D "cn=orcladmin" -q
```

6. On OIDHOST2, use the `opmnctl` command to start the Oracle Internet Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl start (if OPMN is not running)
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=oid1
```

### 8.3.5.2 Performing an Oracle RAC Failover

Follow these steps to perform an Oracle RAC failover:

1. Use the `srvctl` command to stop a database instance:

```
srvctl stop instance -d db_unique_name -i inst_name_list
```

2. Use the `srvctl` command to check the status of the database:

```
srvctl status database -d db_unique_name -v
```

3. Check the status of Oracle Internet Directory:

---

**Note:** See the "Configuring Your Environment" section of *Oracle Fusion Middleware Reference for Oracle Identity Management* for a list of the environment variables you must set before using the `ldapbind` command.

---

```
ldapbind -h oid_host -p 389 -D "cn=orcladmin" -q
ldapbind -h oid_host -p 389 -D "cn=orcladmin" -q
ldapbind -h oid.mycompany.com -p 389 -D "cn=orcladmin" -q
```

---

**Note:** The `-q` option above prompts the user for a password. LDAP tools have been modified to disable the options `-w password` and `-P password` when the environment variable `LDAP_PASSWORD_PROMPTONLY` is set to `TRUE` or `1`. Use this feature whenever possible.

---

4. Use the `srvctl` command to start the database instance:

```
srvctl start instance -d db_unique_name -i inst_name_list
```

### 8.3.6 Troubleshooting Oracle Internet Directory High Availability

This section provides information that can help you troubleshoot Oracle Internet Directory high availability issues:

- Log files for Oracle Internet Directory are found in the following directory:

```
ORACLE_INSTANCE/diagnostics/log/OID
```

- The order in which log files should be examined when troubleshooting is:

1. `oidmon-xxx.log`
2. `oidldapd01-xxxx.log`
3. `oidldapd01s-xxxx.log`

- This section shows some of the error messages that may be related to high availability, and their meaning:

**Error:** ORA-3112, ORA-3113 errors in the log file

**Cause:** one of the database node is down, OID connects again to surviving node.

**Action:** See why database node went down or Oracle process got killed

**Error:** Failing Over...Please stand by in the log file

**Cause:** OID server received a notification from the Oracle process that one of the database node is down. OID will connect to the surviving node.

If the failover is successful you would see this message:

Failover ended...resuming services.

If the failover was not successful, you would see these errors:

- a. Tried 10 times, now quitting from failover function...
- b. Bad Failover Event:



c. Forcing Failover abort as setting of DB parameters for the session failed

If high availability event notification is enabled, you would see a message similar to the following:

```
HA Callback Event
Thread Id: 8
Event type: 0
HA Source: OCI_HA_INSTANCE
Host name: dbhost1
Database name: orcl
Instance name: orcl1
Timestamp: 14-MAY-09 03.25.24 PM -07:00
Service name: orcl.us.oracle.com
HA status: DOWN - TAF Capable
```

If TAF is disabled, HA status will be shown as "DOWN."

**Action:** See why database node went down.

**Error:** Time Difference of at least 250 sec found between node1 and node2.

**Cause:** There is time difference between the two nodes

**Action:** Synchronize the system time.

**Error:** Node=% did not respond for configured %d times, Failing over...

**Cause:** One of the OID nodes (oidmon) is not responding.

**Action:** See if the node is alive or OIDMON process is running.

For more information about troubleshooting Oracle Internet Directory, see *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

## 8.3.7 Additional Oracle Internet Directory High Availability Issues

This section describes issues for Oracle Internet Directory in a high availability environment.

### 8.3.7.1 Changing the Password of the ODS Schema Used by Oracle Internet Directory

You can change the Oracle Internet Directory database schema password (that is, the password of the ODS user in the database) using the Oracle Internet Directory Database Password Utility (`oidpasswd`) from any of the Oracle Internet Directory nodes. However, since the ODS schema password is stored in a password wallet under the `ORACLE_INSTANCE` of each OID instance, the password wallet must be updated in each Oracle Internet Directory node.

To change the ODS database user password, invoke the following command on one of the Oracle Internet Directory nodes:

```
oidpasswd connect=database-connection-string change_oiddb_pwd=true
```

On all other Oracle Internet Directory nodes, invoke the following command to synchronize the password wallet:

```
oidpasswd connect=database-connection-string create_wallet=true
```

If you change the ODS password on one Oracle RAC node by using the OID Database Password Utility (`oidpasswd`), then you must do one of the following to update the wallet `ORACLE_HOME/ldap/admin/oidpwd11ldap1` on the other Oracle RAC nodes:

- Invoke the OID Database Password Utility on all the other nodes to update the wallet file only. This applies to replication password changes also, but for replication password changes you would use the Replication Environment Management Tool to update the password instead of the OID Database Password Utility.
- Copy the changed wallet to the other nodes.

If you run the `oidpasswd` command on one node only, and do not update the wallet on all the Oracle RAC nodes, the Oracle Internet Directory instance on the second node will not be able to start on the other nodes. You will see this error in the `OIDMON` log file:

```
[gsdsiConnect] ORA-1017, ORA-01017: invalid username/password; logon denied.
```

As mentioned above, the fix is to copy the `oidpwdlldap1` file to the other Oracle RAC nodes, or to invoke the `oidpasswd` tool with the `create_wallet=true` option on the other nodes.

## 8.4 Oracle Virtual Directory High Availability

This section provides an introduction to Oracle Virtual Directory and describes how to design and deploy a high availability environment for Oracle Virtual Directory.


This section includes the following topics:

- [Section 8.4.1, "Oracle Virtual Directory Component Architecture"](#)
- [Section 8.4.2, "Oracle Virtual Directory High Availability Concepts"](#)
- [Section 8.4.3, "Oracle Virtual Directory High Availability Configuration Steps"](#)
- [Section 8.4.4, "Validating Oracle Virtual Directory High Availability"](#)
- [Section 8.4.5, "Oracle Virtual Directory Failover and Expected Behavior"](#)
- [Section 8.4.6, "Troubleshooting Oracle Virtual Directory High Availability"](#)

### 8.4.1 Oracle Virtual Directory Component Architecture

Oracle Virtual Directory is an LDAP version 3 enabled service that provides virtualized abstraction of one or more enterprise data sources into a single directory view. Oracle Virtual Directory provides the ability to integrate LDAP-aware applications into diverse directory environments while minimizing or eliminating the need to change either the infrastructure or the applications. Oracle Virtual Directory supports a diverse set of clients, such as Web Applications and portals, and it can connect to directories, databases and Web Services as shown in [Figure 8-4](#).

[Figure 8-4](#) shows Oracle Virtual Directory in a non-high availability architecture.

**Figure 8–4 Oracle Virtual Directory in a Non-High Availability Architecture**

The Oracle Virtual Directory server is written in Java and internally it is organized into multiple layers. These layers are logical layers—Oracle Virtual Directory appears as a single complete service to the administrator and to clients.

#### 8.4.1.1 Oracle Virtual Directory Runtime Considerations

OPMN is used to start, monitor, and manage the Oracle Virtual Directory process, and to restart the Oracle Virtual Directory process if it goes down. For information on using OPMNCTL to start and stop Oracle Virtual Directory instances, see *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

OPMN invokes the JVM to start the VDEServer process with the required parameters. JVM parameters are configured in `opmn.xml` (`oracle.security.jps.config` is used for the JPS Config File Location, `vde.soTimeoutBackend` is used to control orphan server connections).

You can also use the Oracle Enterprise Manager Fusion Middleware Control to start and stop Oracle Virtual Directory instances. For information, see *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

Except for JPS, which is installed when Oracle Virtual Directory is installed, Oracle Virtual Directory does not have external dependencies. It can run by itself.

Oracle Virtual Directory can be configured to store LDAP objects in the local file system. This feature can be used by JPS and other components.

Oracle Virtual Directory provides two types of listeners: LDAP and HTTP. Both listeners support SSL/TLS on top of their basic protocols. The LDAP layer also provides the ability to support LDAP-SASL to support digital certificate authentication.

The LDAP(S) protocols provide LDAPv2/v3 based services, and the HTTP(S) protocols provide one or more services such as DSMLv2, or basic white page functions provided by an XSLT enabled Web Gateway.

Based on the nature of the operation, client connections can either be persistent or short-lived.

#### 8.4.1.2 Oracle Virtual Directory Component Characteristics

This section describes the various configuration artifacts for Oracle Virtual Directory. The following Oracle Virtual Directory configuration files are located under `ORACLE_INSTANCE/config/OVD/OVDComponentName`:

- `server.os_xml`:

Oracle Virtual Directory provides the ability to regulate items such as the number of entries the server can return for an anonymous user or for an authenticated user. You can also limit inbound transaction traffic, which can be used to protect proxied sources from Denial Of Service attacks or to limit LDAP traffic to control access to a limited directory infrastructure resource. These properties and others are configured in `server.os_xml`.

- `listeners.os_xml`:

Oracle Virtual Directory provides services to clients through connections known as Listeners. Oracle Virtual Directory supports two types of Listeners: LDAP and HTTP. An Oracle Virtual Directory configuration can have any number of listeners or it can even have zero Listeners, thus restricting access to only the administrative gateway. Most Oracle Virtual Directory deployments will need no more than two HTTP Listeners and two LDAP Listeners, where one Listener is for SSL and one for non-SSL for each protocol. The Listener configuration file is `listeners.os_xml`.

- `adapters.os_xml`:

To present the single virtual directory view of data in multiple and various data repositories, Oracle Virtual Directory must connect to those repositories so it can virtualize the data and route data to and from the repositories. Oracle Virtual Directory uses adapters to connect to its underlying data repositories. Each adapter manages a namespace in the directory identified by a specific parent distinguished name (DN). There is no limit to how many adapters you can configure. You can also combine and overlap adapters to present a customized directory tree. The adapters configuration file is `adapters.os_xml`.

---

**Note:** For information on configuring a No-Authorization SSL connection between Oracle Virtual Directory and a proxy LDAP directory, see *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*. The procedure described in that manual can be used for any proxy LDAP directory configured to support anonymous ciphers.

---

- `acls.os_xml`

Oracle Virtual Directory provides granular access controls that can be applied uniformly across all connected data stores and which are compliant with the Internet Engineering Task Force's RFC 2820, Access Control Requirements for LDAP. The access control rules are modeled on the IETF's internet draft titles LDAP Access Control Model for LDAPv3, (March 2, 2001 draft).

Oracle Virtual Directory provides virtualized abstraction of one or more enterprise data sources into a single directory view. Accordingly, Access Control Lists (ACLs) and adapter namespaces are independent of each other. Removing all entries in a namespace, or changing the root value of an adapter, will not effect ACLs

automatically. ACLs and adapter namespaces must be configured independently of each other. The ACL configuration file is `acls.os_xml`.

Oracle Virtual Directory instance-specific data is stored in the `ORACLE_INSTANCE` home. The wallet is also stored in the instance home.

If a single Oracle Virtual Directory instance fails, use OPMN to restart the instance.

**8.4.1.2.1 Oracle Virtual Directory Log File** The log files for an Oracle Virtual Directory instance are stored in the following directory in the instance home:

```
ORACLE_INSTANCE/diagnostics/logs/OVD/OVDComponentName/
```

For more information on using the Oracle Virtual Directory log files to troubleshoot Oracle Virtual Directory issues, see [Section 8.4.6, "Troubleshooting Oracle Virtual Directory High Availability"](#).


## 8.4.2 Oracle Virtual Directory High Availability Concepts

This section provides conceptual information about using Oracle Virtual Directory in a high availability two-node cluster. See [Section 8.4.2.2, "Oracle Virtual Directory Prerequisites"](#) for prerequisites and [Section 8.4.3, "Oracle Virtual Directory High Availability Configuration Steps"](#) for specific steps for setting up the two-node cluster.

### 8.4.2.1 Oracle Virtual Directory High Availability Architecture

[Figure 8–5](#) shows the Oracle Virtual Directory high availability architecture in an active-active configuration.

**Figure 8–5 Oracle Virtual Directory in a High Availability Environment**



[Figure 8–5](#) shows Oracle Virtual Directory in the directory tier in a high availability architecture. The two-node cluster is set up at installation time. The load balancing router routes requests to the two Oracle Virtual Directory instances that are clustered on OVDHOST1 and OVDHOST2. Fast Connection Failover (FCF) is used for notification when an Oracle RAC instance becomes unavailable.

The two computers have the same Oracle Virtual Directory configuration. The Oracle Virtual Directory configuration for each instance is stored in its `ORACLE_INSTANCE`. Each Oracle Virtual Directory Instance configuration must be updated separately by using Oracle Directory Services Manager to connect to each Oracle Virtual Directory instance one at a time. The alternate approach is to update the configuration of one Oracle Virtual Directory instance and then use cloning to copy the configuration to the other Oracle Virtual Directory instance or instances. See the "Cloning Oracle Fusion Middleware" chapter in *Oracle Fusion Middleware Administrator's Guide* for more information about cloning.

---

**Note:** Oracle Directory Services Manager should be used to manage Oracle Virtual Directory configuration. Oracle Directory Services Manager should not connect to Oracle Virtual Directory through the load balancer to perform configuration updates to an Oracle Virtual Directory instance; instead, it should connect explicitly to a physical Oracle Virtual Directory instance to perform a configuration update for that instance.

---

OPMN is used to start and stop Oracle Virtual Directory instances. When a cluster that includes multiple Oracle Virtual Directory instances is started, there is no impact on the individual Oracle Virtual Directory instances in the cluster.

The load balancing router detects a hang or failure of an Oracle Virtual Directory instance and routes LDAP and HTTP traffic to surviving instances. When the instance becomes available again, the load balancing router detects this and routes traffic to all the surviving instances.

If an instance fails in the middle of a transaction, the transaction is not committed to the back end.

If one Oracle Virtual Directory instance in the two-node Oracle Virtual Directory cluster fails, the load balancing router detects this and reroutes the LDAP client traffic to the surviving instance or instances in the cluster. When the Oracle Virtual Directory instance comes up again, the load balancing router detects this and reroutes the LDAP client traffic instance to the surviving instance or instances in the cluster.

**8.4.2.1.1 Oracle Virtual Directory High Availability Connect Features** Oracle Virtual Directory offers multiple high availability capabilities, including:

- Fault Tolerance and Failover: Oracle Virtual Directories provide fault tolerance in two forms:
  - They can be configured in fault tolerant configurations.
  - They can manage flow to fault tolerant proxied sources.

Multiple Oracle Virtual Directories can be quickly deployed simply by copying, or even sharing configuration files. When combined with round-robin DNS, redirector, or cluster technology, Oracle Virtual Directory provides a complete fault-tolerant solution.

For each proxied directory source, Oracle Virtual Directory can be configured to access multiple hosts (replicas) for any particular source. It intelligently fails over between hosts and spreads the load between them. Flexible configuration options allow administrators to control percentages of a load to be directed towards specific replica nodes and to indicate whether a particular host is a read-only replica or a read-write server (master). This avoids unnecessary referrals resulting from attempts to write to a read-only replica.

- Load Balancing: Oracle Virtual Directory was designed with powerful load balancing features that allow it to spread load and manage failures between its proxied LDAP directory sources.

Oracle Virtual Directory's virtual directory tree capability allows large sets of directory information to be broken up into multiple distinct directory servers. Oracle Virtual Directory is able to recombine the separated data sets back into one virtual tree by combining the separate directory tree branches.

If you have multiple LDAP servers for a particular source, the Oracle Virtual Directory LDAP Adapter can load balance and failover for these servers on its own. This load balancing and failover happens transparently to the client and does not require any additional hardware or changes to the client connecting to Oracle Virtual Directory.

The Database adapter supports load balancing and failover if the underlying JDBC driver provides this functionality. Additionally, Oracle Virtual Directory is certified for use with Oracle Real Application Clusters (Oracle RAC). See [Section 4.1.6, "JDBC Clients"](#) for more information about using Oracle Virtual Directory with Oracle RAC.

Oracle Virtual Directory Routing also provides load balancing capabilities. Routing allows search filters to be included in addition to the search base to determine optimized search targets. In this load balancing approach, Oracle Virtual Directory automatically routes queries to the appropriate virtualized directory sources enabling the ability to work with many millions of directory entries.

---

**Note:** Oracle Virtual Directory's value is as a virtualization and proxy service, not as a directory store. If you need a highly available directory storage system, Oracle recommends using Oracle Internet Directory.

---

The log files for an Oracle Virtual Directory instance are stored in the following directory in the instance home:

`ORACLE_INSTANCE/diagnostics/logs/OVD/OVDComponentName/`

For more information on using the Oracle Virtual Directory log files to troubleshoot Oracle Virtual Directory issues, see [Section 8.4.6, "Troubleshooting Oracle Virtual Directory High Availability"](#).

### 8.4.2.2 Oracle Virtual Directory Prerequisites

This section describes prerequisites for Oracle Virtual Directory.

---

**Note:** Oracle requires that you synchronize the system clocks on the cluster nodes when you are using Oracle Virtual Directory in a high availability deployment.

---

**8.4.2.2.1 Load Balancer Virtual Server Names for Oracle Virtual Directory** When Oracle Virtual Directory is deployed in a high availability configuration, it is recommended to use an external load balancer to front-end the Oracle Virtual Directory instances and load balance the LDAP requests between the various Oracle Virtual Directory instances.

Refer to [Section 8.2.5.4, "Configuring Virtual Server Names and Ports for the Load Balancer"](#) for details.

## 8.4.3 Oracle Virtual Directory High Availability Configuration Steps

Oracle Virtual Directory can be deployed in a highly available configuration either in a standalone mode or as a part of a WebLogic Server domain.

Since Oracle Directory Services Manager and Oracle Enterprise Manager Fusion Middleware Control are required to manage Oracle Virtual Directory effectively, it is recommended that Oracle Virtual Directory be deployed as part of an Oracle WebLogic Server domain. For information on configuring Oracle Virtual Directory and Oracle Directory Services Manager in a collocated configuration, see [Section 8.7, "Collocated Architecture High Availability."](#)

In a high availability environment, it is recommended that Oracle Virtual Directory be set up in a clustered deployment, where the clustered Oracle Virtual Directory instances access the same Oracle RAC database repository or LDAP repository.

#### 8.4.3.1 Configuring Oracle Virtual Directory Without a WebLogic Domain

This section describes the steps to deploy Oracle Virtual Directory without an Oracle WebLogic Server domain.

---

**Note:** When Oracle Virtual Directory is installed on a host using the **Configure without a Domain** option in the Oracle Identity Management 11g installer, by default the installer uses ovd1 as the component name for the Oracle Virtual Directory instance.

During a **Configure without a Domain** installation, the installer cannot detect if another Oracle Virtual Directory instance with a component name of ovd1 already exists on the host and is registered with a domain.

The Oracle Virtual Directory instances installed on OVDHOST1 and OVDHOST2 in [Section 8.4.3.1.1, "Configuring Oracle Virtual Directory on OVDHOST1"](#) and [Section 8.4.3.1.2, "Configuring Oracle Virtual Directory on OVDHOST2"](#) are installed using the **Configure without a Domain** installation option.

---

**8.4.3.1.1 Configuring Oracle Virtual Directory on OVDHOST1** Follow these steps to configure the Oracle Virtual Directory instance on OVDHOST1:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OVDHOST1 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Ensure that ports 6501 and 7501 are not in use by any service on the computer by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":6501"
```

```
netstat -an | grep LISTEN | grep ":7501"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":6501"
```

```
netstat -an | findstr "LISTEN" | findstr ":7501"
```



4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 6501 and 7501 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Virtual Directory):

```
The Non-SSL LDAP port for OVD
Oracle Virtual Directory (Non-SSL) Port No for LDAP= 6501
The SSL LDAP Port for OVD
Oracle Virtual Directory (SSL) Port No for LDAP = 7501
```

7. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Configure without a Domain** and then click **Next**.
10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.
  - **Oracle Home Directory:**

```
/u01/app/oracle/product/fmw/idm_1
```
  - **Oracle Instance Location:**

```
/u01/app/oracle/admin/ora_inst1
```
  - **Oracle Instance Name:**

```
ora_inst1
```

---

**Note:** Ensure that the Oracle Home Location directory path for OVDHOST1 is the same as the Oracle Home Location directory path for OVDHOST2. For example, if the Oracle Home Location directory path for OVDHOST1 is:

```
/u01/app/oracle/product/fmw/idm_1
```

then the Oracle Home Location directory path for OVDHOST2 must be:

```
/u01/app/oracle/product/fmw/idm_1
```

---

Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Virtual Directory** and click **Next**.

13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

14. Specify the following values on the Specify Virtual Directory screen:

In the Client Listeners section, enter:

- **LDAP v3 Name Space:** Enter the name space for Oracle Virtual Directory. The default value is `dc=us,dc=mycompany,dc=com`.

`dc=us,dc=mycompany,dc=com`

- **HTTP Web Gateway:** Select this option to enable the Oracle Virtual Directory HTTP Web Gateway.
- **Secure:** Select this option if you enabled the HTTP Web Gateway and you want to secure it using SSL.

In the OVD Administrator section, enter:

- **Administrator User Name:** Enter the user name for the Oracle Virtual Directory administrator, for example: `cn=orcladmin`
- **Password:** Enter the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Confirm Password:** Confirm the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Configure the Administrative Server in secure mode:** Select this option to secure the Oracle Virtual Directory Administrative Listener using SSL. This option is selected by default. Oracle recommends selecting this option.

Click **Next**.

15. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.

16. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

17. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
18. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

**8.4.3.1.2 Configuring Oracle Virtual Directory on OVDHOST2** Follow these steps to configure the Oracle Virtual Directory instance on OVDHOST2:

---

**Note:** The instructions in this section can also be used to scale out Oracle Virtual Directory in your 11g Oracle Identity Management high availability configuration.

---

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OVDHOST2 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. On OVDHOST1, ports 6501 and 7501 were used for Oracle Virtual Directory. The same ports should be used for the Oracle Virtual Directory instance on OVDHOST2. Therefore, ensure that ports 6501 and 7501 are not in use by any service on OVDHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":6501"
```

```
netstat -an | grep LISTEN | grep ":7501"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":6501"
```

```
netstat -an | findstr "LISTEN" | findstr ":7501"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 6501 and 7501 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Virtual Directory):

```
The Non-SSL LDAP port for OVD
Oracle Virtual Directory (Non-SSL) Port No for LDAP = 6501
The SSL LDAP Port for OVD
```

Oracle Virtual Directory (SSL) Port No for LDAP = 7501

7. Start the Oracle Identity Management 11g Configuration Assistant located under the *ORACLE\_HOME/bin* directory as follows:

On UNIX, issue this command: **./config.sh**

On Windows, double-click **config.exe**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Configure without a Domain** and then click **Next**.
10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.
  - **Oracle Home Directory:**  
*/u01/app/oracle/product/fmw/idm\_1*
  - **Oracle Instance Location:**  
*/u01/app/oracle/admin/ora\_inst2*
  - **Oracle Instance Name:**  
*ora\_inst2*

---

**Note:** Ensure that the Oracle Home Location directory path for OVDHOST2 is the same as the Oracle Home Location directory path for OVDHOST1. For example, if the Oracle Home Location directory path for OVDHOST1 is:

*/u01/app/oracle/product/fmw/idm\_1*

then the Oracle Home Location directory path for OVDHOST2 must be:

*/u01/app/oracle/product/fmw/idm\_1*

---

Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:
  - **Email Address:** Provide the email address for your My Oracle Support account.
  - **Oracle Support Password:** Provide the password for your My Oracle Support account.
  - Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Virtual Directory** and click **Next**.
13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the *staticports.ini* file that you copied to the temporary directory.

Click **Next**.

**14. Specify the following values on the Specify Virtual Directory screen:**

In the Client Listeners section, enter:

- **LDAP v3 Name Space:** Enter the name space for Oracle Virtual Directory. The default value is `dc=us,dc=mycompany,dc=com`.

`dc=us,dc=mycompany,dc=com`

- **HTTP Web Gateway:** Select this option to enable the Oracle Virtual Directory HTTP Web Gateway.
- **Secure:** Select this option if you enabled the HTTP Web Gateway and you want to secure it using SSL.

In the OVD Administrator section, enter:

- **Administrator User Name:** Enter the user name for the Oracle Virtual Directory administrator, for example: `cn=orcladmin`
- **Password:** Enter the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Confirm Password:** Confirm the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Configure the Administrative Server in secure mode:** Select this option to secure the Oracle Virtual Directory Administrative Listener using SSL. This option is selected by default. Oracle recommends selecting this option.

Click **Next**.

**15. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.**

**16. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.**

Click **Next**.

**17. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.**

**18. On the Installation Complete screen, click **Finish** to confirm your choice to exit.**

**8.4.3.1.3 Registering Oracle Virtual Directory with a WebLogic Domain** It is recommended that you manage your Oracle Virtual Directory component with Oracle Enterprise Manager Fusion Middleware Control. To be able to manage Oracle Virtual Directory using Oracle Enterprise Manager Fusion Middleware Control, you must register the component and the Oracle Fusion Middleware instance that contains it with an Oracle WebLogic Server domain. You can register an Oracle Fusion Middleware instance with a WebLogic domain during installation or Oracle instance creation, but you are not required to do so. If an Oracle Fusion Middleware instance was not previously registered with a WebLogic domain, you can register it by using `opmnctl registerinstance`.

Before using the `opmnctl registerinstance` command to register an Oracle Virtual Directory instance with an Oracle WebLogic Server domain, ensure that the WebLogic Server is installed.

Then execute the `opmnctl registerinstance` command in this format (note that the `ORACLE_HOME` and `ORACLE_INSTANCE` variables have to be set for each installation before executing this command in the home directory for the Oracle Virtual Directory instance):

```
opmnctl registerinstance -adminHost WLSTHostName -adminPort WLSPort
-adminUsername adminUserName
```

For example:

```
opmnctl registerinstance -adminHost idmhost1.mycompany.com -adminPort 7001
-adminUsername weblogic
```

```
Command requires login to weblogic admin server (idmhost1.mycompany.com)
Username: weblogic
Password: *****
```

For additional details on registering Oracle Virtual Directory components with a WebLogic domain, see the "Registering an Oracle Instance Using OPMNCTL" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

### 8.4.3.2 Configuring Oracle Virtual Directory With a WebLogic Domain

This section describes the steps to deploy Oracle Virtual Directory in a high availability configuration as part of a WebLogic Server domain.

In this configuration, Oracle Virtual Directory and a WebLogic Server domain is configured on the first host, and only Oracle Virtual Directory is configured on the second host. The Oracle Virtual Directory instance on the second host joins the domain created on the first host.

**8.4.3.2.1 Configuring Oracle Virtual Directory on OVDHOST1** Follow these steps to configure the Oracle Virtual Directory instance on OVDHOST1:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OVDHOST1 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Ensure that ports 6501 and 7501 are not in use by any service on OVDHOST1 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":6501"
```

```
netstat -an | grep LISTEN | grep ":7501"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":6501"
```

```
netstat -an | findstr "LISTEN" | findstr ":7501"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 6501 and 7501 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Virtual Directory):

```
The Non-SSL LDAP port for OVD
Oracle Virtual Directory (Non-SSL) Port No for LDAP = 6501
The SSL LDAP port for OVD
Oracle Virtual Directory (SSL) Port No for LDAP = 7501
```

7. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: `./config.sh`

On Windows, double-click `config.exe`

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Create New Domain** and specify the following values:
  - **User Name:** weblogic
  - **Password:** <password for the weblogic user>
  - **Confirm Password:** <Confirm the password for the weblogic user>
  - **Domain Name:** IDMDomain
10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:**  
`/u01/app/oracle/product/fmw`
  - **Oracle Home Directory:** idm
  - **WebLogic Server Directory:**  
`/u01/app/oracle/product/fmw/wlserver_10.3`
  - **Oracle Instance Location:**  
`/u01/app/oracle/admin/ovd_inst1`
  - **Oracle Instance Name:**  
`ovd_inst1`

---

---

**Note:** Ensure that the Oracle Home Location directory path for OVDHOST1 is the same as the Oracle Home Location directory path for OVDHOST2. For example, if the Oracle Home Location directory path for OVDHOST1 is:

`/u01/app/oracle/product/fmw/idm`

then the Oracle Home Location directory path for OVDHOST2 must be:

`/u01/app/oracle/product/fmw/idm`

---

---

Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Virtual Directory** and click **Next**.

---

---

**Note:** The Oracle Directory Services Manager and Fusion Middleware Control management components are automatically selected for this installation. You cannot deselect these choices.

---

---

13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

14. Specify the following values on the Specify Virtual Directory screen:

In the Client Listeners section, enter:

- **LDAP v3 Name Space:** Enter the name space for Oracle Virtual Directory. The default value is `dc=us,dc=mycompany,dc=com`.

`dc=us,dc=mycompany,dc=com`

- **HTTP Web Gateway:** Select this option to enable the Oracle Virtual Directory HTTP Web Gateway.
- **Secure:** Select this option if you enabled the HTTP Web Gateway and you want to secure it using SSL.

In the OVD Administrator section, enter:

- **Administrator User Name:** Enter the user name for the Oracle Virtual Directory administrator, for example: `cn=orcladmin`
- **Password:** Enter the password for the Oracle Virtual Directory administrator. For example: `*****`



- **Confirm Password:** Confirm the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Configure the Administrative Server in secure mode:** Select this option to secure the Oracle Virtual Directory Administrative Listener using SSL. This option is selected by default. Oracle recommends selecting this option.

Click **Next**.

15. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.
16. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.  
Click **Next**.
17. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
18. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

**8.4.3.2.2 Creating boot.properties for the Administration Server on OVDHOST1** This section describes how to create a `boot.properties` file for the Administration Server on OVDHOST1. The `boot.properties` file enables the Administration Server to start without prompting for the administrator username and password.

Follow these steps to create the `boot.properties` file:

1. On OVDHOST1, go the following directory:

```
MW_HOME/user_projects/domains/domainName/servers/AdminServer/security
```

For example:

```
cd /u01/app/oracle/product/fmw/user_
projects/domains/IDMDomain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` under the `security` directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

3. Stop the Administration Server if it is running.

See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.

4. Start the Administration Server on OVDHOST1 using the `startWebLogic.sh` script located under the `MW_HOME/user_projects/domains/domainName/bin` directory.
5. Validate that the changes were successful by opening a web browser and accessing the following pages:
  - WebLogic Server Administration Console at:  
`http://oidhost1.mycompany.com:7001/console`
  - Oracle Enterprise Manager Fusion Middleware Control at:  
`http://oidhost1.mycompany.com:7001/em`

Log into these consoles using the `weblogic` user credentials.

**8.4.3.2.3 Configuring Oracle Virtual Directory on OVDHOST2** Follow these steps to configure the Oracle Virtual Directory instance on OVDHOST2:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OVDHOST2 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. On OVDHOST1, ports 6501 and 7501 were used for Oracle Virtual Directory. The same ports should be used for the Oracle Virtual Directory instance on OVDHOST2. Therefore, ensure that ports 6501 and 7501 are not in use by any service on OVDHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":6501"
```

```
netstat -an | grep LISTEN | grep ":7501"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":6501"
```

```
netstat -an | findstr "LISTEN" | findstr ":7501"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for ports 6501 and 7501 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.

6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the lines where you specify the port numbers for Oracle Virtual Directory):

```
The Non-SSL LDAP port for OVD
Oracle Virtual Directory (Non-SSL) Port No for LDAP = 6501
The SSL LDAP port for OVD
Oracle Virtual Directory (SSL) Port No for LDAP = 7501
```

7. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Extend Existing Domain** and specify the following values:
  - **HostName:** `ovdhost1.mycompany.com` (This is the host where the Oracle WebLogic Administration Server is running.)
  - **Port:** `7001` (This is the Administration Server port.)
  - **User Name:** `weblogic`
  - **Password:** `<password for the weblogic user>`
10. On the Specify Installation Location screen, specify the following values:

- **Oracle Middleware Home Location:**

```
/u01/app/oracle/product/fmw
```

- **Oracle Home Directory:** `idm`

- **WebLogic Server Directory:**

```
/u01/app/oracle/product/fmw/wlserver_10.3
```

- **Oracle Instance Location:**

```
/u01/app/oracle/admin/ovd_inst2
```

- **Oracle Instance Name:**

```
ovd_inst2
```

---

**Note:** Ensure that the Oracle Home Location directory path for OVDHOST1 is identical to the Oracle Home Location directory path for OVDHOST2. For example, if the Oracle Home Location directory path for OVDHOST1 is:

```
/u01/app/oracle/product/fmw/idm
```

then the Oracle Home Location directory path for OVDHOST2 must be:

```
/u01/app/oracle/product/fmw/idm
```

---

Click **Next**.

11. On the Specify Email for Security Updates screen, specify these values:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Virtual Directory** and click **Next**.

---

**Note:** The Oracle Directory Services Manager and Fusion Middleware Control management components are automatically selected for the installation.

---

13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

14. Specify the following values on the Specify Virtual Directory screen:

In the Client Listeners section, enter:

- **LDAP v3 Name Space:** Enter the name space for Oracle Virtual Directory. The default value is `dc=us,dc=mycompany,dc=com`.

`dc=us,dc=mycompany,dc=com`

- **HTTP Web Gateway:** Select this option to enable the Oracle Virtual Directory HTTP Web Gateway.
- **Secure:** Select this option if you enabled the HTTP Web Gateway and you want to secure it using SSL.

In the OVD Administrator section, enter:

- **Administrator User Name:** Enter the user name for the Oracle Virtual Directory administrator, for example: `cn=orcladmin`
- **Password:** Enter the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Confirm Password:** Confirm the password for the Oracle Virtual Directory administrator. For example: `*****`
- **Configure the Administrative Server in secure mode:** Select this option to secure the Oracle Virtual Directory Administrative Listener using SSL. This option is selected by default. Oracle recommends selecting this option.

Click **Next**.

15. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.

16. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

17. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.

18. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

### 8.4.3.3 Configuring Oracle Virtual Directory with Highly Available Data Sources

Oracle Virtual Directory can be configured to use either an Oracle RAC database repository or a highly available LDAP repository as a data source.

This section describes how to configure Oracle Virtual Directory in a high availability environment with an Oracle RAC database repository and with an LDAP repository.

---

**Note:** When configuring Oracle Virtual Directory in a high availability topology, the configuration steps must be completed on all the nodes individually.

---

**8.4.3.3.1 Configuring Oracle Virtual Directory with an Oracle RAC Database** In an Oracle Virtual Directory high availability environment with an Oracle RAC database as the repository, you must create and configure an Oracle Virtual Directory Database Adapter.

For introductory information about Oracle Virtual Directory Adapters, see the "Understanding Oracle Virtual Directory Adapters" section of *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

For information about creating Database Adapters for Oracle RAC databases, see the "Creating Database Adapters for RAC Databases" section of *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

For information about configuring Database Adapters, see the "Configuring Database Adapters" section of *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

**8.4.3.3.2 Configuring Oracle Virtual Directory with LDAP** Oracle Virtual Directory connects to a LDAP repository, such as Oracle Internet Directory, through an LDAP adapter. For high availability environments, the LDAP adapter can be configured by adding the host and port information of the nodes with a load balancing algorithm or by adding the load balancer URL of the LDAP repository.

For introductory information about Oracle Virtual Directory Adapters, see the "Understanding Oracle Virtual Directory Adapters" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

For information about creating and configuring LDAP Adapters, see the "Configuring LDAP Adapters" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*.

## 8.4.4 Validating Oracle Virtual Directory High Availability

Oracle Virtual Directory instances are front-ended by a hardware load balancer in high availability deployments. Requests are load balanced between the Oracle Virtual Directory instances, and in case of a failure (either an instance failure or a host failure),

the hardware load balancer detects the failure and routes all requests to the surviving instance or instances.

To test Oracle Virtual Directory high availability, stop one node at a time and connect to the Oracle Virtual Directory instances through the load balancer URL using a tool such as `ldapbind`. The connection should succeed all the time as long as one node is running and the topology is configured correctly.

Use the `ldapbind` command line tool to connect to the Oracle Virtual Directory instances through the load balancer virtual host. The `ldapbind` tool enables you to determine whether you can authenticate a client to a server.

Follow the steps below to validate the high availability setup of Oracle Virtual Directory on OVDHOST1 and OVDHOST2:

1. Configure your environment by following the steps in the "Configuring Your Environment" section of *Oracle Fusion Middleware Reference for Oracle Identity Management*. That section has a list of the environment variables you must set before using the `ldapbind` command.
2. On OVDHOST1, use the `opmnctl` command to stop the Oracle Virtual Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=ovd1
```

3. On OVDHOST2, check the status of Oracle Virtual Directory by binding to the load balancing virtual address using `ldapbind`:

Non-SSL:

```
ldapbind -h ovd.mycompany.com -p 6501 -D "cn=orcladmin" -q
```

---

**Note:** The `-q` option above prompts the user for a password. LDAP tools have been modified to disable the options `-w password` and `-P password` when the environment variable `LDAP_PASSWORD_PROMPTONLY` is set to `TRUE` or `1`. Use this feature whenever possible.

---

4. A successful bind shows that the load balancer is routing all traffic to OVDHOST2.
5. On OVDHOST1, use the `opmnctl` command to start the Oracle Virtual Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl start (if OPMN is not running)
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=ovd1
```

6. On OVDHOST2, use the `opmnctl` command to stop the Oracle Virtual Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=ovd1
```

7. On OVDHOST1, check the status of Oracle Virtual Directory by binding to the load balancing virtual address using `ldapbind`:

Non-SSL:

```
ldapbind -h ovd.mycompany.com -p 6501 -D "cn=orcladmin" -q
```

8. On OVDHOST2, use the `opmnctl` command to start the Oracle Virtual Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl start (if OPMN is not running)
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=ovd1
```

For more information about the `ldapbind` command, see the `ldapbind` section in *Oracle Fusion Middleware Reference for Oracle Identity Management*.

#### 8.4.4.1 Validating Oracle Virtual Directory High Availability Using SSL

Oracle Virtual Directory is configured to use the SSL Server Authentication Only Mode by default. When using command line tools like `ldapbind` to validate a connection using connection secured by SSL Server Authentication mode, the server certificate must be stored in an Oracle Wallet. Follow the steps below to perform this task:

1. Create an Oracle Wallet by executing the following command:

```
ORACLE_HOME/bin/orapki wallet create -wallet DIRECTORY_FOR_SSL_WALLET -pwd
WALLET_PASSWORD
```

2. Export the Oracle Virtual Directory server certificate by executing the following command:

```
ORACLE_HOME/jdk/jre/bin/keytool -exportcert -keystore OVD_KEystore_FILE
-storepass PASSWORD -alias OVD_SERVER_CERT_ALIAS -rfc -file
OVD_SERVER_CERT_FILE
```

3. Add the Oracle Virtual Directory server certificate to the Oracle Wallet by executing the following command:

```
ORACLE_HOME/bin/orapki wallet add -wallet DIRECTORY_FOR_SSL_WALLET
-trusted_cert -cert OVD_SERVER_CERT_FILE -pwd WALLET_PASSWORD
```

4. Follow the instructions shown in [Section 8.4.4, "Validating Oracle Virtual Directory High Availability"](#) but use the `ldapbind` command shown below to validate the high availability setup of Oracle Virtual Directory on OVDHOST1 and OVDHOST2. Use the Oracle Wallet from step 3 while executing the following command:

```
ORACLE_HOME/bin/ldapbind -D cn=orcladmin -q -U 2 -h HOST -p SSL_PORT -W
"file://DIRECTORY_FOR_SSL_WALLET" -q
```

5. When an Oracle Virtual Directory high availability deployment is front ended by a hardware load balancer, the wallets on all the Oracle Virtual Directory nodes must contain the client certificates of all the Oracle Virtual Directory instances that are a part of that topology. Add the client certificates of all the Oracle Virtual Directory instances in the topology to the wallets on all the nodes in the topology. This ensures that a valid connection request made through the load balancer URL does not fail.

---

**Note:** If you are using default settings after installing 11g Release 1 (11.1.1), you can use the following values for the variables described in this section:

- For `OVD_KEystore_FILE`, use:
 

```
ORACLE_INSTANCE/config/OVD/ovd1/keystores/keys.jks
```
  - For `OVD_SERVER_CERT_ALIAS`, use `serverselfsigned`
  - For `PASSWORD` used for the `-storepass` option, use the `orcladmin` account password.
-

## 8.4.5 Oracle Virtual Directory Failover and Expected Behavior

This section includes steps for performing a failover of Oracle Virtual Directory and for performing a failover of Oracle RAC.

### 8.4.5.1 Performing an Oracle Virtual Directory Failover

Follow these steps to perform a failover of an Oracle Virtual Directory instance and to check the status of Oracle Virtual Directory:

1. Use the `opmnctl` command to stop the Oracle Virtual Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=ovd1
```

2. Use the `opmnctl` command to start the Oracle Virtual Directory instance:

```
ORACLE_INSTANCE/bin/opmnctl start (if OPMN is not running)
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=ovd1
```

3. Check the status of Oracle Virtual Directory:

```
ORACLE_INSTANCE/bin/opmnctl status ias-component=ovd1
```

### 8.4.5.2 Performing an Oracle RAC Failover

Follow these steps to perform an Oracle RAC failover:

1. Use the `srvctl` command to stop a database instance:

```
srvctl stop instance -d db_unique_name -i inst_name_list
```

2. Use the `srvctl` command to check the status of the database:

```
srvctl status database -d db_unique_name -v
```

3. Check the status of Oracle Virtual Directory:

```
ORACLE_INSTANCE/bin/opmnctl status ias-component=ovd1
```

4. Use the `srvctl` command to start the database instance:

```
srvctl start instance -d db_unique_name -i inst_name_list
```

---

---

**Note:** When Oracle Virtual Directory is configured with a Database Adapter against an Oracle RAC database, the first search performed after the Oracle RAC failover fails. However, subsequent search requests succeed without any issues.

---

---

## 8.4.6 Troubleshooting Oracle Virtual Directory High Availability

Information in the Oracle Virtual Directory log files can be helpful in troubleshooting Oracle Virtual Directory issues. Log files for Oracle Virtual Directory are found under the following directory:

```
ORACLE_INSTANCE/diagnostics/log/OVD/<OVDDComponentName>
```

The order in which log files should be examined when troubleshooting is:

1. `diagnostic.log`: This file captures diagnostic messages.
2. `http-errors.log`: This file captures HTTP errors.



3. `access.log`: This file captures information about processes and clients that access the Oracle Virtual Directory instance.

#### 8.4.6.1 Troubleshooting LDAP Adapter Creation

When creating an LDAP adapter, you specify the host name and port number of LDAP server in the Connection page of adapter creation wizard. If the LDAP server is listening in **SSL Server-Only Auth** or **Mutual Authentication** mode, ODSM imports the server certificate into Oracle Virtual Directory's trust store. However, if you specify the load balancer name that front-ends more than one LDAP server, it imports only one of the LDAP server's certificates. This causes a problem when the Oracle Virtual Directory server's request is routed to the LDAP server, whose certificate is not trusted.

To avoid this problem, during LDAP adapter creation, in addition to specifying the load balancer host and port details, specify the host and port details of LDAP servers front-ended by the load balancer, so that certificates of all LDAP servers are imported. After the adapter is created, you can edit adapter settings to remove host and port details of physical LDAP servers, or their weight can be set to zero.

## 8.5 Oracle Directory Integration Platform High Availability

This section provides an introduction to Oracle Directory Integration Platform and describes how to design and deploy a high availability environment for Oracle Directory Integration Platform and Oracle Directory Services Manager.

See [Section 8.6, "Oracle Directory Services Manager High Availability"](#) for more information about Oracle Directory Services Manager.

This section includes the following topics:

- [Section 8.5.1, "Oracle Directory Integration Platform Component Architecture"](#)
- [Section 8.5.2, "Oracle Directory Integration Platform High Availability Concepts"](#)
- [Section 8.5.3, "Oracle Directory Integration Platform and Oracle Directory Services Manager High Availability Configuration Steps"](#)
- [Section 8.5.4, "Oracle Directory Integration Platform Failover and Expected Behavior"](#)
- [Section 8.5.5, "Troubleshooting Oracle Directory Integration Platform High Availability"](#)

### 8.5.1 Oracle Directory Integration Platform Component Architecture

Oracle Directory Integration Platform is a J2EE application that enables you to integrate your applications and directories, including third-party LDAP directories, with Oracle Internet Directory.

Oracle Directory Integration Platform includes services and interfaces that allow you to deploy synchronization solutions with other enterprise repositories. It can also be used to provide Oracle Internet Directory interoperability with third party metadirectory solutions.

Oracle Directory Integration Platform provides two distinct services depending on the type of integration needed:

- Synchronization through the Oracle Directory Integration Platform Synchronization Service, which keeps connected directories consistent with the central Oracle Internet Directory.

- Notification through Oracle Directory Integration Platform Provisioning Service, which sends notifications to target applications periodically to reflect changes made to a user's status or information.

---

**Note:** Throughout the rest of this chapter, these terms will be used:

- "Synchronization Service" for Oracle Directory Integration Platform Synchronization Service
  - "Provisioning Service" for Oracle Directory Integration Platform Provisioning Service
- 

Figure 8–6 shows the Oracle Directory Integration Platform architecture.

**Figure 8–6 Oracle Directory Integration Platform Architecture**




Figure 8–6 shows the various components of Oracle Directory Integration Platform. Oracle Internet Directory is synchronized with connected directories using Oracle Directory Integration Platform's Synchronization Enterprise JavaBeans (EJB) and the Quartz Scheduler.

Similarly, changes in Oracle Internet Directory are sent to various applications using Oracle Directory Integration Platform's Provisioning Enterprise JavaBeans (EJB) and the Quartz Scheduler.

#### 8.5.1.1 Oracle Directory Integration Platform Component Characteristics

Oracle Directory Integration Platform is a J2EE application that enables you to integrate your applications and directories, including third-party LDAP directories, with Oracle Internet Directory.

Oracle Directory Integration Server provides synchronization services through the Synchronization Service and integration services through the Provisioning Service, as described below:

- The Synchronization Service provides:
  - Scheduling: Processing a synchronization profile based on a predefined schedule
  - Mapping: Executing rules for converting data between connected directories and Oracle Internet Directory
  - Data propagation: Exchanging data with connected directories by using a connector
  - Error handling
- The Provisioning Service provides:
  - Data propagation: Exchanging data with connected directories by using a connector
  - Event notification: Notifying an application of a relevant change to the user or group data stored in Oracle Internet Directory
  - Error handling

Oracle Directory Integration Platform is deployed on Oracle WebLogic Server. By default, Oracle Directory Integration Platform is installed as part of the Oracle Identity Management software stack; however, depending on your architectural requirements, it can also be installed as a standalone application.

#### 8.5.1.1.1 Runtime Processes The Oracle Directory Integration Server provides:

- The Synchronization Service using the Synchronization Enterprise JavaBeans (EJB) and the Quartz Scheduler. The Synchronization EJB is stateless in nature.  
 The Quartz scheduler invokes the Synchronization EJBs that synchronize the appropriate change to all the connected directories.  
 The Synchronization Service supplies these changes using the interface and format required by the connected directory. Synchronization through the Oracle Directory Integration Platform connectors ensures that Oracle Internet Directory remains up-to-date with all the information that Oracle Internet Directory clients need.
- The Provisioning Service uses the Provisioning Enterprise JavaBeans (EJB) and the Quartz Scheduler. The Provisioning EJB is stateless in nature.  
 The Quartz scheduler invokes the Provisioning EJBs that in turn notify each of the provisioned application of the changes.
- UpdateJob EJB periodically polls the Oracle Internet Directory changelog for changes in the Synchronization or Provisioning Profiles. When a profile change is detected, the Quartz scheduler jobs are updated accordingly.

**8.5.1.1.2 Process Lifecycle** Oracle Directory Integration Platform deploys to an Oracle WebLogic Server as an externally managed application. By default, the Oracle Directory Integration Platform application leverages the underlying WebLogic Server infrastructure for startup, shutdown, monitoring and other lifecycle events. You can configure WebLogic Node Manager to monitor the server process and restart it in case of failure.

Oracle Directory Integration Platform initializes when the Managed Server it is deployed on starts up. After the Oracle Directory Integration Platform application

starts up, the initialization code creates an initial set of jobs for the existing directory integration profiles and then invokes the UpdateJobBean EJB.

UpdateJobBean updates the jobs submitted to the Quartz Scheduler.

Depending upon the job, the Quartz scheduler invokes either the Provisioning EJB or the Synchronization EJB. If the Provisioning EJB is invoked, it, in turn, notifies each of the provisioned application of Oracle Internet Directory changes. If the Synchronization EJB is invoked, it, in turn, synchronizes the appropriate change to all the connected directories.

### Starting and Stopping

Manage Oracle Directory Integration Platform lifecycle events using one or more of the command line tools and consoles listed below:

- Oracle WebLogic Scripting Tool (WLST)
- Oracle Enterprise Manager Fusion Middleware Control
- Oracle WebLogic Server Administration Console
- Oracle WebLogic Node Manager

**8.5.1.1.3 Request Flow** Oracle Directory Integration Platform does not service any external requests. Its main functionality is to support synchronization or provisioning based on the profiles created.

This section describes the information flow during synchronization and provisioning.

### Oracle Directory Integration Platform Synchronization Service Flow

Oracle Directory Integration Platform relies on the directory synchronization profiles to synchronize changes between Oracle Internet Directory and the connected directories.

Depending on where the changes are made, synchronization can occur:

- From a connected directory to Oracle Internet Directory
- From Oracle Internet Directory to a connected directory
- In both directions

### Synchronizing from Oracle Internet Directory to a Connected Directory

Oracle Internet Directory maintains a change log in which it stores incremental changes made to directory objects. It stores these changes sequentially based on the change log number.

Synchronization from Oracle Internet Directory to a connected directory makes use of this change log. Consequently, when running the Oracle Directory Integration Platform, you must start Oracle Internet Directory with the default setting in which change logging is enabled.

Each time the Oracle Directory Integration Platform Synchronization Service processes a synchronization profile, it:

1. Retrieves the latest change log number up to which all changes have been applied.
2. Checks each change log entry more recent than that number.
3. Selects changes to be synchronized with the connected directory by using the filtering rules in the profile.

4. Applies the mapping rules to the entry and makes the corresponding changes in the connected directory.

The appropriate entries or attributes are then updated in that connected directory. If the connected directory does not use DB, LDAP, tagged, or LDIF formats directly, then the agent identified in its profile is invoked. The number of the last change successfully used is then stored in the profile.

Periodically, Oracle Internet Directory purges the change log after all profiles have used what they need, and identifies where subsequent synchronization should begin.

### **Synchronizing from a Connected Directory to Oracle Internet Directory**

When a connected directory uses DB, LDAP, tagged, or LDIF formats directly, changes to its entries or attributes can be automatically synchronized by the Oracle Directory Integration Platform Synchronization Service. Otherwise, the connector has an agent in its synchronization profile, which writes the changes to a file in the LDIF or tagged format. The Oracle Directory Integration Platform Synchronization Service then uses this file of connected directory data to update Oracle Internet Directory.

### **Provisioning Flow**

Provisioning refers to the process of providing users, groups, and other objects with access to applications and other resources that are available within an environment. A provisioning-integrated application refers to an application that has registered for provisioning events and registered a provisioning-integration profile in Oracle Internet Directory.

An application can be provisioned with Oracle Directory Integration Platform Provisioning in one of the methods below:

- Synchronous provisioning
- Asynchronous provisioning
- Provisioning data flow

Each of these provisioning methods is described in its own section below.

### **Synchronous Provisioning**

Applications that maintain user information in Oracle Internet Directory and use the Data Access Java plug-in to create, modify, and delete user entries whenever the change occurs in Oracle Internet Directory are said to be provisioned synchronously, since the Data Access Java plug-in can be invoked directly from Oracle Identity Management, including the Provisioning Console and command-line LDAP tools.

Synchronous provisioning with the Oracle Directory Integration Platform Provisioning Service can be invoked from Oracle Identity Management Tools (such as the Provisioning Console in the Oracle Fusion Middleware Control and bulkprov), synchronization with third-party directories, or command-line LDAP tools.

Synchronous provisioning with the Oracle Directory Integration Platform Provisioning Service from Oracle Identity Management Tools such as through the Provisioning Console screen in the Oracle Fusion Middleware Control, bulk provisioning with the provProfileBulkProv command, and from third-party directories follows this process:

1. A new user entry is created in Oracle Internet Directory.
2. The Oracle Identity Management component that created the new user entry invokes the Data Access Java plug-in.
3. The Data Access Java plug-in provisions the new user account in the application.

Synchronous provisioning from command line LDAP tools follows this process:

1. A command-line LDAP tool creates a new user entry in Oracle Internet Directory.
2. At the next scheduled synchronization interval, the Oracle Directory Integration Platform identifies new user entries in Oracle Internet Directory that require provisioning.
3. The Oracle Directory Integration Platform invokes the Data Access Java plug-in.
4. The Data Access Java plug-in provisions the new user accounts in the application.

### **Asynchronous Provisioning**

In asynchronous provisioning, the provisioning is handled by a PL/SQL plug-in and not by any component of Oracle Identity Management

The Oracle Directory Integration Platform propagates PL/SQL events to a provisioning-integrated application, which then executes a PL/SQL plug-in to process the events. Execution of a PL/SQL plug-in occurs within the application repository and not within the address space of any Oracle Identity Management component. Because provisioning is handled by a PL/SQL plug-in and not by any component of Oracle Identity Management, provisioning-integrated applications that implement a PL/SQL plug-in are provisioned asynchronously.

Asynchronous provisioning with the Oracle Directory Integration Platform Provisioning Service can be invoked from Oracle Identity Management Tools (such as Provisioning Console and bulkprov), synchronization with third-party directories, or command-line LDAP tools.

Asynchronous provisioning from Oracle Identity Management Tools such as the Provisioning Console, bulk provisioning with the provProfileBulkProv command, and third-party directories follows this process:

1. A new user entry and an associated entry containing application-specific user preferences are created in Oracle Internet Directory.
2. At the next scheduled synchronization interval, the Oracle Directory Integration Platform identifies new user entries in Oracle Internet Directory that require provisioning.
3. Provisioning events are sent from the Oracle Directory Integration Platform to the PL/SQL plug-in.

Asynchronous provisioning using command line LDAP tools follows this process:

1. A new user entry is created in Oracle Internet Directory using a command-line LDAP tool.
2. At the next scheduled synchronization interval, the Oracle Directory Integration Platform identifies new user entries in Oracle Internet Directory that require provisioning, and creates an associated entry containing application-specific user preferences.
3. Provisioning events are sent from the Oracle Directory Integration Platform to the PL/SQL plug-in.

### **Provisioning Data Flow**

Regardless of whether it is provisioned synchronously or asynchronously, an application can invoke the Pre-Data Entry and Post-Data Entry plug-ins to enhance provisioning intelligence and implement business policies. Both plug-ins are invoked by Oracle Identity Management components such as the Oracle Internet Directory Provisioning Console and bulk provisioning with the provProfileBulkProv command.

The Pre-Data Entry plug-in populates fields according to provisioning policies. The primary purpose of this plug-in is to determine whether a user should be provisioned in an application. For example, if an organization has a policy where only managers are provisioned for a financial application, the Pre-Data Entry plug-in can be used to identify which user entries to provision. Common user attributes are already populated when this plug-in is invoked, so it should have adequate information to make provisioning decisions.

The Post-Data Entry plug-in primarily validates data entered by users for common attributes and application-specific attributes. The validation for the plug-in must be successful for provisioning to continue.

The provisioning data flow follows this process:

1. Base user information is created.
2. The Pre-Data Entry plug-in is invoked, which populates fields according to policies.
3. The Post-Data Entry plug-in is invoked, which validates data entered by the user.
4. Depending on the provisioning approach, either asynchronous or synchronous provisioning procedures are invoked.

If provisioning is performed with the Provisioning Console, then after the Pre-Data Entry Plug-in is invoked, but before the Post-Data Entry plug-in is invoked, an administrator can modify the application attributes.

**8.5.1.1.4 Configuration Artifacts** Oracle Directory Integration Platform-related configuration details are maintained in the `dip-config.xml` file. This configuration file is packaged as part of the Oracle Directory Integration Platform application. The default location for the `dip-config.xml` file is:

```
MW_HOME/user_projects/domains/domainName/config/
fmwconfig/servers/serverName/applications/DIP_11.1.1.2.0/configuration
```

Where *domainName* is the name of your domain *serverName* is the name of your managed server.

Config MBeans manages parameters. [Table 8–6](#) shows the configuration parameters required in `dip-config.xml` to start the Oracle Directory Integration Platform:

**Table 8–6 Configuration Parameters Required to Start Directory Integration Platform**

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                        |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OID Host       | Host name of the Oracle Internet Directory to which Oracle Directory Integration Platform needs to connect.                                                                                                                                                                                                                                        |
| OID Port       | The Oracle Internet Directory port                                                                                                                                                                                                                                                                                                                 |
| SSL Mode       | The SSL mode used to connect to Oracle Internet Directory. Valid values are: <ul style="list-style-type: none"> <li>■ 0: non-SSL</li> <li>■ 1: SSL with no authentication (handshake only). This is the default mode</li> <li>■ 2: SSL with server-only authentication enabled. Authentication is based on the Trust Point Certificate.</li> </ul> |
| JKS Location   | File system location to store the Java Keystore (JKS)                                                                                                                                                                                                                                                                                              |
| Quartz Threads | Maximum number of threads that can be used by Quartz for scheduling the processes.                                                                                                                                                                                                                                                                 |

Once the Oracle Directory Integration Platform server is up and running, it reads further details from Oracle Internet Directory for handling its synchronization and provisioning functions.

For information on creating synchronization profiles and provisioning profiles, see:

- "Creating Synchronization Profiles" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform*.
- "Managing Provisioning Profiles Using oidprovtool" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform*.

**8.5.1.1.5 External Dependencies** Oracle Directory Integration Platform uses an Oracle Internet Directory to store its metadata. The Quartz Scheduler uses the ODSSM schema to store its scheduling information in the database. The same database is used by Oracle Internet Directory and Oracle Directory Integration Platform. The ODSSM schema required for Oracle Directory Integration Platform is created as part of Oracle Internet Directory schema creation.

Oracle Directory Integration Platform is also dependent on the Oracle Credential Store Framework (CSF), a secure framework provided by Oracle and the Java Keystore (JKS) to store wallets and credentials used to connect to Oracle Internet Directory and third party LDAP stores over SSL.

Oracle Directory Integration Platform is also dependent on the Oracle Fusion Middleware Common Audit Framework, which is installed by default.

**8.5.1.1.6 Oracle Directory Integration Platform Log File** Oracle Directory Integration Platform is a J2EE application deployed on top of Oracle WebLogic Server. All log messages are logged in the server log file of the Oracle WebLogic Server that the application is deployed on. The default location of the server log is:

```
WEBLOGIC_SERVER_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

## 8.5.2 Oracle Directory Integration Platform High Availability Concepts

This section provides conceptual information about using Oracle Directory Integration Platform in a high availability configuration.


In the Oracle Directory Integration Platform high availability configuration described in this section, Oracle Directory Integration Platform and Oracle Directory Services Manager are installed and configured on two hosts in a two-node high availability active-active configuration.

### 8.5.2.1 Oracle Directory Integration Platform High Availability Architecture

[Figure 8–7](#) shows the Oracle Directory Integration Platform and Oracle Directory Services Manager high availability architecture in an active-active configuration.



**Figure 8–7 Oracle Directory Integration Platform and Oracle Directory Services Manager in a High Availability Architecture**



In [Figure 8–7](#), the application tier includes the IDMHOST1 and IDMHOST2 computers.

On IDMHOST1, the following installations have been performed:

- An Oracle Directory Integration Platform instance and Oracle Directory Services Manager instance have been installed on the WLS\_ODS1 Managed Server. The Oracle RAC database has been configured in a JDBC multi data source to protect the instances from Oracle RAC node failure.
- A WebLogic Administration Server has been installed. Under normal operations, this is the active Administration Server.

On IDMHOST2, the following installations have been performed:

- An Oracle Directory Integration Platform instance and Oracle Directory Services Manager instance have been installed in the WLS\_ODS2 Managed Server. The

Oracle RAC database has been configured in a JDBC multi data source to protect the instances from Oracle RAC node failure.

The instances in the WLS\_ODS2 Managed Server on IDMHOST2 and the instances in the WLS\_ODS1 Managed Server on IDMHOST1 are configured as the CLUSTER\_ODS cluster.

- A WebLogic Administration Server has been installed. Under normal operations, this is the passive Administration Server. You will make this Administration Server active if the Administration Server on IDMHOST1 becomes unavailable.

**8.5.2.1.1 Starting and Stopping the Cluster** In a high availability architecture, Oracle Directory Integration Platform and Oracle Directory Services Manager are deployed on an Oracle WebLogic Cluster that has at least two servers as a part of the cluster.

By default, the WebLogic Server starts, stops and monitors the applications. By default, both the Oracle Directory Integration Platform and Oracle Directory Services Manager applications leverage the high availability features of the underlying WebLogic Clusters. In case of hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

In a high availability environment, WebLogic Node Manager is configured to monitor the WebLogic servers. In case of failure, Node Manager restarts the WebLogic Server. If Node Manager cannot restart the server, then the front-ending load balancing router detects failure of a WebLogic instance in the Cluster and routes traffic to surviving instances.

**8.5.2.1.2 Cluster-Wide Configuration Changes** When Oracle Internet Directory is deployed in an active-active high availability configuration, all the Oracle Internet Directory instances belonging to the cluster share the same database. Any changes made to Oracle Directory Integration Platform on one Oracle Internet Directory node would automatically be propagated to all the Oracle Internet Directory instances in the cluster.

The following subsections describe configuration changes made to the Oracle Directory Integration Platform application in an Oracle Internet Directory multimaster replication deployment. In a multimaster replication deployment, configuration changes need to be applied to all the nodes in the cluster manually, as described below.

### Directory Integration Profiles

Changes made to directory integration profiles on one Oracle Internet Directory node are not automatically replicated to other Oracle Internet Directory nodes in a default multimaster Oracle Internet Directory replication environment. They need to be manually copied over from the primary node to the secondary nodes on a periodic basis. This allows a directory synchronization profile to execute on a secondary node if a problem occurs on the primary node.

One of the parameters used by Oracle Directory Integration Platform is `orcllastappliedchangenumber`. The value assigned to the `lastchangenumber` attribute in a directory synchronization profile depends on the directory server on which Oracle Directory Integration Platform is running. In an active-active Oracle Directory Integration Platform configuration, you must manually update the `lastchangenumber` attribute in all instances.

The next section details the steps to copy the synchronization profiles and the provisioning profiles from the primary Oracle Internet Directory to the secondary Oracle Internet Directory in a multimaster replication deployment.

## Directory Synchronization Profiles

After you copy an export profile to a target node, you must update the `lastchangenumber` attribute with the target node value. To update the value:

1. Disable the synchronization profile.
2. Get the value of the `lastchangenumber` attribute on the target node using the `ldapsearch` command.
3. Use `ldapsearch` to get the LDIF dump of the profile entry.
4. Use `ldapadd` to add the profile to the other Oracle Internet Directory instance.
5. Use the `updatechgnum` operation of the `manageSyncProfiles` command to update the `lastchangenumber` attribute in the export profile you copied to the target node with the value you obtained in Step 2.
6. Enable the synchronization profile.

## Directory Provisioning Profiles

In a default multimaster Oracle Internet Directory replication environment, the Oracle Directory Integration Platform is installed in the same location as the primary Oracle Internet Directory. The information and steps in this section are applicable only when multimaster replication is set up.

If the primary node fails, event propagation stops for all profiles located on the node. Although the events are queued and not lost while the primary node is stopped, the events will not be propagated to any applications that expect them. To ensure that events continue to be propagated even when the primary node is down for the Version 1.0 and 2.0 profiles, the directory provisioning profiles must be copied to other secondary nodes.

However, directory provisioning profiles should only be copied from the primary node to any secondary nodes immediately after an application is installed and before any user changes are made in Oracle Internet Directory.

To synchronize the directory provisioning profiles between a primary node and any secondary nodes, you must do the following:

1. On the primary node, use the `ldifwrite` command to create an LDIF dump of the entries from this container:

```
cn=provisioning profiles,cn=changelog subscriber,cn=oracle internet directory
```

2. Copy the LDIF dump to the secondary node.
3. Use the `ldapadd` command to add the profiles on the secondary node.

### 8.5.2.2 Protection from Failures and Expected Behavior

This section discusses protection from different types of failure in an Oracle Directory Integration Platform active-active cluster.

**8.5.2.2.1 Process Failure** In a high availability environment, the Oracle Directory Integration Platform application is deployed on an Oracle WebLogic Server cluster comprised of at least two WebLogic instances.

By default, the Oracle Directory Integration Platform application leverages the high availability features of the underlying WebLogic Clusters. When Oracle Directory Integration Platform is deployed, the Quartz scheduler is started with a clustering option. When started with the clustering option, depending on the load on the node, the scheduler runs the job on any of the available nodes in the cluster. In case of

hardware or other failures on one or more nodes, the scheduler runs the jobs on the available nodes.

In addition, in a high availability environment, WebLogic Node Manager is configured to monitor the WebLogic servers. In case of failure, Node Manager restarts the WebLogic Server.

Within the Oracle Directory Integration Platform application, the Quartz Scheduler invokes the Provisioning or Synchronization EJBs that do the actual work. As soon as the Quartz scheduler invokes an EJB, it tags that EJB as executing the job. In case the EJB fails, the Quartz scheduler marks the job as failed and reschedules it to be executed later by another EJB.

**8.5.2.2.2 Expected Client Application Behavior When Failure Occurs** Oracle Directory Integration Platform failover is not transparent to end users.

Configuration changes made on a Oracle Directory Integration Platform instance deployed in a high availability topology are not propagated automatically to all the Oracle Directory Integration Platform instances in the topology.

It is recommended that you use the `manageDIPServerConfig` tool to make the same configuration changes on all Oracle Directory Integration Platform instances in the topology. This ensures that the configuration across all the Oracle Directory Integration Platform instances in the topology is uniform.

See *Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform* for more information about the `manageDIPServerConfig` tool.

**8.5.2.2.3 External Dependency Failure** Oracle Directory Integration Platform requires the back-end repository, Oracle Internet Directory, Credential Store Framework and the WebLogic Managed Server to be available during startup. Oracle Directory Integration Platform fails to start if any of these is not available.

### 8.5.2.3 Oracle Directory Integration Platform Prerequisites

This section describes prerequisites for setting up Oracle Directory Integration Platform in the high availability architecture.

---

**Note:** Oracle Directory Integration Platform uses Quartz to maintain its jobs and schedules in the database. For the Quartz jobs to be run on multiple Oracle Directory Integration Platform nodes in a cluster, it is required that the system clocks on the cluster nodes be synchronized.

---

Oracle Internet Directory should be installed and configured by following the instructions in these sections:

- [Section 8.3.2.3, "Oracle Internet Directory Prerequisites"](#)
- [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps"](#)

## 8.5.3 Oracle Directory Integration Platform and Oracle Directory Services Manager High Availability Configuration Steps

In a high availability environment, it is recommended that Oracle WebLogic Server utilities be used for clustering, load balancing, and failover of Oracle Directory Integration Platform instances and Oracle Directory Services Manager.

This section describes how to perform the following installation and configuration on IDMHOST1 and IDMHOST2:

- [Section 8.5.3.1, "Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST1"](#)
- [Section 8.5.3.2, "Creating boot.properties for the Administration Server on IDMHOST1"](#)
- [Section 8.5.3.3, "Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST2"](#)
- [Section 8.5.3.4, "Post-Installation Steps for Oracle Directory Integration Platform and Oracle Directory Services Manager"](#)
- [Section 8.5.3.5, "Installing Oracle Fusion Middleware Components on WEBHOST1 and WEBHOST2"](#)
- [Section 8.5.3.6, "Configuring Oracle HTTP Server on WEBHOST1 and WEBHOST2"](#)

It is not required to install Oracle Directory Integration Platform and Oracle Directory Services Manager in the same installation session, on the same Managed Server, or on the same host. The installation and configuration steps in this section are one example of installing Oracle Directory Integration Platform and Oracle Directory Services Manager in a two-node active-active cluster.

### 8.5.3.1 Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST1

Follow these steps to configure Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST1:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on IDMHOST1 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Ensure that port 7006 is not in use by any service on IDMHOST1 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":7006"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7006"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for port 7006 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port numbers for Oracle Directory Services Manager):

```
The port for ODSM server port
ODS Server Port No = 7006
```

7. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Create New Domain** and enter the domain details:
  - **User Name:** `weblogic`
  - **User Password:** `*****`
  - **Confirm Password:** `*****`
  - **Domain Name:** `IDMDomain`

Click **Next**.

10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be changed.
  - **Oracle Home Directory:** This value is prefilled and cannot be changed)

`ods`

- **WebLogic Server Directory:**

`/u01/app/oracle/product/fmw/wlserver_10.3`

- **Oracle Instance Location:**

`/u01/app/oracle/admin/ods_instance1`

- **Oracle Instance Name:**

`ods_instance1`

Click **Next**.

11. On the Specify Oracle Configuration Manager Details screen, specify these values:
  - **Email Address:** Provide the email address for your My Oracle Support account.
  - **Oracle Support Password:** Provide the password for your My Oracle Support account.
  - Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, select **Oracle Directory Integration Platform, Management Components - Oracle Directory Services Manager**. Deselect all the other components. Ignore the warning that says Oracle Internet Directory needs to be configured as a prerequisite.

Select the **Clustered** check box.

Click **Next**.

---

**Note:** The default Oracle WebLogic Server clustering mode set by the installer is unicast (not multicast).

---

13. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the filename for the `staticports.ini` file that you copied to the temporary directory.

Click **Next**.

14. On the Specify OID Details screen, specify the following:

- **Hostname:** `oid.mycompany.com`
- **Port:** `636`
- **Username:** `cn=orcladmin`
- **Password:** `*****`

Click **Next**.

15. On the Specify Schema Database screen, select **Use Existing Schema** and specify the following values:

- **Connect String:**

```
infradbhost1-vip.mycompany.com:1521:oiddb1^infradbhost2-vip.mycompany.com:1521:oiddb2@oid.mycompany.com
```

---

**Note:** The Oracle RAC database connect string information needs to be provided in the format  
*host1:port1:instance1^host2:port2:instance2@servicename.*

During this installation, it is not required for all the Oracle RAC instances to be up. If one Oracle RAC instance is up, the installation can proceed.

It is required that the information provided above is complete and accurate. Specifically, the correct host, port, and instance name must be provided for each Oracle RAC instance, and the service name provided must be configured for all the specified Oracle RAC instances.

Any incorrect information entered in the Oracle RAC database connect string has to be corrected manually after the installation.

---

- **User Name:** `ODSSM`
- **Password:** `*****`

Click **Next**.

16. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.
17. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

18. On the Configuration screen, multiple configuration assistants are launched in succession; this process can be lengthy. When it completes, click **Next**.
19. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

### 8.5.3.2 Creating `boot.properties` for the Administration Server on IDMHOST1

This section describes how to create a `boot.properties` file for the Administration Server on IDMHOST1. The `boot.properties` file enables the Administration Server to start without prompting for the administrator username and password.

Follow these steps to create the `boot.properties` file:

1. On IDMHOST1, go the following directory:

```
MW_HOME/user_projects/domains/domainName/servers/AdminServer/security
```

For example:

```
cd /u01/app/oracle/product/fmw/user_
projects/domains/IDMDomain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` under the `security` directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

---

---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

---

3. Stop the Administration Server if it is running.  
See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.
4. Start the Administration Server on IDMHOST1 using the `startWebLogic.sh` script located under the `MW_HOME/user_projects/domains/domainName/bin` directory.
5. Validate that the changes were successful by opening a web browser and accessing the following pages:
  - WebLogic Server Administration Console at:



`http://oidhost1.mycompany.com:7001/console`

- Oracle Enterprise Manager Fusion Middleware Control at:

`http://oidhost1.mycompany.com:7001/em`

Log into these consoles using the `weblogic` user credentials.

### 8.5.3.3 Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST2

To configure Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST2:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on IDMHOST2 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: `./config.sh`

On Windows, double-click `config.exe`

4. On the Welcome screen, click **Next**.
5. On the Select Installation Type screen, select **Install & Configure** then click **Next**.
6. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, please fix them and restart your installation.

Click **Next**.

7. On the Select Domain screen, select the **Expand Cluster** option and specify the values shown in the example below:

- **HostName:** `idmhost1.mycompany.com`
- **Port:** `7001`
- **UserName:** `weblogic`
- **User Password:** `<password for the weblogic user>`

Click **Next**.

---

**Note:** Oracle WebLogic Administration Server runs on IDMHOST1.

---

8. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** Prefilled and cannot be changed.  
`/u01/app/oracle/product/fmw`
  - **Oracle Home Directory:** Prefilled and cannot be changed.  
`ods`

- **WebLogic Server Directory:**

```
/u01/app/oracle/product/fmw/wlserver_10.3
```

- **Oracle Instance Location:**

```
/u01/app/oracle/admin/ods_instance2
```

- **Oracle Instance Name:**

```
ods_instance2
```

Click **Next**.

9. On the Configure Components screen, de-select all products except **Oracle DIP and Management Components**.

Click **Next**.

10. On the Configure Ports screen, select **Specify Ports Using Configuration File** and enter the full pathname to the `staticports.ini` file that you edited in the temporary directory.

Click **Next**.

11. On the Installation Summary screen, review the choices you made. If you must make any changes, click **Back**. If you made the correct selections, click **Install**.
12. On the Installation Progress screen, view the progress of the installation.
13. On UNIX systems, once the installation is done, the `oracleRoot.sh` confirmation dialog box opens. It explains that a configuration script needs to be run as root before the installation can proceed.

Leave the confirmation dialog box open, open another shell window, log in as root, and run this script file:

```
/u01/app/oracle/product/fmw/ods/oracleRoot.sh
```

14. On the Configuration Progress screen, view the progress of the configuration.
15. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

### 8.5.3.4 Post-Installation Steps for Oracle Directory Integration Platform and Oracle Directory Services Manager

In the previous section, the installer created a second Managed Server, `wls_ods2` on `IDMHOST2`. However, the Oracle Directory Integration Platform application is not deployed on `IDMHOST2` and the newly created Managed Server is not automatically started. Also, the WebLogic Administration Console shows the state of the `wls_ods2` Managed Server on `IDMHOST2` as `UNKNOWN`.

Follow the post-installation steps in this section to complete the installation and configuration of the Oracle Directory Integration Platform and Oracle Directory Services Manager applications on `IDMHOST2`.

#### 8.5.3.4.1 Copy the Oracle Directory Integration Platform Configuration from `IDMHOST1` to `IDMHOST2`

To copy the Oracle Directory Integration Platform directly application from `IDMHOST1` to `IDMHOST2`:

Copy the following directory on `IDMHOST1`: `MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_ods1/applications` to the following location on `IDMHOST2`: `MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_ods2`.

For example, from IDMHOST1, execute the following command:

```
scp -rp MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_ods1/applications/odsm user@IDMHOST2:/MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_ods2
```

**8.5.3.4.2 Restart the Managed Server on IDMHOST2 in a Cluster** To restart the newly-created wls\_ods2 Managed Server in a cluster on IDMHOST2:

1. Use this URL to access the Oracle WebLogic Server Administration Console:

`http://idmhost1.mycompany.com:7001/console`

Log into the console using the administrator's credentials.

2. In the left pane, expand **Environment** and select **Clusters**.

See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.

3. Click on the link for the cluster (cluster\_ods) containing the Managed Server (wls\_ods2) you want to stop.
4. Select **Control**.
5. Under **Managed Server Instances in this Cluster**, select the check box next to the Managed Server (wls\_ods2) you want to start and click **Start**.
6. On the Cluster Life Cycle Assistant page, click **Yes** to confirm.
7. Node Manager starts the server on the target machine. When the Node Manager finishes its start sequence, the server's state is indicated in the **State** column in the Servers status table. The state should be RUNNING.

### 8.5.3.5 Installing Oracle Fusion Middleware Components on WEBHOST1 and WEBHOST2

This section describes how to install the required binaries for the Oracle Home (ORACLE\_HOME) for Oracle HTTP Server.

Oracle strongly recommends that you read the release notes for any additional installation and deployment considerations prior to starting the setup process.

**8.5.3.5.1 Installing Oracle HTTP Server for the Web Tier** This section describes how to install Oracle HTTP Server.

Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.

On Linux platforms, if the `/etc/oraInst.loc` file exists, verify that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

Start the installer for Oracle Web Tier components.

```
HOST1> ./runInstaller
```

Then perform these installation steps:

1. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
2. On the Specify Installation Location screen, enter the following values:

**MW\_HOME:** Enter the value of the MW\_HOME, for example:

`/u01/app/product/fmw`

Select the previously installed Middleware Home from the drop-down list. For the Oracle HTTP Server Oracle Home (OHS\_ORACLE\_HOME) directory, enter the directory name `WEB`.

Click **Next**.

3. On the Summary screen, click **Install**.

When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the `root` user.

4. On the Installation Complete screen, click **Finish**.

**8.5.3.5.2 Upgrading the Oracle HTTP Server Oracle Home to Patch Set 3** To upgrade your Oracle HTTP Server software to 11.1.1.4 (Patch Set 3):

1. Start the Web Tier Upgrade Installer by running `./runinstaller`.
2. On the Welcome screen, click **Next**.
3. On the Prerequisite Checks screen, click **Next**.
4. On the Specify Install Location screen, provide the path to the Oracle Middleware Home and the name of the Oracle HTTP Server Oracle Home directory.
5. On the Installation Summary screen, validate your selections, and then click **Install**.
6. The Installation Progress screen shows the progress of the install.

Once the installation is done, the `oracleRoot.sh` confirmation dialog box appears. This dialog box advises you that a configuration script needs to be run as `root` before the installation can proceed. Leaving the confirmation dialog box open, open another shell window, log in as `root`, and run this script file:  
`/u01/app/oracle/product/fmw/id/oracleRoot.sh`. After the script completes, click **OK** on the Confirmation Dialog box.

7. On the Installation Complete screen, click **Finish** to exit.

### 8.5.3.6 Configuring Oracle HTTP Server on WEBHOST1 and WEBHOST2

This section describes how to configure Oracle HTTP Server on WEBHOST1 and WEBHOST2. Oracle HTTP Server is not required if Oracle Directory Integration Platform is the only component being installed.

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle HTTP Server software has been installed and upgraded on WEBHOST1 and WEBHOST2 as described in [Section 8.5.3.5, "Installing Oracle Fusion Middleware Components on WEBHOST1 and WEBHOST2."](#)

3. Ensure that port 7777 is not in use by any service on WEBHOST1 or WEBHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":7777"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr ":7777"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for port 7777 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port number for Oracle HTTP Server):

```
The http_main port for ohs component
OHS Port = 7777
```

7. Start the Configuration Assistant for Oracle Fusion Middleware 11g Web Tier utilities located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.

9. On the Configure Components screen:

- Select **Oracle HTTP Server**.
- Select **Associate Selected Components with Weblogic Domain**.

Click **Next**.

10. On the Specify WebLogic Domain screen:

Enter the location where you installed Oracle WebLogic Server. Note that the Administration Server must be running.

- **Domain Host Name:** IDMHOST1
- **Domain Port No:** 7001
- **User Name:** weblogic
- **Password:** \*\*\*\*\*

Click **Next**.

11. On the Specify Component Details screen:

- Enter the following values for WEBHOST1:

- **Instance Home Location:** /u01/app/oracle/admin/ohs\_inst1
- **Instance Name:** ohs\_inst1
- **OHS Component Name:** ohs1
- Enter the following values for WEBHOST2:
  - **Instance Home Location:** /u01/app/oracle/admin/ohs\_inst2
  - **Instance Name:** ohs\_inst2
  - **OHS Component Name:** ohs2

Click **Next**.

**12.** On the Specify Webtier Port Details screen:

- Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
- Enter the Oracle HTTP Server port, for example, 7777.

Click **Next**.

**13.** On the Oracle Configuration Manager screen, enter the following:

- **Email Address:** Provide the email address for your My Oracle Support account
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- **I wish to receive security updates via My Oracle Support:** Click this checkbox.

**14.** On the Installation Summary screen, ensure that the selections are correct. If they are not, click **Back** and make the necessary fixes. After ensuring that the selections are correct, click **Next**.

**15.** On the Installation Progress screen on UNIX systems, a dialog appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

**16.** On the Configuration Progress screen, multiple configuration assistants are launched in succession; this process can be lengthy.

**17.** On the Configuration Completed screen, click **Finish** to exit.

#### 8.5.3.6.1 Configuring Oracle HTTP Server for Oracle Directory Services Manager High Availability

Follow the instructions in this section to configure Oracle HTTP Server for Oracle Directory Services Manager high availability.

**1.** Configure Oracle HTTP Server to use the load balancing router virtual hosts.

The Oracle HTTP Server instances on WEBHOST1 and WEBHOST2 should be configured to use the virtual hosts set up in the load balancer. Refer to [Section 8.2.5.4, "Configuring Virtual Server Names and Ports for the Load Balancer"](#) for more information about the virtual hosts.

To configure the Oracle HTTP Server instances to use the load balancing router virtual hosts, edit the `httpd.conf` file to define the virtual host directives as follows:

```
NameVirtualHost *:7777
```

```
<VirtualHost *:7777>
 ServerName admin.mycompany.com:7777
 ServerAdmin you@your.address
 RewriteEngine On
 RewriteOptions inherit
</VirtualHost>
```

The `httpd.conf` file is located under the `ORACLE_  
INSTANCE/config/OHS/componentName` directory on both `WEBHOST1` and `WEBHOST2`.

2. Configure Oracle HTTP Server to route to Oracle Directory Services Manager, Oracle Enterprise Manager Fusion Middleware Control, and the Oracle WebLogic Server Administration Console.

To enable the Oracle HTTP Server instances to route to the Oracle Directory Services Manager applications on `IDMHOST1` and `IDMHOST2`, add the directives below to the `mod_wl_ohs.conf` file located under the `ORACLE_  
INSTANCE/config/OHS/componentName` directory on `WEBHOST1` and `WEBHOST2`:

```
<IfModule weblogic.module>
WebLogicHost IDMHOST1.MYCOMPANY.COM
WebLogicPort PORT
</IfModule>

<Location /odsm>
SetHandler weblogic-handler
WebLogicCluster IDMHOST1.MYCOMPANY.COM:<PORT>, IDMHOST2.MYCOMPANY.COM:<PORT>
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

#AdminServer and EM
<Location /console>
SetHandler weblogic-handler
WebLogicHost VIP1
WebLogicPort 7001
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

<Location /consolehelp>
SetHandler weblogic-handler
WebLogicHost VIP1
WebLogicPort 7001
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

<Location /em>
SetHandler weblogic-handler
WebLogicHost VIP1
WebLogicPort 7001
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

<Location /odsm-config>
SetHandler weblogic-handler
```

```
WebLogicCluster IDMH0ST1.MYCOMPANY.COM:<PORT>, IDMH0ST2.MYCOMPANY.COM:<PORT>
</Location>
```

3. Stop and start Oracle HTTP Server using the `opmnctl` command:

```
ORACLE_INSTANCE/bin/opmnctl stopall
ORACLE_INSTANCE/bin/opmnctl startall
```

4. Validate that you can access the consoles using the load balancing router virtual host:

Oracle Directory Services Manager Console:

`http://admin.mycompany.com:7777/odsm`

Oracle WebLogic Server Administration Console:

`http://idmhost1.mycompany.com:7001/console`

Oracle Enterprise Manager Fusion Middleware Control:

`http://idmhost1.mycompany.com:7001/em`

---

**Note:** When Oracle Directory Integration Platform deployed in a high availability configuration is enabled for SSL with server-only Authentication (that is, SSL Mode2), it is required to use a common key store that contains the certificates of all the Oracle Internet Directory instances. You can create a common keystore by importing the certificates from all the Oracle Internet Directory instances and then copying the keystore to all nodes where Oracle Directory Integration Platform is running.

For more information about using SSL with Oracle Directory Integration Platform, refer to the *Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform*.

---

## 8.5.4 Oracle Directory Integration Platform Failover and Expected Behavior

In a high availability environment, the Oracle Directory Integration Platform application is deployed on an Oracle WebLogic Server cluster comprised of at least two WebLogic instances.

By default, the Oracle Directory Integration Platform application leverages the high availability features of the underlying WebLogic Clusters. In case of hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

In addition, in a high availability environment, WebLogic Node Manager is configured to monitor the WebLogic servers. In case of failure, Node Manager restarts the WebLogic Server. In case of a database instance failure, the surviving Oracle RAC node takes over any remaining processes. There may be innocuous errors in the Managed Servers logs during an Oracle RAC failover, as discussed in [Section 8.5.5, "Troubleshooting Oracle Directory Integration Platform High Availability."](#)

## 8.5.5 Troubleshooting Oracle Directory Integration Platform High Availability

This section describes how to deal with issues involving Oracle Directory Integration Platform high availability.



### 8.5.5.1 Managed Server Log File Exceptions Received for Oracle Directory Integration Platform During an Oracle RAC Failover

During an Oracle RAC failover, exceptions similar to the ones below are seen in the Managed Server log files running the Oracle Directory Integration Platform application. These errors are thrown when the multi data sources configured on the WebLogic Server platform try to verify the health of the Oracle RAC database instances during failover. These are innocuous errors that you can ignore. The Oracle Directory Integration Platform application will recover and begin to operate normally after a lag of one or two minutes. During an Oracle RAC failover, there will be no Oracle Directory Integration Platform down time if one Oracle RAC instance is running at all times.

```
RuntimeException:
[2008-11-21T00:11:10.915-08:00] [wls_ods] [ERROR] []
[org.quartz.impl.jdbcjobstore.JobStoreTX] [tid: 25] [userId: <anonymous>]
[ecid: 0000Hqy69UiFW7V6u3FCEH199aj0000009,0] [APP: DIP] ClusterManager: Error
managing cluster: Failed to obtain DB connection from data source
'schedulerDS': java.sql.SQLException: Could not retrieve datasource via JNDI
url 'jdbc/schedulerDS' java.sql.SQLException: Cannot obtain connection:
driverURL = jdbc:weblogic:pool:schedulerDS, props =
{EmulateTwoPhaseCommit=false, connectionPoolID=schedulerDS,
jdbcTxDataSource=true, LoggingLastResource=false,
dataSourceName=schedulerDS}.[]
Nested Exception: java.lang.RuntimeException: Failed to setAutoCommit to true
for pool connection
```

```
AuthenticationException while connecting to OID:
[2008-11-21T00:12:08.812-08:00] [wls_ods] [ERROR] [DIP-10581] [oracle.dip]
[tid: 11] [userId: <anonymous>] [ecid: 0000Hqy6m54FW7V6u3FCEH199ap0000000,0]
[APP: DIP] DIP was not able to get the context with the given details {}[]
javax.naming.AuthenticationException: [LDAP: error code 49 - Invalid
Credentials]
```

Most of the exceptions will be related to the scheduler or LDAP, for example:

1. Could not retrieve datasource via JNDI url 'jdbc/schedulerDS'  
java.sql.SQLException.
2. javax.naming.AuthenticationException: [LDAP: error code 49 - Invalid Credentials]

### 8.5.5.2 Dealing with Error Messages Received After Starting WebLogic Node Manager

If you see the following error message after starting WebLogic Node Manager, follow the steps that appear below after the error message:

```
<Dec 15, 2008 8:40:05 PM> <Warning> <Uncaught exception in server handler:
javax.net.ssl.SSLKeyException: [Security:090482]BAD_CERTIFICATE alert was
received from stbee21.us.oracle.com - 152.68.64.2155. Check the peer to
determine why it rejected the certificate chain (trusted CA configuration,
hostname verification). SSL debug tracing may be required to determine the
exact reason the certificate was rejected.> javax.net.ssl.SSLKeyException:
[Security:090482]BAD_CERTIFICATE alert was received from stbee21.us.oracle.com -
152.68.64.215. Check the peer to determine why it rejected the certificate chain
(trusted CA configuration, hostname verification). SSL debug tracing may be
required to determine the exact reason the certificate was rejected.
```

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.

2. In the left pane of the Console, expand **Servers** and AdminServer (admin).
3. Select the **Configuration > SSL > Advanced Link**.
4. Select **None** for **Hostname Verification**.
5. Click **Save** to save the setting.
6. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.
7. Restart all servers.

If the Managed Server is started in Admin mode, perform these steps:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Servers** and the name of the server that is running in ADMIN mode.
3. Select the **Control > Start/Stop** tab.
4. Select the name of the server.
5. Click **Resume**.
6. Select **Yes** to resume servers.

#### 8.5.5.3 If WebLogic Node Manager Fails to Start

If WebLogic Node Manager fails to start, ensure that you copied the following domains file from IDMHOST1 to IDMHOST2:

```
WL_HOME/common/nodemanager/nodemanager.domains
```

#### 8.5.5.4 Configuration Changes Do Not Automatically Propagate to All Oracle Directory Integration Platform Instances in a Highly Available Topology

When you change the configuration of an Oracle Directory Integration Platform instance in a high availability topology, the configuration change does not propagate automatically to all the Oracle Directory Integration Platform instances in the topology.

Instead, use the manageDIPServerConfig tool to make the same configuration change for each Oracle Directory Integration Platform instance in the topology, which will ensure the same configuration across all the Oracle Directory Integration Platform instances in the topology. See *Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform* for more information about the manageDIPServerConfig tool.

#### 8.5.5.5 Operation Cannot Be Completed for Unknown Errors Message

Sometimes the following error message appears intermittently when you use the manageSyncProfiles command:

```
OPERATION CANNOT BE COMPLETED FOR UNKNOWN ERRORS
```

When you see this error message, start and stop the Managed Server (wls\_ods1 or wls\_ods2). If the problem persists, repeat the copy method on the second node.

## 8.6 Oracle Directory Services Manager High Availability

This section introduces Oracle Directory Services Manager and describes how to design and deploy a high availability environment for Oracle Directory Integration Platform and Oracle Directory Services Manager.

See [Section 8.5, "Oracle Directory Integration Platform High Availability"](#) for more information about Oracle Directory Integration Platform.

This section includes the following topics:

- [Section 8.6.1, "Oracle Directory Services Manager Component Architecture"](#)
- [Section 8.6.2, "Oracle Directory Services Manager High Availability Concepts"](#)
- [Section 8.6.3, "Oracle Directory Services Manager High Availability Configuration Steps"](#)
- [Section 8.6.4, "Validating Oracle Directory Services Manager High Availability"](#)
- [Section 8.6.5, "Oracle Directory Services Manager Failover and Expected Behavior"](#)
- [Section 8.6.6, "Troubleshooting Oracle Directory Services Manager"](#)
- [Section 8.6.7, "Additional Considerations for Oracle Directory Services Manager High Availability"](#)


### 8.6.1 Oracle Directory Services Manager Component Architecture

Oracle Directory Services Manager is a unified graphical user interface (GUI) for managing instances of Oracle Internet Directory and Oracle Virtual Directory. It is a replacement for Oracle Directory Manager, which is now deprecated. Oracle Directory Services Manager enables you to configure the structure of the directory, define objects in the directory, add and configure users, groups, and other entries.

Oracle Directory Services Manager is the interface used by end users to manage their directory entries, schemas, security, and other features.

[Figure 8–8](#) shows the Oracle Directory Services Manager architecture.

**Figure 8–8 Oracle Directory Services Manager Architecture**



[Figure 8–8](#) shows Oracle Directory Services Manager deployed in non-high availability architecture. Oracle Directory Services Manager is deployed on the Oracle WebLogic Server. Oracle Directory Services Manager is configured to communicate with the Oracle Virtual Directory and the Oracle Internet Directory instances it manages.

Oracle Directory Services Manager uses the HTTP(s) protocol to communicate with client browsers. It uses the LDAP(s) protocol to communicate with Oracle Internet Directory and over web services to communicate with Oracle Virtual Directory.

### 8.6.1.1 Oracle Directory Services Manager Component Characteristics

Oracle Directory Services Manager is an Oracle Application Development Framework (ADF)-based J2EE application that is deployed on top of the Oracle WebLogic Server. By default, Oracle Directory Services Manager is deployed to the Administration Server within the WebLogic domain, but depending on the requirements in your environment, it can also be deployed to a Managed Server.

Oracle Directory Services Manager can be deployed on the same node with Oracle Internet Directory or on a separate node.

Oracle Directory Services Manager can be invoked directly or from Oracle Enterprise Manager Fusion Middleware Control. Supported browsers include Firefox 2, Firefox 3, and Internet Explorer 7.

To invoke Oracle Directory Services Manager directly, enter the following URL into your browser's address field:

```
http://host:port/odsm/faces/odsm.jspx
```

where:

- *host* is the name of the Managed Server where Oracle Directory Services Manager is running.
- *port* is the Managed Server port number.

To invoke Oracle Directory Services Manager from Oracle Enterprise Manager Fusion Middleware Control:

- Select **Directory Services Manager** from the **Oracle Internet Directory** menu in the Oracle Internet Directory target, then **Data Browser**, **Schema**, **Security**, or **Advanced**.
- From the **Oracle Virtual Directory** menu in the Oracle Virtual Directory target, select **Directory Services Manager**, then **Data Browser**, **Schema**, **Adapter**, **Extensions**, or **Quick configuration wizard**.

Oracle Directory Services Manager is deployed to an Oracle WebLogic Server as an externally staged application. The WebLogic server manages the startup, shutdown, monitoring of the Oracle Directory Services Manager application. Oracle Directory Services Manager initializes when the Managed Server starts up. Oracle Node Manager is configured to monitor the server process and restarts it in case of failure.

**8.6.1.1.1 Lifecycle Management** The lifecycle events for the Oracle Directory Services Manager application can be managed using one or more of the command line tools and consoles listed below:

The following tools can be used to start and stop Oracle Directory Services Manager processes:

- Oracle WebLogic Server Scripting Tool (WLST)
- Oracle Enterprise Manager Fusion Middleware Control
- WebLogic Server Administration Console
- WebLogic Node Manager

**8.6.1.1.2 Oracle Directory Services Manager Log File** Oracle Directory Services Manager messages are logged in the server log file of the Oracle WebLogic Server where it is running. The default location of the server log is:

```
WEBLAGIC_SERVER_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

## 8.6.2 Oracle Directory Services Manager High Availability Concepts


This section provides conceptual information about using Oracle Directory Services Manager in a high availability configuration.

In the Oracle Directory Services Manager high availability configuration described in this section, Oracle Directory Services Manager and Oracle Directory Integration Platform are installed and configured on two hosts in a two-node high availability active-active configuration.

### 8.6.2.1 Oracle Directory Services Manager High Availability Architecture

[Figure 8–9](#) shows the Oracle Directory Integration Platform and Oracle Directory Services Manager high availability architecture in an active-active configuration.

**Figure 8–9 Oracle Directory Services Manager and Oracle Directory Integration Platform in a High Availability Architecture**



In [Figure 8–9](#), the application tier includes the IDMHOST1 and IDMHOST2 computers.

On IDMHOST1, the following installations have been performed:

- An Oracle Directory Integration Platform instance and Oracle Directory Services Manager instance have been installed on the WLS\_ODS1 Managed Server. The Oracle RAC database has been configured in a JDBC multi data source to protect the instances from Oracle RAC node failure.
- A WebLogic Administration Server has been installed. Under normal operations, this is the active Administration Server.

On IDMHOST2, the following installations have been performed:

- An Oracle Directory Integration Platform instance and Oracle Directory Services Manager instance have been installed in the WLS\_ODS2 Managed Server. The

Oracle RAC database has been configured in a JDBC multi data source to protect the instances from Oracle RAC node failure.

The instances in the WLS\_ODS2 Managed Server on IDMHOST2 and the instances in the WLS\_ODS1 Managed Server on IDMHOST1 are configured as the CLUSTER\_ODS cluster.

- A WebLogic Administration Server has been installed. Under normal operations, this is the passive Administration Server. You will make this Administration Server active if the Administration Server on IDMHOST1 becomes unavailable.

**8.6.2.1.1 Starting and Stopping the Cluster** In a high availability architecture, Oracle Directory Integration Platform and Oracle Directory Services Manager are deployed on an Oracle WebLogic Cluster that has at least two servers as a part of the cluster.

By default, the WebLogic Server starts, stops and monitors the applications. By default, both the Oracle Directory Integration Platform and Oracle Directory Services Manager applications leverage the high availability features of the underlying WebLogic Clusters. In case of hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

In a high availability environment, WebLogic Node Manager is configured to monitor the WebLogic servers. In case of failure, Node Manager restarts the WebLogic Server. If Node Manager cannot restart the server, then the front-ending load balancing router detects failure of a WebLogic instance in the Cluster and routes traffic to surviving instances.

### 8.6.2.2 Protection from Failures and Expected Behaviors

This section describes protection from different types of failure in an Oracle Directory Services Manager active-active cluster.

**8.6.2.2.1 Process Failure** In a high availability environment, the Oracle Directory Services Manager applications are deployed on an Oracle WebLogic Server cluster comprised of at least two WebLogic instances.

By default, the Oracle Directory Services Manager applications leverage the high availability features of the underlying WebLogic Clusters. In case of hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

In addition, in a high availability environment, WebLogic Node Manager is configured to monitor the WebLogic servers. In case of failure, Node Manager restarts the WebLogic Server. If Node Manager cannot restart the server, then `mod_wl_ohs`, which is configured as a part of Oracle HTTP Server, routes the request to the surviving WebLogic instance.

OPMN monitors the Oracle HTTP Server processes and restarts the process in case of failure. If OPMN is unable to restart the HTTP process, the front-ending load balancing router detects the failure of an Oracle HTTP Server instance and routes traffic to surviving instances.

Oracle Directory Services Manager maintains a session state, but in case of failure, the session state information is not carried over to the surviving node.

**8.6.2.2.2 Expected Client Application Behavior When Failure Occurs** Oracle Directory Services Manager failover is not transparent. You have to reestablish the connection during an Oracle WebLogic Server instance failover using Oracle Directory Services Manager.

**8.6.2.2.3 Expected Dependency Failure** Oracle Directory Services Manager requires the WebLogic Managed Server to be available during startup. If it is not available, Oracle Directory Services Manager fails to start.

### 8.6.2.3 Oracle Directory Services Manager Prerequisites

This section describes prerequisites for setting up Oracle Directory Services Manager in the high availability architecture.

---

**Note:** Oracle requires that you synchronize the system clocks on the cluster nodes when you are using Oracle Directory Services Manager in a high availability deployment.

---

Oracle recommends that you install and configure the components listed below before you install Oracle Directory Services Manager. See the sections below for the prerequisites as well as installation and configuration steps:

- Oracle database  
Install and configure an Oracle Real Application Clusters database by following the instructions in these sections:
  - [Section 8.2, "Prerequisites for Oracle Identity Management High Availability Configuration"](#)
  - [Section 8.2.2, "Database Prerequisites"](#)
  - [Section 8.2.3, "Installing and Configuring the Database Repository"](#)
- Oracle Repository Creation Utility  
Install the Oracle Repository Creation Utility to create and load the required schemas by following the instructions in these sections:
  - [Section 8.2.4, "Obtaining the Repository Creation Utility Software"](#)
  - [Section 8.2.5, "Configuring the Database for Oracle Fusion Middleware 11g Metadata"](#)
- Load balancers and virtual hosts  
See [Section 8.2.5.4, "Configuring Virtual Server Names and Ports for the Load Balancer"](#) for information about configuring the load balancer and creating the required virtual server.
- Oracle Internet Directory  
Install and configure Oracle Internet Directory by following the instructions in these sections:
  - [Section 8.3.2.3, "Oracle Internet Directory Prerequisites"](#)
  - [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps"](#)

## 8.6.3 Oracle Directory Services Manager High Availability Configuration Steps

See [Section 8.5.3, "Oracle Directory Integration Platform and Oracle Directory Services Manager High Availability Configuration Steps"](#) for the steps for installing and configuring the high availability active-active configuration shown in [Figure 8–9](#).



## 8.6.4 Validating Oracle Directory Services Manager High Availability

This section describes how to validate Oracle Directory Services Manager in a high availability configuration.

### 8.6.4.1 Performing a WebLogic Server Instance Failover

While you are accessing Oracle Directory Services Manager through Oracle HTTP Server, you can use the steps in this section to fail over a WebLogic Server instance and validate that Oracle Directory Services Manager is still accessible.

The Oracle HTTP Server virtual server name in this example is:

`http://admin.mycompany.com`

Perform these steps:

1. Access Oracle Directory Services Manager through the Oracle HTTP Server virtual server name:

`http://admin.mycompany.com/odsm/faces/odsm.jspx`

2. When the Welcome to Oracle Directory Services Manager screen is displayed, open a web browser and enter the following URL:

`http://hostname:port/console`

where `hostname` is the DNS name of the Administration Server and `port` is the address of the port on which the Administration Server is listening for requests (7001 by default).

3. On the login page, enter the user name and the password you used to start the Administration Server and click **Log In**.
4. Shut down one of the Managed Servers where Oracle Directory Services Manager is deployed by following these steps:
  - a. In the left pane of the Oracle WebLogic Server Administration Console, expand **Environment** and select **Servers**.  
See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.
  - b. Click the name of the Managed Server that you want to shut down. For example, `wls_ods1`.
  - c. In the settings for `wls_ods1` screen, select the **Control --> Start/Stop** tab.
  - d. Select the check box next to the name of the Managed Server (`wls_ods1`) that you want to shut down and click **Shutdown > When work completes**.
5. Check the status of the Managed Server (`wls_ods1`):
  - a. In the left pane of the Console, expand **Environment** and select **Servers**.
  - b. The State for the Managed Server (`wls_ods1`) should be **SHUTDOWN**.
6. Access Oracle Directory Services Manager through the Oracle HTTP Server virtual server name:

`http://admin.mycompany.com/odsm/faces/odsm.jspx`

The Welcome to Oracle Directory Services Manager screen is displayed.

### 8.6.4.2 Performing an Oracle RAC Database Failover

While you are accessing Oracle Directory Services Manager through the load balancing router, you can follow the steps in this section to fail over one Oracle RAC database instance at a time and ensure that Oracle Directory Services Manager is still functional. This example checks Oracle Directory Services Manager as well as Oracle Internet Directory access to the database.

The Oracle HTTP Server virtual server name in this example is:

`http://admin.mycompany.com`

The LDAP virtual server name in this example is:

`oid.mycompany.com:389`

Perform these steps:

1. Access Oracle Directory Services Manager through the Oracle HTTP Server virtual server name:

`http://admin.mycompany.com/odsm/faces/odsm.jspx`

The Welcome to Oracle Directory Services Manager screen is displayed.

2. Verify that you can connect to Oracle Internet Directory using the LDAP virtual server:

- a. Select the **Connect to a directory --> Create A New Connection** link in the upper right hand corner.

- b. In the New Connection screen, fill in the connection information below and click **Connect**:

- \* **Directory Type:** OID
- \* **Name:** OIDHA
- \* **Server:** oid.mycompany.com
- \* **Port:** 389
- \* **SSL Enabled:** Leave blank
- \* **User Name:** cn=orcladmin
- \* **Password:** \*\*\*\*\*
- \* **Start Page:** Home (default)

3. Shut down the Oracle Internet Directory Schema database instance on INFRADBHOST1-VIP by running the following `srvctl` commands (in this example, the database name is INFRADB):

```
ORACLE_HOME/bin/srvctl stop instance -d infradb -i infradb1
ORACLE_HOME/bin/srvctl status database
```

4. Refresh the Oracle Directory Services Manager screen or navigate to one of the tabs (Home, Data Browser, Schema, Security, Advanced). You should still be able to access the Oracle Internet Directory information.

5. Restart the Oracle Internet Directory Schema database instance on INFRADBHOST1-VIP by running the following `srvctl` commands:

```
ORACLE_HOME/bin/srvctl start instance -d infradb -i infradb1
ORACLE_HOME/bin/srvctl status database
```

6. Repeat steps 3, 4, and 5 for INFRADBHOST2-VIP.

## 8.6.5 Oracle Directory Services Manager Failover and Expected Behavior

This section provides steps for using Oracle Directory Services Manager to validate a failover of a Managed Server, to validate a failover of an Oracle Internet Directory instance, and to validate a failover of an Oracle RAC database.

### 8.6.5.1 Using Oracle Directory Services Manager to Validate a Failover of a Managed Server

To use Oracle Directory Services Manager to validate a failover of a Managed Server:

1. In a web browser, launch the Oracle Directory Services Manager page using the Oracle HTTP Server host and port. For example:  
`http://WEBHOST1:PORT/odsm/faces/odsm.jspx`
2. Make a connection to Oracle Internet Directory.
3. Go to the Administration Console and stop the `wls_ods1` Managed Server.
4. Go back to the Oracle Directory Services Manager page and continue your work.
5. The Oracle Directory Services Manager page will be disconnected.
6. Re-launch the Oracle Directory Services Manager page using the Oracle HTTP Server host and port again from the same browser.
7. Reestablish a new connection.

The current behavior is the expected behavior. Oracle Directory Services Manager failover is not transparent. You must reestablish the connection during a WebLogic Server instance failover using Oracle Directory Services Manager.

### 8.6.5.2 Using Oracle Directory Services Manager to Validate a Failover of an Oracle Internet Directory Instance

To use Oracle Directory Services Manager to validate a failover of an Oracle Internet Directory instance:

1. In a web browser, launch the Oracle Directory Services Manager page using the Oracle HTTP Server host and port. For example:  
`http://WEBHOST1:PORT/odsm/faces/odsm.jspx`
2. Make a connection to the Oracle Internet Directory hardware load balancer.
3. Shut down one Oracle Internet Directory instance at a time.
4. During the failover, go back to the Oracle Directory Services Manager page and continue your work.

It is expected behavior when Oracle Directory Services Manager opens a window with a message that Oracle Internet Directory is down. Oracle Directory Services Manager reconnects to Oracle Internet Directory but the connection may not persist during the failover. For additional information, see [Section 8.6.6.4, "Oracle Directory Services Manager Displays "LDAP Server is down" Message During Oracle Internet Directory Failover."](#)

While accessing Oracle Directory Services Manager, fail over the Oracle Internet Directory instances one at a time and ensure the LDAP store is still accessible. Oracle

Directory Services Manager might not have a persistent connection to Oracle Internet Directory.

### 8.6.5.3 Using Oracle Directory Services Manager to Validate an Oracle RAC Failover

To use Oracle Directory Services Manager to validate an Oracle RAC failover:

1. In a web browser, launch the Oracle Directory Services Manager page using the Oracle HTTP Server host and port. For example:

```
http://WEBHOST1:PORT/odsm/faces/odsm.jspx
```

2. Connect to the Oracle Internet Directory from the Oracle Directory Services Manager page.
3. Shut down one Oracle RAC database instance at a time.
4. Go back to the Oracle Directory Services Manager page and continue your work from the Oracle Internet Directory connection established in Step 2.

It is expected behavior for Oracle Directory Services Manager to temporarily lose its connection during an Oracle RAC failover. See [Section 8.6.6.5, "Oracle Directory Services Manager Temporarily Loses Its Connection During Oracle RAC Failover"](#) for more information about the error message that opens in Oracle Directory Services Manager during an Oracle RAC failover.

While accessing Oracle Directory Services Manager through the hardware load balancer, perform a failover on one Oracle RAC database instance at a time and ensure that Oracle Directory Services Manager is still functional. This will check Oracle Directory Services Manager as well as Oracle Internet Directory access to the Oracle RAC database.

## 8.6.6 Troubleshooting Oracle Directory Services Manager

This section describes how to deal with issues involving Oracle Directory Services Manager high availability.

### 8.6.6.1 Dealing with Error Messages Received After Starting WebLogic Node Manager

If the following error message appears after you start WebLogic Node Manager, follow the steps that appear below after the error message:

```
<Dec 15, 2008 8:40:05 PM> <Warning> <Uncaught exception in server handler:
javax.net.ssl.SSLKeyException: [Security:090482]BAD_CERTIFICATE alert was
received from stbee21.us.oracle.com - 152.68.64.2155. Check the peer to
determine why it rejected the certificate chain (trusted CA configuration,
hostname verification). SSL debug tracing may be required to determine the
exact reason the certificate was rejected.> javax.net.ssl.SSLKeyException:
[Security:090482]BAD_CERTIFICATE alert was received from stbee21.us.oracle.com -
152.68.64.215. Check the peer to determine why it rejected the certificate chain
(trusted CA configuration, hostname verification). SSL debug tracing may be
required to determine the exact reason the certificate was rejected.
```

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Servers** and AdminServer (admin).
3. Select the **Configuration > SSL > Advanced Link**.

4. Select **None** for **Hostname Verification**.
5. Click **Save**.
6. Click **Activate Changes** in the Change Center of the Administration Console.
7. Restart all servers.

If the Managed Server is started in Admin mode, perform these steps:

1. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
2. In the left pane of the Console, expand **Servers** and the name of the server that is running in ADMIN mode.
3. Select the **Control > Start/Stop** tab.
4. Select the name of the server. Click **Resume**.
5. Select **Yes** to resume servers.

#### 8.6.6.2 If WebLogic Node Manager Fails to Start

If WebLogic Node Manager fails to start, ensure that you copied the following domains file from IDMHOST1 to IDMHOST2:

```
WL_HOME/common/nodemanager/nodemanager.domains
```

#### 8.6.6.3 Oracle Directory Services Manager Failover Using Oracle HTTP Server is Not Transparent

When you perform an Oracle Directory Services Manager failover using Oracle HTTP Server, the failover is not transparent. You see this behavior when you perform the following steps:

1. Oracle Directory Services Manager is deployed in a high availability active-active configuration using Oracle HTTP Server.
2. Open an Oracle Directory Services Manager page using the Oracle HTTP Server name and port number.
3. Make a connection to an LDAP server, for example, an Oracle Internet Directory server or an Oracle Virtual Directory server.
4. Work with the LDAP server using the current Oracle Directory Services Manager Oracle HTTP Server host and port.
5. Shut down one Managed Server at a time using the Oracle WebLogic Server Administration Console.
6. Go back to the Oracle Directory Services Manager page and port, and the connection which was established earlier with Oracle Internet Directory or Oracle Virtual Directory. When you do, a message is displayed advising you to reestablish a new connection to the Oracle Directory Services Manager page.

In this situation, you must do the following:

1. In your web browser, exit the current Oracle Directory Services Manager page.
2. Launch a new web browser page and specify the same Oracle Directory Services Manager Oracle HTTP Server name and port.
3. Reestablish a new connection to the LDAP server you were working with earlier (Oracle Internet Directory or Oracle Virtual Directory).

#### **8.6.6.4 Oracle Directory Services Manager Displays "LDAP Server is down" Message During Oracle Internet Directory Failover**

In a high availability configuration where Oracle Directory Services Manager is connected to Oracle Internet Directory through a load balancer, Oracle Directory Services Manager opens the `LDAP Server is down` message during failover from one instance of Oracle Internet Directory to another.

The connection reestablishes less than a minute after the Oracle Internet Directory failover is complete and you can continue without logging in again.

#### **8.6.6.5 Oracle Directory Services Manager Temporarily Loses Its Connection During Oracle RAC Failover**

Oracle Directory Services Manager temporarily loses its connection to an Oracle Internet Directory instance that is using an Oracle RAC database during an Oracle RAC failover. Oracle Directory Services Manager may display the message `Failure accessing Oracle database (oracle errcode=errcode)`, where `errcode` is one of the following values: 3113, 3114, 1092, 28, 1041, or 1012.

The connection will be reestablished in less than a minute after the Oracle RAC failover is complete, and you will be able to continue without logging in again.

### **8.6.7 Additional Considerations for Oracle Directory Services Manager High Availability**

When using Oracle Directory Services Manager to manage an Oracle Internet Directory cluster, use the load balancer virtual address as the connect string. However, when using Oracle Directory Services Manager to manage an Oracle Virtual Directory cluster, you must specify the host name and port for a specific Oracle Virtual Directory instance.

## **8.7 Collocated Architecture High Availability**

This section describes how to design and deploy Oracle Internet Directory, Oracle Directory Integration Platform, Oracle Virtual Directory, and Oracle Directory Services Manager in collocated high availability environments. The introduction to these components is provided in previous sections of this manual.

This section describes the collocated architecture for Oracle Internet Directory, Oracle Directory Integration Platform, Oracle Virtual Directory and Oracle Directory Services Manager when deployed in a high availability configuration.

This section includes the following topics:

- [Section 8.7.1, "Collocated Architecture Overview"](#)
- [Section 8.7.2, "Collocated Architecture High Availability Deployment"](#)
- [Section 8.7.3, "Validating the Collocated Components High Availability"](#)
- [Section 8.7.4, "Troubleshooting Collocated Components Manager High Availability"](#)
- [Section 8.7.5, "Additional Considerations for Collocated Components High Availability"](#)


### **8.7.1 Collocated Architecture Overview**

See the sections below for an architecture overview of each component in the collocated architectures described in this section:

- Oracle Internet Directory: [Section 8.3.1, "Oracle Internet Directory Component Architecture"](#)
- Oracle Virtual Directory: [Section 8.4.1, "Oracle Virtual Directory Component Architecture"](#)
- Oracle Directory Integration Platform: [Section 8.5.1, "Oracle Directory Integration Platform Component Architecture"](#)
- Oracle Directory Services Manager: [Section 8.6.1, "Oracle Directory Services Manager Component Architecture"](#)

Figure 8–10 shows Oracle Internet Directory, Oracle Directory Integration Platform, Oracle Virtual Directory, and Oracle Directory Services Manager collocated on a single host and deployed in a non-high availability architecture.


**Figure 8–10 Collocated Components Architecture**



All the components in Figure 8–10 are deployed on the same host, but have separate Oracle homes and Oracle instances. Oracle Internet Directory uses a standalone Oracle database as the security metadata repository.

### 8.7.2 Collocated Architecture High Availability Deployment

Figure 8–11 shows Oracle Internet Directory, Oracle Virtual Directory, Oracle Directory Integration Platform, and Oracle Directory Services Manager collocated on IDMHOST1 and IDMHOST2 and deployed in a high availability architecture.

**Figure 8–11 Collocated Components in a High Availability Architecture**

### 8.7.2.1 Collocated Architecture Prerequisites

See the sections below for the prerequisites of each component in the collocated architectures described in this section:

- Oracle Internet Directory: [Section 8.3.2.3, "Oracle Internet Directory Prerequisites"](#)
- Oracle Virtual Directory: [Section 8.4.2.2, "Oracle Virtual Directory Prerequisites"](#)
- Oracle Directory Integration Platform: [Section 8.5.2.3, "Oracle Directory Integration Platform Prerequisites"](#)
- Oracle Directory Services Manager: [Section 8.6.2.3, "Oracle Directory Services Manager Prerequisites"](#)
- Oracle Identity Federation: [Section 8.13.2.3, "Oracle Identity Federation Prerequisites"](#)

### 8.7.2.2 Configuring Collocated Components for High Availability

This section provides the steps to install and configure Oracle Internet Directory, Oracle Directory Integration Platform, Oracle Virtual Directory and Oracle Directory Services Manager on IDMHOST1 and IDMHOST2 in a high availability architecture:



---

**Note:** In a collocated environment, the Oracle Identity Management components should be installed in separate Oracle Homes. They can share the same MW\_HOME.

For each component, ensure that the Oracle Home Location directory path for the first instance is the same as the Oracle Home Location directory path for the second instance.

For example, if the Oracle Home Location directory path for the first Oracle Internet Directory instance on OIDHOST1 is  
 /u01/app/oracle/product/fmw/idm, then the Oracle Home Location directory path for the second Oracle Internet Directory instance on OIDHOST2 must also be  
 /u01/app/oracle/product/fmw/idm.

---

1. Install the database. For more information, see [Section 8.2.3, "Installing and Configuring the Database Repository."](#)
2. Install RCU. For more information, see [Section 8.2.4, "Obtaining the Repository Creation Utility Software."](#)
3. Configure the database. For more information, see [Section 8.2.5, "Configuring the Database for Oracle Fusion Middleware 11g Metadata."](#)
4. Run RCU to install the required schemas for Oracle Internet Directory and Oracle Identity Federation. For more information, see [Section 8.3.2.3.2, "Using RCU to Create Oracle Internet Directory Schemas in the Repository"](#) and [Section 8.13.2.3.1, "Using RCU to Create Oracle Identity Federation Schemas in the Repository."](#)
5. Install and configure Oracle Internet Directory on the first host. For more information, see [Section 8.3.3.2.1, "Configuring Oracle Internet Directory on OIDHOST1"](#) or [Section 8.3.3.3.1, "Configuring Oracle Internet Directory on OIDHOST1."](#)
6. Install and configure Oracle Internet Directory on the second host. For more information, see [Section 8.3.3.2.3, "Configuring Oracle Internet Directory on OIDHOST2"](#) or [Section 8.3.3.3.3, "Configuring Oracle Internet Directory on OIDHOST2."](#)
7. Install and configure Oracle Virtual Directory on the first host. For more information, see [Section 8.4.3.1.1, "Configuring Oracle Virtual Directory on OVDHOST1"](#) or [Section 8.4.3.2.1, "Configuring Oracle Virtual Directory on OVDHOST1."](#)
8. Install and configure Oracle Virtual Directory on the second host. For more information, see [Section 8.4.3.1.2, "Configuring Oracle Virtual Directory on OVDHOST2"](#) or [Section 8.4.3.2.3, "Configuring Oracle Virtual Directory on OVDHOST2."](#)
9. Install and configure Oracle Directory Integration Platform and Oracle Directory Services Manager on the first host. For more information, see [Section 8.5.3.1, "Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST1."](#)
10. Install and configure Oracle Directory Integration Platform and Oracle Directory Services Manager on the second host. For more information, see [Section 8.5.3.3, "Configuring Oracle Directory Integration Platform and Oracle Directory Services Manager on IDMHOST2."](#)

## 8.7.3 Validating the Collocated Components High Availability

See the following sections for information about validating components in the collocated high availability architectures and for information about how to failover the components and Oracle RAC.

### 8.7.3.1 Validation Tests

See the sections below for information on validating the following components in the collocated high availability architectures:

- Oracle Internet Directory: [Section 8.3.4, "Validating Oracle Internet Directory High Availability"](#)
- Oracle Virtual Directory: [Section 8.4.4, "Validating Oracle Virtual Directory High Availability"](#)
- Oracle Directory Integration Platform: [Section 8.6.4, "Validating Oracle Directory Services Manager High Availability"](#)
- Oracle Directory Services Manager: [Section 8.6.4, "Validating Oracle Directory Services Manager High Availability"](#)
- Oracle Identity Federation: [Section 8.13.3.6, "Validating Oracle Identity Federation High Availability"](#)

### 8.7.3.2 Failures and Expected Behaviors

See the sections below for information on failures and expected behaviors for the following components in the collocated high availability architectures:

- Oracle Internet Directory: [Section 8.3.5, "Oracle Internet Directory Failover and Expected Behavior"](#)
- Oracle Virtual Directory: [Section 8.4.5, "Oracle Virtual Directory Failover and Expected Behavior"](#)
- Oracle Directory Integration Platform: [Section 8.5.4, "Oracle Directory Integration Platform Failover and Expected Behavior"](#)
- Oracle Directory Services Manager: [Section 8.6.5, "Oracle Directory Services Manager Failover and Expected Behavior"](#)
- Oracle Identity Federation: [Section 8.13.4, "Oracle Identity Federation Failover and Expected Behavior"](#)

## 8.7.4 Troubleshooting Collocated Components Manager High Availability

See the sections below for information on troubleshooting the following components in the collocated high availability architectures:

- Oracle Internet Directory: [Section 8.3.6, "Troubleshooting Oracle Internet Directory High Availability"](#)
- Oracle Virtual Directory: [Section 8.4.6, "Troubleshooting Oracle Virtual Directory High Availability"](#)
- Oracle Directory Integration Platform: [Section 8.5.5, "Troubleshooting Oracle Directory Integration Platform High Availability"](#)
- Oracle Directory Services Manager: [Section 8.6.6, "Troubleshooting Oracle Directory Services Manager"](#)

- Oracle Identity Federation: [Section 8.13.5, "Troubleshooting Oracle Identity Federation High Availability"](#)

### 8.7.5 Additional Considerations for Collocated Components High Availability

See the sections below for information on additional considerations for the following components in the collocated high availability architectures:

- Oracle Internet Directory: [Section 8.3.7, "Additional Oracle Internet Directory High Availability Issues"](#)
- Oracle Directory Services Manager: [Section 8.6.7, "Additional Considerations for Oracle Directory Services Manager High Availability"](#)

## 8.8 Oracle Access Manager High Availability

This section provides an introduction to Oracle Access Manager 11gR1 and describes how to design and deploy a high availability environment for Oracle Access Manager 11gR1.

Oracle Access Manager 11gR1 is the successor product to both Oracle Access Manager 10g (access only) and Oracle Single Sign-On 10g. Oracle Access Manager 11gR1 provides a single authoritative source for all authentication and authorization services. The core service provided is the checking of valid session tokens, the requesting of credentials if the session token is invalid or missing, and the issuing of session tokens, intercepting resource requests and evaluating access control policies to control access to resources. Oracle Access Manager 11gR1 features a pure Java server while continuing to use Oracle Single Sign-On 10g and Oracle Access Manager 10g agent components. The main new feature for 11gR1 for the agent is the Shared Secret Key Per WebGate (SSKPWG) feature.

Oracle Access Manager provides Single Sign-On features, thus preventing the user from re-logging in every time after authenticating once. It accomplishes this by managing the user session life cycle, which also involves facilitating global logout by orchestrating logout across all relying parties in the valid user session. Oracle Access Manager also ensures that resource access by users is authorized subject to the specified authorization policy.

Unlike Oracle Access Manager 10g, Oracle Access Manager 11gR1 no longer has any identity service and is a first class consumer of identity information from other identity management services such as native LDAP and Oracle Identity Manager.

Oracle Access Manager 11gR1 is an architecture evolution in the area of Access Management that delivers unified product architecture for enterprise and web Single Sign-On. Components of Oracle Access Manager leverage the Java EE middleware platform with a common underlying data model, shared underlying functionality, consistent semantics with a focus of interoperability and seamless integration. Oracle Access Manager 11gR1 offers co-existence and incremental migration paths for existing Single Sign-On deployments.

This section includes the following topics:

- [Section 8.8.1, "Oracle Access Manager Component Architecture"](#)
- [Section 8.8.2, "Oracle Access Manager High Availability Concepts"](#)
- [Section 8.8.3, "Oracle Security Token Service High Availability"](#)
- [Section 8.8.4, "Oracle Access Manager High Availability Configuration Steps"](#)

## 8.8.1 Oracle Access Manager Component Architecture

Figure 8–12 shows the Oracle Access Manager 11gR1 component architecture.

**Figure 8–12 Oracle Access Manager Single Instance Architecture**




Figure 8–12 shows the following components:

- **User agents:** These include web browsers, Java applications, and Web services applications. The user agents access the Access Server and the administration and configuration tools using HTTP.
- **Protected resources:** A protected resource is an application or web page to which access is restricted. Access to protected resources is controlled by WebGates or Custom Agents.
- **Administration and configuration tools:** Oracle Access Manager can be administered and configured by the Oracle Access Manager console, the Oracle Enterprise Manager Fusion Middleware Control and the Oracle Enterprise Manager Grid Control, and the WebLogic Scripting Tool (WLST).
- **Access Server:** The Access Server includes the Credential Collector, OSSO Proxy, and OAM Proxy components. The Coherence Distributed Object Cache is used to propagate configuration file changes between Access Servers

### 8.8.1.1 Oracle Access Manager Component Characteristics

An Oracle Access Manager 11gR1 deployment is composed of the following system entities:

- **Oracle Access Manager Agents -** Oracle Access Manager agents are extensions of the Access Server that are responsible for ensuring that access is controlled as per the policies managed in the Access Server. Agents are the clients/programs (such as WebGate, Java Agents, custom Agents, and Oracle Single Sign-On Apache Modules) used by end-users to access various Web resources that are protected by Oracle Access Manager 11gR1. The most commonly used user agents are web browsers. Oracle Access Manager 11gR1 is a Web Single Sign-On solution that is primarily focused on controlling resources accessible via HTTP (resources identified by a URL) and custom resource types.

Agents require the Access Server component to perform their functions. If the Access Server is unavailable, no access to protected servers will be permitted (users will be locked out of the system).

Oracle Access Manager 11gR1 agents connect to the Access Server over the front channel and back channel. The front channel communication takes place over HTTP(S) when the user needs to be authenticated. Front channel communications tend to be short-lived.

When an Oracle Access Manager agent communicates with the Access Server using the front channel HTTP binding, it needs to communicate through a load balancing router. This information is passed to the agent and configured as the Challenge Redirect URL in the authentication scheme.

Back channel communication takes place using a proprietary protocol called Oracle Access Protocol (OAP), which takes place over a TCP connection. Agents use back channel communications with the Access Server for every resource access request in order to make an authorization decision, thus these back-channel connections are persistent and long-lived.

When an Oracle Access Manager agent communicates with the Access Server using back channel OAP binding, it is configured to use a primary / secondary model.

Oracle Access Manager agents are externally staged (deployed in the web tier) because this provides the best scalability.

WebGate caches information about resource requests and authentication schemes. The WebGate cache is flushed based on a configured timeout or a server-initiated cache purge.

WebGates refresh their configuration by polling the server every 60 seconds. When configuration changes are detected, they persist immediately. Existing connections terminate and new connections are created if a connection information change occurs.

WebGate configuration includes information about the agent identity, agent credentials, agent-server security context, and connection parameters.

- Protected Resources - These are the applications that are protected by Oracle Access Manager 11gR1 (also referred to as partner applications). Access to these resources is subject to the access control policies in Oracle Access Manager 11gR1 and is enforced by Oracle Access Manager agents that are deployed in the access path of the protected resource (for example, Oracle Access Manager agents deployed in the Web Server, J2EE agents deployed in the Application Server). Agents are entities that control access to protected applications based on security policies. Agents present in the resource access path intercept every resource access request to enforce the security policy that protects the resources.
- Access Server - This is the server side component that provides the core runtime access management services. It has an event-based design pattern that implements the core Oracle Access Manager services. An Access Server is a J2EE application that is packaged as an EAR file and is composed of Servlets and JSPs in addition to Java classes. It provides various Identity Provider (IDP) services. The Access Server in Oracle Access Manager 11gR1 provides Single Sign-On, Authentication, and Authorization services.
- JMX Mbeans - Runtime Mbeans are packaged as part of the Access Server package. Config Mbeans are packaged as standalone WAR files.

- WebLogic 11g SSPI providers are composed of Java classes that implement the SSPI interface along with the Access Java Access JDK. AccessGates are built using the pure Java Access JDK.
- Administration Console - The Oracle Access Manager Administration Console is a J2EE application that hosts the Administration Console and provides services like Administration/Configuration to manage the Oracle Access Manager 11gR1 deployment. In Oracle Access Manager 11gR1, this component must be deployed to the WebLogic Administration Server.
- WebLogic Scripting Tool (WLST) is composed of Java classes, which are included in the Access Server package. Limited administration of the Oracle Access Manager 11gR1 deployment is supported via the command line.
- Fusion Middleware Control and Enterprise Manager Grid Control - Oracle Access Manager 11gR1 integrates with the Enterprise Manager Grid Control to display performance metrics and deployment topology.
- Coherence Distributed Object Cache - Oracle Access Manager 11gR1 components rely on this infrastructure for real time change propagation.
- External credential collectors are a set of JSPs.
- The Oracle Access Manager Proxy is a customized version of the Apache MINA server (based on the JCA architecture), which includes MessageDrivenBeans and ResourceAdapters in addition to Java Server classes. It is included in the Access Server package.
- The Oracle Single Sign-On (OSSO) Proxy comprises Java classes, which are included in the Access Server package.
- Data Repositories - Oracle Access Manager 11gR1 handles different types of information including Identity, Policy, Partner, Session and Transient data:
  - LDAP for Identity data
  - Files for Configuration and Partner data
  - Coherence in-memory for Session and Transient Data
  - Policy data will be stored in files or in an RDBMS
- Oracle Access Manager 10g WebGates are C-based agents that are intended to be deployed in web servers.
- Oracle Single Sign-On Apache modules are C-based agents that are intended to be deployed in Oracle HTTP Server web servers.
- Oracle Access Manager 11g WebGates are C-based agents that are intended to be deployed in web servers.

**8.8.1.1.1 Oracle Access Manager State Information** Authenticated user session information is persisted via the Coherence Distributed Object Cache. Use the Coherence Distributed Object Cache in-memory mode for Oracle Access Manager 11gR1.

Oracle Access Manager may create a transient state for unauthenticated users during the login processing. This state is generally not replicated among Oracle Access Manager nodes. To protect against effects of node failures during the login processing, the state may be optionally stored in an encrypted client cookie.

To store the transient state for unauthenticated users during login processing, change the Oracle Access Manager server parameter RequestCacheType from BASIC to COOKIE by following these steps:

1. Set up the environment for WLST by running this command:

```
DOMAIN_HOME/bin/setDomainEnv.sh
```

2. Start WLST by issuing this command:

```
Start WLST by issuing this command:
ORACLE_HOME/common/bin/wlst.sh
```

3. Connect to your domain:

```
wls:/IDM_Domain/serverConfig> connect()
```

4. Enter the WebLogic Administration username and password, and enter the URL for the Administration Server in the format:

```
t3://OAMHOST1.mycompany.com:7001
```

5. Issue this command:

```
wls:/IDM_Domain/serverConfig> configRequestCacheType(type="COOKIE")
```

6. Check that the command worked by issuing this command:

```
wls:/IDM_Domain/serverConfig> displayRequestCacheType()
```

7. Restart the Oracle Access Manager managed servers.

**8.8.1.1.2 Oracle Access Manager Request Flow** The following list shows the steps in an Oracle Access Manager request flow:

1. The user tries to access a Oracle Access Manager 11gR1 protected Web Resource using his web browser.
2. The Oracle Access Manager agent<sup>1</sup> intercepts the request and tries to ascertain if the user has an authenticated session.
3. Since this is the user's first access, the user is redirected to the Oracle Access Manager 11gR1 Access Server for authentication.
4. Access Server's credential collector<sup>2</sup> component displays a Login Form.<sup>3</sup> The user submits his credentials to the Access Server.
5. Access Server validates the user's credentials and generates a security token. The user is redirected to the resource he tried to access in Step 1.
6. The Oracle Access Manager agent intercepts the request and extracts the security token (cookie).
7. The Oracle Access Manager agent then makes a back channel call<sup>4</sup> to the Access Server (OAP over TCP) to validate the session and authorize the request.
8. Oracle Access Manager authenticates the user from the LDAP repository.
9. Access server verifies the user's permissions against the configured policy for the web resource.

<sup>1</sup> The agent in use is specific to a deployment and different types of agents (with different features) can be used in a deployment.

<sup>2</sup> In addition to the built-in Credential Collector, Oracle Access Manager is capable of supporting external credential collectors.

<sup>3</sup> The credential collection will be different for non-username and password authentication schemes.

<sup>4</sup> Only WebGates support back channel communication.

10. Access server responds to the WebGate request indicating that access is allowed.
11. The Oracle Access Manager agent allows the request to go through.<sup>5</sup>
12. The user is now able to access the web resource he tried to access in Step 1.

**8.8.1.1.3 Oracle Access Manager Process Lifecycle** You can start Access Server and Administration Console from the user interface and command line tools that WebLogic Server provides.

The Access Server supports a health check request (a ping request over HTTP) that the load balancer can use for death detection.

Oracle Access Manager agents are native applications that reside in the protected application environment. No tools are provided as part of OAM 11gR1 but it is expected that environment specific tooling, where available, will be leveraged for the above purpose.

Oracle Access Manager 11gR1 is instrumented for server side metrics (using DMS) and this information is published to the Administration Console. Using DMS metrics collection, you can monitor the agent and server component metrics as a proxy for component monitoring. In addition, Oracle Access Manager 11gR1 supports fine-grained real time activity tracing, which can also serve as a proxy for component monitoring.

On startup, Access Server initializes connections to the backend repositories. If the repository is not reachable, the Access Server retries the connections to the repositories, using a timeout that grows exponentially with a configurable upper bound.

Access Server will provide continuity of service based on locally cached data if the backend connections go down. This will continue until the caches grow stale or the backend connections become alive again.

**8.8.1.1.4 Oracle Access Manager Configuration Artifacts** The Oracle Access Manager configuration artifacts include these files:

- `DOMAIN_HOME/config/fmwconfig/oam-configuration.xml`  
The configuration file, which contains instance specific information.
- `DOMAIN_HOME/config/fmwconfig/oam-policy.xml`
- `DOMAIN_HOME/config/fmwconfig/.oamkeystore`  
This is used for storing symmetric and asymmetric keys.
- `DOMAIN_HOME/config/fmwconfig/component_events.xml`  
Used for audit definition.
- `DOMAIN_HOME/config/fmwconfig/jazn-data.xml`  
Used for Administration Console permissions
- `DOMAIN_HOME/config/fmwconfig/servers/instanceName/logging.xml`  
Used for logging configuration.
- `DOMAIN_HOME/config/fmwconfig/servers/instanceName/dms_config.xml`  
Used for tracing logging.

---

<sup>5</sup> The agent may perform some housekeeping tasks, such as session refresh, before allowing the request to go through to the web resource.



- `DOMAIN_HOME/config/fmwconfig/cwallet.sso`  
Used for passwords

**8.8.1.1.5 Oracle Access Manager External Dependencies** Oracle Access Manager has external runtime dependencies on the:

- LDAP based Identity Store
  - User Identity Repository
  - LDAP access abstracted by User/Role API.

---

**Note:** Oracle Access Manager always connects to one Identity store, which can be a physical server or a load balancer IP. If the primary is down, Oracle Access Manager reconnects and expects the load balancer to connect it to the secondary.

---

- OCSP Responder Service
  - Real-time X.509 Certification Validation
- RDBMS Policy Store
  - Policy (Authentication and Authorization) Repository
  - RDBMS access abstracted by the OAM policy engine
- Oracle Identity Manager (when OIM based password management is enabled)
  - Oracle Identity Manager is used to provide Password Management Services and replaces the Oracle Access Manager 10g Identity Server
- Oracle Identity Manager Policy Store (when Oracle Identity Manager-based password management is enabled)
  - LDAP Repository containing Oblix Schema elements that are used to store Configuration, Metadata, and so on.
- Oracle Adaptive Access Manager (when Oracle Adaptive Access Manager Advanced Authentication Scheme is selected)
- Oracle Identity Federation (when Oracle Identity Federation Authentication Scheme is selected)

**8.8.1.1.6 Oracle Access Manager Log File Location** Oracle Access Manager is a J2EE application deployed on WebLogic Server. All log messages are logged in the server log file of the WebLogic Server that the application is deployed on. The default location of the server log is:


```
WL_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

## 8.8.2 Oracle Access Manager High Availability Concepts

This section provides conceptual information about using Oracle Access Manager in a high availability two-node cluster.

### 8.8.2.1 Oracle Access Manager High Availability Architecture

Figure 8–13 shows an Oracle Access Manager high availability architecture:

**Figure 8–13 Oracle Access Manager High Availability Architecture**

In Figure 8–13, incoming authentication requests are received by the hardware load balancer, which routes them to WEBHOST1 or WEBHOST2 in the web tier. These hosts

have Oracle HTTP Server installed. The Oracle HTTP Server then forwards requests on to the WebLogic managed servers using the WebLogic plugin `mod_wl_ohs`.

The load balancing router should use session stickiness for HTTP traffic only. OAP traffic does not use a load balancing router, so session stickiness is not required for OAP traffic.

Applications which are accessed by other Oracle HTTP Servers whose resources have restricted access must also have a WebGate, Oracle Single Sign-On Server agent (`mod_osso` agent), or custom agent configured. The WebGate on WEBHOST3 communicates with the Access Servers on OAMHOST1 and OAMHOST2 in the application tier using OAP. WEBHOST3 is an application web server, and for authentication, HTTP redirect is used to route requests to the load balancer and to WEBHOST1 and WEBHOST2. For a high availability deployment, you can optionally configure another host (for example, WEBHOST4) with the same components as WEBHOST3.

OAMHOST1 and OAMHOST2 deploy managed servers which host the Oracle Access Server application. These managed servers are configured in a cluster which allows the Access Servers to work in an active-active manner.

The WebLogic Administration Server runs on OAMHOST1 and deploys the WebLogic Administration Console, Oracle Enterprise Manager Fusion Middleware Control, and the Oracle Access Manager Console. The Administration Server can be configured to run in active-passive mode on OAMHOST2, which means that if OAMHOST1 becomes unavailable, then Administration Server can be manually started on OAMHOST2.

In the directory tier, the virtual IP `ovd.mycompany.com` is set up to route Oracle Virtual Directory requests to OVDHOST1 and OVDHOST2, which comprise an active-active Oracle Virtual Directory cluster. The virtual IP `oid.mycompany.com` is set up to route Oracle Internet Directory requests to OIDHOST1 and OIDHOST2, which comprise an active-active Oracle Internet Directory cluster.

An Oracle RAC database provides high availability in the data tier.

In Oracle Access Manager 11gR1, only one Oracle Access Manager cluster is supported per WebLogic Server domain. In addition, Oracle Access Manager clusters cannot span WebLogic Server domains.

A single instance Oracle Access Manager 11gR1 deployment satisfies the following high availability requirements:

- Load handling
- External connection management and monitoring
- Recovery
- Fault containment
- Fault diagnostics
- Administration Server offline

A multiple instance Oracle Access Manager 11gR1 deployment satisfies the following additional high availability requirements:

- Redundancy
- Client connection failover/continuity
- Client load balancing
- State management

Use of an external load balancing router is recommended for inbound HTTP connections. Outbound external connections to LDAP Servers (or OAM policy engine PDP/PIP) are load balanced with support for connection failover. Oracle Access Manager agents can load balance connections across multiple Access Servers.

Oracle Access Manager agents open persistent TCP connections to the Access Servers. This requires firewall connection timeouts to be sufficiently large to avoid premature termination of TCP connections.

The Access Server and Oracle Access Manager Administration Console interface with the OAM policy engine for policy evaluation and management. The OAM policy engine internally depends on a database as the policy repository. The database interactions are encapsulated within the OAM policy engine, with only the connectivity configuration information managed by Oracle Access Manager. The high availability characteristics of the interaction between Oracle Access Manager and the OAM policy engine are:

- The database connection information is configured in the Oracle Access Manager configuration file synchronized among the Oracle Access Manager instances. Should the database connection information change at runtime, Access Server instances will re-initialize OES to complete the change activation.
- Database communication is managed within the OAM policy engine, and generally decoupled from Oracle Access Manager and OAM policy engine interactions. The very first startup of an Oracle Access Manager server instance will fail, however, if the database is unreachable. An OAM policy engine bootstrap failure is treated as fatal by Oracle Access Manager, and the startup operation is aborted.
- Transient database unavailability is transparently tolerated by OAM policy engine policy evaluation services, allowing Oracle Access Manager server runtimes to continue functioning uninterrupted. After the initial OAM policy engine bootstrap, the Oracle Access Manager instances may even be restarted while the database is unreachable -- the OAM policy engine will continue operating against its locally cached policies.
- Oracle Access Manager policy management interfaces (in the Oracle Access Manager Administration Console and the CLI tool) will fail if the database is unreachable, as seen by the OAM policy engine management service interfaces. The operation may be retried at a later point in time, but no automated retry is provided for management operations.
- Following a successful policy modification in the database repository, the OAM policy engine layer in the Oracle Access Manager server runtimes retrieves and activates the changes within a configurable OAM policy engine database poll interval (configured through Oracle Access Manager configuration). A positive acknowledgement of a policy change must be received from each Oracle Access Manager server runtime, otherwise the policy change cannot be considered successfully activated. The administrator can use the Oracle Access Manager Administration Console to remove any Oracle Access Manager instance with a policy activation failure from service.

**8.8.2.1.1 Starting and Stopping the Cluster** In a high availability architecture, Oracle Access Manager server is deployed on an Oracle WebLogic Cluster, which has at least two servers as a part of the cluster.

By default, Oracle WebLogic Server starts, stops, monitors and manages the various lifecycle events for the application. The Oracle Access Manager application leverages the high availability features of the underlying Oracle WebLogic Clusters. In case of

hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

In a high availability environment, WebLogic Node Manager is configured to monitor the Oracle WebLogic Servers. In case of failure, Node Manager restarts the WebLogic Server.

In a high availability environment, a hardware load balancer is used to load balance requests between the multiple Oracle Access Manager instances. If one of the Oracle Access Manager instances fails, the load balancer detects the failure and routes requests to the surviving instances.

**8.8.2.1.2 Cluster-Wide Configuration Changes** The standard Java EE artifacts that Oracle Access Manager uses are configured as part of the Oracle WebLogic domain in which Oracle Access Manager is installed. Oracle WebLogic Clusters provide automatic configuration synchronization for artifacts, such as data sources, across the WebLogic Server domain. At the same time, the WebLogic Server cluster synchronizes the deployments and libraries used by the Oracle Access Manager components.

Additionally, Oracle Access Manager application level configuration is stored in the Oracle Access Manager repository. Propagation of Oracle Access Manager configuration changes to all the cluster members is based on a distribution mechanism that leverages the Coherence Distributed Object Cache. All Oracle Access Manager components are notified of change events from the coherence layer, which are then taken up by the components. To ensure atomicity of the change, Oracle Access Manager components reload their entire configuration every time a change happens.

Oracle Access Manager configuration applies to all instances in a cluster. The only exceptions to the above (instance-specific configuration) supported in Oracle Access Manager 11gR1 are the Oracle Access Manager proxy host, Oracle Access Manager proxy port, and the instance-specific Coherence configuration (when Well Known Addresses (WKA) is used). The IP address of the proxy host and proxy port are stored in a configuration file. The Oracle Access Manager proxy port is the endpoint for OAP requests from agents. The IP address of the Coherence WKA is also stored in a configuration file. The Coherence WKA is used to determine the Coherence nodes that are authorized to receive Oracle Access Manager-specific traffic. The `oam-configuration.xml` file is the configuration file that stores this configuration information.

It is possible to configure clients of the Oracle Access Manager proxy to access the service using this virtual/logical IP.

The Oracle Access Manager proxy can be deployed (and its clients still able to access the service) when the logical IP and the component instance is migrated to any other physically different machine configured similarly.

Adding and removing Access Server instances is transparent to other Oracle Access Manager Access Server instances in the cluster. However, take care to ensure that the removal of a specific Access Server does not affect the load balancing and failover characteristics of the agents.

Restarting an Oracle Access Manager Access Server has no impact on any other running components or members of the cluster.

Online application redeployment does not cause any problems.

### 8.8.2.2 Protection from Failures and Expected Behaviors

This section describes how an Oracle Identity Management high availability cluster deployment protects components from failure. This section also describes expected behavior if component failure occurs.

The WebLogic Server infrastructure protects the Identity Management Service Infrastructure system from all process failures.

The following features protect the Oracle Access Manager high availability configuration from failure:

- Session state is maintained in the Coherence Distributed Object Cache, which provides replication and failover for all session state information. Since Coherence is used in the in-memory mode, this information will not survive a simultaneous restart of all Oracle Access Manager components.
- If an Access Server crashes, a WebGate with a persistent connection to that server waits for the connection to timeout, then it switches over to the secondary (backup) Access Server. Any outstanding requests are failed over to the secondary server.
- Oracle Access Manager Access Servers support a heart beat check - a ping request over HTTP. In addition, the WebLogic Node Manager on the Managed Server can monitor the application and restart it if it is not running.
- When a WebLogic Server node fails, external connection failover is based on the configuration and is based on the retry timeout as well as the number of retries. Oracle Access Manager Agent-Access Server failover is based on a timeout.
- When the load balancing router or proxy server detects a WebLogic Server node failure, subsequent client connections are routed to the active instance, which picks up the session state from the Coherence Distributed Object Cache and carries on with the processing.
- Back channel OAP bindings use a primary/secondary model for failover. Front Channel HTTP bindings use a load balancing router for failover.
- When the lifetime of a connection expires, pending requests are allowed to complete before the connection is terminated. The connection object is returned to the pool.
- When it receives an exception from another service, Oracle Access Manager retries external connection requests. You can configure the number of retries or set to "no retries."

**8.8.2.2.1 WebLogic Server Crash** When a user first requests access to a protected resource, the Oracle HTTP server establishes a connection with an Oracle Access Server and, following initial user authentication, the Oracle Access Manager agent (for example, WebGate) places entries in a browser cookie for two access servers. From this moment, the client application communicates directly with one of the access servers identified in the cookie. If that access server becomes unavailable, the application communicates with the alternate access server identified in the cookie. If both servers become unavailable, the user must reauthenticate.

If a WLS\_OAMx server crashes, Node Manager attempts to restart it locally.

Ongoing requests from the HTTP Server timeout and new requests are directed to the other WLS\_OAMx server. Once the server's restart completes on the failed node, Oracle HTTP Server resumes routing any incoming requests to the server.

---

**Note:** Oracle Access Manager servers also support a heartbeat check. This heartbeat determines if the access server can service its requests. It checks things such as:

- Whether the LDAP store can be accessed
- Whether the policy store can be accessed

If the heartbeat succeeds, then the Access Server is capable of servicing requests, and requests will be sent to it. If the heartbeat fails, however, then requests will not be routed to the Access Server with the issue.

---

**8.8.2.2.2 Node Failure** Node failures are treated in the same way as WebLogic Server crashes.

**8.8.2.2.3 Database Failure** The Oracle Access Manager service Infrastructure is protected against failures in the database by using multi data sources. These multi-data sources are typically configured during the initial set up of the system (Oracle Fusion Middleware Configuration Wizard allows you to define these multi-pools directly at installation time). They guarantee that when an Oracle RAC database instance fails, the connections are reestablished with available database instances. The multi data source allows you to configure connections to multiple instances in an Oracle RAC database.

For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

### 8.8.3 Oracle Security Token Service High Availability

This section describes Oracle Security Token Service (OSTS) high availability.

Oracle Security Token Service is a shared Web Service (JAX-WS) that provides a standards-based consolidated mechanism of trust brokerage between different identity domains and infrastructure tiers. Oracle Security Token Service brokers trust between a Web Service Consumer (WSC) and a Web Service Provider (WSP) and provides security token lifecycle management services to providers and consumers. Oracle Security Token Service can help simplify the effort needed to bridge access to various systems using a standardized set of interfaces.

An Oracle Security Token Service high availability deployment depends on OAM; the Oracle Access Manager (OAM) application contains the Oracle Security Token Service server runtime. An Oracle Security Token Service high availability deployment cannot occur independent of OAM. Oracle Access Manager and Oracle Security Token Service are bundled together in the OAM J2EE Application EAR file, installed together, and deployed on the same managed server in a WebLogic domain. Oracle Security Token Service is also integrated with the Oracle Access Manager Console.

For more details on administering and configuring Oracle Security Token Service, see one of the following topics in *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

- Oracle Security Token Service Key Terms and Concepts
- About Oracle Security Token Service with Oracle Access Manager
- About Oracle Security Token Service Deployments
- About Oracle Security Token Service Administration

- Oracle Security Token Service Implementation Scenarios
- Managing Oracle Security Token Service Settings and Set Up
- Managing Oracle Security Token Service Certificates and Keys
- Managing Templates, Endpoints, and Policies
- Managing Token Service Partners and Partner Profiles

For information on patching, see Migrating Oracle Access Manager 11.1.1.3.0 to 11.1.1.6.0


This section includes the following topics:

- [Section 8.8.3.1, "Oracle Security Token Service High Availability Architecture"](#)
- [Section 8.8.3.2, "Oracle Security Token Service Component Characteristics"](#)
- [Section 8.8.3.3, "Oracle Security Token Service High Availability Configuration Steps"](#)
- [Section 8.8.3.4, "Validating Oracle Security Token Service High Availability"](#)
- [Section 8.8.3.5, "Oracle Security Token Service Failover and Expected Behavior"](#)
- [Section 8.8.3.6, "Disabling and Enabling Oracle Security Token Service"](#)
- [Section 8.8.3.7, "Troubleshooting Oracle Security Token Service"](#)
- [Section 8.8.3.8, "Log File Location"](#)

### 8.8.3.1 Oracle Security Token Service High Availability Architecture

The following figure shows Oracle Security Token Service in a high availability architecture:

**Figure 8–14 Oracle Security Token Service High Availability Architecture**





This figure shows a two-node deployment of Oracle Access Manager/Oracle Security Token Service components. This section provides details about an Oracle Security Token Service high availability deployment. For more details about the overall Oracle Access Manager high availability architecture, which is deployed as part of, see [Section 8.8.2.1, "Oracle Access Manager High Availability Architecture"](#).

Security Token Service is the server side component that implements the WS-Trust protocol support.

The load balancer receives token requests and routes them to the Security Token Service (STS).

The Oracle Access Manager Administration console is a J2EE application that provides services to manage the Oracle Security Token Service deployment. As part of the OAM deployment, Administration Console must deploy to the Weblogic AdminServer.

In Oracle Security Token Service, each WLS domain supports one Oracle Security Token Service cluster. OAM/Oracle Security Token Service clusters cannot span WebLogic Server domains.

Oracle Security Token Service is primarily based on the OASIS WS-Trust protocol. However, Oracle Security Token Service delegates the processing of other WS-\* protocols present in the SOAP message to the JAX-WS stack.

Oracle recommends using external load balancers for inbound HTTP/SOAP connections. Outbound external connections to LDAP servers are load balanced with support for connection failover.

**8.8.3.1.1 Clients and Client Connections** Web Service clients that implement the WS-Trust protocol interact with Oracle Security Token Service to issue or validate tokens. Clients designed to interact with an STS server, such as OWSM Client, as part of a Web Service call to a Relying Party can invoke Oracle Security Token Service.

The client connection process is as follows:

1. The Web Service client sends a SOAP message over http or https.

The WSS protocol protects the message. The payload contains a WS-Trust request (RST) indicating the operation to perform, which kind of token to issue or validate, and additional information about the token characteristics.

2. The server processes the request and sends a response over the same channel the server received it on.

The WSS protocol protects the message. The payload contains a WS-Trust response (RSTRC) if the processing was successful or a SOAP fault if an error occurs.

**8.8.3.1.2 Cluster Wide Configuration Changes** Each Oracle Security Token Service Access Server instance is a peer of other instances with no inter-instance communication. Because all initialization happens before the Server is ready to receive requests combined with built in throttling capabilities, the WebLogic Server handles surge conditions gracefully without any significant effect on Oracle Security Token Service Access Server performance characteristics.

When the cluster stops, the Oracle Security Token Service denies new requests and permits existing requests to complete before the Access Server application shuts down. If a forced shutdown occurs, Oracle Security Token Service can recover for any corrupted/invalid data that the shutdown causes.

Propagation of configuration changes to all the cluster members is based on a distribution mechanism that leverages the Coherence distributed object cache. The coherence layer notifies all Oracle Security Token Service components of change

events. The components then uptake these change events. OAM components reload their entire configuration every time a change happens.

---

**Note:** The addition/removal of Access Server instance(s) is transparent to other Oracle Security Token Service instances in the cluster. Verify that removing a specific Oracle Security Token Service server does not affect the load.

---

### 8.8.3.2 Oracle Security Token Service Component Characteristics

This section describes Oracle Security Token Service component characteristics and includes the following topics:

- [Section 8.8.3.2.1, "Oracle Security Token Service Component Lifecycle"](#)
- [Section 8.8.3.2.2, "Runtime Processes"](#)
- [Section 8.8.3.2.3, "Starting and Stopping Oracle Security Token Service"](#)
- [Section 8.8.3.2.4, "J2EE Components and Subcomponents"](#)
- [Section 8.8.3.2.5, "Session State Information"](#)
- [Section 8.8.3.2.6, "Configuration Artifacts"](#)
- [Section 8.8.3.2.7, "External Dependencies"](#)

**8.8.3.2.1 Oracle Security Token Service Component Lifecycle** On startup, the OAM/Oracle Security Token Service Server initializes connections to the backend repositories. If the repository is not reachable, the OAM/Oracle Security Token Service server retries the connections to the repositories using a timeout that grows exponentially with a configurable upper limit.

OAM/Oracle Security Token Service Server provide continuity of service based on locally cached data if the backend connections go down. Service continues until the caches grow stale or the backend connections come up again.

**8.8.3.2.2 Runtime Processes** The following graphic shows the Oracle Security Token Service runtime process.

**Figure 8–15 Oracle Security Token Service Runtime Process**



The Oracle Security Token Service runtime process works as described below:

1. A Web Service Consumer (WSC) sends a Web Services-Trust Request Security Token (RST) message for a security token type that the WSP requires. Authentication of the client occurs by using the transport layer authentication, or by binding the WSS Token to the RST message.
2. The Security Token Service (STS) validates the RST message, authenticates the request, then authorizes the requested operation.
3. The appropriate security token is generated in accordance with the metadata that the RST message specifies. For the policy driven exchange use-case, the STS looks

up the appropriate token generation policy to generate the appropriate security token.

4. STS generates an RST message that contains the generated security token; it sends the message to the WSC as a response.

---

**Note:** WSP validation of the security token depends on the token type. When STS acts as a trust intermediary only, validation is performed against the underlying security infrastructure, such as Kerberos.

---

**8.8.3.2.3 Starting and Stopping Oracle Security Token Service** Because they are J2EE applications, you can start the Access Server (where Oracle Security Token Service is deployed) and Admin Console from the user interface and Command Line tool that the Application Server provides.

**8.8.3.2.4 J2EE Components and Subcomponents** J2EE Components and sub-components include the following:

- STS - An event based design pattern that implements the core Oracle Security Token Service 11gR1-PS1. It is packaged as a WAR application in the OAM EAR file and comprises a WS Provider Servlet and Java classes. The STS Web Application is bound to the /sts root path
- Admin Console - A stand-alone console based on ADF/IDM Shell, and packaged as an EAR file.
- JMX Mbeans - Packaged with the Access Server package. Config Mbeans are packaged as standalone JAR files.
- WSLT Command - Consists of Java classes that are in the OAM/Oracle Security Token Service package.
- OWSM Agent - Web Service interceptor providing support for WSS protocol, part of JRF.
- ORAProvider - JRF Web Service Provider

**8.8.3.2.5 Session State Information** Oracle Security Token Service is a stateless J2EE application with the exception of the Nonce caching for Username Tokens, where OSTs will keep track of presented username tokens when the nonce is present, in order to prevent replay attacks.

**8.8.3.2.6 Configuration Artifacts** Oracle Access Manager and Oracle Security Token Service are built together and use the same modules for configuration, logging, and other processes. The Oracle Security Token Service configuration artifacts include the following files.

- DOMAIN\_HOME/config/fmwconfig/oam-config.xml — The configuration file, which contains instance-specific information.
- DOMAIN\_HOME/config/fmwconfig/oam-policy.xml — Present only when OES Micro SM is not being used.
- DOMAIN\_HOME/config/fmwconfig/servers/instanceName/logging.xml — Logging config
- DOMAIN\_HOME/config/fmwconfig/cwallet.sso — Passwords

- DOMAIN\_HOME/config/fmwconfig/oamkeystore — keystore containing keys and certificates OAM/Oracle Security Token Service owns
- DOMAIN\_HOME/config/fmwconfig/amtruststore — keystore containing the trust anchors used for X509 cert validation
- DOMAIN\_HOME/config/fmwconfig/amcrl.jar — zip file containing CRLs used for certificate revocation
- DOMAIN\_HOME/config/fmwconfig/default-keystore.jks — OWSM keystore used to store keys and certificates used by the OWSM Agent, as well as trusted anchors used to validate X.509 certificates for WSS operations

**8.8.3.2.7 External Dependencies** Oracle Security Token Service has external dependencies on the:

- LDAP based Identity Store
  - User Identity Repository
  - LDAP access abstracted by User/Role API.

---

**Note:** Oracle Access Manager always connects to one Identity store, which can be a physical server or a load balancer IP. If the primary is down, Oracle Access Manager reconnects and expects the load balancer to connect it to the secondary.

---

- OCSP Responder Service
  - Real-time X.509 Certification Validation
- RDBMS Policy Store
  - Policy (Authentication and Authorization) Repository
  - RDBMS access abstracted by the OAM policy engine

### 8.8.3.3 Oracle Security Token Service High Availability Configuration Steps

Oracle Security Token Service High Availability is configured as part of Oracle Access Manager. All Oracle Security Token Service system configuration is done using the Oracle Access Manager Console. See [Section 8.8.4, "Oracle Access Manager High Availability Configuration Steps"](#) for steps to configure OAM.

### 8.8.3.4 Validating Oracle Security Token Service High Availability

You can verify that Oracle Security Token Service endpoints are up and running on the different Oracle Security Token Service servers. To do so, access the WSDL document of an Oracle Security Token Service endpoint directly:

`http(s)://[hostname:port]/sts/[ENDPOINT]?WSDL`

Replace [ENDPOINT] with the existing published endpoint.

### 8.8.3.5 Oracle Security Token Service Failover and Expected Behavior

This section describes Oracle Security Token Service failover characteristics in a high availability environment.

Oracle Access Manager Access Servers support a heartbeat check--a ping request over HTTP. In addition, the WebLogic Node Manager on the Managed Server can monitor the application and restart it if necessary.

If a WebLogic Server node fails, external connection failover is based on the configuration, the retry timeout, and the number of retries. Oracle Access Manager Agent-Access Server failover is based on a timeout.

When the load balancing router or proxy server detects a WebLogic Server node failure, subsequent client connections route to the active instance, which picks up the session state from the Coherence Distributed Object Cache and continues processing.

Front Channel HTTP bindings use a load balancing router for failover

When it receives an exception from another service, Oracle Access Manager retries external connection requests. The number of retries is configurable and includes a no retries option.

See the following topics for more information:

- [Section 8.8.2.2.1, "WebLogic Server Crash"](#)
- [Section 8.8.2.2.2, "Node Failure"](#)
- [Section 8.8.2.2.3, "Database Failure"](#)

**8.8.3.5.1 Death Detection and Restart** OAM/Oracle Security Token Service Access Servers support a heartbeat check in the form of a ping request sent over HTTP. Also, the WebLogic Node Manager on the managed server can monitor the application and restart if the event isn't running. Restarting an OAM Access Server does not affect any other cluster components or members.

**8.8.3.5.2 Node Failure** External Connection failover is based on the configuration, retry timeout, and the number of retries. The LBR or Proxy Server detects node failure and subsequent client connections are routed to the active instance, which picks up the session state from the Coherence DOC and continues with the processing.

### 8.8.3.6 Disabling and Enabling Oracle Security Token Service

Oracle Security Token Service is enabled by default. To disable Oracle Security Token Service, you use the Oracle Access Manager Console. See Enabling or Disabling Available Services in the Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service.

### 8.8.3.7 Troubleshooting Oracle Security Token Service

Oracle Security Token Service logs are logged to the Managed Servers log files. However, you can edit the logging.xml so that it logs OSTS information to a separate log file, diagnostic.log, in the folder

`<DomainHome>/config/fmwconfig/servers/<servername>/sts/log/.`

To create an Oracle Security Token Service log file to troubleshoot Oracle Security Token Service:

1. Open the file  
`<DomainHome>/config/fmwconfig/servers/<servername>/logging.xml`
2. Add the following in the appropriate sections:
 

```
<log_handler name='sts-handler'
class='oracle.core.ojdl.logging.ODLHandlerFactory'>
 <property name='path' value='sts/log' />
 <property name='maxFileSize' value='10485760' />
 <property name='maxLogSize' value='104857600' />
</log_handler>
```

```
<logger name='oracle.security.fed' level='TRACE:32'>
 <handler name='sts-handler' />
</logger>
```

### 8.8.3.8 Log File Location

All log messages go to the server log file of the WebLogic Server that the application is deployed on. The default location of the server log is:

```
WL_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

### 8.8.3.9 Additional Considerations

The Oracle Security Token Service server can detect fake requests, such as replay attacks, that can occur if a user tries to steal token data from a request and send another request with the same token. In this case, the server detects the second fake request. The second issuance request with the same token in <Env: Body> goes to the Oracle Security Token Service server. The server denies the request after checking its UNT token cache, which indicates a replay attack.

## 8.8.4 Oracle Access Manager High Availability Configuration Steps

This section provides high-level instructions for setting up a high availability deployment for Oracle Access Manager. This deployment includes two Oracle HTTP Servers, which distribute requests to two OAM servers. These OAM servers interact with an Oracle Real Application Clusters (Oracle RAC) database and optionally an external LDAP store. If any single component fails, then the remaining components will continue to function.

This section includes the following topics:

- [Section 8.8.4.1, "Prerequisites for Oracle Access Manager Configuration"](#)
- [Section 8.8.4.2, "Run the Repository Creation Utility to Create the OAM Schemas in a Database"](#)
- [Section 8.8.4.3, "Installing Oracle WebLogic Server"](#)
- [Section 8.8.4.4, "Install and Configure the Oracle Access Manager Application Tier"](#)
- [Section 8.8.4.5, "Creating boot.properties for the Administration Server on OAMHOST1"](#)
- [Section 8.8.4.6, "Start OAMHOST1"](#)
- [Section 8.8.4.7, "Validating OAMHOST1"](#)
- [Section 8.8.4.8, "Configure OAM on OAMHOST2"](#)
- [Section 8.8.4.9, "Start OAMHOST2"](#)
- [Section 8.8.4.10, "Validating OAMHOST2"](#)
- [Section 8.8.4.11, "Configure Oracle Access Manager to Work with Oracle HTTP Server"](#)
- [Section 8.8.4.11, "Configure Oracle Access Manager to Work with Oracle HTTP Server"](#)
- [Section 8.8.4.12, "Configure Oracle Access Manager to use an External LDAP Store"](#)
- [Section 8.8.4.13, "Validating the Oracle Access Manager Configuration"](#)

- [Section 8.8.4.14, "Configuring Oracle Coherence to Keep Configuration Files in Sync"](#)
- [Section 8.8.4.15, "Scaling Up and Scaling Out the Oracle Access Manager Topology"](#)

#### 8.8.4.1 Prerequisites for Oracle Access Manager Configuration

Before you configuring Oracle Access Manager for high availability, you must:

- Run the Repository Creation Utility to create the OAM schemas in a database.  
See [Section 8.8.4.2, "Run the Repository Creation Utility to Create the OAM Schemas in a Database"](#) for instructions on running the Repository Creation Utility to create the OAM schemas.
- Install Oracle WebLogic Server on OAMHOST1 and OAMHOST2.  
Follow the steps in [Section 8.8.4.3, "Installing Oracle WebLogic Server"](#) to install Oracle WebLogic Server on OAMHOST1 and OAMHOST2.
- Install the Oracle Identity Management executables on OAMHOST1 and OAMHOST2.  
Follow the steps in [Section 8.8.4.4, "Install and Configure the Oracle Access Manager Application Tier"](#) to install the Oracle Identity Management executables on OAMHOST1 and OAMHOST2.
- Ensure that a highly available LDAP implementation is available. For information about installing and configuring Oracle Internet Directory, see [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps."](#) For information about installing and configuring Oracle Virtual Directory, see [Section 8.4.3, "Oracle Virtual Directory High Availability Configuration Steps."](#)

#### 8.8.4.2 Run the Repository Creation Utility to Create the OAM Schemas in a Database

See [Section 8.2.4.1, "Executing the Repository Creation Utility"](#) for instructions on running the Repository Creation Utility to create the OAM schemas in your database repository.

#### 8.8.4.3 Installing Oracle WebLogic Server

Prior to installing the Oracle WebLogic Server, ensure that your machines meet the system, patch, kernel, and other requirements as specified in *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

Start the installer, then proceed as follows:

1. On the Welcome screen, click **Next**.
2. On the Choose Middleware Home Directory screen, select **Create a New Middleware Home**.

For Middleware Home Directory, enter:

`ORACLE_BASE/product/fmw`

---

**Note:** `ORACLE_BASE` is the base directory under which Oracle products are installed. The recommended value is `/u01/app/oracle`.

---

Click **Next**.

3. On the Register for Security Updates screen, enter your contact information so that you can be notified of security updates.

Click **Next**.

4. On the Choose Install Type screen, select **Custom**.

Click **Next**.

5. On the Choose Products and Components screen, select only **Oracle JRockit SDK**, and click **Next**.

6. On the Choose Product Installation Directories screen, accept the directory `ORACLE_BASE/product/fmw/wlserver_10.3`.

Click **Next**.

7. On the Installation Summary screen, click **Next**.

8. On the Installation Complete screen, deselect **Run Quickstart**.

Click **Done**.

#### 8.8.4.4 Install and Configure the Oracle Access Manager Application Tier

This section describes how to install Oracle Fusion Middleware components on OAMHOST1 and OAMHOST2.

**8.8.4.4.1 Install Oracle Fusion Middleware for Identity Management** This section includes the steps for installing the Oracle Identity Management software into the previously created Middleware Home directory. The steps should be performed on OAMHOST1 and OAMHOST2.

On Linux platforms, if the `/etc/oraInst.loc` file exists, verify that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

Start the installer for Oracle Fusion Middleware as follows:

```
OAMHOST1> runInstaller
```

When the installer prompts you for a JRE/JDK location, enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `ORACLE_BASE/product/fmw/jrockit_160_<version>`.

Then proceed as follows:

1. On the Welcome screen, click **Next**.
2. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
3. On the Specify Installation Location screen, enter the following values:

- Oracle Middleware Home: Select the previously installed Middleware home from the list for `MW_HOME`, for example:

```
/u01/app/oracle/product/fmw
```

- Oracle Home Directory: Enter `idm`.

Click **Next**.



4. On the Installation Summary screen, click **Install**.

When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the root user.

5. On the Installation Complete screen, click **Finish**.

#### 8.8.4.4.2 Configure Oracle Identity Management on OAMHOST1

This section creates the Oracle Identity Management domain on OAMHOST1.

Start the configuration wizard by running the command:

```
MW_HOME/oracle_common/common/bin/config.sh
```

Then proceed as follows:

1. In the Welcome screen, select **Create a New WebLogic Domain**, and then click **Next**.
2. In the Select Domain Source Screen:
 

Select **Generate a domain configured automatically to support the following products**:

And select the following products:

  - **Oracle Enterprise Manager**
  - **Oracle JRF** (selected by default)
  - **Oracle Access Manager with Database Policy Store**

Click **Next**.
3. In the Specify Domain and Location screen enter:
  - **Domain name:** `IDM_Domain`
  - **Domain Location:** Accept the default.
  - **Application Directory:** Accept the default.

Click **Next**.
4. In the Configure Administrator Username and Password screen, enter the username and password to be used for the domain's administrator, and click **Next**.
5. In the Configure Server Start Mode and JDK screen, make the following selections:
  - **WebLogic Domain Startup Mode:** Select **Production Mode**.
  - **JDK Selection:** Select **JROCKIT SDK1.6.0\_<version>**.
6. In the Configure JDBC Component Schema screen, select all of the data sources, then select **Configure selected data sources as RAC multi data sources**.
 

Click **Next**.
7. In the Configure RAC Multi Data Source Component Schema screen, select the first data source, **OAM Admin Server**, and enter the following:
  - **Data source:** `OAM`
  - **Service Name:** `oam.mycompany.com`
  - **User Name:** `OAM_OAM` (assuming OAM was used as the RCU prefix)
  - **Password:** The password for above account

In the top right box, click **Add** to add an Oracle RAC host. Enter the following information:

- **Host Name:** OAMDBHOST1
- **Instance Name:** oamdb1
- **Port:** 1521

Click **Add** again to add the second database host and enter the following information:

- **Host Name:** OAMDBHOST2
- **Instance Name:** oamdb2
- **Port:** 1521

Click **Next**.

8. In the Test Component Schema screen, the configuration wizard attempts to validate the data source.

If the data source validation succeeds, click **Next**.

If it fails, click **Previous**, correct the issue, and try again.

9. In the Select Optional Configuration screen, select:

- **Administration Server**
- **Managed Server Clusters and Machines**

Click **Next**.

10. In the Customize Server and Cluster Configuration screen, select **Yes**, and click **Next**.

11. In the Configure the Administration Server screen, enter the following values:

- **Name:** AdminServer
- **Listen Address:** OAMHOST1.mycompany.com
- **Listen Port:** 7001
- **SSL listen port:** Not applicable
- **SSL enabled:** leave unchecked

Click **Next**.

12. On the Configure Managed Servers screen, create an entry for each OAMHOST in the topology, that is, one for the OAM Server running on OAMHOST1 and one for the OAM Server running on OAMHOST2.

Select the OAM\_SERVER entry and change the entry to the following values:

- **Name:** WLS\_OAM1
- **Listen Address:** OAMHOST1.mycompany.com
- **Listen Port:** 14100

For the second OAM\_SERVER, click **Add** and supply the following information:

- **Name:** WLS\_OAM2
- **Listen Address:** OAMHOST2.mycompany.com
- **Listen Port:** 14100

Click **Next**.

13. In the Configure Clusters screen, create a cluster by clicking **Add**.

Enter name: OAM\_Cluster

Leave all other fields at the default settings.

Click **Next**.

14. On the Assign Servers to Clusters screen, associate the managed servers with the cluster, as follows:

- Click the cluster name **OAM\_Cluster** in the right window.
- Click the managed server **WLS\_OAM1**, then click the arrow to assign it to the cluster.
- Repeat for managed server **WLS\_OAM2**.

Click **Next**.

15. On the Configure Machines screen, create a machine for each host in the topology. Click the UNIX tab if your hosts use a UNIX-based operating system. Otherwise, click the Machines tab. Supply the following information:

- **Name:** Name of the host. The best practice is to use the DNS name (OAMHOST1)
- **Node Manager Listen Address:** The DNS name of the machine (OAMHOST1.mycompany.com)
- **Node Manager Port:** A port for Node Manager to use.

Repeat the steps for OAMHOST2:

- **Name:** Name of the host. The best practice is to use the DNS name (OAMHOST2)
- **Node Manager Listen Address:** The DNS name of the machine (OAMHOST2.mycompany.com)
- **Node Manager Port:** A port for Node Manager to use.

Click **Next**.

16. In the Assign Servers to Machines screen, indicate which managed servers will run on the machines just created.

- Click the machine **OAMHOST1** in the right window.
- Click the managed server **WLS\_OAM1** in the left window.
- Click the arrow to assign the managed server to the host **OAMHOST1**.
- Click the machine **OAMHOST2** in the right window.
- Click the managed server **WLS\_OAM2** in the left window.
- Click the arrow to assign the managed server to the host **OAMHOST2**.

Click **Next**.

17. On the Configuration Summary screen, click **Create** to begin the creation process.

When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the root user.

---

**Note:** You cannot run the `config.sh` script twice to make configuration changes. You must use another tool to make additional configuration changes such as using the MBeans Browser in Fusion Middleware Control

---

#### 8.8.4.5 Creating `boot.properties` for the Administration Server on OAMHOST1

This section describes how to create a `boot.properties` file for the Administration Server on OAMHOST1. The `boot.properties` file enables the Administration Server to start without prompting for the administrator username and password.

Follow these steps to create the `boot.properties` file:

1. On OAMHOST1, go the following directory:

```
MW_HOME/user_projects/domains/domainName/servers/AdminServer/security
```

For example:

```
cd /u01/app/oracle/product/fmw/user_
projects/domains/IDMDomain/servers/AdminServer/security
```

2. Use a text editor to create a file called `boot.properties` under the `security` directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

3. Stop the Administration Server if it is running.

See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.

4. Start the Administration Server on OAMHOST1 using the `startWebLogic.sh` script located under the `MW_HOME/user_projects/domains/domainName/bin` directory.
5. Validate that the changes were successful by opening a web browser and accessing the following pages:

- WebLogic Server Administration Console at:

```
http://oamhost1.mycompany.com:7001/console
```

- Oracle Enterprise Manager Fusion Middleware Control at:

```
http://oamhost1.mycompany.com:7001/em
```

Log into these consoles using the `weblogic` user credentials.

### 8.8.4.6 Start OAMHOST1

Now you will start OAMHOST1.

This section describes the steps for starting OAMHOST1.

**8.8.4.6.1 Create the Node Manager Properties File on OAMHOST1** Before you can start managed servers from the console, you must create a Node Manager property file. You do this by running the script `setNMProps.sh`, which is located in the `MW_HOME/oracle_common/common/bin` directory. For example:

```
OAMHOST1> $MW_HOME/oracle_common/common/bin/setNMProps.sh
```

**8.8.4.6.2 Start Node Manager** Start Node Manager by issuing the following command:

```
OAMHOST1> $MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
```

**8.8.4.6.3 Start Oracle Access Manager on OAMHOST1** To start Oracle Access Manager on OAMHOST1, follow these steps:

1. Log into the WebLogic Administration Console using this URL:  
`http://oamhost1.mycompany.com:7001/console`
2. Supply the WebLogic administrator username and password.
3. Select **Environment - Servers** from the **Domain Structure** menu.
4. Click the Control tab.
5. Click the server **WLS\_OAM1**.
6. Click **Start**.
7. Click **OK** to confirm that you want to start the server.

### 8.8.4.7 Validating OAMHOST1

Validate the implementation by connecting to the OAM Console at the following URL:

```
http://OAMHOST1.mycompany.com:7001/oamconsole
```

The implementation is valid if the OAM Admin console login page is displayed and you can login using the WebLogic administrator account.

### 8.8.4.8 Configure OAM on OAMHOST2

Once the configuration has succeeded on OAMHOST1, you can propagate it to OAMHOST2. You do this by packing the domain using the `pack` script on OAMHOST1, and unpacking the domain using the `unpack` script on OAMHOST2.

Both scripts reside in the `MW_HOME/oracle_common/common/bin` directory.

On OAMHOST1, enter:

```
pack.sh -domain=$MW_HOME/user_projects/domains/IDM_Domain \
 -template=/tmp/idm_domain.jar -template_name="OAM Domain" -managed=true
```

This creates a file called `idm_domain.jar` in the `/tmp` directory. Copy this file to OAMHOST2.

On OAMHOST2, enter:

```
unpack.sh -domain=$MW_HOME/user_projects/domains/IDM_Domain \
```

```
-template=/tmp/idm_domain.jar
```

#### 8.8.4.9 Start OAMHOST2

Now you will start OAMHOST2.

This section describes the steps for starting OAMHOST2.

**8.8.4.9.1 Create the Node Manager Properties File on OAMHOST2** Before you can start managed servers from the console, you must create a Node Manager property file. You do this by running the script `setNMProps.sh`, which is located in the `MW_HOME/oracle_common/common/bin` directory. For example:

```
OAMHOST1> $MW_HOME/oracle_common/common/bin/setNMProps.sh
```

**8.8.4.9.2 Start Node Manager** Start Node Manager by issuing the following command:

```
OAMHOST2> $MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
```

**8.8.4.9.3 Start Oracle Access Manager on OAMHOST2** To start Oracle Access Manager on OAMHOST2, follow these steps:

1. Log into the WebLogic Administration Console using this URL:

```
http://oamhost2.mycompany.com:7001/console
```

2. Supply the WebLogic administrator username and password.
3. Select **Environment - Servers** from the **Domain Structure** menu.
4. Click the Control tab.
5. Click the server **WLS\_OAM2**.
6. Click **Start**.
7. Click **OK** to confirm that you want to start the server.

#### 8.8.4.10 Validating OAMHOST2

Validate the implementation by connecting to the OAM server using the following URL:

```
http://OAMHOST2.mycompany.com:14100/oam
```

The implementation is valid if the OAM login page is displayed. Note at this point it will show an "Action failed" error on the page. This is normal because you are accessing the page directly rather than as a login request.

#### 8.8.4.11 Configure Oracle Access Manager to Work with Oracle HTTP Server

This section provides steps for configuring Oracle Access Manager to work with the Oracle HTTP Server.

**8.8.4.11.1 Update Oracle HTTP Server Configuration** On WEBHOST1 and WEBHOST2, create a file named `oam.conf` in the following directory:

```
ORACLE_INSTANCE/config/OHS/ohs1/moduleconf
```

Create the file with the following lines:

```

NameVirtualHost *:7777
<VirtualHost *:7777>

 ServerName sso.mycompany.com:7777
 ServerAdmin you@your.address
 RewriteEngine On
 RewriteOptions inherit

 <Location /oam>
 SetHandler weblogic-handler
 WebLogicCluster oamhost1.mycompany.com:14100,oamhost2.mycompany.com:14100
 </Location>

</VirtualHost>

```

#### 8.8.4.11.2 Restart Oracle HTTP Server

Restart the Oracle HTTP Server on WEBHOST1 and WEBHOST2:

```

WEBHOST1>opmnctl stopall
WEBHOST1>opmnctl startall

WEBHOST2>opmnctl stopall
WEBHOST2>opmnctl startall

```

#### 8.8.4.11.3 Make Oracle Access Manager Server Aware of the Load Balancer

By default, Oracle Access Manager sends requests to the login page located on the local server. In a high availability deployment, you must change this setup so that login page requests go to the load balancer.

To make Oracle Access Manager aware of the load balancer:

1. Log into the Oracle Access Manager Console at this URL as the `weblogic` user:  
<http://IDMHOST1.mycompany.com/oamconsole>
2. Click on the **System Configuration** tab.
3. Expand **Access Manager Settings**.
4. Double click **Access Manager Settings**.
5. Enter the following information:
  - OAM Server Host: `sso.mycompany.com`
  - OAM Server Port: `7777`
  - OAM Server Protocol: `http`
6. Click **Apply**.
7. Restart managed servers `WLS_OAM1` and `WLS_OAM2`.

#### 8.8.4.12 Configure Oracle Access Manager to use an External LDAP Store

By default, Oracle Access Manager uses its own built-in LDAP server. In a highly available configuration, it is recommended that an external LDAP directory be used as the directory store. This section describes how to set up an external LDAP store.

---

**Note:** It is recommended that you back up the environment and LDAP store before performing any of the subsequent steps in this section.

---

**8.8.4.12.1 Extending Directory Schema for Oracle Access Manager** Pre-configuring the Identity Store extends the schema in the backend directory regardless of directory type.

To extend the directory schema for Oracle Access Manager, perform the following tasks on OAMHOST1:

1. Set the Environment Variables: MW\_HOME, JAVA\_HOME, IDM\_HOME and ORACLE\_HOME.

Set IDM\_HOME to IDM\_ORACLE\_HOME

Set ORACLE\_HOME to IAM\_ORACLE\_HOME

2. Create a properties file `extend.props` that contains the following:

```
IDSTORE_HOST : idstore.mycompany.com
IDSTORE_PORT : 389
IDSTORE_BINDDN : cn=orcladmin
IDSTORE_USERNAMEATTRIBUTE: cn
IDSTORE_LOGINATTRIBUTE: uid
IDSTORE_USERSEARCHBASE:cn=Users,dc=mycompany,dc=com
IDSTORE_GROUPSEARCHBASE: cn=Groups,dc=us,dc=oracle,dc=com
IDSTORE_SEARCHBASE: dc=mycompany,dc=com
IDSTORE_SYSTEMIDBASE: cn=systemids,dc=mycompany,dc=com
```

Where:

- IDSTORE\_HOST and IDSTORE\_PORT are, respectively, the host and port of your Identity Store directory. If you are using a non-OID directory, then specify the Oracle Virtual Directory host which should be `IDSTORE.mycompany.com`.)
  - IDSTORE\_BINDDN Is an administrative user in the Identity Store Directory
  - IDSTORE\_USERSEARCHBASE is the location in your Identity Store where users are placed.
  - IDSTORE\_GROUPSEARCHBASE is the location in your Identity Store where groups are placed.
  - IDSTORE\_SEARCHBASE is the location in the directory where Users and Groups are stored.
  - IDSTORE\_SYSTEMIDBASE is the location in your directory where the Oracle Identity Manager reconciliation user are placed.
  - IDSTORE\_SYSTEMIDBASE is the location of a container in the directory where users can be placed when you do not want them in the main user container. This happens rarely but one example is the Oracle Identity Manager reconciliation user which is also used for the bind DN user in Oracle Virtual Directory adapters.
3. Configure the Identity Store by using the command `idmConfigTool`, which is located at: `IAM_ORACLE_HOME/idmtools/bin`.

The Linux command syntax is:



```
idmConfigTool.sh -preConfigIDStore input_file=configfile
```

The Windows command syntax is:

```
idmConfigTool.bat -prepareIDStore input_file=configfile
```

For example:

```
idmConfigTool.sh -preConfigIDStore input_file=extend.props
```

When the command runs, the system prompts you for the account password with which you are connecting to the Identity Store.

Sample command output:

```
./preconfig_id.sh
Enter ID Store Bind DN password :
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/idm_idstore_groups_
template.ldif
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/idm_idstore_groups_acl_
template.ldif
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/systemid_pwdpolicy.ldif
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFileINFO: ->
LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/idstore_tuning.ldifApr
5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oid_schema_extn.ldif
The tool has completed its operation. Details have been logged to
automation.log
```

4. Check the log file for errors and warnings and correct them.

**8.8.4.12.2 Create Users and Groups in LDAP** To add users that Oracle Access Manager requires to the Identity Store, follow these steps:

1. Set the Environment Variables MW\_HOME, JAVA\_HOME, IDM\_HOME, and ORACLE\_HOME.
  - Set IDM\_HOME to IDM\_ORACLE\_HOME.
  - Set ORACLE\_HOME to IAM\_ORACLE\_HOME.
2. Create a properties file oam.props that contains the following:

```
IDSTORE_HOST : idstore.mycompany.com
IDSTORE_PORT : 389
IDSTORE_BINDDN : cn=orcladmin
IDSTORE_USERNAMEATTRIBUTE: cn
IDSTORE_LOGINATTRIBUTE: uid
IDSTORE_USERSEARCHBASE: cn=Users,dc=mycompany,dc=com
IDSTORE_GROUPSEARCHBASE: cn=Groups,dc=mycompany,dc=com
IDSTORE_SEARCHBASE: dc=mycompany,dc=com
```

```
POLICYSTORE_SHARES_IDSTORE: true
OAM11G_IDSTORE_ROLE_SECURITY_ADMIN: OAMAdministrators
IDSTORE_OAMSOFTWAREUSER: oamLDAP
IDSTORE_OAMADMINUSER: oamadmin
```

Where:

- IDSTORE\_HOST and IDSTORE\_PORT are, respectively, the host and port of your Identity Store directory.
  - IDSTORE\_BINDDN is an administrative user in the Identity Store Directory.
  - IDSTORE\_USERSEARCHBASE is the location in the directory where Users are Stored.
  - IDSTORE\_GROUPSEARCHBASE is the location in the directory where Groups are Stored.
  - IDSTORE\_SEARCHBASE is the location in the directory where Users and Groups are stored.
  - POLICYSTORE\_SHARES\_IDSTORE is set to true if your Policy and Identity Stores are in the same directory. If not, it is set to false.
  - IDSTORE\_OAMADMINUSER is the name of the user you want to create as your Oracle Access Manager Administrator.
  - IDSTORE\_OAMSOFTWAREUSER is a user that gets created in LDAP that is used when Oracle Access Manager is running to connect to the LDAP server.
3. Configure the Identity Store using the command `idmConfigTool` which is located at `IAM_ORACLE_HOME/idmtools/bin`.

The Linux command syntax is:

```
idmConfigTool.sh -prepareIDStore mode=OAM input_
file=configfile
```

The Windows command syntax:

```
idmConfigTool.bat -prepareIDStore mode=OAM input_
file=configfile
```

For example:

```
idmConfigTool.sh -prepareIDStore mode=OAM input_file=oam.props
```

After the command runs, the system prompts you to enter the password for the account with which you are connecting to the ID Store.

Sample command output:

```
Enter ID Store Bind DN password :
Apr 5, 2011 3:53:28 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/oam/server/oim-intg/schema/OID_oblix_schema_
add.ldif
Apr 5, 2011 3:54:12 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/oam/server/oim-intg/schema/OID_oblix_schema_
index_add.ldif
Apr 5, 2011 3:55:10 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:
```

```

/u01/app/oracle/product/fmw/IAM/oam/server/oim-intg/schema/OID_oblix_pwd_
schema_add.ldif
Apr 5, 2011 3:55:11 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/oam/server/oim-intg/schema/OID_oim_pwd_schema_
add.ldif
*** Creation of Oblix Anonymous User ***
Apr 5, 2011 3:55:11 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oam_10g_anonymous_user_
template.ldif
Enter User Password for oblixanonymous:
Confirm User Password for oblixanonymous:
Apr 5, 2011 3:55:53 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oam_group_member_
template.ldif
The tool has completed its operation. Details have been logged to
automation.log

```

4. Check the log file for any errors or warnings and correct them.

See Oracle Fusion Middleware Integration Overview for Oracle Identity Management Suite for more information about the `idmConfigTool` command.

#### 8.8.4.12.3 Create a User Identity Store To create a user identity store:

1. Go to the Oracle Access Manager Console at the URL:

`http://adminvhn.mycompany.com:7001/oamconsole`

2. Log in using the WebLogic administration user.
3. Click **Add User Identity Store** and enter the following information:
  - **Store Name:** LDAP\_DIR
  - **Store Type:** OVD
  - **Description:** Enter a description of the Directory Store
  - **Enable SSL:** Select this if you communicate with your directory over SSL
  - **Location:** Enter the location, for example `ovd.mycompany.com:389`
  - **Bind DN:** Enter the user permitted to search the LDAP store. For example, `cn=orcladmin`
  - **Password:** Enter the `orcladmin` password
  - **User Name Attribute:** For example: `uid`
  - **User Search Base:** Enter the location of users in the LDAP store. For example, `cn=Users,dc=mycompany,dc=com`
  - **Group Name Attribute:** For example: `orclguid`
  - **Group Search Base:** Enter the location of groups in the LDAP store. For example, `cn=Groups,dc=mycompany,dc=com`
  - **OAM Administrator Role:** OAMAdministrator

4. Click **Apply**.
5. Click **Test Connection** to validate the connection to the LDAP server.

**8.8.4.12.4 Set LDAP to System and Default Store** Now that you have defined the LDAP identity store, you must set it as the primary authentication store. To do this, follow these steps in the Oracle Access Manager Console:

1. Click the **System Configuration** tab.
2. Select **Data Sources - User Identity Stores** from the navigation pane.
3. Click **LDAP\_DIR**.
4. Select **Open** from the **Actions** menu.
5. Click **Set as Default Store**.
6. Click **Set as System Store**.
7. Click the Add [+] icon in **Access System Administrators**.
8. Enter **OAM\*** in the search name field and click **Search**.
9. Select **OAMAdministrator** from the search results and click **Add Selected**.
10. Click **Apply**.
11. In the Validate System Administrator window, enter the username and password of the OAM administrator, for example, oamadmin.
12. Click **Validate**.
13. Test the connection by clicking **Test Connection**.

**8.8.4.12.5 Set Authentication to Use External LDAP** By default, Oracle Access Manager uses the integrated LDAP store for user validation. You must update the LDAP authentication module so that it can validate users against the new external LDAP store.

To update the LDAP authentication module to use external LDAP:

1. Click the **System Configuration** tab.
2. Select **Access Manager Settings - Authentication Modules - LDAP Authentication Modules**.
3. Click **LDAP**.
4. Select **Open** from the **Actions** menu.
5. Set **User Identity Store** to **LDAP\_DIR**.
6. Click **Apply**.
7. Restart the managed servers Admin Server, WLS\_OAM1 and WLS\_OAM2.

### 8.8.4.13 Validating the Oracle Access Manager Configuration

Validate the configuration by logging into the Oracle Access Manager Console at <http://oamhost1.mycompany.com:7001/oamconsole> as oamadmin.

### 8.8.4.14 Configuring Oracle Coherence to Keep Configuration Files in Sync

In a highly available environment Oracle uses Oracle Coherence to keep configuration files in sync. Oracle Coherence uses port 9095 by default, but this can be changed through the Oracle Access Manager Console.

Log in to the console using the url `http://admin.us.oracle.com/oamconsole`, then follow these steps:

1. Click on the **System Configuration** tab.
2. Expand **Servers** in the navigator.
3. Double-click on the Managed Server whose port you wish to change.
4. Click on the **Coherence** tab.
5. Change the value of **Local Port** to the desired value.
6. Click **Apply**.
7. Restart the Administration Server and all the managed servers residing in the same cluster as the managed server that has been updated.

#### 8.8.4.15 Scaling Up and Scaling Out the Oracle Access Manager Topology

This section describes how to scale up and scale out an Oracle Access Manager high availability topology. Perform a scale up operation to add a new Oracle Access Manager managed server instance is added to a node already running one or more server instances. Perform a scale out operation to add a new Oracle Access Manager managed server instance to a new node.

##### 8.8.4.15.1 Scaling Up Oracle Access Manager Scale up OAM as follows:

Log in to the Oracle WebLogic Server Administration Console at `http://oamhost1.mycompany.com:7001/console`.

1. From the Domain Structure window of the Oracle WebLogic Server Administration Console, expand the **Environment** node and then **Servers**. The Summary of Servers page appears.
2. In the Change Center, click **Lock & Edit**.
3. Select an existing server on the host you wish to extend, for example: `WLS_OAM1`.
4. Click **Clone**.
5. Enter the following information:
  - **Server Name:** A new name for the server, for example: `WLS_OAM3`.
  - **Server Listen Address:** The name of the host on which the managed server will run.
  - **Server Listen Port:** The port the new managed server will use, this port must be unique within the host.
6. Click **OK**.
7. Click on the newly created server **WLS\_OAM3**
8. Set the SSL listen port. This should be unique on the host that the managed server will be running on.
9. Click **Save**.
10. Disable host name verification for the new managed server. Before starting and verifying the `WLS_OAM3` managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in `OAMHOSTn`.

If the source server from which the new one was cloned had already disabled hostname verification, these steps are not required, as the hostname verification settings were propagated to the cloned server. To disable host name verification:

- a. In Oracle Enterprise Manager Fusion Middleware Control, select **Oracle WebLogic Server Administration Console**.
  - b. Expand the **Environment** node in the Domain Structure window.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Select **WLS\_OAM3** in the Names column of the table. The Settings page for server appears.
  - e. Click the **SSL** tab.
  - f. Click **Advanced**.
  - g. Set **Hostname Verification** to **None**.
  - h. Click **Save**.
11. Click **Activate configuration** from the Change Center menu.

Register the new managed server with OAM. You now need to configure the new managed server now as an OAM server. You do this from the Oracle OAM console. Proceed as follows:

1. Log in to the OAM console at  
`http://oamhost1.mycompany.com:7001/oamconsole` as the `oamadmin` user.
2. Click the **System Configuration** tab.
3. Click **Server Instances**.
4. Select **Create** from the Actions menu.
5. Enter the following information:
  - **Server Name:** `WLS_OAM3`
  - **Host:** Host that the server will run on
  - **Port:** Listen port that was assigned when the managed server was created
  - **OAM Proxy Port:** Port you want the OAM proxy to run on. This is unique for the host
  - **Proxy Server ID:** `AccessServerConfigProxy`
  - **Mode:** `Open`
6. Click **Coherence** tab.  
Set **Local Port** to a unique value on the host.  
Click **Apply**.
7. Click **Apply**.

You can now start the Access server. To use the Access server, you must inform any Webgates of its existence. You do that as follows:

1. Log in to the OAM console at  
`http://oamhost1.mycompany.com:7001/oamconsole` as the `oamadmin` user.
2. Click the **System Configuration** tab.

3. Expand **Agents** -> **OAM Agents** -> **10g Agents** (for 10g WebGates) or **Agents** > **OAM Agents** > **11g Agents** (for 11g WebGates).
4. Double click the Webgate you want to change.
5. Add the new server to either the primary or secondary server list by clicking the Add + icon.
6. Select the server name from the list.
7. Click **Apply**

**8.8.4.15.2 Scaling Out Oracle Access Manager** Scale out is very similar to scale up, but first requires the software to be installed on the new node.

1. Install Oracle WebLogic Server on the new host as described in [Section 8.8.4.3, "Installing Oracle WebLogic Server."](#)
2. Install Oracle Fusion Middleware Identity Management components on the new host as described in [Section 8.8.4.4, "Install and Configure the Oracle Access Manager Application Tier."](#)
3. Log in to the Oracle WebLogic Server Administration Console at `http://oamhost1.mycompany.com:7001/oamconsole`.
4. From the Domain Structure window of the Oracle WebLogic Server Administration Console, expand the **Environment** node and then **Servers**. The Summary of Servers page appears.
5. In the Change Center, click **Lock & Edit**.
6. Select an existing server on the host you want to extend, for example: `WLS_OAM1`.
7. Click **Clone**.
8. Enter the following information:
  - **Server Name:** A new name for the server, for example: `WLS_OAM3`.
  - **Server Listen Address:** The name of the host on which the managed server will run.
  - **Server Listen Port:** The port the new managed server will use. This port must be unique within the host.
9. Click **OK**.
10. Click the newly created server `WLS_OAM3`.
11. Set the SSL listen port. This should be unique on the host that the managed server will run on.
12. Click **Save**.
13. Disable host name verification for the new managed server. Before starting and verifying the `WLS_OAM3` managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in `OAMHOSTn`.

If the source server from which the new one was cloned had already disabled hostname verification, these steps are not required, as the hostname verification settings was propagated to the cloned server. To disable host name verification, proceed as follows:

- a. In Oracle Enterprise Manager Fusion Middleware Control, select Oracle WebLogic Server Administration Console.
- b. Expand the **Environment** node in the Domain Structure pane.
- c. Click **Servers**. The Summary of Servers page appears.
- d. Select **WLS\_OAM3** in the Names column of the table. The Settings page for server appears.
- e. Click the **SSL** tab.
- f. Click **Advanced**.
- g. Set **Hostname Verification** to **None**.
- h. Click **Save**.

14. Click **Activate Configuration** from the Change Center menu.

15. Pack the domain on OAMHOST1 using the command:

```
pack.sh -domain=ORACLE_BASE/admin/IDM_Domain/aserver/IDM_Domain -template
=/tmp/idm_domain.jar -template_name="OAM Domain" -managed=true
```

The `pack.sh` script is located in `MW_HOME/oracle_common/common/bin`.

16. Unpack the domain on the new host using the command:

```
unpack.sh -domain=ORACLE_BASE/admin/IDM_Domain/mserver/IDM_Domain
-template=/tmp/idm_domain.jar -template_name="OAM Domain" -app_dir=ORACLE_
BASE/admin/IDM_Domain/mserver/applications
```

The `unpack.sh` script is located in `MW_HOME/oracle_common/common/bin`.

17. Before you can start managed servers from the console, you must create a node manager properties file on OAMHOST3. You do this by running the script `setNMProps.sh`, which is located in `MW_HOME/oracle_common/common/bin`. Type:

```
MW_HOME/oracle_common/common/bin/setNMProps.sh
```

Register the new managed server with OAM. The new managed server now needs to be configured as an OAM server. You do this from the Oracle OAM console, as follows:

1. Log in to the OAM console at `http://oamhost1.mycompany.com:7001/oamconsole` as the `oamadmin` user.
2. Click the **System Configuration** tab.
3. Click **Server Instances**.
4. Select **Create** from the Actions menu.
5. Enter the following information:
  - **Server Name:** `WLS_OAM3`
  - **Host:** Host that the server will be running on, `OAMHOST3`.
  - **Port:** Listen port that was assigned when the managed server was created.
  - **OAM Proxy Port:** Port you want the OAM proxy to run on. This is unique for the host.
  - **Proxy Server ID:** `AccessServerConfigProxy`



- **Mode:** Open

## 6. Click **Apply**.

You can now start the Access Server. In order for the server to be used, however, you must inform any Webgates of its existence. You do that as follows:

1. Log in to the OAM console at `http://oamhost1.mycompany.com:7001/oamconsole` as the `oamadmin` user.
2. Click the **System Configuration** tab.
3. Expand **Agents** -> **OAM Agents** -> **10g Agents**.
4. Double click the Webgate you want to change.
5. Add the new server to either the primary or secondary server list by clicking the Add [+] icon.
6. Select the server name from the list.
7. Click **Apply**.

Update the Web Tier. Now that the new managed server has been created and started, the web tier will start to direct requests to it. Best practice, however, is to inform the web server that the new managed server has been created.

You do this by updating the file `OAM.conf` on each of the web tiers. This file resides in the directory: `ORACLE_INSTANCE/config/OHS/component_name/moduleconf`.

Add the new server to the `WebLogicCluster` directive in the file, for example, change:

```
<Location /OAM_admin>
 SetHandler weblogic-handler
 WebLogicCluster
 OAMhost1.mycompany.com:14200,OAMhost2.mycompany.com:14200
</Location>
```

to:

```
<Location /OAM_admin>
 SetHandler weblogic-handler
 WebLogicCluster

OAMhost1.mycompany.com:14200,OAMhost2.mycompany.com:14200,OAMhsot3.mycompany.com:14300
</Location>
```

## 8.9 Oracle Identity Manager High Availability

This section introduces Oracle Identity Manager and describes how to design and deploy a high availability environment for Oracle Identity Manager.

Oracle Identity Manager is a user provisioning and administration solution that automates the process of adding, updating, and deleting user accounts from applications and directories. It also improves regulatory compliance by providing granular reports that attest to who has access to what. Oracle Identity Manager is available as a stand-alone product or as part of Oracle Identity and Access Management Suite.

Automating user identity provisioning can reduce Information Technology (IT) administration costs and improve security. Provisioning also plays an important role in regulatory compliance. Key features of Oracle Identity Manager include password management, workflow and policy management, identity reconciliation, reporting and auditing, and extensibility through adapters.

Oracle Identity Manager provides the following key functionalities:

- User Administration
- Workflow and Policy
- Password Management
- Audit and Compliance Management
- User Provisioning
- Organization and Role Management


For details about Oracle Identity Manager, see the *Oracle Fusion Middleware System Administrator's Guide for Oracle Identity Manager*.

This section includes the following topics:

- [Section 8.9.1, "Oracle Identity Manager Component Architecture"](#)
- [Section 8.9.2, "Oracle Identity Manager High Availability Concepts"](#)
- [Section 8.9.3, "Oracle Identity Manager High Availability Configuration Steps"](#)

## 8.9.1 Oracle Identity Manager Component Architecture

[Figure 8–16](#) shows the Oracle Identity Manager architecture:

**Figure 8–16 Oracle Identity Manager Component Architecture**

### 8.9.1.1 Oracle Identity Manager Component Characteristics

Oracle Identity Manager Server is Oracle's self-contained, standalone identity management solution, based on Java EE standards. It provides User Administration, Workflow and Policy, Password Management, Audit and Compliance Management, User Provisioning and Organization and Role Management functionalities.

Oracle Identity Manager is a standard Java EE application that is deployed on Oracle WebLogic Server and uses a database to store runtime and configuration data. The MDS schema contains configuration information; the runtime and user information is stored in the OIM schema.

Oracle Identity Manager connects to the SOA managed servers over RMI to invoke the SOA EJBs.

Oracle Identity Manager uses the human workflow module of the Oracle SOA Suite for managing its request workflow. Oracle Identity Manager connects to SOA using SOA Soapurl for connecting to SOA web services. This is the front end URL for SOA; this should be the load balancer or webserver URL in case of clustered SOA servers. When the workflow is completed, SOA calls back Oracle Identity Manager web services using OimFrontEndURL. Oracle SOA is deployed along with the Oracle Identity Manager.

Several Oracle Identity Manager modules use JMS queues. Each queue is processed by a separate Message Driven Bean (MDB), which is also part of the Oracle Identity

Manager application. Message producers are also part of the Oracle Identity Manager application.

Oracle Identity Manager uses embedded Oracle Entitlements Server (microkernel), which is also part of the Oracle Identity Manager engine. Oracle Entitlements Server (OES) is used for authorization checks inside Oracle Identity Manager. For example, one of the policy constraints determines that only users with certain roles are allowed create users. This is defined using the Oracle Identity Manager user interface.

Oracle Identity Manager uses a Quartz based scheduler for scheduled activities. There are various scheduled activities that happen in the background. For example, one of the scheduled tasks is to disable users after the end date of the users.

Oracle Identity Manager simply links to Oracle BI Publisher for all the reporting features. BI Publisher is expected to be in a different domain or same domain, so the integration is only a simple static URL integration. There is no interaction between BI Publisher and Oracle Identity Manager runtime components. BI Publisher is configured to use the same OIM database schema for reporting purposes.

When you enable LDAPSync to communicate directly with external Directory Servers such as Oracle Internet Directory, ODSEE, and Microsoft Active Directory, support for high availability/failover features requires that you configure the Identity Virtualization Library (libOVD).

To configure libOVD, use the WLST command `addLDAPHost`. To manage libOVD, see *Managing Identity Virtualization Library (libOVD) Adapters* in the guide *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Manager* for a list of WLST commands.

### **8.9.1.2 Runtime Processes**

Oracle Identity Manager is a Java EE application that is deployed on Oracle WebLogic Server as a no stage application. The Oracle Identity Manager server is initialized when the WebLogic Server it is deployed on starts up. As part of the application initialization, the quartz based scheduler is also started. Once initialization is done, the system is ready to receive requests from clients.

Remote Manager and Design Console need to be started as standalone utilities separately.

### **8.9.1.3 Component and Process Lifecycle**

Oracle Identity Manager is deployed to an Oracle WebLogic Server as an externally managed application. By default, the Oracle WebLogic Server starts, stops, monitors and manages other lifecycle events for the Oracle Identity Manager application.

Oracle Identity Manager is a standard Java EE application, and it starts up after the application server components have been started up. Also Oracle Identity Manager uses the authenticator which is part of the Oracle Identity Manager component mechanism; it starts up before the WebLogic JNDI are initialized and the application is started. This is loaded from the OIM ORACLE\_HOME.

Oracle Identity Manager uses a Quartz technology-based scheduler. Quartz starts the scheduler thread on all the WebLogic Server instances. It uses the database as the centralized storage for picking and executing the scheduled activities. If one of the scheduler instances picks up a job, the other instances will not pick up that same job.

Oracle Identity Manager caches certain system configuration values in the cache in memory of the server instance from the database. These caches are independently loaded and not shared among the servers. Any changes to the system configuration triggers the cache cleanup; this process notifies all the servers in the cluster. Oracle

Identity Manager uses oscache, jgroups for this purpose. Jgroups uses multicast addresses. A valid multicast address is randomly generated during installation and seeded to the Oracle Identity Manager metadata repository file.

WebLogic Node Manager can be configured to monitor the server process and restart it in case of failure.

The Oracle Enterprise Manager Fusion Middleware Control is used to monitor as well as to modify the configuration of the application.

#### 8.9.1.4 Starting and Stopping Oracle Identity Manager

Oracle Identity Manager lifecycle events can be managed using one or more of the following command line tools and consoles:

- Oracle WebLogic Scripting Tool (WLST)
- Oracle WebLogic Server Administration Console
- Oracle Enterprise Manager Fusion Middleware Control
- Oracle WebLogic Node Manager

#### 8.9.1.5 Configuration Artifacts

The configuration for the Oracle Identity Manager server is stored in the MDS repository and is managed using MBeans. The oim-config.xml file is the main configuration file. The OIM configuration can be managed using the MBean browser through the Oracle Enterprise Manager Fusion Middleware Control or through the command line MDS utilities. The oim-config.xml file is stored in the /db/oim-config.xml location of the MDS Repository.

For more information about the MDS utilities, see the MDS utilities section of *Oracle Fusion Middleware Developer's Guide for Oracle Identity Manager*.

JMS is configured out of the box by the installer. All the necessary JMS queues, connection pools, data sources and so on are configured on the WebLogic application servers. The following queues are created when Oracle Identity Manager is deployed:

- oimAttestationQueue
- oimAuditQueue
- oimDefaultQueue
- oimKernelQueue
- oimProcessQueue
- oimReconQueue
- oimSODQueue

The Design Console and Remote Manager configuration is stored in the xlconfig.xml file.

#### 8.9.1.6 External Dependencies

Oracle Identity Manager is a Java EE application that is deployed on an Oracle WebLogic Managed Server. Oracle Identity Manager uses the Worklist and Human workflow modules of the Oracle SOA Suite for request flow management. Oracle Identity Manager interacts with external repositories to store configuration and runtime data. Oracle Identity Manager requires these repositories to be available during initialization and during runtime. All Oracle Identity Manager credentials are

stored in the OIM repository. The external components required for the Oracle Identity Manager server to function are listed below:

- Oracle WebLogic Server
  - Administration Server
  - Managed Server
- Data Repositories
  - Configuration Repository (MDS Schema)
  - Runtime Repository (OIM Schema)
  - User Repository (OIM Schema)
  - SOA Repository (SOA Schema)
- External LDAP Stores (when using LDAP Sync)
- BI Publisher

The Design Console is a tool used by the administrator for development and customization. The Design Console communicates directly with the Oracle Identity Manager engine, so it relies on the same components that the Oracle Identity Manager server relies on.

Remote Manager is an optional independent standalone application, which calls the custom APIs on the local system. So it needs the JAR files for those custom APIs in its classpath.

#### **8.9.1.7 Oracle Identity Manager Log File Locations**

Oracle Identity Manager is a Java EE application deployed on Oracle WebLogic Server. All server related logs messages are logged in the server log file and all Oracle Identity Manager specific messages are logged into the diagnostic log file of the Oracle WebLogic Server where the application is deployed.

The Oracle WebLogic Server log files are located under the following directory:

*DOMAIN\_HOME/servers/serverName/logs*

The three main log files are *serverName.log*, *serverName.out*, and *serverName-diagnostic.log*, where *serverName* is the name of the Oracle WebLogic Server. For example, if the Oracle WebLogic Server name is *wls\_OIM1*, then the diagnostic log file name would be *wls\_OIM1-diagnostic.log*.

You can view the log files using the Oracle Enterprise Manager Fusion Middleware Control.

### **8.9.2 Oracle Identity Manager High Availability Concepts**

This section provides an introduction to Oracle Identity Management high availability concepts and describes how to design and deploy a high availability environment for Oracle Identity Manager.


---

**Note:** Be aware of the following when you deploy Oracle Identity Manager in a high availability configuration:

- Oracle Identity Manager can be deployed on an Oracle RAC database, but Oracle RAC failover is not transparent for Oracle Identity Manager in this release. If an Oracle RAC failover occurs, end users may have to resubmit their requests.
  - Oracle Identity Manager always requires that at least one of the nodes in the SOA cluster be available. If the SOA cluster is not available, end user requests will fail. Oracle Identity Manager does not retry for a failed SOA call. Therefore, the end user must retry when a SOA call fails.
- 

### 8.9.2.1 Oracle Identity Manager High Availability Architecture

[Figure 8–17](#) shows Oracle Identity Manager deployed in a high availability architecture in an active-active configuration.

**Figure 8–17 Oracle Identity Manager High Availability Architecture**

On OIMHOST1, the following installations have been performed:

- An Oracle Identity Manager instance has been installed in the WLS\_OIM1 Managed Server and a SOA instance has been installed in the WLS\_SOA1 Managed Server.



- The Oracle RAC database has been configured in a JDBC multi data source to protect the instance from Oracle RAC node failure.
- A WebLogic Server Administration Server has been installed. Under normal operations, this is the active Administration Server.

On OIMHOST2, the following installations have been performed:

- An Oracle Identity Manager instance has been installed in the WLS\_OIM2 Managed Server and a SOA instance has been installed in the WLS\_SOA2 Managed Server.
- The Oracle RAC database has been configured in a JDBC multi data source to protect the instance from Oracle RAC node failure.
- The instances in the WLS\_OIM1 and WLS\_OIM2 Managed Servers on OIMHOST1 and OIMHOST2 are configured as the OIM\_Cluster cluster.
- The instances in the WLS\_SOA1 and WLS\_SOA2 Managed Servers on OIMHOST1 and OIMHOST2 are configured as the SOA\_Cluster cluster.
- A WebLogic Server Administration Server has been installed. Under normal operations, this is the passive Administration Server. You make this Administration Server active if the Administration Server on OIMHOST1 becomes unavailable.

The following virtual host names are used in the Oracle Identity Manager high availability configuration in [Figure 8–17](#):

- OIMVHN1 is the virtual host name that maps to the listen address for the WLS\_OIM1 managed server, and it fails over with server migration of the WLS\_OIM1 managed server. It is enabled on the node where the WLS\_OIM1 managed server is running (OIMHOST1 by default).
- OIMVHN2 is the virtual host name that maps to the listen address for the WLS\_OIM2 managed server, and it fails over with server migration of the WLS\_OIM2 managed server. It is enabled on the node where the WLS\_OIM2 managed server is running (OIMHOST2 by default).
- SOAVHN1 is the virtual host name that is the listen address for the WLS\_SOA1 managed server, and it fails over with server migration of the WLS\_SOA1 managed server. It is enabled on the node where the WLS\_SOA1 managed server is running (OIMHOST1 by default).
- SOAVHN2 is the virtual host name that is the listen address for the WLS\_SOA2 managed server, and it fails over with server migration of the WLS\_SOA2 managed server. It is enabled on the node where the WLS\_SOA2 managed server is running (OIMHOST2 by default).
- VHN refers to the virtual IP addresses for the Oracle Real Application Clusters (Oracle RAC) database hosts.

### 8.9.2.2 Starting and Stopping the Oracle Identity Manager Cluster

In a high availability architecture, Oracle Identity Manager is deployed on an Oracle WebLogic cluster that has at least two servers as members of the cluster.

By default, Oracle WebLogic Server starts, stops, monitors, and manages the various lifecycle events for the application. The Oracle Identity Manager application leverages the high availability features of the underlying Oracle WebLogic clusters. In case of hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

You can use one or more of the following command line tools and consoles to manage Oracle Identity Manager lifecycle events:

- Oracle WebLogic Server Administration Console
- Oracle Enterprise Manager Fusion Middleware Control
- Oracle WebLogic Scripting Tool (WLST)

### **8.9.2.3 Cluster-Wide Configuration Changes**

For high availability environments, changing the configuration of one Oracle Identity Manager instance changes the configuration of all the other instances, because all the Oracle Identity Manager instances share the same configuration repository.

### **8.9.2.4 Considerations for Synchronizing with LDAP**

Synchronization information between LDAP and the Oracle Identity Manager database is handled by a process called Reconciliation, which is a scheduled process that runs in the background primarily. It is also possible to manually run this process.

If an LDAP outage occurs during the Synchronization process, the data which did not get into Oracle Identity Manager will be picked up during the next run of the reconciliation task.

## **8.9.3 Oracle Identity Manager High Availability Configuration Steps**

This section provides high-level instructions for setting up a high availability deployment for Oracle Identity Manager.

This section includes the following topics about configuring Oracle Identity Management for high availability:

- [Section 8.9.3.1, "Prerequisites for Oracle Identity Manager Configuration"](#)
- [Section 8.9.3.2, "Creating and Configuring the WebLogic Domain for OIM and SOA on OIMHOST1"](#)
- [Section 8.9.3.3, "Post-Installation Steps on OIMHOST1"](#)
- [Section 8.9.3.4, "Configuring Oracle Identity Manager on OIMHOST1"](#)
- [Section 8.9.3.5, "Post-Configuration Steps for the Managed Servers"](#)
- [Section 8.9.3.6, "Validate the Oracle Identity Manager Instance on OIMHOST1"](#)
- [Section 8.9.3.7, "Propagating Oracle Identity Manager to OIMHOST2"](#)
- [Section 8.9.3.8, "Post-Installation Steps on OIMHOST2"](#)
- [Section 8.9.3.9, "Validate the Oracle Identity Manager Instance on OIMHOST2"](#)
- [Section 8.9.3.10, "Configuring Oracle Internet Directory using the LDAP Configuration Post-setup Script"](#)
- [Section 8.9.3.11, "Configuring Server Migration for the OIM and SOA Managed Servers"](#)
- [Section 8.9.3.12, "Configuring a Shared JMS Persistence Store"](#)
- [Section 8.9.3.13, "Configuring a Default Persistence Store for Transaction Recovery"](#)
- [Section 8.9.3.14, "Install Oracle HTTP Server on WEBHOST1 and WEBHOST2"](#)
- [Section 8.9.3.15, "Configuring Oracle Identity Manager to Work with the Web Tier"](#)

- [Section 8.9.3.16, "Validate the Oracle HTTP Server Configuration"](#)
- [Section 8.9.3.17, "Oracle Identity Manager Failover and Expected Behavior"](#)
- [Section 8.9.3.18, "Troubleshooting Oracle Identity Manager High Availability"](#)
- [Section 8.9.3.19, "Scaling Up and Scaling Out the Oracle Identity Manager Topology"](#)

### 8.9.3.1 Prerequisites for Oracle Identity Manager Configuration

Before you configure Oracle Identity Manager for high availability, you must:

- Run the Repository Creation Utility to create the OIM schemas in a database.  
See [Section 8.9.3.1.1, "Running RCU to Create the OIM Schemas in a Database"](#) for instructions on running the Repository Creation Utility to create the OIM schemas.
- Install Oracle WebLogic Server on OIMHOST1 and OIMHOST2  
Follow the steps in [Section 8.9.3.1.2, "Installing Oracle WebLogic Server"](#) to install Oracle WebLogic Server on OIMHOST1 and OIMHOST2.
- Install the Oracle SOA Suite on OIMHOST1 and OIMHOST2.  
Follow the steps in [Section 8.9.3.1.3, "Installing the Oracle SOA Suite on OIMHOST1 and OIMHOST2"](#) to install the Oracle SOA Suite software on OIMHOST1 and OIMHOST2.
- Install the Oracle Identity Management software on OIMHOST1 and OIMHOST2:  
Follow the steps in [Section 8.9.3.1.4, "Installing the Oracle Identity Manager on OIMHOST1 and OIMHOST2"](#) to install the Oracle Identity Management software on OIMHOST1 and OIMHOST2.
- Ensure that a highly available LDAP implementation is available.

---

**Note:** This section is required only for LDAPSvc-enabled Oracle Identity Manager installations and for Oracle Identity Manager installations that integrate with Oracle Access Manager.

If you are not planning to enable the LDAPSvc option or to integrate with Oracle Access Manager, you can skip this section.

---

For information about installing and configuring Oracle Internet Directory, see [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps."](#) For information about installing and configuring Oracle Virtual Directory, see [Section 8.4.3, "Oracle Virtual Directory High Availability Configuration Steps."](#) Note that Oracle Identity Manager does not communicate directly with Oracle Internet Directory. It communicates with Oracle Virtual Directory, which communicates with Oracle Internet Directory.

- Create the `wlfullclient.jar` file on OIMHOST1 and OIMHOST2.  
Oracle Identity Manager uses the `wlfullclient.jar` library for certain operations. Oracle does not ship this library, so you must create this library manually. Oracle recommends creating this library under the `MW_HOME/wlserver_10.3/server/lib` directory on all the machines in the application tier of your environment. You do not need to create this library on directory tier machines such as OIHOST1, OIHOST2, OVDHOST1 and OVDHOST2.

Follow these steps to create the `wlfullclient.jar` file:

1. Navigate to the `MW_HOME/wlserver_10.3/server/lib` directory.
2. Set your `JAVA_HOME` to `MW_HOME/jdk160_18` and ensure that your `JAVA_HOME/bin` directory is in your path.
3. Create the `wlfullclient.jar` file by running:

```
java -jar wljarbuilder.jar
```

**8.9.3.1.1 Running RCU to Create the OIM Schemas in a Database** Before you can install the Oracle Identity Manager and SOA instances on OIMHOST1 and OIMHOST2, you must use the Repository Creation Utility (RCU) to create the collection of schemas used by Oracle Identity Manager.

RCU ships on its own CD as part of the Oracle Fusion Middleware 11g kit.

Follow these steps to run RCU and create the Oracle Identity Manager schemas in an Oracle RAC database repository:

1. Issue this command:

```
prompt> RCU_HOME/bin/rcu &
```

2. On the Welcome screen, click **Next**.
3. On the Create Repository screen, select the **Create** operation to load the component schemas into an existing database.

Click **Next**.

4. On the Database Connection Details screen, enter connection information for the existing database as follows:

- **Database Type:** Oracle Database
- **Host Name:** Name of the computer on which the database is running. For an Oracle RAC database, specify the VIP name or one node name. Example: OIMDBHOST1-VIP or OIMDBHOST2-VIP
- **Port:** The port number for the database. For example: 1521
- **Service Name:** The service name of the database. For example: oim.mycompany.com
- **Username:** SYS
- **Password:** The SYS user password
- **Role:** SYSDBA

Click **Next**.

5. Click **OK** on the Checking Prerequisites screen after the Global Prerequisites complete successfully.
6. On the Select Components screen, create a new prefix and select the components to be associated with this deployment:

**Create a New Prefix:** ha

**Components:**

- Under **Identity Management:**
  - Oracle Identity Manager - OIM

- Note that Metadata Services - MDS is selected by default.
- Under **SOA and BAM Infrastructure**:
  - SOA Infrastructure - SOAINFRA is selected by default.
  - User Messaging Service - ORASDPM is selected by default.

Click **Next**.

7. Click **OK** on the Checking Prerequisites screen after the Component Prerequisites complete successfully.
8. On the Schema Passwords screen, enter the passwords for the mail and additional (auxiliary) schema users.

Click **Next**.

9. On Map Tablespaces screen, select the tablespaces for the components.
10. On the Summary screen, click **Create**.
11. On the Completion Summary screen, click **Close**.

**8.9.3.1.2 Installing Oracle WebLogic Server** Prior to installing the Oracle WebLogic Server, ensure that your machines meet the system, patch, kernel, and other requirements as specified in *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

Start the installer, then perform these steps on OIMHOST1 and OIMHOST2:

1. On the Welcome screen, click **Next**.
2. On the Choose Middleware Home Directory screen, select **Create a New Middleware Home**.

For Middleware Home Directory, enter:

`ORACLE_BASE/product/fmw`

---

**Note:** `ORACLE_BASE` is the base directory under which Oracle products are installed. The recommended value is `/u01/app/oracle`.

---

Click **Next**.

3. On the Register for Security Updates screen, enter your contact information so that you can be notified of security updates.

Click **Next**.

4. On the Choose Install Type screen, select **Custom**.

Click **Next**.

5. On the Choose Products and Components screen, select only **Oracle JRockit SDK**, and click **Next**.

6. On the Choose Product Installation Directories screen, accept the directory `ORACLE_BASE/product/fmw/wlserver_10.3`.

Click **Next**.

7. On the Installation Summary screen, click **Next**.
8. On the Installation Complete screen, deselect **Run Quickstart**.

Click **Done**.

**8.9.3.1.3 Installing the Oracle SOA Suite on OIMHOST1 and OIMHOST2** Perform these steps to install Oracle Fusion Middleware Components on OIMHOST1 and OIMHOST2.

On Linux platforms, if the `/etc/oraInst.loc` file exists, verify that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

Start the installer for Oracle Fusion Middleware Component as follows:

```
HOST1> runInstaller
```

When the installer prompts you for a JRE/JDK location, enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `ORACLE_BASE/product/fmw/jrockit_160_14_<version>`.

Then proceed as follows:

1. On the Welcome screen, click **Next**.
2. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
3. On the Specify Installation Location screen, enter the following values:
  - Oracle Middleware Home: Select the previously installed Middleware home from the list for `MW_HOME`, for example  
`/u01/app/oracle/product/fmw`
  - Oracle Home Directory:
    - Enter `soa` as the Oracle home directory name when installing the Oracle SOA Suite in the `ORACLE_HOME`

Click **Next**.

4. On the Installation Summary screen, click **Install**.

When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the `root` user.

5. On the Installation Complete screen, click **Finish**.

**8.9.3.1.4 Installing the Oracle Identity Manager on OIMHOST1 and OIMHOST2** Perform these steps to install Oracle Fusion Middleware Components on OIMHOST1 and OIMHOST2.

On Linux platforms, if the `/etc/oraInst.loc` file exists, verify that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

Start the installer for Oracle Fusion Middleware Component as follows:

```
HOST1> runInstaller
```

When the installer prompts you for a JRE/JDK location, enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `ORACLE_BASE/product/fmw/jrockit_160_<version>`.

Then proceed as follows:

1. On the Welcome screen, click **Next**.

2. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
3. On the Specify Installation Location screen, enter the following values:
  - Oracle Middleware Home: Select the previously installed Middleware home from the list for MW\_HOME, for example  
`/u01/app/oracle/product/fmw`
  - Oracle Home Directory:
    - Enter `iam` as the Oracle home directory name when installing the Oracle Identity and Access Management Suite in the ORACLE\_HOME
 Click **Next**.
4. On the Installation Summary screen, click **Install**.  
 When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the root user.
5. On the Installation Complete screen, click **Finish**.

### 8.9.3.2 Creating and Configuring the WebLogic Domain for OIM and SOA on OIMHOST1

The domain needs to be created on OIMHOST1. Follow these steps:

1. Start the Configuration Wizard by executing this command:  
`MW_HOME/oracle_common/common/bin/config.sh`
2. On the Welcome screen, select **Create a WebLogic Domain**.  
 Click **Next**.
3. On the Select Domain Source screen, select **Generate a domain configured automatically to support the following added products**.  
 From the list, select:
  - **Oracle Identity Manager**  
 Note that Oracle SOA Suite and Oracle WSM Policy Manager are selected automatically.
  - **Oracle Enterprise Manager**  
 Note that Oracle JRF is selected automatically.
4. On the Specify Domain Name and Location screen, enter the name and location for the domain and all its applications.  
 Provide the following:
  - **Domain Name:** IDMDomain
  - **Domain Location:** Accept the default.
  - **Application Location:** Accept the default.
 Click **Next**.
5. On the Configure Administration Server Username and Password screen, provide the following:
  - **Name:** weblogic

- **User Password:** Enter the password for the `weblogic` user.
- **Confirm User Password:** Enter the password for the `weblogic` user.
- **Description:** Provide a description for the user.

Click **Next**.

6. On the Configure Server Start Mode and JDK screen, select **Production Mode** and **JRockit SDK 1.6.n**.

Click **Next**.

7. On the Configure JDBC Component Schemas screen, select the Component Schemas shown below:

- SOA Infrastructure
- User Messaging Service
- OIM MDS Schema
- OWSM MDS Schema
- SOA MDS Schema
- OIM Schema

Select the check box next to **Configure selected component schemas as RAC multi data source schemas in the next panel**.

Click **Next**.

8. On the Configure RAC Multi Data Source Component Schemas screen, select all the Multi Data Source schemas and enter the following:

- **Service Name:** `oim.mycompany.com`
- For the first RAC node:
  - **Host Name:** `OIMDBHOST1-VIP.mycompany.com`
  - **Instance Name:** `oimdb1`
  - **Port:** 1521
- For the second RAC node:
  - **Host Name:** `OIMDBHOST2-VIP.mycompany.com`
  - **Instance Name:** `oimdb2`
  - **Port:** 1521

Select each schema individually and enter the schema's username and password, as shown in [Table 8-7](#):

**Table 8-7 Entering the Schema Owner and Password for Each multi data Source Schema**

| Schema Name            | Schema Owner | Password             |
|------------------------|--------------|----------------------|
| SOA Infrastructure     | HA_SOAINFRA  | <enter the password> |
| User Messaging Service | HA_ORASDPM   | <enter the password> |
| OIM MDS Schema         | HA_MDS       | <enter the password> |
| OWSM MDS Schema        | HA_MDS       | <enter the password> |
| SOA MDS Schema         | HA_MDS       | <enter the password> |



**Table 8–7 (Cont.) Entering the Schema Owner and Password for Each multi data Source Schema**

| Schema Name | Schema Owner | Password             |
|-------------|--------------|----------------------|
| OIM Schema  | HA_OIM       | <enter the password> |

Click **Next**.

9. On the Test Component Schema screen, select **All the Schemas** and then click **Test Connections**. Validate that the test for all the schemas completed successfully.

Click **Next**.

10. On the Select Optional Configuration screen, select:

- Administration Server
- JMS Distributed Destination
- Managed Server Clusters and Machines

Click **Next**.

11. In the Configure the Administration Server screen, enter the following values:

- **Name:** AdminServer
- **Listen Address:** oimhost1.mycompany.com
- **Listen Port:** 7001
- **SSL listen port:** Not applicable
- **SSL enabled:** Leave unchecked

Click **Next**.

12. On the Select JMS Distributed Destination Type screen, ensure that all JMS System Resources listed are Uniform Distributed Destinations. If they are not, select **UDD** from the drop down box. Validate that the entries look like those in [Table 8–8](#):

**Table 8–8 Values to Choose for JMS System Resources**

| JMS System Resource  | Uniform/Weighted Distributed Destination |
|----------------------|------------------------------------------|
| UMSJMSSystemResource | UDD                                      |
| SOAJMSModule         | UDD                                      |
| OIMJMSModule         | UDD                                      |
| BPMJMSModule         | UDD                                      |

Click **Next**.

An Override Warning box with the following message is displayed:

"CFGFWK-40915: At least one JMS system resource has been selected for conversion to a Uniform Distributed Destination (UDD). This conversion will take place only if the JMS System resource is assigned to a cluster."

Click **OK** on the Warning box.

13. When you first enter the Configure Managed Servers screen, the configuration wizard will have created two default managed servers (oim\_server1 and soa\_server1) for you. Change the details of the default managed servers and then add the second managed server. Follow the steps below:

For the `oim_server1` entry, change the entry to the following values:

- **Name:** `WLS_OIM1`
- **Listen Address:** `OIMVHN1`
- **Listen Port:** 14000

For the `soa_server1` entry, change the entry to the following values:

- **Name:** `WLS_SOA1`
- **Listen Address:** `SOAVHN1`
- **Listen Port:** 8001

For the second OIM Server, click **Add** and supply the following information:

- **Name:** `WLS_OIM2`
- **Listen Address:** `OIMVHN2`
- **Listen Port:** 14000

For the second SOA Server, click **Add** and supply the following information:

- **Name:** `WLS_SOA2`
- **Listen Address:** `SOAVHN2`
- **Listen Port:** 8001

Click **Next**.

---

---

**Note:** Follow the steps for adding the second managed server to add additional managed servers.

---

---

14. On the Configure Clusters screen, create two clusters by clicking **Add**.

Supply the following information for the OIM Cluster:

- **Name:** `oim_cluster`
- **Cluster Messaging Mode:** unicast

Supply the following information for the SOA Cluster:

- **Name:** `soa_cluster`
- **Cluster Messaging Mode:** unicast

Leave all the other fields at the default settings and click **Next**.

15. On the Assign Servers to Clusters screen, associate the managed servers with the cluster. Click on the cluster name in the right window.

Click on the managed server under **servers**, then click on the arrow to assign it to the cluster.

Assign the **WLS\_OIM1** and **WLS\_OIM2** managed servers to be members of the **oim\_cluster**.

Assign the **WLS\_SOA1** and **WLS\_SOA2** managed servers to be members of the **soa\_cluster**.

Click **Next**.

16. On the Configure Machines screen, create a machine for each host in the topology.

Click on the **Unix** tab if your hosts use some type of Unix operating system. Otherwise, click on **Machines**.

Provide the following information:

- **Name:** Name of the host. A good practice is to use the DNS name here.
- **Node Manager Listen Address:** Enter the DNS name of the machine here.
- **Node Manager Port:** Supply a port for Node Manager to use.

Click **Next**.

---

**Note:** On UNIX, delete the default local machine entry under the **Machines** tab.

---

17. On the Assign Servers to Machines screen, you will assign the managed servers that will run on the machines you just created. Follow these steps:

Click on a machine in the right hand window.

Click on the managed servers you want to run on that machine in the left window.

Click on the arrow to assign the managed servers to the machine.

Repeat these steps until all the managed servers are assigned to the appropriate machine.

A typical example would be:

- Host1: **Admin Server, WLS\_OIM1, and WLS\_SOA1**
- Host2: **WLS\_OIM2 and WLS\_SOA2**

Click **Next**.

18. On the Configuration Summary screen, click **Create** to create the domain.

### 8.9.3.3 Post-Installation Steps on OIMHOST1

This section describes the post-installation steps to perform on OIMHOST1. It includes these sections:

- [Section 8.9.3.3.1, "Creating boot.properties for the Administration Server on OIMHOST1"](#)
- [Section 8.9.3.3.2, "Update Node Manager on OIMHOST1"](#)
- [Section 8.9.3.3.3, "Start Node Manager on OIMHOST1"](#)
- [Section 8.9.3.3.4, "Start the Administration Server on OIMHOST1"](#)

**8.9.3.3.1 Creating boot.properties for the Administration Server on OIMHOST1** This section describes how to create a boot.properties file for the Administration Server on OIMHOST1. The boot.properties file enables the Administration Server to start without prompting for the administrator username and password.

Follow these steps to create the boot.properties file:

1. On OIMHOST1, create the following directory:

```
MW_HOME/user_projects/domains/domainName/servers/AdminServer/security
```

For example:

```
$ mkdir -p
```

*MW\_HOME/user\_projects/domains/domainName/servers/AdminServer/security*

2. Use a text editor to create a file called `boot.properties` under the security directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

**8.9.3.3.2 Update Node Manager on OIMHOST1** Before the managed servers can be started via the WebLogic Administration Console, Node Manager requires that the `StartScriptEnabled` property be set to true.

To do this, run the `setNMProps.sh` script located under the following directory:

*MW\_HOME/oracle\_common/common/bin*

**8.9.3.3.3 Start Node Manager on OIMHOST1** Start the Node Manager on OIMHOST1 using the `startNodeManager.sh` script located under the following directory:

*MW\_HOME/wlserver\_10.3/server/bin*

**8.9.3.3.4 Start the Administration Server on OIMHOST1** Follow these steps to start the Administration Server and validate its startup:

1. Start the Administration Server on OIMHOST1 by issuing the command:

*DOMAIN\_HOME/bin/startWebLogic.sh*

2. Validate that the Administration Server started up successfully by opening a web browser and accessing the following pages:

- WebLogic Server Administration Console at:

`http://oimhost1.mycompany.com:7001/console`

- Oracle Enterprise Manager Fusion Middleware Control at:

`http://oimhost1.mycompany.com:7001/em`

Log into these consoles using the `weblogic` user credentials.

### 8.9.3.4 Configuring Oracle Identity Manager on OIMHOST1

This section describes how to configure the Oracle Identity Manager and SOA managed servers before starting them.

This section includes the following topics:

- [Section 8.9.3.4.1, "Prerequisites for Configuring Oracle Identity Manager"](#)
- [Section 8.9.3.4.2, "Running the Oracle Identity Management Configuration Wizard"](#)

**8.9.3.4.1 Prerequisites for Configuring Oracle Identity Manager** Before configuring Oracle Identity Manager, ensure that the following tasks have been performed:

---

**Note:** This section is required only for LDAPSync-enabled Oracle Identity Manager installations and for Oracle Identity Manager installations that integrate with Oracle Access Manager.

If you are not planning to enable the LDAPSync option or to integrate with Oracle Access Manager, you can skip this section.

---

1. ["Extending the Directory Schema for Oracle Identity Manager"](#)
2. ["Creating Users and Groups for Oracle Identity Manager"](#)
3. ["Creating Adapters in Oracle Virtual Directory"](#)

### Extending the Directory Schema for Oracle Identity Manager

Pre-configuring the Identity Store extends the schema in the back end directory regardless of directory type.

To pre-configure the Identity Store, perform these steps on OIMHOST1:

1. Set the environment variables MW\_HOME, JAVA\_HOME and ORACLE\_HOME.

Set ORACLE\_HOME to IAM\_ORACLE\_HOME.

2. Create a properties file `extend.props` that contains the following:

```
IDSTORE_HOST : idstore.mycompany.com
IDSTORE_PORT : 389
IDSTORE_BINDDN: cn=orcladmin
IDSTORE_USERNAMEATTRIBUTE: cn
IDSTORE_LOGINATTRIBUTE: uid
IDSTORE_USERSEARCHBASE: cn=Users,dc=mycompany,dc=com
IDSTORE_GROUPSEARCHBASE: cn=Groups,dc=mycompany,dc=com
IDSTORE_SEARCHBASE: dc=mycompany,dc=com
IDSTORE_SYSTEMIDBASE: cn=systemids,dc=mycompany,dc=com
```

Where:

- IDSTORE\_HOST and IDSTORE\_PORT are, respectively, the host and port of your Identity Store directory. If you are using a non-OID directory, then specify the Oracle Virtual Directory host (which should be IDSTORE.mycompany.com.)
- IDSTORE\_BINDDN is an administrative user in the Identity Store Directory
- IDSTORE\_USERSEARCHBASE is the location in the directory where Users are Stored.
- IDSTORE\_GROUPSEARCHBASE is the location in the directory where Groups are Stored.
- IDSTORE\_SEARCHBASE is the location in the directory where Users and Groups are stored.
- IDSTORE\_SYSTEMIDBASE is the location of a container in the directory where users can be placed when you do not want them in the main user container. This happens rarely but one example is the Oracle Identity Manager reconciliation user which is also used for the bind DN user in Oracle Virtual Directory adapters.

3. Configure the Identity Store using the command `idmConfigTool` which is located at `IAM_ORACLE_HOME/idmtools/bin`.

The command syntax for Linux:

```
idmConfigTool.sh -preConfigIDStore input_file=configfile
```

The command syntax for Windows:

```
idmConfigTool.bat -preConfigIDStore input_file=configfile
```

For example:

```
idmConfigTool.sh -preConfigIDStore input_file=extend.props
```

After the command runs, the system prompts you to enter the password of the account with which you are connecting to the ID Store.

Sample command output:

```
./preconfig_id.sh
Enter ID Store Bind DN password :
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/idm_idstore_groups_
template.ldif
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/idm_idstore_groups_acl_
template.ldif
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/systemid_pwdpolicy.ldif
Apr 5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFileINFO: ->
LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/idstore_tuning.ldifApr
5, 2011 3:39:25 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oid_schema_extn.ldif
The tool has completed its operation. Details have been logged to
automation.log
```

4. Check the log file for any errors or warnings and correct them.

### Creating Users and Groups for Oracle Identity Manager

Follow the steps in the procedure to add the `oimadmin` user to the Identity Store and assign it to an Oracle Identity Manager administrative group. You must also create a user outside of the standard `cn=Users` location to perform reconciliation. Oracle recommends that you select this user as the bind DN when connecting to directories with Oracle Virtual Directory.

---

**Note:** This command also creates a container in your Identity Store for reservations.

---

To add the `xelsysadm` user to the Identity Store and assign it to an administrative group, perform the following tasks on `OIMHOST1`:

1. Set the Environment Variables: `MW_HOME`, `JAVA_HOME`, `IDM_HOME`, and `ORACLE_HOME`

Set `IDM_HOME` to `IDM_ORACLE_HOME`

Set `ORACLE_HOME` to `IAM_ORACLE_HOME`

2. Create a properties file `oim.props` that contains the following:

```
IDSTORE_HOST : idstore.mycompany.com
IDSTORE_PORT : 389
IDSTORE_BINDDN : cn=orcladmin
IDSTORE_USERNAMEATTRIBUTE: cn
IDSTORE_LOGINATTRIBUTE: uid
IDSTORE_USERSEARCHBASE:cn=Users,dc=mycompany,dc=com
IDSTORE_GROUPSEARCHBASE: cn=Groups,dc=us,dc=oracle,dc=com
IDSTORE_SEARCHBASE: dc=mycompany,dc=com
POLICYSTORE_SHARES_IDSTORE: true
IDSTORE_SYSTEMIDBASE: cn=systemids,dc=mycompany,dc=com
IDSTORE_OIMADMINUSER: oimadmin
IDSTORE_OIMADMINGROUP:OIMAdministrators
```

Where:

- `IDSTORE_HOST` and `IDSTORE_PORT` are, respectively, the host and port of your Identity Store directory.
  - `IDSTORE_BINDDN` is an administrative user in the Identity Store Directory
  - `IDSTORE_OIMADMINUSER` is the name of the administration user you would like to use to log in to the Oracle Identity Manager console.
  - `IDSTORE_OIMADMINGROUP` is the name of the group you want to create to hold your Oracle Identity Manager administrative users.
  - `IDSTORE_USERSEARCHBASE` is the location in your Identity Store where users are placed.
  - `IDSTORE_GROUPSEARCHBASE` is the location in your Identity Store where groups are placed.
  - `IDSTORE_SYSTEMIDBASE` is the location in your directory where the Oracle Identity Manager reconciliation user are placed.
  - `POLICYSTORE_SHARES_IDSTORE` is set to true if your Policy and Identity stores are in the same directory. If not, it is set to false.
3. Configure the Identity Store by using the command `idmConfigTool`, located at `IAM_ORACLE_HOME/idmtools/bin`

The Linux command syntax is:

```
idmConfigTool.sh -prepareIDStore mode=OIM input_
file=configfile
```

The Windows command syntax is:

```
idmConfigTool.bat -prepareIDStore mode=OIM input_
file=configfile
```

For example:

```
idmConfigTool.sh -prepareIDStore mode=OIM input_file=oim.props
```

When the command runs, the system prompts you for the account password with which you are connecting to the Identity Store. The system also requests the passwords you want to assign to the accounts:

```
IDSTORE_OIMADMINUSER
oimadmin
```

Oracle recommends that you set `oimadmin` to the same value as the account you create as part of the Oracle Identity Manager configuration.

Sample command output:

```
Enter ID Store Bind DN password :
*** Creation of oimadmin ***
Apr 5, 2011 4:58:51 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oim_user_template.ldif
Enter User Password for oimadmin:
Confirm User Password for oimadmin:
Apr 5, 2011 4:59:01 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oim_group_template.ldif
Apr 5, 2011 4:59:01 AM oracle.ldap.util.LDIFLoader loadOneLdifFileINFO: ->
LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oim_group_member_
template.ldif
Apr 5, 2011 4:59:01 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oim_groups_acl_
template.ldif
Apr 5, 2011 4:59:01 AM oracle.ldap.util.LDIFLoader loadOneLdifFile
INFO: -> LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oim_reserve_
template.ldif
*** Creation of Xel Sys Admin User ***
Apr 5, 2011 4:59:01 AM oracle.ldap.util.LDIFLoader loadOneLdifFileINFO: ->
LOADING:

/u01/app/oracle/product/fmw/IAM/idmtools/templates/oid/oam_user_template.ldif
Enter User Password for xelsysadm:
Confirm User Password for xelsysadm:
The tool has completed its operation. Details have been logged to
/home/oracle/idmtools/oim.log
```

4. Check the log file for errors and warnings and correct them.

### Creating Adapters in Oracle Virtual Directory

Oracle Identity Manager uses Oracle Virtual Directory to connect to external LDAP stores. You must create a user adapter and a change log adapter in Oracle Virtual Directory to enable Oracle Identity Manager to connect to the external LDAP store, such as Oracle Internet Directory. Follow these steps to create the adapters.



---

**Note:** Creating adapters in Oracle Virtual Directory is not required if your implementation is Oracle Internet Directory only.

---

### User Adapter

Create the user adapter on the Oracle Virtual Directory instances running on OVDHOST1 and OVDHOST2 individually.

To create the User Adapter in Oracle Virtual Directory using Oracle Directory Services Manager:

1. Open a browser and bring up the ODSM console at <http://admin.mycompany.com/odsm>.

---

**Note:** Although Oracle Directory Services Manager is not shown in [Figure 8–17](#), it is required to managed Oracle Internet Directory and Oracle Virtual Directory. It is assumed that Oracle Directory Services Manager exists in your environment.

---

2. Create connections to each of the Oracle Virtual Directory instances running on OVDHOST1 and OVDHOST2, if they do not already exist
3. Connect to each Oracle Virtual Directory instance by using the appropriate connection entry.
4. On the Home page, click the **Adapter** tab.
5. Start the New Adapter Wizard by clicking **Create Adapter** at the top of the adapter window.
6. Create a new adapter using the New Adapter Wizard, with the following parameters:

---

**Note:** If you created a User Adapter previously, skip the steps to create the Adapter and follow the steps to edit the Adapter.

---

| Screen     | Field                | Value/Step                                                                                                                                                |
|------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type       | Adapter Type         | LDAP                                                                                                                                                      |
|            | Adapter Name         | User Adapter                                                                                                                                              |
|            | Adapter Template     | User_OID                                                                                                                                                  |
| Connection | Use DNS Setting      | No                                                                                                                                                        |
|            | Host                 | oid.mycompany.com                                                                                                                                         |
|            | Port                 | 389                                                                                                                                                       |
|            | Server Proxy Bind DN | cn=oimadmin,cn=system<br>ids,dc=mycompany,dc=c<br>om                                                                                                      |
|            | Proxy Password       | oimadmin password. This<br>is same as the password<br>provided in " <a href="#">Extending the<br/>Directory Schema for Oracle<br/>Identity Manager</a> ". |

| Screen          | Field                   | Value/Step                                                        |
|-----------------|-------------------------|-------------------------------------------------------------------|
| Connection Test |                         | Validate that the test succeeds.                                  |
| Namespace       | <b>Remote Base</b>      | dc=mycompany,dc=com                                               |
|                 | <b>Mapped Namespace</b> | dc=mycompany,dc=com                                               |
| Summary         |                         | Verify that the summary is correct and then click <b>Finish</b> . |

7. Edit the User Adapter as follows:

- a. Select the OIM User Adapter.
- b. Click the **Plug-ins** Tab.
- c. Click the **User Management** Plug-in, then click **Edit** in the plug-ins table. The plug-in editing window appears.
- d. In the Parameters table, update the parameter values as follows:

| Parameter     | value |
|---------------|-------|
| directoryType | oid   |
| pwdMaxFailure | 10    |
| oamEnabled    | true  |

- e. Click **OK**.
- f. Click **Apply**.

### Change Log Adapter

Create the change log adapter on the Oracle Virtual Directory instances running on OVDHOST1 and OVDHOST2 individually. Follow these steps to create the Change Log Adapter in Oracle Virtual Directory using Oracle Directory Services Manager.

- Open a browser and bring up the ODSM console at <http://admin.mycompany.com/odsm>.
- Create connections to each of the Oracle Virtual Directory instances running on OVDHOST1 and OVDHOST2, if they do not already exist.
- Connect to an Oracle Virtual Directory instance by using the appropriate connection entry.
- On the Home page click on the **Adapter** tab.
- Start the New Adapter Wizard by clicking **Create Adapter** at the top of the adapter window.
- Create a new adapter using the New Adapter Wizard, with the following parameters:

| Screen | Field               | Value/Step             |
|--------|---------------------|------------------------|
| Type   | <b>Adapter Type</b> | LDAP                   |
|        | <b>Adapter Name</b> | OIM Change Log Adapter |

| Screen           | Field                       | Value/Step                                                                                                                              |
|------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Connection       | <b>Adapter Template</b>     | Changelog_OID                                                                                                                           |
|                  | <b>Use DNS Setting</b>      | No                                                                                                                                      |
|                  | <b>Host</b>                 | oid.mycompany.com                                                                                                                       |
|                  | <b>Port</b>                 | 389                                                                                                                                     |
|                  | <b>Server Proxy Bind DN</b> | cn=oimAdminUser,cn=systemids,dc=mycompany,dc=com                                                                                        |
|                  | <b>Proxy Password</b>       | oimadmin password. This is the same password provided in <a href="#">"Extending the Directory Schema for Oracle Identity Manager"</a> . |
| Connection Test  |                             | Validate that the test succeeds.                                                                                                        |
| Naming Space     | <b>Remote Base</b>          | cn=changelog                                                                                                                            |
| Mapped Namespace |                             | cn=changelog                                                                                                                            |
| Summary          |                             | Verify that the summary is correct then click <b>Finish</b> .                                                                           |

7. To edit the change adapter, follow these steps.

- a. Select the OIM Change Log Adapter.
- b. Click the **Plug-ins** tab.
- c. In the Deployed Plus-ins table, click the **changelog** plug-in, then click **Edit** in the plug-ins table. The plug-in editing window appears.
- d. In the Parameters table, update the parameter values.
- e. Click **OK**.
- f. Click **Apply**.

Edit the Change Log Adapter to either add or modify the properties so that they match the values shown in the following table. You must add the `mapObjectclass`, `modifierDNFilter`, `sizeLimit`, and `targetDNFilter` proprieties to the adapter.

| Parameter                | Value                                                    |
|--------------------------|----------------------------------------------------------|
| <b>directoryType</b>     | oid                                                      |
| <b>mapAttribute</b>      | targetGUID=orclguid                                      |
| <b>requiredAttribute</b> | orclguid                                                 |
| <b>modifierDNFilter</b>  | cn=oimadmin,cn=users,cn=OIM,cn=Products,cn=OracleContext |
| <b>sizeLimit</b>         | 1000                                                     |

| Parameter      | Value                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| targetDNFilter | dc=mycompany, dc=com<br><br>Search based from which reconciliation needs to happen. This value must be the same as the LDAP SearchDN that is specified during OIM installation. |
| mapUserState   | true                                                                                                                                                                            |
| oamEnabled     | true                                                                                                                                                                            |

## Stopping and Starting Oracle Internet Directory and Oracle Virtual Directory

Stop and Start:

- The Oracle Virtual Directory instances running on both OVDHOST1 and OVDHOST2.
- The Oracle Internet Directory instances running on both OIDHOST1 and OIDHOST2.

as described in [Section 8.14, "Starting and Stopping Oracle Identity Management Components."](#)

**8.9.3.4.2 Running the Oracle Identity Management Configuration Wizard** You must configure the Oracle Identity Manager server instances before you can start the Oracle Identity Manager and SOA managed servers. These configuration steps need to be performed only once, for example, during the initial creation of the domain. The Oracle Identity Management Configuration Wizard loads the OIM metadata into the database and configures the instance.

Before proceeding, ensure that the following are true:

- The administration server is up and running.
- The environment variables DOMAIN\_HOME and WL\_HOME are *not* set in the current shell.

The Oracle Identity Management Configuration Wizard is located under the Identity Management Oracle home. Type:

```
IAM_ORACLE_HOME/bin/config.sh
```

Proceed as follows:

1. On the Welcome screen, click **Next**
2. On the Components to Configure screen, select **OIM Server**. Select **OIM Remote Manager**, if required in your topology.  
Click **Next**.
3. On the Database screen, provide the following values:
  - **Connect String:** The connect string for the OIM database. For example:  
oimdbhost1-vip.mycompany.com:1521:oimdb1^oimdbhost2-vip.mycompany.com:1521:oimdb2@oim.mycompany.com
  - **OIM Schema User Name:** HA\_OIM
  - **OIM Schema password:** *password*
  - **MDS Schema User Name:** HA\_MDS

- **MDS Schema Password:** *password*

Select **Next**.

4. On the WebLogic Administration Server screen, provide the following details for the WebLogic Admin Server:

- **URL:** The URL to connect to the WebLogic Administration Server. For example: `t3://oimhost1.mycompany.com:7001`
- **UserName:** `weblogic`
- **Password:** Password for the `weblogic` user

Click **Next**.

5. On the OIM Server screen, provide the following values:

- **OIM Administrator Password:** Password for the OIM Administrator. This is the password for the `xelsysadm` user, the same password you entered earlier for `idmconfigtool`.
- **Confirm Password:** Confirm the password.
- **OIM HTTP URL:** Proxy URL for the OIM Server. This is the URL for the Hardware load balancer that is front ending the OHS servers for OIM. For example: `http://oiminternal.mycompany.com:80`.
- **Key Store Password:** Key store password. The password must have an uppercase letter and a number. For example: `MyPassword1`

Click **Next**.

6. On the LDAP Sync and OAM screen, select **Configure BI Publisher** and provide the **BI Publisher URL**, if required in your environment. Enter the URL to connect to the BI Publisher in your environment.

Select **Enable LDAP Sync**.

---



---

**Notes:**

- Do not select **Enable Identity Administration Integration with OAM**.
  - BI Publisher is not a part of the *IDMDomain*. For information about BI Publisher high availability, see [Section 15.2, "High Availability for Oracle Business Intelligence Publisher."](#)
- 
- 

Click **Next**.

7. On the LDAP Server screen, provide the following LDAP server details:

---



---

**Note:** This release of Oracle Identity Manager LDAPSvc module supports connecting to Oracle Internet Directory only through the Oracle Virtual Directory.

The LDAP server that this step and the next step refer to is an Oracle Virtual Directory.

---



---

- **Directory Server Type:** The directory server type. Select `OID`, `ACTIVE_DIRECTORY`, `IPLANET`, or `OVD`. The default is `OID`.

- **Directory Server ID:** The directory server ID.
- **Server URL:** The URL to access the LDAP server. For example:  
ldap://ovd.mycompany.com:389 if you use the Oracle Virtual Directory Server, ldap://oid.mycompany.com:389 if you use Oracle Internet Directory.
- **Server User:** The username to connect to the server. For example:  
cn=orcladmin.
- **Server Password:** The password to connect to the LDAP server.
- **Server SearchDN:** The Search DN. For example: dc=mycompany, dc=com.

Click **Next**.

8. On the LDAP Server Continued screen, enter the following LDAP server details:

- **LDAP Role Container:** The DN for the Role Container. This is the container where the OIM roles are stored. For example:  
cn=Groups, dc=mycompany, dc=com.
- **LDAP User Container:** The DN for the User Container. This is the container where the OIM users are stored. For example:  
cn=Users, dc=mycompany, dc=com.
- **User Reservation Container:** The DN for the User Reservation Container.

---

**Note:** Use the same container DN Values that `idmconfigtool` creates during the procedure "Creating Users and Groups for Oracle Identity Manager."

---

Click **Next**.

9. On the Remote Manager screen, provide the following values:

---

**Note:** This screen appears only if you selected the Remote Manager utility in step 2.

---

- **Service Name:** HA\_RManager
  - **RMI Registry Port:** 12345
  - **Listen Port (SSL):** 12346
10. On the Configuration Summary screen, verify the summary information.  
Click **Configure** to configure the Oracle Identity Manager instance.
  11. On the Configuration Progress screen, once the configuration completes successfully, click **Next**.
  12. On the Configuration Complete screen, view the details of the Oracle Identity Manager Instance configured.  
Click **Finish** to exit the Configuration Assistant.
  13. Stop the WebLogic Administration Server, as described in [Section 8.14, "Starting and Stopping Oracle Identity Management Components."](#)
  14. Start the WebLogic Administration Server, as described in [Section 8.14, "Starting and Stopping Oracle Identity Management Components."](#)

### 8.9.3.5 Post-Configuration Steps for the Managed Servers

This section describes post-configuration steps for the managed servers:

- [Section 8.9.3.5.1, "Start the WLS\\_SOA1 and WLS\\_OIM1 Managed Servers on OIMHOST1"](#)
- [Section 8.9.3.5.2, "Updating the Coherence Configuration for the Coherence Cluster"](#)

**8.9.3.5.1 Start the WLS\_SOA1 and WLS\_OIM1 Managed Servers on OIMHOST1** Follow these steps to start the WLS\_SOA1 and WLS\_OIM1 managed servers on OIMHOST1:

1. Stop the WebLogic Administration Server on OIMHOST1. Use the WebLogic Administration Console to stop the Administration Server.
2. Start the WebLogic Administration Server on OIMHOST1 using the `startWebLogic.sh` script under the `$DOMAIN_HOME/bin` directory. For example:

```
/u01/app/oracle/admin/OIM/bin/startWebLogic.sh > /tmp/admin.out 2>1&
```

3. Validate that the WebLogic Administration Server started up successfully by bringing up the WebLogic Administration Console.
4. Start the WLS\_SOA1 managed server using the WebLogic Administration Console.
5. Start the WLS\_OIM1 managed server using the WebLogic Administration Console.

**8.9.3.5.2 Updating the Coherence Configuration for the Coherence Cluster** Follow the steps below to update the Coherence configuration for the SOA managed servers:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node.
3. Click **Servers**. The Summary of Servers page appears.
4. Click the name of the server (represented as a hyperlink) in the **Name** column of the table. The settings page for the selected server appears.
5. Click the Server Start tab.
6. Enter the following for WLS\_SOA1 and WLS\_SOA2 into the **Arguments** field.

For WLS\_SOA1, enter the following (on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1 -Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost1vhn1
```

For WLS\_SOA2, enter the following (on a single line, without a carriage return):

```
-Dtangosol.coherence.wka1=soahost1vhn1 -Dtangosol.coherence.wka2=soahost2vhn1
-Dtangosol.coherence.localhost=soahost2vhn1
```

7. Click **Save** and activate the changes.
8. This change requires the SOA servers to be restarted.

### 8.9.3.6 Validate the Oracle Identity Manager Instance on OIMHOST1

Validate the Oracle Identity Manager Server instance on OIMHOST1 by bringing up the Oracle Identity Manager Console using a web browser.

The URL for the Oracle Identity Manager Console is:

`http://oimvhn1.mycompany.com:14000/oim`

Log in using the `xelsysadm` password.

### 8.9.3.7 Propagating Oracle Identity Manager to OIMHOST2

Once the configuration has succeeded on OIMHOST1, the configuration can be propagated to OIMHOST2. This is done by packing the domain on OIMHOST1 and then unpacking it on OIMHOST2.

Follow these steps to pack the domain on OIMHOST1 and unpack it on OIMHOST2:

1. On OIMHOST1, invoke the pack utility in the `MW_HOME/oracle_common/common/bin` directory:

```
pack.sh -domain=MW_HOME/user_projects/domains/OIM_Domain -
template=/u01/app/oracle/admin/templates/oim_domain.jar -
template_name="OIM Domain" -managed=true
```

2. The previous step created the `oim_domain.jar` file in the following directory:

```
/u01/app/oracle/admin/templates
```

Copy the `oim_domain.jar` file from OIMHOST1 to a temporary directory on OIMHOST2.

3. On OIMHOST2, invoke the unpack utility in the `MW_HOME/oracle_common/common/bin` directory and specify the location of the `oim_domain.jar` file in its temporary directory:

```
unpack.sh -domain=MW_HOME/user_projects/domains/OIM_Domain -
template=tmp/oim_domain.jar
```

### 8.9.3.8 Post-Installation Steps on OIMHOST2

This section describes the post-installation steps to perform on OIMHOST2. It includes these sections:

- [Section 8.9.3.8.1, "Update Node Manager on OIMHOST2"](#)
- [Section 8.9.3.8.2, "Start Node Manager on OIMHOST2"](#)
- [Section 8.9.3.8.3, "Start the WLS\\_SOA2 and WLS\\_OIM2 Managed Servers on OIMHOST2"](#)

**8.9.3.8.1 Update Node Manager on OIMHOST2** Before managed servers can be started via the WebLogic Administration Console, Node Manager requires that the `StartScriptEnabled` property be set to true.

To do this, run the `setNMProps.sh` script located under the following directory:

```
MW_HOME/oracle_common/common/bin
```

**8.9.3.8.2 Start Node Manager on OIMHOST2** Start the Node Manager on OIMHOST2 using the `startNodeManager.sh` script located under the following directory:

```
MW_HOME/wlserver_10.3/server/bin
```



**8.9.3.8.3 Start the WLS\_SOA2 and WLS\_OIM2 Managed Servers on OIMHOST2** Follow these steps to start the WLS\_SOA2 and WLS\_OIM2 managed servers on OIMHOST2:

1. Validate that the WebLogic Administration Server started up successfully by bringing up the WebLogic Administration Console.
2. Start the WLS\_SOA2 managed server using the WebLogic Administration Console.
3. Start the WLS\_OIM2 managed server using the WebLogic Administration Console. The WLS\_OIM2 managed server must be started after the WLS\_SOA2 managed server is started.

### 8.9.3.9 Validate the Oracle Identity Manager Instance on OIMHOST2

Validate the Oracle Identity Manager Server instance on OIMHOST2 by bringing up the Oracle Identity Manager Console using a web browser.

The URL for the Oracle Identity Manager Console is:

`http://oimvhn2.mycompany.com:14000/oim`

Log in using the `xelsysadm` password.

### 8.9.3.10 Configuring Oracle Internet Directory using the LDAP Configuration Post-setup Script

---

**Note:** This section is required only for LDAPSync-enabled Oracle Identity Manager installations and for Oracle Identity Manager installations that integrate with Oracle Access Manager.

If you are not planning to enable the LDAP-Sync option or to integrate with Oracle Access Manager, you can skip this section.

---

The Oracle Identity Manager LDAP configuration post-setup script updates the Oracle Identity Manager LDAP Sync scheduled jobs with the last change number from Oracle Internet Directory. The LDAP configuration post-setup script is located under the `IAM_ORACLE_HOME/server/ldap_config_util` directory. To run the script, follow these steps:

1. Ensure that the `wlfullclient.jar` file is created as described in [Section 8.9.3.1, "Prerequisites for Oracle Identity Manager Configuration."](#)
2. Edit the `ldapconfig.props` file located under the `IAM_ORACLE_HOME/server/ldap_config_util` directory and provide the following values:
  - **OIMProviderURL:**  
`t3://oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000`
  - **OIDURL:** `ldap://oid.mycompany.com:389`
  - **OIDAdminUsername:** `cn=orcladmin`
  - **OIDSearchBase:** `dc=mycompany,dc=com`
  - **UserContainerName:** `cn=Users`
  - **RoleContainerName:** `cn=Groups`
  - **ReservationContainerName:** `cn=Reserve`

3. Save the file.
4. Set the WL\_HOME and JAVA\_HOME environment variables.
5. Run LDAPConfigPostSetup.sh. The script prompts for the OID Admin Password and the OIM Admin Password. For example:

```
Prompt> ./LDAPConfigPostSetup.sh
[Enter OID admin password:]
[Enter OIM admin password:]
```

where `OID admin password` is the password for the Administrative account for managing Oracle Internet Directory and `OIM admin password` is the password for the Administrative LDAP account to be created in the LDAP repository for use by Oracle Identity Manager for LDAP-Sync required actions.

---

**Note:** `xelsysadm` is the top-level administrator for OIM; it is the OIM admin password.

---

### 8.9.3.11 Configuring Server Migration for the OIM and SOA Managed Servers

For this high availability topology, Oracle recommends that you configure server migration for the WLS\_OIM1, WLS\_SOA1, WLS\_OIM2, and WLS\_SOA2 managed servers.

- The WLS\_OIM1 and WLS\_SOA1 managed servers on OIMHOST1 are configured to restart automatically on OIMHOST2 if a failure occurs on OIMHOST1.
- The WLS\_OIM2 and WLS\_SOA2 managed servers on OIMHOST2 are configured to restart automatically on OIMHOST1 if a failure occur on OIMHOST2.

In this configuration, the WLS\_OIM1, WLS\_SOA1, WLS\_OIM2 and WLS\_SOA2 servers listen on specific floating IPs that WebLogic Server Migration fails over.

The following steps enable server migration for the WLS\_OIM1, WLS\_SOA1, WLS\_OIM2, and WLS\_SOA2 managed servers, which in turn enables a managed server to fail over to another node if a server or process failure occurs:

- Step 1: [Setting Up a User and Tablespace for the Server Migration Leasing Table](#)
- Step 2: [Creating a Multi Data Source Using the Oracle WebLogic Administration Console](#)
- Step 3: [Editing Node Manager's Properties File](#)
- Step 4: [Setting Environment and Superuser Privileges for the wlsifconfig.sh Script](#)
- Step 5: [Configuring Server Migration Targets](#)
- Step 6: [Testing the Server Migration](#)

**8.9.3.11.1 Setting Up a User and Tablespace for the Server Migration Leasing Table** The first step to set up a user and tablespace for the server migration leasing table:

---

**Note:** If other servers in the same domain are already configured with server migration, you can use the same tablespace and data sources. In this case, you don't need to recreate the data sources and multi data source for database leasing, however, you must retarget them to the clusters you're configuring for server migration.

---

1. Create a tablespace named `leasing`. For example, log on to SQL\*Plus as the `sysdba` user and run the following command:

```
SQL> create tablespace leasing logging datafile 'DB_
HOME/oradata/orcl/leasing.dbf' size 32m autoextend on next 32m maxsize 2048m
extent management local;
```

2. Create a user named '`leasing`' and assign to it the `leasing` tablespace:

```
SQL> create user leasing identified by welcome1;
SQL> grant create table to leasing;
SQL> grant create session to leasing;
SQL> alter user leasing default tablespace leasing;
SQL> alter user leasing quota unlimited on LEASING;
```

3. Create the `leasing` table using the `leasing.ddl` script:

- a. Copy the `leasing.ddl` file located in either the `WL_HOME/server/db/oracle/817` or the `WL_HOME/server/db/oracle/920` directory to your database node.
- b. Connect to the database as the **leasing** user.
- c. Run the `leasing.ddl` script in SQL\*Plus:

```
SQL> @Copy_Location/leasing.ddl;
```

**8.9.3.11.2 Creating a Multi Data Source Using the Oracle WebLogic Administration Console** You create a data source to each of the Oracle RAC database instances during the process of setting up the multi data source, both for these data sources and the global leasing multi data source. When you create a data source:

- Ensure that this is a non-XA data source.
- The names of the multi data sources are in the format of `<MultiDS>-rac0`, `<MultiDS>-rac1`, and so on.
- Use Oracle's Driver (Thin) Version 9.0.1, 9.2.0, 10, 11.
- Data sources do not require support for global transactions. Therefore, do *not* use any type of distributed transaction emulation/participation algorithm for the data source (do not choose the **Supports Global Transactions** option, or the **Logging Last Resource, Emulate Two-Phase Commit**, or **One-Phase Commit** options of the **Supports Global Transactions** option), and specify a service name for your database.
- Target these data sources to the cluster(s).
- Make sure the data source's connection pool initial capacity is set to 0 (zero). To do this, select **Services**, **JDBC** and then **Datasources**. In the Datasources screen, click the **Datasource Name**, click the **Connection Pool** tab, and enter 0 (zero) in the **Initial Capacity** field.

### Creating a Multi Data Source

To create a multi data source:

1. Log into the Oracle WebLogic Server Administration Console using the Admin credentials.
2. In the Domain Structure window, expand the **Services** node then expand the **DataSource** node.
3. Click **Lock and Edit** in the Change Center.

4. Click **New** then click **Multi Data Sources**.
5. Enter `leasing` as the name.
6. Enter `jdbc/leasing` as the JNDI name.
7. Select **Failover as algorithm** (default). Click **Next**.
8. Select the target cluster(s). Click **Next**.
9. Select **non-XA driver** (the default). Click **Next**.
10. Click **Create New Data Source**.
11. Enter `leasing-rac0` as the name. Enter `jdbc/leasing-rac0` as the JNDI name. Enter `oracle` as the database type. For the driver type, select Oracle Driver (Thin) for RAC server-Instance connection Version 10,11.

---

**Note:** When you create multi data sources for the leasing table, enter names in the format `<MultiDS>-rac0`, `<MultiDS>-rac1`, and so on.

---

12. Click **Next**.
13. Deselect **Supports Global Transactions**. Click **Next**.
14. Enter the service name, database name (the RAC Node instance name such as `racdb1`, `racdb2`), host port, and password for your leasing schema. Click **Next**.
15. Click **Test Configuration** and verify that the connection works. Click **Next**.

---

**Note:** Do not click **Finish** until you are done creating all of the data sources that you want to make. If you do not see the data sources that you created, this indicates that you clicked **Finish** too soon. When you click **Finish**, the Administration Console does not present the screen in which you target the data sources.

---

16. Target the data source to the cluster(s).
17. Select the data source and add it to the right screen.
18. Click **Create a New Data Source** for the second instance of your Oracle RAC database, target it to the cluster(s), repeating the steps for the second instance of your Oracle RAC database.
19. Add the second data source to your multi data source.
20. Save and click **Activate Changes**.

**8.9.3.11.3 Editing Node Manager's Properties File** You must edit the `nodemanager.properties` file to add the following properties for each node where you configure server migration:

```
Interface=eth0
NetMask=255.255.255.0
UseMACBroadcast=true
```

- **Interface:** Specifies the interface name for the floating IP (such as `eth0`).

---

**Note:** Do not specify the sub interface, such as `eth0:1` or `eth0:2`. This interface is to be used without the `:0`, or `:1`. The Node Manager's scripts traverse the different `:X` enabled IPs to determine which to add or remove. For example, the valid values in Linux environments are `eth0`, `eth1`, or, `eth2`, `eth3`, `ethn`, depending on the number of interfaces configured.

---



---

**Note:** In Windows, the Interface should be set to the Network Interface Name. For example: `Interface="Local Area Connection"`.

---

- **NetMask:** Net mask for the interface for the floating IP. The net mask should be the same as the net mask on the interface; `255.255.255.0` is an example. The actual value depends on your network.
- **UseMACBroadcast:** Specifies whether or not to use a node's MAC address when sending ARP packets, that is, whether or not to use the `-b` flag in the `arping` command.

Verify in Node Manager's output (shell where Node Manager starts) that these properties are being used or problems may arise during migration. You should see an entry similar to the following in Node Manager's output:

```
...
StateCheckInterval=500
Interface=eth0
NetMask=255.255.255.0
...
```

---

**Note:** The following steps are not required if the server properties (start properties) are properly set and Node Manager can start the servers remotely.

---

1. Set the following property in the `nodemanager.properties` file:
  - **StartScriptEnabled:** Set this property to `true` to enable Node Manager to start the managed servers.
2. Start Node Manager on `OIMHOST1` and `OIMHOST2` by running the `startNodeManager.sh` script in the `WL_HOME/server/bin` directory.

---

**Note:** When running Node Manager from a shared storage installation, multiple nodes start using the same `nodemanager.properties` file. However, each node may require different `NetMask` or `Interface` properties. In this case, specify individual parameters on a per-node basis using environment variables. For example, to use a different interface (`eth3`) in `HOSTn`, use the `Interface` environment variable as follows: `HOSTn> export JAVA_OPTIONS=-DInterface=eth3` and start Node Manager after the variable is set in the shell.

---

**8.9.3.11.4 Setting Environment and Superuser Privileges for the wlsifconfig.sh Script** To set environment and superuser privileges for the `wlsifconfig.sh` script for each node where you configure server migration:

---

**Note:** On Windows, the script is `wlsifconfig.cmd` and users with the Administrator privilege can run it.

---

1. Ensure that your PATH environment variable includes these files:

**Table 8–9 Files Required for the PATH Environment Variable**

| File                             | Located in this directory                     |
|----------------------------------|-----------------------------------------------|
| <code>wlsifconfig.sh</code>      | <code>DOMAIN_HOME/bin/server_migration</code> |
| <code>wlscontrol.sh</code>       | <code>WL_HOME/common/bin</code>               |
| <code>nodemanager.domains</code> | <code>WL_HOME/common</code>                   |

2. Grant sudo configuration for the `wlsifconfig.sh` script.
  - Configure sudo to work without a password prompt.
  - For security reasons, Oracle recommends restricting to the subset of commands required to run the `wlsifconfig.sh` script. For example, perform the following steps to set the environment and superuser privileges for the `wlsifconfig.sh` script:
  - Grant sudo privilege to the WebLogic user (`oracle`) with no password restriction, and grant execute privilege on the `/sbin/ifconfig` and `/sbin/arping` binaries.
  - Ensure that the script is executable by the WebLogic user. The following is an example of an entry inside `/etc/sudoers` granting sudo execution privilege for `oracle` and also over `ifconfig` and `arping`:

```
oracle ALL=NOPASSWD: /sbin/ifconfig,/sbin/arping
```

---

**Note:** Ask the system administrator for the sudo and system rights as appropriate to this step.

---

**8.9.3.11.5 Configuring Server Migration Targets** You first assign all the available nodes for the cluster's members and then specify candidate machines (in order of preference) for each server that is configured with server migration. To configure cluster migration in a migration in a cluster:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand **Environment** and select **Clusters**.
3. Click the cluster you want to configure migration for in the Name column.
4. Click the **Migration** tab.
5. Click **Lock and Edit**.
6. In the **Available** field, select the machine to which to allow migration and click the right arrow.

7. Select the data source to use for automatic migration. In this case, select the leasing data source.
8. Click **Save**.
9. Click **Activate Changes**.
10. Set the candidate machines for server migration. You must perform this task for all managed servers as follows:
  - a. In the Domain Structure window of the Oracle WebLogic Server Administration Console, expand **Environment** and select **Servers**.
 

**Tip:** Click **Customize this table** in the Summary of Servers page and move Current Machine from the Available window to the Chosen window to view the machine that the server runs on. This will be different from the configuration if the server migrates automatically.
  - b. Select the server that you want to configure migration for.
  - c. Click the **Migration** tab.
  - d. In the **Available** field, located in the Migration Configuration section, select the machines you want to allow migration to and click the right arrow.
  - e. Select **Automatic Server Migration Enabled**. This enables Node Manager to start a failed server on the target node automatically.
  - f. Click **Save** then Click **Activate Changes**.
  - g. Repeat the steps above for any additional managed servers.
  - h. Restart the administration server, Node Managers, and the servers for which server migration has been configured.

**8.9.3.11.6 Testing the Server Migration** To verify that server migration works properly:

**From OIMHOST1:**

1. Stop the WLS\_OIM1 managed server. To do this, run this command:

```
OIMHOST1> kill -9 pid
```

where *pid* specifies the process ID of the managed server. You can identify the pid in the node by running this command:

```
OIMHOST1> ps -ef | grep WLS_OIM1
```

2. Watch the Node Manager console. You should see a message indicating that WLS\_OIM1's floating IP has been disabled.
3. Wait for Node Manager to try a second restart of WLS\_OIM1. It waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Node Manager should now log a message indicating that the server will not be restarted again locally.

**From OIMHOST2:**

1. Watch the local Node Manager console. After 30 seconds since the last try to restart WLS\_OIM1 on OIMHOST1, Node Manager on OIMHOST2 should prompt that the floating IP for WLS\_OIM1 is being brought up and that the server is being restarted in this node.

2. Access the soa-infra console in the same IP.

Follow the steps above to test server migration for the WLS\_OIM2, WLS\_SOA1, and WLS\_SOA2 managed servers.

Table 8–10 shows the managed servers and the hosts they migrate to in case of a failure.

**Table 8–10 WLS\_OIM1, WLS\_OIM2, WLS\_SOA1, WLS\_SOA2 Server Migration**

| Managed Server | Migrated From | Migrated To |
|----------------|---------------|-------------|
| WLS_OIM1       | OIMHOST1      | OIMHOST2    |
| WLS_OIM2       | OIMHOST2      | OIMHOST1    |
| WLS_SOA1       | OIMHOST1      | OIMHOST2    |
| WLS_SOA2       | OIMHOST2      | OIMHOST1    |

### Verification From the Administration Console

Migration can also be verified in the Administration Console:

1. Log into the Oracle WebLogic Server Administration Console at <http://oimhost1.mycompany.com:7001/console> using the Admin credentials.
2. Click **Domain** on the left console.
3. Click the **Monitoring** tab and then the **Migration** sub tab.

The Migration Status table provides information on the status of the migration.

---

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

---

### 8.9.3.12 Configuring a Shared JMS Persistence Store

Configure the location for all of the persistence stores as a shared directory that is visible from both OIMHOST1 and OIMHOST2. As JMS messages are persisted in the file system for each server's local file system, shared storage is necessary for the JMS persistence store for WebLogic server migration. Without shared storage, a migrated server cannot access pending JMS messages.

---

**Note:** See [Considerations for Using File Stores on NFS](#) for important information about JMS messages and transaction logs, as well as releasing locks on file stores.

---

To change all persistent stores to use a shared base directory:

1. Log into the Oracle WebLogic Server Administration Console at <http://oimhost1.mycompany.com:7001/console> using the Admin credentials.
2. In the Domain Structure window, expand the **Services** node and then click the **Persistence Stores** node. The Summary of Persistence Stores page is displayed.
3. Select the hyperlink(s) for the persistence store from the **Name** column of the table: BPM, SOA, OIM, and UMS. The Settings page for the persistence store opens.



4. On the Configuration tab, in the **Directory** field, enter the location of a persistent storage solution (such as NAS or SAN) that is available to other servers in the cluster. Specifying this location enables pending JMS messages to be sent.
5. The location should follow this directory structure:

- For the WLS\_SOA1 and WLS\_SOA2 servers, use a directory structure similar to:

`ORACLE_BASE/admin/domainName/soaClusterName/jms`

- For the WLS\_OIM1 and WLS\_OIM2 servers, use a directory structure similar to:

`ORACLE_BASE/admin/domainName/oimClusterName/jms`

---

**Note:** The WLS\_OIM1 and WLS\_OIM2 servers must be able to access this directory.

The WLS\_SOA1 and WLS\_SOA2 servers must be able to access this directory.

This directory must exist *before* you restart the server.

---

6. Click **Save and Activate**.
7. Restart the servers to make the change in the persistent stores take effect.

### 8.9.3.13 Configuring a Default Persistence Store for Transaction Recovery

Each Oracle WebLogic Managed Server has a transaction log that stores information about inflight transactions that are coordinated by the server that may not have been completed. The WebLogic Server uses this transaction log for recovery from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to all the servers in the cluster. Without shared storage, other servers in the cluster cannot do a transaction recovery in the case of a server failure, so the operation may need to be retried.

---

**Note:** Preferably, this location should be a dual-ported SCSI disk or on a Storage Area Network (SAN).

---

Perform these steps to set the location for the default persistence stores for the Oracle Identity Manager and SOA Servers:

1. Log into the Oracle WebLogic Server Administration Console at `http://oimhost1.mycompany.com:7001/console` using the Admin credentials.
2. In the Domain Structure window, expand the **Environment** node and then click the **Servers** node. The Summary of Servers page is displayed.
3. Select the name of the server (represented as a hyperlink) in the **Name** column of the table. The Settings page for the selected server is displayed, and it defaults to the Configuration tab.
4. Click the Services tab.

5. In the Default Store section of the page, enter the path to the folder where the default persistent stores will store their data files. The directory structure of the path should be:

- For the WLS\_SOA1 and WLS\_SOA2 servers, use a directory structure similar to:

*ORACLE\_BASE/admin/domainName/soaClusterName/tlogs*

- For the WLS\_OIM1 and WLS\_OIM2 servers, use a directory structure similar to:

*ORACLE\_BASE/admin/domainName/oimClusterName/tlogs*

6. Click **Save**.

---

**Note:** To enable migration of the Transaction Recovery Service, specify a location on a persistent storage solution that is available to the managed servers in the cluster. WLS\_SOA1, WLS\_SOA2, WLS\_OIM1, and WLS\_OIM2 must be able to access this directory.

---

#### 8.9.3.14 Install Oracle HTTP Server on WEBHOST1 and WEBHOST2

For instructions on installing Oracle HTTP Server on WEBHOST1 and WEBHOST2, see [Section 8.5.3.5.1, "Installing Oracle HTTP Server for the Web Tier."](#)

#### 8.9.3.15 Configuring Oracle Identity Manager to Work with the Web Tier

This section describes how to configure Oracle Identity Manager to work with the Oracle Web Tier.

##### 8.9.3.15.1 Prerequisites

Verify that the following tasks have been performed:

1. Oracle Web Tier has been installed on WEBHOST1 and WEBHOST2.
2. Oracle Identity Manager has been installed and configured on OIMHOST1 and OIMHOST2.
3. The load balancer has been configured with a virtual hostname (*sso.mycompany.com*) pointing to the web servers on WEBHOST1 and WEBHOST2.
4. The load balancer has been configured with a virtual hostname (*oiminternal.mycompany.com*) pointing to web servers WEBHOST1 and WEBHOST2.
5. For details on configuring the VIPs on the load balancer, see [Section 8.2.5.4, "Configuring Virtual Server Names and Ports for the Load Balancer."](#)

##### 8.9.3.15.2 Configuring Oracle HTTP Servers to Front End the OIM and SOA Managed Servers

1. On each of the web servers on WEBHOST1 and WEBHOST2, create a file called *oim.conf* in the directory *ORACLE\_INSTANCE/config/OHS/component/moduleconf*. This file must contain the following information:

```
oim admin console(idmshell based)
<Location /admin>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
```

```

 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassthrough ON
 </Location>

oim self and advanced admin webapp consoles(canonic webapp)

<Location /oim>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassthrough ON
</Location>

SOA Callback webservice for SOD - Provide the SOA Managed Server Ports
<Location /sodcheck>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:8001,oimvhn2.mycompany.com:8001
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassthrough ON
</Location>

Callback webservice for SOA. SOA calls this when a request is
approved/rejected
Provide the SOA Managed Server Port
<Location /workflowservice>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassthrough ON
</Location>

xlWebApp - Legacy 9.x webapp (struts based)
<Location /xlWebApp>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassthrough ON
</Location>

Nexaweb WebApp - used for workflow designer and DM
<Location /Nexaweb>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassthrough ON
</Location>

used for FA Callback service.
<Location /callbackResponseService>

```

```

 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassThrough ON
 </Location>

spml xsd profile
<Location /spml-xsd>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassThrough ON
</Location>

<Location /HTTPlnt>
 SetHandler weblogic-handler
 WLCookieName oimjsessionid
 WebLogicCluster oimvhn1.mycompany.com:14000,oimvhn2.mycompany.com:14000
 WLLogFile "${ORACLE_INSTANCE}/diagnostics/logs/mod_wl/oim_component.log"
 WLProxySSL ON
 WLProxySSLPassThrough ON
</Location>

```

2. Create a file called `virtual_hosts.conf` in `ORACLE_INSTANCE/config/OHS/component/moduleconf`. The file must contain the following information:

```

NameVirtualHost *:7777
<VirtualHost *:7777>

 ServerName http://sso.mycompany.com:7777
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>

<VirtualHost *:7777>
 ServerName http://oiminternal.mycompany.com:80
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>

```

3. Save the file on both `WEBHOST1` and `WEBHOST2`.
4. Stop and start the Oracle HTTP Server instances on both `WEBHOST1` and `WEBHOST2` as described in [Section 8.14, "Starting and Stopping Oracle Identity Management Components."](#)

### 8.9.3.16 Validate the Oracle HTTP Server Configuration

To validate that Oracle HTTP Server is configured properly, follow these steps:

1. In a web browser, enter the following URL for the Oracle Identity Manager Console:

```
http://sso.mycompany.com:7777/oim
```

The Oracle Identity Manager Console login page should display.

2. Log into the Oracle Identity Manager Console using the credentials for the xelsysadm user.

### 8.9.3.17 Oracle Identity Manager Failover and Expected Behavior

In a high availability environment, WebLogic Node Manager is configured to monitor the Oracle WebLogic Servers. In case of failure, Node Manager restarts the WebLogic Server.

In a high availability environment, a hardware load balancer is used to load balance requests between the multiple Oracle Identity Manager instances. If one of the Oracle Identity Manager instances fails, the load balancer detects the failure and routes requests to the surviving instances.

In a high availability environment, the state and configuration information is stored in a database that is shared by all the members of the cluster.

The surviving Oracle Identity Manager instances will continue to seamlessly process any unfinished transactions started on the failed instance, since the state information is in the shared database and is available to all the members in the cluster.

When an Oracle Identity Manager instance fails, its database and LDAP connections are released. The surviving instances in the active-active deployment make their own connections to continue processing unfinished transactions on the failed instance.

Be aware of the following when you deploy Oracle Identity Manager in a high availability configuration:

- Oracle Identity Manager can be deployed on an Oracle RAC database, but Oracle RAC failover is not transparent for Oracle Identity Manager in this release. If an Oracle RAC failover occurs, end users may have to resubmit their requests.
- Oracle Identity Manager always requires that at least one of the nodes in the SOA cluster be available. If the SOA cluster is not available, end user requests will fail. Oracle Identity Manager does not retry for a failed SOA call. Therefore, the end user must retry when a SOA call fails.

### 8.9.3.18 Troubleshooting Oracle Identity Manager High Availability

If you are creating a user in Oracle Identity Manager (by logging into Oracle Identity Manager, clicking the Administration tab, clicking the **Create User** link, entering the required information in the fields, and clicking **Save**) in an active-active Oracle Identity Manager configuration, and the Oracle Identity Manager server that is handling the request fails, you may see a ResourceConnectionValidationException in the Oracle Identity Manager log file, similar to:

```
[2010-06-14T15:14:48.738-07:00] [oim_server2] [ERROR] [] [XELLERATE.SERVER]
[tid: [ACTIVE].ExecuteThread: '0' for queue: 'weblogic.kernel.Default
(self-tuning)'] [userId: xelsysadm] [ecid:
004YGJGmYrtEkJV6u3M6UH00073A0005EI,0:1] [APP: oim#11.1.1.3.0] [dcid:
12eb0f9c6e8796f4:-785b18b3:12938857792:-7ffd-0000000000000037] [URI:
/admin/faces/pages/Admin.jspx] Class/Method:
PooledResourceConnection/heartbeat encounter some problems: Operation timed
out[[
com.oracle.oim.gcp.exceptions.ResourceConnectionValidationException: Operation
timed out
 at
oracle.iam.ldapsync.impl.repository.LDAPConnection.heartbeat(LDAPConnection.ja
va:162)
```

```

 at
com.oracle.oim.gcp.ucp.PooledResourceConnection.heartbeat(PooledResourceConnec
tion.java:52)
 .
 .
 .

```

Although this exception is received, the user is created fine.

### 8.9.3.19 Scaling Up and Scaling Out the Oracle Identity Manager Topology

You can scale out or scale up the Oracle Identity Manager high availability topology. When you scale up the topology, you add new managed servers to nodes that are already running one or more managed servers. When you scale out the topology, you add new managed servers to new nodes.

**8.9.3.19.1 Scaling Up Oracle Identity Manager** In this case, you already have a node that runs a managed server configured with SOA components. The node contains a Middleware home, an Oracle HOME (SOA) and a domain directory for existing managed servers.

You can use the existing installations (the Middleware home, and domain directories) for creating new WLS\_OIM and WLS\_SOA servers. There is no need to install the OIM and SOA binaries in a new location, or to run pack and unpack.

Follow these steps for scaling up the topology:

1. Using the Administration Console, clone either the WLS\_OIM1 or the WLS\_SOA1 into a new managed server. The source managed server to clone should be one that already exists on the node where you want to run the new managed server.

To clone a managed server:

- a. Select **Environment -> Servers** from the Administration Console.
- b. Select the managed server that you want to clone (for example, WLS\_OIM1 or WLS\_SOA1).
- c. Select **Clone**.
- d. Name the new managed server WLS\_OIMn or WLS\_SOAn, where n is a number to identity the new managed server.

The rest of the steps assume that you are adding a new server to OIMHOST1, which is already running WLS\_SOA1 and WLS\_OIM1.

2. For the listen address, assign the host name or IP to use for this new managed server. If you are planning to use server migration as recommended for this server, this should be the VIP (also called a floating IP) to enable it to move to another node. The VIP should be different from the one used by the managed server that is already running.
3. Create JMS Servers for SOA, OIM, BPM, and UMS on the new managed server.
  - a. Use the Oracle WebLogic Server Administration Console to create a new persistent store for the new SOAJMServer and name it, for example, SOAJMSFileStore\_n. Specify the path for the store. This should be a directory on shared storage.

```
ORACLE_BASE/admin/domain_name/cluster_name/jms
```

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

- b. Create a new JMS Server for SOA, for example, SOAJMSServer\_n. Use the SOAJMSFileStore\_n for this JMSServer. Target the SOAJMSServer\_n Server to the recently created Managed Server (WLS\_SOAn).
- c. Create a new persistence store for the new UMSJMSServer, for example, UMSJMSFileStore\_n. Specify the path for the store. This should be a directory on shared storage.

*ORACLE\_BASE/admin/domain\_name/cluster\_name/jms/UMSJMSFileStore\_n*

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

You can also assign SOAJMSFileStore\_n as store for the new UMS JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

---

- d. Create a new JMS Server for UMS, for example, UMSJMSServer\_n. Use the UMSJMSFileStore\_n for this JMSServer. Target the UMSJMSServer\_n Server to the recently created managed server (WLS\_SOAn).
- e. Create a new persistence store for the new OIMJMSServer, for example, OIMJMSFileStore\_n. Specify the path for the store. This should be a directory on shared storage.

*ORACLE\_BASE/admin/domain\_name/cluster\_name/jms/OIMJMSFileStore\_n*

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

You can also assign SOAJMSFileStore\_n as store for the new OIM JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

---

- f. Create a new JMS Server for OIM, for example, OIMJMSServer\_n. Use the OIMJMSFileStore\_n for this JMSServer. Target the OIMJMSServer\_n Server to the recently created Managed Server (WLS\_OIMn).
- g. Update the SubDeployment targets for the SOA JMS Module to include the recently created SOA JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **SOAJMSModule** (represented as a hyperlink in the **Names** column of the table). The Settings page for SOAJMSModule appears. Click the **SubDeployments** tab. The subdeployment module for SOAJMS appears.

---

**Note:** This subdeployment module name is a random name in the form of SOAJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS\_SOA1 and WLS\_SOA2).

---

- h. Click the **SOAJMSServerXXXXXX** subdeployment. Add the new JMS Server for SOA called SOAJMSServer\_n to this subdeployment. Click **Save**.
- i. Update the SubDeployment targets for the UMSJMSSystemResource to include the recently created UMS JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **UMSJMSSystemResource** (represented as a hyperlink in the Names column of the table). The Settings page for UMSJMSSystemResource appears. Click the SubDeployments tab. The subdeployment module for UMSJMS appears.

---

**Note:** This subdeployment module name is a random name in the form of UCMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS\_SOA1 and WLS\_SOA2).

---

- j. Click the **UMSJMSServerXXXXXX** subdeployment. Add the new JMS Server for UMS called UMSJMSServer\_n to this subdeployment. Click **Save**.
- k. Update the SubDeployment targets for OIMJMSModule to include the recently created OIM JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **OIMJMSSystemResource** (represented as a hyperlink in the **Names** column of the table). The Settings page for OIMJMSSystemResource appears. Click the SubDeployments tab. The subdeployment module for OIMJMS appears.

---

**Note:** This subdeployment module name is a random name in the form of OIMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS\_OIM1 and WLS\_OIM2).

---

- l. Click the **OIMJMSServerXXXXXX** subdeployment. Add the new JMS Server for OIM called OIMJMSServer\_n to this subdeployment. Click **Save**.
4. Configure Oracle Coherence for deploying composites, as described in [Section 5.13.8, "Configuring Oracle Coherence for Deploying Composites."](#)

---

**Note:** Only the localhost field needs to be changed for the server. Replace the localhost with the listen address of the new server added, for example:

Dtangosol.coherence.localhost=SOAHOST1VHNn

---



5. Configure the transaction persistent store for the new server. This should be a shared storage location visible from other nodes.

From the Administration Console, select **Server\_name** > **Services** tab. Under Default Store, in Directory, enter the path to the folder where you want the default persistent store to store its data files.

6. Disable host name verification for the new managed server. Before starting and verifying the WLS\_SOA<sub>n</sub> managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in SOAHOST<sub>n</sub>. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required (the hostname verification settings is propagated to the cloned server).

To disable host name verification:

- a. In the Oracle Enterprise Manager Console, select Oracle WebLogic Server Administration Console.
  - b. Expand the **Environment** node in the Domain Structure window.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Select WLS\_SOA<sub>n</sub> in the **Names** column of the table. The Settings page for the server appears.
  - e. Click the SSL tab.
  - f. Click **Advanced**.
  - g. Set **Hostname Verification** to **None**.
  - h. Click **Save**.
7. Start and test the new managed server from the Administration Console.
    - a. Shut down the existing managed servers in the cluster.
    - b. Ensure that the newly created managed server, WLS\_SOA<sub>n</sub>, is up.
    - c. Access the application on the newly created managed server (<http://vip:port/soa-infra>). The application should be functional
  8. Configure Server Migration for the new managed server.

---

**Note:** Since this is a scale up operation, the node should already contain a Node Manager and environment configured for server migration. The floating IP for the new SOA managed server should also be already present.

---

Configure server migration by following these steps:

- a. Log into the Administration Console.
- b. In the left pane, expand **Environment** and select **Servers**.
- c. Select the server (represented as a hyperlink) for which you want to configure migration. The Settings page for that server appears.
- d. Click the Migration tab.

- e. In the Available field, in the Migration Configuration section, select the machines to which to allow migration and click the right arrow. Select the same migration targets as for the servers that already exist on the node.  
  
For example, for new managed servers on SOAHOST1, which is already running WLS\_SOA1, select SOAHOST2. For new managed servers on SOAHOST2, which is already running WLS\_SOA2, select SOAHOST1.  
  
Make sure the appropriate resources are available to run the managed servers concurrently during migration.
  - f. Select the **Automatic Server Migration Enabled** option. This enables the Node Manager to start a failed server on the target node automatically.
  - g. Click **Save**.
  - h. Restart the Administration Server, managed servers, and Node Manager.
  - i. Repeat the steps in this list for configuring server migration for the newly created WLS\_OIMn managed server.
9. Test server migration for this new server. Follow these steps from the node where you added the new server:
    - a. Stop the WLS\_SOAn managed server.  
  
To do this, run "kill -9 *pid*" on the PID of the managed server. You can identify the PID of the node using "ps -ef | grep WLS\_SOAn".
    - b. Watch the Node Manager Console: you should see a message indicating that WLS\_SOA1's floating IP has been disabled.
    - c. Wait for the Node Manager to try a second restart of WLS\_SOAn. Node Manager waits for a fence period of 30 seconds before trying this restart.
    - d. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

**8.9.3.19.2 Scaling Out Oracle Identity Manager** When you scale out the topology, you add new managed servers configured with SOA to new nodes.

Before performing the steps in this section, check that you meet these requirements:

- There must be existing nodes running managed servers configured with SOA within the topology.
- The new node can access the existing home directories for WebLogic Server and SOA. (Use the existing installations in shared storage for creating a new WLS\_SOA or WLS\_WSM managed server. You do not need to install WebLogic Server or SOA binaries in a new location but you do need to run pack and unpack to bootstrap the domain configuration in the new node.)

---

**Note:** If there is no existing installation in shared storage, installing WebLogic Server and SOA in the new nodes is required as described in [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines."](#)

---

---

**Note:** When an *ORACLE\_HOME* or *WL\_HOME* is shared by multiple servers in different nodes, Oracle recommends keeping the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. To update the oraInventory in a node and "attach" an installation in a shared storage to it, use *ORACLE\_HOME*/oui/bin/attachHome.sh. To update the Middleware home list to add or remove a *WL\_HOME*, edit the user\_home/boa/beahomelist file. See the following steps.

---

Follow these steps for scaling out the topology:

1. On the new node, mount the existing Middleware home, which should include the SOA installation and the domain directory, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To attach *ORACLE\_HOME* in shared storage to the local Oracle Inventory, execute the following commands:

```
SOAHOSTn>cd ORACLE_BASE/product/fmw/soa/
SOAHOSTn>./attachHome.sh -jreLoc ORACLE_BASE/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the *MW\_HOME*/boa/beahomelist file and add *ORACLE\_BASE*/product/fmw to it.

3. Log in to the Oracle WebLogic Administration Console.
4. Create a new machine for the new node that will be used, and add the machine to the domain.
5. Update the machine's Node Manager's address to map the IP of the node that is being used for scale out.
6. Use the Oracle WebLogic Server Administration Console to clone WLS\_SOA1 into a new managed server. Name it WLS\_SOAn, where n is a number.

---

**Note:** These steps assume that you are adding a new server to node n, where no managed server was running previously.

---

7. Assign the host name or IP to use for the new managed server for the listen address of the managed server.

If you are planning to use server migration for this server (which Oracle recommends), this should be the VIP (also called a floating IP) for the server. This VIP should be different from the one used for the existing managed server.

8. Create JMS servers for SOA, OIM (if applicable), and UMS on the new managed server.
  - a. Use the Oracle WebLogic Server Administration Console to create a new persistent store for the new SOAJMSServer and name it, for example, SOAJMSFileStore\_n. Specify the path for the store. This should be a directory on shared storage.

```
ORACLE_BASE/admin/domain_name/cluster_name/jms/SOAJMSFileStore_n
```

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

- b. Create a new JMS Server for SOA, for example, SOAJMSServer\_n. Use the SOAJMSFileStore\_n for this JMSServer. Target the SOAJMSServer\_n Server to the recently created Managed Server (WLS\_SOAn).
- c. Create a new persistence store for the new UMSJMSServer, for example, UMSJMSFileStore\_n. Specify the path for the store. This should be a directory on shared storage.

*ORACLE\_BASE/admin/domain\_name/cluster\_name/jms/UMSJMSFileStore\_n*

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

You can also assign SOAJMSFileStore\_n as store for the new UMS JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

---

- d. Create a new JMS Server for UMS, for example, UMSJMSServer\_n. Use the UMSJMSFileStore\_n for this JMSServer. Target the UMSJMSServer\_n Server to the recently created managed server (WLS\_SOAn).
- e. Create a new persistence store for the new OIMJMSServer, for example, OIMJMSFileStore\_n. Specify the path for the store. This should be a directory on shared storage.

*ORACLE\_BASE/admin/domain\_name/cluster\_name/jms/OIMJMSFileStore\_n*

---

**Note:** This directory must exist before the managed server is started or the start operation fails.

---

You can also assign SOAJMSFileStore\_n as store for the new OIM JMS Servers. For the purpose of clarity and isolation, individual persistent stores are used in the following steps.

---

- f. Create a new JMS Server for OIM, for example, OIMJMSServer\_n. Use the OIMJMSFileStore\_n for this JMSServer. Target the OIMJMSServer\_n Server to the recently created Managed Server (WLS\_OIMn).
- g. Update the SubDeployment targets for the SOA JMS Module to include the recently created SOA JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **SOAJMSModule** (represented as a hyperlink in the **Names** column of the table). The Settings page for SOAJMSModule appears. Click the **SubDeployments** tab. The subdeployment module for SOAJMS appears.

---

**Note:** This subdeployment module name is a random name in the form of SOAJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS\_SOA1 and WLS\_SOA2).

---

- h. Click the **SOAJMSServerXXXXXX** subdeployment. Add the new JMS Server for SOA called SOAJMSServer\_n to this subdeployment. Click **Save**.
- i. Update the SubDeployment targets for the UMSJMSSystemResource to include the recently created UMS JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **UMSJMSSystemResource** (represented as a hyperlink in the Names column of the table). The Settings page for UMSJMSSystemResource appears. Click the SubDeployments tab. The subdeployment module for UMSJMS appears.

---

**Note:** This subdeployment module name is a random name in the form of UCMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS\_SOA1 and WLS\_SOA2).

---

- j. Click the **UMSJMSServerXXXXXX** subdeployment. Add the new JMS Server for UMS called UMSJMSServer\_n to this subdeployment. Click **Save**.
- k. Update the SubDeployment targets for OIMJMSModule to include the recently created OIM JMS Server. To do this, expand the **Services** node and then expand the **Messaging** node. Choose **JMS Modules** from the Domain Structure window of the Oracle WebLogic Server Administration Console. The JMS Modules page appears. Click **OIMJMSSystemResource** (represented as a hyperlink in the **Names** column of the table). The Settings page for OIMJMSSystemResource appears. Click the SubDeployments tab. The subdeployment module for OIMJMS appears.

---

**Note:** This subdeployment module name is a random name in the form of OIMJMSServerXXXXXX resulting from the Config Wizard JMS configuration for the first two servers (WLS\_OIM1 and WLS\_OIM2).

---

- l. Click the **OIMJMSServerXXXXXX** subdeployment. Add the new JMS Server for OIM called OIMJMSServer\_n to this subdeployment. Click **Save**.
9. Run the pack command on SOAHOST1 to create a template pack as follows:

```
SOAHOST1> cd ORACLE_COMMON_HOME/common/bin
SOAHOST1> ./pack.sh -managed=true/
-domain=MW_HOME/user_projects/domains/soadomain/
-template=soadomaintemplateScale.jar -template_name=soa_domain_templateScale
```

Run the following command on SOAHOST1 to copy the template file created to SOAHOSTn:

```
SOAHOST1> scp soadomaintemplateScale.jar oracle@SOAHOSTN:/
ORACLE_BASE/product/fmw/soa/common/bin
```

Run the unpack command on SOAHOSTn to unpack the template in the managed server domain directory as follows:

```
SOAHOSTn> cd ORACLE_BASE/product/fmw/soa/common/bin
SOAHOSTN> ./unpack.sh /
-domain=ORACLE_BASE/product/fmw/user_projects/domains/soadomain/
-template=soadomaintemplateScale.jar
```

10. Configure Oracle Coherence for deploying composites, as described in [Section 5.13.8, "Configuring Oracle Coherence for Deploying Composites."](#)

---

**Note:** Only the localhost field needs to be changed for the server. Replace the localhost with the listen address of the new server added, for example:

Dtangosol.coherence.localhost=SOAHOST1VHNn

---

11. Configure the transaction persistent store for the new server. This should be a shared storage location visible from other nodes.

From the Administration Console, select **Server\_name > Services** tab. Under Default Store, in Directory, enter the path to the folder where you want the default persistent store to store its data files.

12. Disable host name verification for the new managed server. Before starting and verifying the WLS\_SOAn managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in SOAHOSTn. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required (the hostname verification settings is propagated to the cloned server).

To disable host name verification:

- a. In the Oracle Enterprise Manager Console, select Oracle WebLogic Server Administration Console.
- b. Expand the **Environment** node in the Domain Structure window.
- c. Click **Servers**. The Summary of Servers page appears.
- d. Select WLS\_SOAn in the **Names** column of the table. The Settings page for the server appears.
- e. Click the SSL tab.
- f. Click **Advanced**.
- g. Set **Hostname Verification** to **None**.
- h. Click **Save**.
13. Start the Node Manager on the new node. To start the Node Manager, use the installation in shared storage from the existing nodes, and start Node Manager by passing the host name of the new node as a parameter as follows:

```
SOAHOSTn> WL_HOME/server/bin/startNodeManager new_node_ip
```

14. Start and test the new managed server from the Administration Console.
  - a. Shut down the existing managed servers in the cluster.

- b. Ensure that the newly created managed server, WLS\_SOAn, is up.
  - c. Access the application on the newly created managed server (<http://vip:port/soa-infra>). The application should be functional
15. Configure Server Migration for the new managed server.

---

**Note:** Since this new node is using an existing shared storage installation, the node is already using a Node Manager and environment configured for server migration that includes netmask, interface, wlsifconfig script superuser privileges. The floating IP for the new SOA managed server is already present in the new node.

---

Configure server migration by following these steps:

- a. Log into the Administration Console.
- b. In the left pane, expand **Environment** and select **Servers**.
- c. Select the server (represented as a hyperlink) for which you want to configure migration. The Settings page for that server appears.
- d. Click the Migration tab.
- e. In the Available field, in the Migration Configuration section, select the machines to which to allow migration and click the right arrow.

---

**Note:** Specify the least-loaded machine as the migration target for the new server. The required capacity planning must be completed so that this node has enough available resources to sustain an additional managed server.

---

- f. Select the **Automatic Server Migration Enabled** option. This enables the Node Manager to start a failed server on the target node automatically.
  - g. Click **Save**.
  - h. Restart the Administration Server, managed servers, and Node Manager.
16. Test server migration for this new server. Follow these steps from the node where you added the new server:
- a. Stop the WLS\_SOAn managed server.  
To do this, run "kill -9 pid" on the PID of the managed server. You can identify the PID of the node using "ps -ef | grep WLS\_SOAn".
  - b. Watch the Node Manager Console: you should see a message indicating that WLS\_SOA1's floating IP has been disabled.
  - c. Wait for the Node Manager to try a second restart of WLS\_SOAn. Node Manager waits for a fence period of 30 seconds before trying this restart.
  - d. Once Node Manager restarts the server, stop it again. Now Node Manager should log a message indicating that the server will not be restarted again locally.

## 8.10 Authorization Policy Manager High Availability

This section provides an introduction to Authorization Policy Manager 11gR1 and describes how to design and deploy a high availability environment for Authorization Policy Manager.

To use Authorization Policy Manager in a high availability active-passive Cold Failover Cluster, follow the instructions for setting up WebLogic Server in an active-passive Cold Failover Cluster in [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

Authorization Policy Manager is not used in high availability active-active deployments because it is deployed to the WebLogic Administration Server JVM.

Authorization Policy Manager is a graphical interface tool for administering application policies.

The intended users of Authorization Policy Manager are security administrators. With this tool, an administrator can view and manage policies across enterprise applications. Administrators can be specified to manage all applications running in the domain or just a subset of them.

Authorization Policy Manager requires that:

- The domain policy store be LDAP-based; the supported policy store is Oracle Internet Directory.
- The domain identity store be LDAP-based; the supported identity store types are Oracle Internet Directory, Oracle Virtual Directory, Oracle WebLogic Server Embedded LDAP, Sun Java System Directory Service version 6.3, Active Directory 2003 and 2008, Novell eDirectory 8.8, and OpenLDAP 2.2 and 2.4.
- Two particular data sources be set: mds-ApplicationMDSDB and apm-DBDS. These data sources can be configured with the WebLogic Console (under **JDBC > Data Sources**).

For more information about using Authorization Policy Manager, see *Oracle Fusion Middleware Authorization Policy Manager Administrator's Guide*.

## 8.11 Oracle Identity Navigator High Availability

Oracle Identity Management Navigator is an administrative portal designed to act as a launch pad for Oracle Identity Management products. It does not replace the individual product consoles. Rather, it allows you to access the Oracle Identity Management consoles from one site. For more details on Oracle Identity Management Navigator, refer to *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Navigator*.

Oracle Identity Navigator is not used in high availability active-active deployments because it is deployed to the WebLogic Administration Server JVM.

To use Oracle Identity Navigator in a high availability active-passive Cold Failover Cluster, follow the instructions for setting up WebLogic Server in an active-passive Cold Failover Cluster in [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)



## 8.12 Oracle Adaptive Access Manager High Availability

This section provides an introduction to Oracle Adaptive Access Manager and describes how to design and deploy a high availability environment for Oracle Adaptive Access Manager.


Oracle Adaptive Access Manager is built on a J2EE-based, multi-tier deployment architecture that separates the platform's presentation, business logic, and data tiers. Because of this separation of tiers, Oracle Adaptive Access Manager can rapidly scale with the performance needs of the customer. The architecture can leverage the most flexible and supported cross-platform J2EE services available: a combination of Java, XML and object technologies. This architecture makes Oracle Adaptive Access Manager a scalable, fault-tolerant solution.

This section includes the following topics:

- [Section 8.12.1, "Oracle Adaptive Access Manager Component Architecture"](#)
- [Section 8.12.2, "Oracle Adaptive Access Manager High Availability Concepts"](#)
- [Section 8.12.3, "Oracle Adaptive Access Manager High Availability Configuration Steps"](#)

### 8.12.1 Oracle Adaptive Access Manager Component Architecture

[Figure 8–18](#) shows the single instance architecture for Oracle Adaptive Access Manager.

**Figure 8–18 Oracle Adaptive Access Manager Single Instance Architecture**

In the Oracle Adaptive Access Manager single instance architecture shown in [Figure 8–18](#), end users access customer web applications, which communicate with the OAAM Server application and its policies using SOAP. Alternately, an OAAM proxy can be set up so that end users communicate with that machine, which then communicates with the OAAM\_Server application using HTTP(S). If the end user is authorized, he or she will be allowed to access the customer web application. The OAAM\_ADMIN component is used for administration and configuration of the OAAM\_SERVER application. The administrator responsible for administering and configuring the OAAM\_Server application uses a web browser to access the OAAM\_ADMIN application. An Oracle RAC database is used to hold policy and configuration information.

### 8.12.1.1 Oracle Adaptive Access Manager Component Characteristics

Oracle Adaptive Access Manager consists of the following two components:

- **OAAM\_ADMIN:** This component is used for administration and configuration of OAAM\_SERVER application. This component is developed using the Oracle JAVA ADF Framework the Identity Management shell and deployed as Web applications in a J2EE container. It is packaged as an EAR file.
- **OAAM\_SERVER:** This component contains the OAAM Admin and OAAM Server sub-components within a single web application. The OAAM\_SERVER component is packaged as an EAR file and is composed of Servlets and JSPs in addition to Java classes. The subcomponents of OAAM\_SERVER are described below by layer:
  - **Presentation Layer:** typically a Web application serving JSPs, servlets, and so on. The presentation layer provides the strong authenticator functionality; it uses the interfaces provided by the business layer (SOAP or Java native) to access its services.
  - **Business Logic Layer:** this layer contains the core application logic that implements the risk analyzing engine. This layer provides Java and SOAP interfaces for the presentation layer. When the Java interface is used, the business logic layer and presentation layer can be part of a single web application. With the SOAP interface, these layers are deployed as different applications.
  - **Data Access Layer:** contains data access components to connect to the supported relational databases. Oracle Adaptive Access Manager uses Oracle's TopLink, which provides a powerful and flexible framework for storing Java objects in a relational database.

The following components can also be used in an Oracle Adaptive Access Manager 11gR1 deployment:

- **Fusion Middleware Control / Enterprise Manager Grid Control:** Oracle Adaptive Access Manager integrates with the Enterprise Manager Grid Control to display performance metrics and deployment topology. Oracle Adaptive Access Manager uses DMS and Discovery Mbeans to integrate with Enterprise Manager. Enterprise Manager is also used to enhance component tracing and configure auditing.

Enterprise Manager can also be used to view log files for each Managed Server in the domain and increase the tracing to Debug, Trace, and Info levels.

- **Data repositories:** Oracle Adaptive Access Manager uses the RDBMS database as its data store. Oracle Adaptive Access Manager supports and works on the following database technologies:
  - Oracle Real Application Clusters
  - Oracle Data Guard
  - Replication Streams
  - Database Partitioning

Oracle Adaptive Access Manager uses RCU to create its schema in the database.

**8.12.1.1.1 Oracle Adaptive Access Manager State Information** The OAAM\_Server component includes the OAAM Server subcomponent and the OAAM Admin subcomponent.

- **OAAM Server** is a stateful application that stores the state in HTTP session.

- OAAM Admin is a stateful application that stores its session information in the database.

The OAAM\_Admin component is an ADF and Identity Management UI shell-based application. It is a stateless application, and its application state is maintained by the ADF framework.

**8.12.1.1.2 Oracle Adaptive Access Manager Runtime Processes** You can perform the following runtime tasks using the Oracle Adaptive Access Manager Administration Console:

- Customer Care application tasks
- System configuration tasks involving policies, groups, and properties
- Viewing session data information
- Viewing the System Statistics dashboard

For example, you can perform the following administration flows:

1. Recent user query:
  - a. View recent logins and session details.
  - b. Perform a query.
  - c. Click **Session Details**.
  - d. Log out.
2. Manual CSR and Agent Case creation:
  - a. To reset customer care, log in.
  - b. Create a case.
  - c. Reset the customer.
  - d. Log out.

You can also perform runtime processing with the Oracle Adaptive Access Manager Server.

**8.12.1.1.3 Oracle Adaptive Access Manager Process Lifecycle** The following runtime processing occurs with Oracle Adaptive Access Manager Server:

1. Oracle Adaptive Access Manager is deployed and integrated with the customer's application.
2. The user will access the customer's application and enter user credentials.
3. Based on the system and rules configured in OAAM, different login flows will be presented, for example:

User Registration: Registration Flows

- a. Flow R1: Login (New User), enter password, personalize device, skip Questions Registration, log out.
- b. Flow R2: Login, enter password, skip Registration, log out.
- c. Flow R3: Login, enter password, personalize device, continue Questions Registration, log out.
- d. Flow R4: Login, enter password, personalize device, continue Questions Registration, enter invalid answers, validation, log out

- e. Flow R5: Login (New Device and New User), enter password, personalize device, continue Questions Registration, log out.

Login Flow:

- a. Flow L1: Login, enter wrong password, Login screen.
- b. Flow L2: Login, enter correct password, Challenge On may be presented, answer correctly, logged in.
- c. Flow L3: Login, enter correct password, Challenge On presented, answer incorrectly, Challenge On presented again, answer correctly, logged in.
- d. Flow L4: Login, enter correct password, Challenge On presented, answer incorrectly, Challenge On presented again, answer incorrectly, Challenge On presented again, answer correctly, logged in.
- e. Flow L5: Login, enter correct password, Challenge On presented, answer incorrectly, Challenge On presented again, answer incorrectly, login blocked.
- f. Flow L6: Login, enter correct password, login blocked (login screen).
- g. Flow L7: Login, enter correct password, logged in.
- h. Flow LNU1Login as a new user: Login, enter correct password, logged in.
- i. Flow LND1: Existing user, login with a new device, enter correct password, logged in.
- j. Flow LNUD1: New user, login with a new device, enter correct password, logged in.

In Session Transaction Flow: Login, enter password, create transaction, update transaction, log out.

4. OAAM tracks and registers the following data elements:

- a. User login
- b. User names
- c. Devices (cookies, browser headers, and flash data supplied)
- d. IP addresses
- e. Transaction data

5. OAAM will trigger the appropriate policy based on login behavior.

As J2EE applications, the Oracle Adaptive Access Manager Server and Administration Console can be started using the user interface and command line tools provided by WebLogic Server.

The Oracle Adaptive Access Manager Server supports a health check request (a ping request over HTTP) that can be used by a load balancer for death detection.

Oracle Adaptive Access Manager is instrumented for server side metrics (using DMS) and this information is published to the Administration Console. When DMS metrics collection is enabled, monitoring the agent and server component metrics can serve as a proxy for component monitoring. In addition, Oracle Adaptive Access Manager supports fine-grained real time activity tracing, which can also serve as a proxy for component monitoring.

**8.12.1.1.4 Oracle Adaptive Access Manager Configuration Artifacts** Oracle Adaptive Access Manager stores its configuration information in the database. These configuration

properties can be changed using the Oracle Adaptive Access Manager Administration Console.

Oracle Adaptive Access Manager does not store any configuration information on the file system or in the exploded EAR file.

**8.12.1.1.5 Oracle Adaptive Access Manager Deployment Artifacts** Oracle Adaptive Access Manager supports the nostage mode of deployment staging. That is, all deployment files are local.

**8.12.1.1.6 Oracle Adaptive Access Manager External Dependencies** Oracle Adaptive Access Manager has an external dependency on the RDBMS database, where it stores its configuration information.

Oracle Adaptive Access Manager uses WebLogic Server multi data sources for Oracle RAC databases.

Oracle Adaptive Access Manager uses the standard Oracle TopLink object caching mechanism.

Oracle Adaptive Access Manager follows standard session object serialization to maintain the persistent state of an object.

Oracle Adaptive Access Manager is not dependent on any hostname, IP address, or port. It will work on a container-specific port or host name.

**8.12.1.1.7 Oracle Adaptive Access Manager Log File Location** Oracle Adaptive Access Manager is a J2EE application deployed on WebLogic Server. All log messages are logged in the server log file of the WebLogic Server that the application is deployed on. The default location of the server log is:


```
WL_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

## 8.12.2 Oracle Adaptive Access Manager High Availability Concepts

This section provides conceptual information about using Oracle Adaptive Access Manager in a high availability two-node cluster.

### 8.12.2.1 Oracle Adaptive Access Manager High Availability Architecture

[Figure 8–19](#) shows the Oracle Adaptive Access Manager high availability architecture:

**Figure 8–19 Oracle Adaptive Access Manager High Availability Architecture**

In [Figure 8–19](#), incoming requests are received by the hardware load balancer, which routes them to WEBHOST1 or WEBHOST2 in the web tier. These hosts have Oracle HTTP Server installed. The Oracle HTTP Server then forwards requests on to the WebLogic managed servers on OAAMHOST1 and OAAMHOST2 using the WebLogic plugin mod\_wl\_ohs.

OAAMHOST1 and OAAMHOST2 contain managed servers which host the Oracle Adaptive Access Manager Server application and the Oracle Adaptive Access Manager Admin application. These managed servers are configured in a cluster which allows the Access Servers to work in an active-active manner.

The WebLogic Administration Server runs on OAAMHOST1 and contains the WebLogic Administration Console and the Oracle Enterprise Manager Fusion Middleware Control. The Administration Server can be configured to run in active-passive mode on OAAMHOST2, which means that if OAAMHOST1 becomes unavailable, then Administration Server can be manually started on OAAMHOST2.

In the directory tier, the virtual IP `ovd.mycompany.com` is set up to route Oracle Virtual Directory requests to OVDHOST1 and OVDHOST2, which comprise an active-active Oracle Virtual Directory cluster. The virtual IP `oid.mycompany.com` is set up to route Oracle Internet Directory requests to OIDHOST1 and OIDHOST2, which comprise an active-active Oracle Internet Directory cluster.

An Oracle RAC database provides high availability in the data tier.

Figure 8–19 shows the OAAM high availability configuration architecture. In the figure, a load balancer routes requests through two Oracle HTTP Server instances on WEBHOST1 and WEBHOST2 to OAAMHOST1 and OAAMHOST2. An OAAM Administration Server instance and an OAAM Managed Server instance is installed on OAAMHOST1 and on OAAMHOST2, and these installations are configured as an OAAM Server cluster (OAAM\_Cluster) and an OAAM Admin Cluster (OAAM\_Admin\_Cluster). The OAAM Server cluster uses the OAAM Server data source and the OAAM Admin cluster uses the OAAM Admin data source to communicate with the Oracle RAC database.

As shown in Figure 8–19, only one Oracle Adaptive Access Manager Server cluster and one Oracle Adaptive Access Manager Administration cluster is supported per WebLogic Server domain. In addition, Oracle Adaptive Access Manager-related clusters cannot span WebLogic Server domains.

A single instance Oracle Adaptive Access Manager deployment satisfies the following high availability requirements:

- Load handling
- External connection management and monitoring
- Recovery
- Fault containment
- Fault diagnostics
- Oracle Adaptive Access Manager Admin / Server offline

A multiple instance Oracle Adaptive Access Manager deployment satisfies the following additional high availability requirements:

- Redundancy
- Client connection failover / continuity
- Client load balancing
- State management

Use of external load balancing routers is recommended for inbound HTTP connections.

Web sessions open persistent TCP connections to the Oracle Adaptive Access Manager Administration Console and servers. This requires that load balancing router and firewall connection timeouts are sufficiently large to avoid premature termination of TCP connections.

**8.12.2.1.1 Starting and Stopping the Cluster** Each Oracle Adaptive Access Manager Administration Console and Server instance is a peer of other instances. Because all initialization happens before the Server is ready to receive requests and because of built in throttling capabilities, surge conditions are dealt with gracefully without any significant impact of the performance characteristics of the Oracle Adaptive Access Manager 11gR1 Access Server.



When the cluster is stopped, new requests will be denied and existing requests will be allowed to complete before the Oracle Adaptive Access Manager Administration Console and Server application shuts down. If a forced shutdown occurs, Oracle Adaptive Access Manager 11gR1 can recover any corrupted or invalid data that the shutdown causes.

Oracle Adaptive Access Manager components are pure J2EE applications and do not have any start or stop functionality of their own. Instead, they rely on container-specific startup and shutdown functionality.

Oracle Adaptive Access Manager components are deployed to WebLogic Server Managed Server nodes. The components can be restarted using Node Manager.

**8.12.2.1.2 Cluster-Wide Configuration Changes** Since Oracle Adaptive Access Manager stores the entire configuration in database, the propagation of configuration changes to all the cluster members transparent. All Oracle Adaptive Access Manager components are notified of change events from the internal layer, which are then taken up by the components. To ensure atomicity of the change, Oracle Adaptive Access Manager components reload their entire configuration every time a change happens.

Oracle Adaptive Access Manager configuration applies to every instance in a cluster.

Adding and removing Oracle Adaptive Access Manager Administration Console and Server instances is transparent to other Oracle Adaptive Access Manager instances in the cluster.

An Oracle Adaptive Access Manager cluster can have any number of instances. There is no restriction on the number of instances per cluster.

Online application redeployment does not cause any problems.

#### **8.12.2.2 Protection from Failures and Expected Behaviors**

The following features protect an Oracle Adaptive Access Manager high availability configuration from failure:

- Session state for the cluster is maintained in memory, which provides replication and failover for all session state information.
- Oracle Adaptive Access Manager Servers support a heart beat check - a ping request over HTTP. In addition, the WebLogic Node Manager on the Managed Server can monitor the application and restart it if it is not running. Restarting an Oracle Adaptive Access Manager Server has no impact on any other running components or cluster members.
- When a WebLogic Server node fails, external connection failover is based on the configuration and is based on the retry timeout as well as the number of retries.
- When the load balancing router or proxy server detects a WebLogic Server node failure, subsequent client connections are routed to the active instance, which picks up the session state for further processing.
- An Oracle Adaptive Access Manager session does not have a direct impact on an Oracle RAC database node failure, because WebLogic Server maintains the state of its database connections.

### **8.12.3 Oracle Adaptive Access Manager High Availability Configuration Steps**

This section provides high-level instructions for setting up a maximum high availability deployment for Oracle Adaptive Access Manager. This deployment includes two Oracle HTTP Servers, which distribute requests to two OAAM servers.

These OAAM servers interact with an Oracle Real Application Clusters (Oracle RAC) database. If any single component fails, then the remaining components will continue to function.

This section includes the following topics:

- [Section 8.12.3.1, "Prerequisites for Oracle Adaptive Access Manager Configuration"](#)
- [Section 8.12.3.2, "Run the Repository Creation Utility to Create the OAAM Schemas in a Database"](#)
- [Section 8.12.3.3, "Installing Oracle WebLogic Server"](#)
- [Section 8.12.3.4, "Install and Configure the Oracle Adaptive Access Manager Application Tier"](#)
- [Section 8.12.3.5, "Creating boot.properties for the Administration Server on OAAMHOST1"](#)
- [Section 8.12.3.6, "Create the Oracle Adaptive Access Manager Administration User"](#)
- [Section 8.12.3.7, "Start OAAMHOST1"](#)
- [Section 8.12.3.8, "Validating OAAMHOST1"](#)
- [Section 8.12.3.9, "Configure Oracle Adaptive Access Manager on OAAMHOST2"](#)
- [Section 8.12.3.10, "Start OAAMHOST2"](#)
- [Section 8.12.3.11, "Validating OAAMHOST2"](#)
- [Section 8.12.3.12, "Configure Oracle Adaptive Access Manager to Work with Oracle HTTP Server"](#)
- [Section 8.12.3.13, "Validating the Oracle Adaptive Access Manager Configuration"](#)
- [Section 8.12.3.14, "Scaling Up and Scaling Out the Oracle Adaptive Access Manager Topology"](#)

---

**Note:** To install and configure Oracle Adaptive Access Manager offline, see *Installing and Configuring Adaptive Access Manager Offline* in the *Oracle Adaptive Access Manager Installation and Configuration Guide*.

---

### 8.12.3.1 Prerequisites for Oracle Adaptive Access Manager Configuration

Before you configuring Oracle Adaptive Access Manager for high availability, you must:

- Run the Repository Creation Utility to create the OAAM schemas in a database.  
See [Section 8.12.3.2, "Run the Repository Creation Utility to Create the OAAM Schemas in a Database"](#) for instructions on running the Repository Creation Utility to create the OAAM schemas.
- Install Oracle WebLogic Server on OAAMHOST1 and OAAMHOST2.  
Follow the steps in [Section 8.12.3.3, "Installing Oracle WebLogic Server"](#) to install Oracle WebLogic Server on OAMHOST1 and OAMHOST2.
- Install and configure the Oracle Identity Management executables on OAAMHOST1 and OAAMHOST2.

Follow the steps in [Section 8.12.3.4, "Install and Configure the Oracle Adaptive Access Manager Application Tier"](#) to install the Oracle Identity Management executables on OAAMHOST1 and OAAMHOST2.

### 8.12.3.2 Run the Repository Creation Utility to Create the OAAM Schemas in a Database

See [Section 8.2.4.1, "Executing the Repository Creation Utility"](#) for instructions on running the Repository Creation Utility to create the OAAM schemas in your database repository.

### 8.12.3.3 Installing Oracle WebLogic Server

Prior to installing the Oracle WebLogic Server, ensure that your machines meet the system, patch, kernel, and other requirements as specified in *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

Start the installer, then proceed as follows:

1. On the Welcome screen, click **Next**.
2. On the Choose Middleware Home Directory screen, select **Create a New Middleware Home**.

For Middleware Home Directory, enter:

`ORACLE_BASE/product/fmw`

---

**Note:** `ORACLE_BASE` is the base directory under which Oracle products are installed. The recommended value is `/u01/app/oracle`.

---

Click **Next**.

3. On the Register for Security Updates screen, enter your contact information so that you can be notified of security updates.

Click **Next**.

4. On the Choose Install Type screen, select **Custom**.

Click **Next**.

5. On the Choose Products and Components screen, select only **Oracle JRockit SDK**, and click **Next**.

6. On the Choose Product Installation Directories screen, accept the directory `ORACLE_BASE/product/fmw/wlserver_10.3`.

Click **Next**.

7. On the Installation Summary screen, click **Next**.

8. On the Installation Complete screen, deselect **Run Quickstart**.

Click **Done**.

### 8.12.3.4 Install and Configure the Oracle Adaptive Access Manager Application Tier

Perform these steps to install Oracle Fusion Middleware Components on OAAMHOST1 and OAAMHOST2.

On Linux platforms, if the `/etc/oraInst.loc` file exists, verify that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

Start the installer for Oracle Fusion Middleware as follows:

```
OAAMHOST1> runInstaller
```

When the installer prompts you for a JRE/JDK location, enter the Oracle SDK location created in the Oracle WebLogic Server installation, for example, `ORACLE_BASE/product/fmw/jrockit_160_<version>`.

Then proceed as follows:

1. On the Welcome screen, click **Next**.
2. On the Prerequisite Checks screen, verify that the checks complete successfully, then click **Next**.
3. On the Specify Installation Location screen, enter the following values:
  - Oracle Middleware Home: Select the previously installed Middleware home from the list for `MW_HOME`, for example:  
`/u01/app/oracle/product/fmw`
  - Oracle Home Directory:
    - Enter `idm` as the Oracle Home directory name when installing the Oracle Identity Management Suite in the `IDM_ORACLE_HOME`.
    - Enter `iam` as the Oracle Home directory name when installing the Oracle Identity and Access Management Suite in the `IAM_ORACLE_HOME`.
    - Enter `soa` as the Oracle Home directory name when installing the Oracle SOA Suite in the `SOA_ORACLE_HOME`.

Click **Next**.

4. On the Installation Summary screen, click **Install and Configure**.
5. In the Welcome screen, select **Create a New WebLogic Domain**, and then click **Next**.
6. In the Select Domain Source Screen:

Select **Generate a domain configured automatically to support the following products**:

And select the following products:

- **Oracle Enterprise Manager**
- **Oracle JRF** (selected by default)
- **Oracle Adaptive Access Manager - Server**
- **Oracle Adaptive Access Manager Admin Server**

Click **Next**.

7. In the Specify Domain and Location screen enter:
  - **Domain name:** `IDM_Domain`
  - **Domain Directory:** Accept the default.
  - **Application Directory:** Accept the default.

Click **Next**.

8. In the Configure Administrator Username and Password screen, enter the username and password to be used for the domain's administrator, and click **Next**.
9. In the Configure Server Start Mode and JDK screen, make the following selections:
  - **WebLogic Domain Startup Mode:** Select **Production Mode**.
  - **JDK Selection:** Select **JROCKIT SDK1.6.0\_<version>**.
10. In the Configure JDBC Component Schema screen, select all of the data sources, then select **Configure selected data sources as RAC multi data sources**.

Click **Next**.

11. In the Configure RAC Multi Data Source Component Schema screen, select the first data source, **OAAM Admin Server**, and enter the following:
  - **Data source:** OAAM Admin Server
  - **Service Name:** `oaam.mycompany.com`
  - **User Name:** OAAM\_OAAM (assuming OAAM was used as the RCU prefix)
  - **Password:** The password for above account

In the top right box, click **Add** to add an Oracle RAC host. Enter the following information:

- **Host Name:** OAAMDBHOST1
- **Instance Name:** oaamdb1
- **Port:** 1521

Click **Add** again to add the second database host and enter the following information:

- **Host Name:** OAAMDBHOST2
- **Instance Name:** oaamdb2
- **Port:** 1521

Deselect this data source. Select the next data source, **OAAM Admin MDS Schema**, and enter the following information:

- **Data Source:** OAAM Admin MDS Schema
- **Service Name:** `oaam.mycompany.com`
- **User Name:** OAAM\_MDS (assuming OAAM was used as the RCU prefix)
- **Password:** Password for the OAAM\_MDS account.

In the top right box, click **Add** to add an Oracle RAC host. Enter the following information:

- **Host Name:** OAAMDBHOST1
- **Instance Name:** oaamdb1
- **Port:** 1521

Click **Add** again to add the second database host and enter the following information:

- **Host Name:** OAAMDBHOST2
- **Instance Name:** oaamdb2

- **Port:** 1521

Deselect this data source. Select the next data source, **OAAM Server**, and enter the following information:

- **Data Source:** OAAM Server
- **Service Name:** oaam.mycompany.com
- **User Name:** OAAM\_OAAM (assuming OAAM was used as the RCU prefix)
- **Password:** Password for the OAAM\_OAAM account.

In the top right box, click **Add** to add an Oracle RAC host. Enter the following information:

- **Host Name:** OAAMDBHOST1
- **Instance Name:** oaamdb1
- **Port:** 1521

Click **Add** again to add the second database host and enter the following information:

- **Host Name:** OAAMDBHOST2
- **Instance Name:** oaamdb2
- **Port:** 1521

Deselect this data source. Select the next data source, **OWSM MDS Schema**, and enter the following information:

- **Data Source:** OWSM MDS Schema
- **Service Name:** oaam.mycompany.com
- **User Name:** OAAM\_MDS (assuming OAAM was used as the RCU prefix)
- **Password:** Password for the OAAM\_MDS account.

In the top right box, click **Add** to add an Oracle RAC host. Enter the following information:

- **Host Name:** INFRADBHOST1
- **Instance Name:** oaamdb1
- **Port:** 1521

Click **Add** again to add the second database host and enter the following information:

- **Host Name:** INFRADBHOST2
- **Instance Name:** oaamdb2
- **Port:** 1521

Deselect this data source.

Click **Next**.

12. In the Test Component Schema screen, the configuration wizard attempts to validate the data source.

If the data source validation succeeds, click **Next**.

If it fails, click **Previous**, correct the issue, and try again.

13. In the Select Optional Configuration screen, select:

- **Administration Server**
- **Managed Server Clusters and Machines**

Click **Next**.

14. In the Customize Server and Cluster Configuration screen, select **Yes**, and click **Next**.

15. In the Configure the Administration Server screen, enter the following values:

- **Name:** AdminServer
- **Listen Address:** Enter the virtual host name of the Administration Server, for example: OAAMHOST1.mycompany.com
- **Listen Port:** 7001
- **SSL listen port:** Not applicable
- **SSL enabled:** leave unchecked

Click **Next**.

16. On the Configure Managed Servers screen, create two entries for each OAAMHOST in the topology, that is, one for the OAAM Server and one for the OAAM Admin Server.

Select the OAAM\_SERVER entry and change the entry to the following values:

- **Name:** WLS\_OAAM1
- **Listen Address:** OAAMHOST1.mycompany.com
- **Listen Port:** 14300
- **SSL Listen Port:** 14301
- **SSL Enabled:** Selected

For the second OAAM\_SERVER, click **Add** and supply the following information:

- **Name:** WLS\_OAAM2
- **Listen Address:** OAAMHOST2.mycompany.com
- **Listen Port:** 14300
- **SSL Listen Port:** 14301
- **SSL Enabled:** Selected

Select the OAAM\_ADMIN\_SERVER entry and change the entry to the following values:

- **Name:** WLS\_OAAM\_ADMIN1
- **Listen Address:** OAAMHOST1.mycompany.com
- **Listen Port:** 14200
- **SSL Listen Port:** 14201
- **SSL Enabled:** Selected

For the second OAAM\_ADMIN\_SERVER, click **Add** and supply the following information:

- **Name:** WLS\_OAAM\_ADMIN2

- **Listen Address:** OAAMHOST2.mycompany.com
- **Listen Port:** 14200
- **SSL Listen Port:** 14201
- **SSL Enabled:** Selected

Leave all other fields at the default settings.

Click **Next**.

17. In the Configure Clusters screen, create a cluster by clicking **Add**.

Enter name: OAAM\_Cluster

Create a second cluster by clicking **Add**.

Enter name: OAAM\_Admin\_Cluster

Leave all other fields at the default settings.

Click **Next**.

18. On the Assign Servers to Clusters screen, associate the managed servers with the cluster, as follows:

- Click the cluster name **OAAM\_Cluster** in the right window.
- Click the managed server **WLS\_OAAM1**, then click the arrow to assign it to the cluster.
- Repeat for managed server **WLS\_OAAM2**.

Then:

- Click the cluster name **OAAM\_Admin\_Cluster** in the right window.
- Click the managed server **WLS\_OAAM\_ADMIN1**, then click the arrow to assign it to the cluster.
- Repeat for managed server **WLS\_OAAM\_ADMIN2**.

Click **Next**.

19. On the Configure Machines screen, create a machine for each host in the topology. Click the UNIX tab if your hosts use a UNIX-based operating system. Otherwise, click the Machines tab. Supply the following information:

- **Name:** Name of the host. The best practice is to use the DNS name (OAAMHOST1)
- **Node Manager Listen Address:** The DNS name of the machine (OAAMHOST1.mycompany.com)
- **Node Manager Port:** A port for Node Manager to use.

Click **Next**.

20. In the Assign Servers to Machines screen, indicate which managed servers will run on the machines just created.

For OAAMHOST1:

- Click the machine **OAAMHOST1** in the right window.
- Click the managed server **WLS\_OAAM1** in the left window.
- Click the arrow to assign the managed server to the host **OAAMHOST1**.
- Click the managed server **WLS\_OAAM\_ADMIN1** in the left window.



- Click the arrow to assign the managed server to the host **OAAMHOST1**.

For **OAAMHOST2**:

- Click the machine **OAAMHOST2** in the right window.
- Click the managed server **WLS\_OAAM2** in the left window.
- Click the arrow to assign the managed server to the host **OAAMHOST2**.
- Click the managed server **WLS\_OAAM\_ADMIN2** in the left window.
- Click the arrow to assign the managed server to the host **OAAMHOST2**.

Click **Next**.

21. On the Configuration Summary screen, click **Create** to begin the creation process.

When prompted, on Linux and UNIX installations, execute the script `oracleRoot.sh` as the root user.

### 8.12.3.5 Creating `boot.properties` for the Administration Server on **OAAMHOST1**

This section describes how to create a `boot.properties` file for the Administration Server on **OAAMHOST1**. The `boot.properties` file enables the Administration Server to start without prompting for the administrator username and password.

Follow these steps to create the `boot.properties` file:

1. Start the Admin Server.
2. On **OAAMHOST1**, go the following directory:

```
MW_HOME/user_projects/domains/domainName/servers/AdminServer/security
```

For example:

```
cd /u01/app/oracle/product/fmw/user_
projects/domains/IDMDomain/servers/AdminServer/security
```

3. Use a text editor to create a file called `boot.properties` under the `security` directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```

---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

4. Stop the Administration Server if it is running.

See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.

5. Start the Administration Server on **OAAMHOST1** using the `startWebLogic.sh` script located under the `MW_HOME/user_projects/domains/domainName/bin` directory.

6. Validate that the changes were successful by opening a web browser and accessing the following pages:

- WebLogic Server Administration Console at:  
`http://oaamhost1.mycompany.com:7001/console`
- Oracle Enterprise Manager Fusion Middleware Control at:  
`http://oaamhost1.mycompany.com:7001/em`

Log into these consoles using the `weblogic` user credentials.

#### 8.12.3.6 Create the Oracle Adaptive Access Manager Administration User

Create an Administrative user, which will be used to access the OAAM console. To do this, perform these steps:

1. Log into the WebLogic Administration Console using this URL:  
`http://oaamhost1.mycompany.com:7001/console`
2. From the **Domain Structure** menu, click **Security Realms** and then **myrealm**.
3. Click the Users and Groups tab.
4. Click the Users sub tab.
5. Click **New**.
6. Enter these values:
  - **Name:** Name for the user, for example: `oaam_admin`
  - **Provider:** Default Authenticator
  - **Password/Confirmation:** Password for userClick **OK**.
7. Once the user has been created, click the user name.
8. Click on the Groups sub tab.
9. Assign the user to the following groups:
  - `OAAMCSRGroup`
  - `OAAMCSRManagerGroup`
  - `OAAMInvestigatorGroup`
  - `OAAMInvestigationManagerGroup`
  - `OAAMRuleAdministratorGroup`
  - `OAAMEnvAdminGroup`
  - `OAAMSOAPServicesGroup`Click **Save**.

#### 8.12.3.7 Start OAAMHOST1

Now you will start OAAMHOST1.

This section describes the steps for starting OAAMHOST1.

**8.12.3.7.1 Create the Node Manager Properties File on OAAMHOST1** Before you can start managed servers from the console, you must create a Node Manager property file. You do this by running the script `setNMProps.sh`, which is located in the `MW_HOME/oracle_common/common/bin` directory. For example:

```
OAAMHOST1> $MW_HOME/oracle_common/common/bin/setNMProps.sh
```

**8.12.3.7.2 Start Node Manager** Start Node Manager by issuing the following command:

```
OAAMHOST1> $MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
```

**8.12.3.7.3 Start Oracle Adaptive Access Manager on OAAMHOST1** To start Oracle Adaptive Access Manager on OAAMHOST1, follow these steps:

1. Log into the WebLogic Administration Console using this URL:

```
http://oaamhost1.mycompany.com:7001/console
```

2. Supply the WebLogic administrator username and password.
3. Select **Environment - Servers** from the **Domain Structure** menu.
4. Click the Control tab.
5. Click the servers **WLS\_OAAM1** and **WLS\_OAAM\_ADMIN1**.
6. Click **Start**.
7. Click **OK** to confirm that you want to start the servers.

### 8.12.3.8 Validating OAAMHOST1

Validate the implementation by connecting to the OAAM Administration Server at the following URL:

```
http://OAAMHOST1.mycompany.com:14200/oaam_admin
```

The implementation is valid if the OAAM Admin console login page is displayed and you can login using the `oaamadmin` account you created in [Section 8.12.3.6, "Create the Oracle Adaptive Access Manager Administration User."](#)

Validate the implementation by connecting to the OAAM Server at:

```
http://OAAMHOST1.mycompany.com:14300/oaam_server
```

The implementation is valid if the OAAM Server login page is displayed.

### 8.12.3.9 Configure Oracle Adaptive Access Manager on OAAMHOST2

Once the configuration has succeeded on OAAMHOST1, you can propagate it to OAAMHOST2. You do this by packing the domain using the `pack` script on OAAMHOST1, and unpacking the domain using the `unpack` script on OAAMHOST2.

Both scripts reside in the `MW_HOME/oracle_common/common/bin` directory.

On OAAMHOST1, enter:

```
pack.sh -domain=$MW_HOME/user_projects/domains/IDM_Domain \
 -template=/tmp/idm_domain.jar -template_name="OAAM Domain" -managed=true
```

This creates a file called `idm_domain.jar` in the `/tmp` directory. Copy this file to OAAMHOST2.

On OAAMHOST2, enter:

```
unpack.sh -domain=$MW_HOME/user_projects/domains/IDM_Domain \
-template=/tmp/idm_domain.jar
```

### 8.12.3.10 Start OAAMHOST2

Now you will start OAAMHOST2.

This section describes the steps for starting OAAMHOST2.

**8.12.3.10.1 Create the Node Manager Properties File on OAAMHOST2** Before you can start managed servers from the console, you must create a Node Manager property file. You do this by running the script `setNMProps.sh`, which is located in the `MW_HOME/oracle_common/common/bin` directory. For example:

```
OAAMHOST1> $MW_HOME/oracle_common/common/bin/setNMProps.sh
```

**8.12.3.10.2 Start Node Manager** Start Node Manager by issuing the following command:

```
OAAMHOST2> $MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
```

**8.12.3.10.3 Start Oracle Adaptive Access Manager on OAAMHOST2** To start Oracle Adaptive Access Manager on OAAMHOST2, follow these steps:

1. Log into the WebLogic Administration Console using this URL:  
`http://oaamhost2.mycompany.com:7001/console`
2. Supply the WebLogic administrator username and password.
3. Select **Environment - Servers** from the **Domain Structure** menu.
4. Click the Control tab.
5. Click the servers **WLS\_OAAM2** and **WLS\_OAAM\_ADMIN2**.
6. Click **Start**.
7. Click **OK** to confirm that you want to start the servers.

### 8.12.3.11 Validating OAAMHOST2

Validate the implementation by connecting to the OAAM Administration Server at the following URL:

```
http://OAAMHOST2.mycompany.com:14200/oaam_admin
```

The implementation is valid if OAAM Admin console login page is displayed and you can login using the `oaamadmin` account you created in [Section 8.12.3.6, "Create the Oracle Adaptive Access Manager Administration User."](#)

Validate the implementation by connecting to the OAAM Server at:

```
http://OAAMHOST2.mycompany.com:14300/oaam_server
```

The implementation is valid if the OAAM Server login page is displayed.

### 8.12.3.12 Configure Oracle Adaptive Access Manager to Work with Oracle HTTP Server

This section provides steps for configuring Oracle Adaptive Access Manager to work with the Oracle HTTP Server.

**8.12.3.12.1 Update Oracle HTTP Server Configuration** On WEBHOST1 and WEBHOST2, create a file named `oaam.conf` in the following directory:

`ORACLE_INSTANCE/config/OHS/ohs1/moduleconf`

Create the file with the following lines:

```
NameVirtualHost *:7777
<VirtualHost *:7777>

 ServerName oaam.mycompany.com:7777
 ServerAdmin you@your.address
 RewriteEngine On
 RewriteOptions inherit

 <Location /oaam_server>
 SetHandler weblogic-handler
 WebLogicCluster oaamhost1.mycompany.com:14300,oaamhost2.mycompany.com:14300
 </Location>

 <Location /oaam_admin>
 SetHandler weblogic-handler
 WebLogicCluster oaamhost1.mycompany.com:14200,oaamhost2.mycompany.com:14200
 WebLogicPort 7001
 </Location>

</VirtualHost>
```

**8.12.3.12.2 Restart Oracle HTTP Server** Restart the Oracle HTTP Server on WEBHOST1 and WEBHOST2:

```
WEBHOST1>opmnctl stopall
WEBHOST1>opmnctl startall

WEBHOST2>opmnctl stopall
WEBHOST2>opmnctl startall
```

**8.12.3.12.3 Change Host Assertion in WebLogic** Because the Oracle HTTP Server acts as a proxy for WebLogic, by default certain CGI environment variables are not passed through to WebLogic. These include the host and port. You must tell WebLogic that it is using a virtual site name and port so that it can generate internal URLs appropriately.

To do this, log into the WebLogic Administration Console at:

`http://oaamhost1.mycompany.com:7001/console`

Then perform these steps:

1. Select **Clusters** from the home page, or select **Environment > Clusters** from the **Domain Structure** menu.
2. Click **Lock and Edit** in the Change Center window to enable editing.

3. Click the Cluster Name (**oaam\_cluster**).
4. In the General tab, set WebLogic Plug in to **Enabled** by checking the box in the Advanced Properties section.
5. Click **Save**.
6. Select HTTP and enter the following values:
  - **Frontend Host:** `oaam.mycompany.com`
  - **Frontend HTTP Port:** `7777`
  - **Frontend HTTPS Port:** Set to SSL port on the load balancer or the Oracle HTTP Server SSL port if using SSL communication.

This ensures that any HTTPS URLs created from within WebLogic are directed to port 443 or 80 on the load balancer.
7. Click **Save**.
8. Select **Clusters** from the home page, or select **Environment > Clusters** from the **Domain Structure** menu.
9. Click the Cluster Name (**oaam\_admin\_cluster**).
10. In the General tab, set WebLogic Plug in to **Enabled** by checking the box in the Advanced Properties section.
11. Click **Save**.
12. Select HTTP and enter the following values:
  - **Frontend Host:** `oaam.mycompany.com`
  - **Frontend HTTP Port:** `7777`
  - **Frontend HTTPS Port:** Set to SSL port on the load balancer or the Oracle HTTP Server SSL port if using SSL communication.
13. Click **Save**.
14. Click **Activate Changes** in the Change Center window to save the changes.
15. Restart managed servers WLS\_OAAM1, WLS\_OAAM2, WLS\_OAAM\_ADMIN1 and WLS\_OAAM\_ADMIN2 as follows:
  - a. Log into the WebLogic Administration Console using this URL:  
`http://oaamhost1.mycompany.com:7001/console`
  - b. Supply the WebLogic administrator username and password.
  - c. Select **Environment - Servers** from the **Domain Structure** menu.
  - d. Click the Control tab.
  - e. Click the servers **WLS\_OAAM1**, **WLS\_OAAM2**, **WLS\_OAAM\_ADMIN1**, and **WLS\_OAAM\_ADMIN2**.
  - f. Click **Shutdown - Force shutdown now**.
  - g. Click **Yes** to confirm that you want to stop the servers.
  - h. Select **Environment - Servers** from the **Domain Structure** menu.
  - i. Click the Control tab.
  - j. Click the servers **WLS\_OAAM1**, **WLS\_OAAM2**, **WLS\_OAAM\_ADMIN1**, and **WLS\_OAAM\_ADMIN2**.

- k. Click **Start**.
  - l. Click **Yes** to confirm that you want to start the servers.
16. Restart the Administration Server.

### 8.12.3.13 Validating the Oracle Adaptive Access Manager Configuration

Log into the Oracle Adaptive Access Manager Administration Console at `http://oaam.mycompany.com:7777/oaam_admin` using the `oaamadmin` account you created.

Also, log into the Oracle Adaptive Access Manager server at `http://oaam.mycompany.com:7777/oaam_server` using the `oaamadmin` account and the password test.

### 8.12.3.14 Scaling Up and Scaling Out the Oracle Adaptive Access Manager Topology

This section describes how to scale up and scale out an Oracle Adaptive Access Manager high availability topology. Perform a scale up operation to add a new Oracle Adaptive Access Manager managed server instance to a node already running one or more server instances. Perform a scale out operation to add a new Oracle Adaptive Access Manager managed server instance to a new node.

**8.12.3.14.1 Scaling Up Oracle Adaptive Access Manager** To scale up OAAM, use the same procedure for both the OAAM server and the OAAM Administration Server.

Log in to the Oracle WebLogic Server console at:

`http://oaamhost1.mycompany.com:7001/console`. Then proceed as follows:

1. From the Domain Structure window of the Oracle WebLogic Server Administration Console, expand the **Environment** node and then **Servers**. The Summary of Servers page appears.
2. Click **Lock & Edit** from the Change Center menu.
3. Select an existing server on the host that you want to extend, for example: `WLS_OAAM1` or `WLS_OAAM_ADMIN1`.
4. Click **Clone**.
5. Enter the following information:
  - **Server Name:** A new name for the server, for example: `WLS_OAAM3`.
  - **Server Listen Address:** The name of the host on which the managed server will run.
  - **Server Listen Port:** The port the new managed server will use. This port must be unique within the host.
6. Click **OK**.
7. Click the newly-created server `WLS_OAAM3`.
8. Set the SSL listen port. This should be unique on the host that the managed server will be running on.
9. Click **Save**.
10. Disable host name verification for the new managed server. Before starting and verifying the `WLS_OAAM3` managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for

the communication between the Oracle WebLogic Administration Server and the Node Manager in OAAMHOST $n$ .

If the source server from which the new one was cloned had already disabled hostname verification, these steps are not required, as the hostname verification settings were propagated to the cloned server. To disable host name verification:

- a. In the Oracle Fusion Middleware Enterprise Manager Console, select **Oracle WebLogic Server Administration Console**.
  - b. Expand the **Environment** node in the Domain Structure pane.
  - c. Click **Servers**. The Summary of Servers page appears.
  - d. Select **WLS\_OAAM3** in the Names column of the table. The Settings page for server appears.
  - e. Click the **SSL** tab.
  - f. Click **Advanced**.
  - g. Set **Hostname Verification** to **None**.
  - h. Click **Save**.
11. Click **Activate configuration** from the Change Center menu.

**8.12.3.14.2 Scaling Out Oracle Adaptive Access Manager** Scale out is very similar to scale up, but first requires the software to be installed on the new node. Proceed as follows:

1. Install Oracle WebLogic Server on the new host as described in [Section 8.12.3.3, "Installing Oracle WebLogic Server."](#)
2. Install Oracle Fusion Middleware Identity Management components on the new host as described in [Section 8.12.3.4, "Install and Configure the Oracle Adaptive Access Manager Application Tier."](#)
3. Log in to the WebLogic console at `http://oaamhost1.mycompany.com:7001/console`.
4. From the Domain Structure pane of the Oracle WebLogic Server Administration Console, expand the **Environment** node and then **Servers**. The Summary of Servers page appears.
5. Click **Lock & Edit** from the Change Center menu.
6. Select an existing server on the host you want to extend, for example: **WLS\_OAAM1** or **WLS\_OAAM\_ADMIN1**.
7. Click **Clone**.
8. Enter the following information:
  - **Server Name:** A new name for the server, for example: **WLS\_OAAM3**
  - **Server Listen Address:** The name of the host on which the managed server will run.
  - **Server Listen Port:** The port the new managed server will use. This port must be unique within the host.
9. Click **OK**.
10. Click the newly-created server **WLS\_OAAM3**.
11. Set the SSL listen port. This should be unique on the host that the managed server will be running on.



**12. Click Save.**

- 13.** Disable host name verification for the new managed server. Before starting and verifying the WLS\_OAAM3 managed server, you must disable host name verification. You can re-enable it after you have configured server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in OAAMHOST $n$ .

If the source server from which the new one was cloned had already disabled hostname verification, these steps are not required, as the hostname verification settings were propagated to the cloned server. To disable host name verification:

- a. In the Oracle Fusion Middleware Enterprise Manager Console, select **Oracle WebLogic Server Administration Console**.
- b. Expand the **Environment** node in the Domain Structure pane.
- c. Click **Servers**. The Summary of Servers page appears.
- d. Select **WLS\_OAAM3** in the Names column of the table. The Settings page for server appears.
- e. Click the **SSL** tab.
- f. Click **Advanced**.
- g. Set **Hostname Verification** to None.
- h. Click **Save**.

**14. Click Activate configuration from the Change Center menu.****15. Pack the domain on OAAMHOST1 using the command:**

```
pack.sh -domain=ORACLE_BASE/admin/IDM_Domain/aserver/IDM_Domain -template
=/tmp/idm_domain.jar -template_name="OAAM Domain" -managed=true
```

The `pack.sh` script is located in `MW_HOME/oracle_common/common/bin`.

**16. Unpack the domain on the new host using the command:**

```
unpack.sh -domain=ORACLE_BASE/admin/IDM_Domain/msserver/IDM_Domain
-template=/tmp/idm_domain.jar -template_name="OAAM Domain" -app_dir=ORACLE_
BASE/admin/IDM_Domain/msserver/applications
```

The `unpack.sh` script is located in `MW_HOME/oracle_common/common/bin`.

- 17.** Before you can start managed servers from the console, you must create a node manager properties file on OAAMHOST2 by running the script `setNMProps.sh`. The `setNMProps.sh` script is located in `MW_HOME/oracle_common/common/bin`. Type:

```
OAAMHOST2> $MW_HOME/oracle_common/common/bin/setNMProps.sh
```

**18. Start Node Manager and the new managed server on the new host**

- 19.** Now that the new managed server has been created and started, the web tier will start to direct requests to it. Best practice, however, is to inform the web server that the new managed server has been created.

You do this by updating the file `oaam.conf` on each of the web tiers. This file resides in the directory: `ORACLE_INSTANCE/config/OHS/component_name/moduleconf`.

Add the new server to the `WebLogicCluster` directive in the file. For example, change:

```
<Location /oaam_admin>
 SetHandler weblogic-handler
 WebLogicCluster oaamhost1.mycompany.com:14200,oaamhost2.mycompany.com:14200
</Location>
```

to:

```
<Location /oaam_admin>
 SetHandler weblogic-handler
 WebLogicCluster
oaamhost1.mycompany.com:14200,oaamhost2.mycompany.com:14200,oaamhost3.mycompany
.com:14300
</Location>
```

## 8.13 Oracle Identity Federation High Availability

This section provides an introduction to Oracle Identity Federation and describes how to design and deploy a high availability environment for Oracle Identity Federation.

### 8.13.1 Oracle Identity Federation Component Architecture

Oracle Identity Federation is a self-contained, standalone federation server that enables single sign-on and authentication in a multiple-domain identity network and supports the broadest set of federation standards. This allows users to federate in heterogeneous environments and business associations, whether or not they have implemented other Oracle Identity Management products in their solution set.

It can be deployed as a multi-protocol hub acting as both an Identity Provider (IdP) and Service Provider (SP).

Acting as an SP, Oracle Identity Federation enables you to manage your resources while off loading actual authentication of users to an IdP, without having to synchronize users across security domains out of band. Once authenticated at the IdP, the SP can allow or deny access to users for the SP's applications depending upon the local access policies.

If a user no longer has an account with the IdP, the federation is terminated and cross-domain single sign-on for that user is automatically disabled. As an IdP, Oracle Identity Federation allows you to manage and authenticate local users based upon federated requests from trusted providers.

Some key features of Oracle Identity Federation include support for:

- Multiple leading federation protocols such as SAML 1.x, SAML 2.0, WS-Federation, SAML 1.x/2.0 attribute sharing functionality, Liberty ID-FF 1.1/ Liberty ID-FF 1.2.
- Cross-protocol single sign-on and sign-out
- Affiliations, by allowing service providers to share their federation information, reducing the number of federations
- X.509 certificate validation
- Integration with Oracle Access Manager and Oracle Single Sign-On
- Integration with Oracle Internet Directory and support for a range of authentication engines, user data repositories and relational databases
- Implementing cross-site access and authentication in an environment containing both identity providers and service providers

- The ability to configure, enable, and disable external sites
- Accessing applications at destination sites using a single sign-on

**Figure 8–20 Oracle Identity Federation Non-High Availability Architecture**




Figure 8–20 shows Oracle Identity Federation deployed on Oracle WebLogic Server in a non-high availability architecture and its relationship to other federation components.

Figure 8–20 shows Oracle Identity Federation as a member of Federation between other identity providers (IdP) and service providers (SP).

The Oracle Identity Federation authentication service can be configured:

- To enable single sign-on access to resources protected by Oracle Single Sign-On (with the Oracle Internet Directory user repository) or Oracle Access Manager (with various repositories). Resources can include Oracle Collaboration Suite, Oracle E-Business Suite, PeopleSoft modules, and so on.
- To interact with Identity and Access Management solutions such as LDAP directories, RDBMS databases, and Oracle Access Manager.

---

**Note:** In an environment where Oracle Identity Federation and Oracle Single Sign-On both protect resources, you can configure either component to serve as the authentication mechanism when a user requests access to a protected resource. Likewise, environments containing both Oracle Identity Federation and Oracle Access Manager provide similar functionality.

---

For more information about Oracle Identity Federation authentication, see *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

---

### 8.13.1.1 Oracle Identity Federation Component Characteristics

Oracle Identity Federation is Oracle's self-contained, standalone federation server, based on J2EE standards. It enables single sign-on and authentication in a

multiple-domain identity network and supports a broad set of federation standards such as SAML 1.x, SAML 2.0, WS-Federation, and SAML 1.x/2.0 attribute sharing functionality.

By default, Oracle Identity Federation is deployed as an externally staged application on Oracle WebLogic Server.

**8.13.1.1.1 Runtime Processes** Oracle Identity Federation is a J2EE application that is deployed on Oracle WebLogic Server as an externally staged application. The Oracle Identity Federation server is initialized when the WebLogic Server it is deployed on starts up.

The Oracle Identity Federation application accesses the Credential Store Framework, to get the credentials to connect to the back-end servers. Once this is complete, the Oracle Identity Federation server is running and is ready to accept and serve requests.

**8.13.1.1.2 Process Lifecycle** Oracle Identity Federation is deployed to an Oracle WebLogic Server as an externally managed application. By default, the Oracle WebLogic Server starts, stops, monitors and manages other lifecycle events for the Oracle Identity Federation application.

The application is initialized when the Oracle WebLogic Server it is deployed to starts up and it stops when the WebLogic Server is shut down.

WebLogic Node Manager can be configured to monitor the server process and restart it in case of failure

The Oracle Enterprise Manager Fusion Middleware Control is used to monitor as well as modify the configuration of the application.

### Starting and Stopping

Oracle Identity Federation lifecycle events can be managed using one or more of the command line tools and consoles listed below:

- Oracle WebLogic Scripting Tool (WLST)
- Oracle WebLogic Server Administration Console
- Oracle Enterprise Manager Fusion Middleware Control
- WebLogic Node Manager

**8.13.1.1.3 Request Flow** Users typically access applications in multiple domains through a corporate portal. For example, Alpha Corporation could have a Portal Server in place to manage Alpha's user logins, page personalization, and so on. The portal server might consist of homegrown logic running within an application server, or it might be a commercial product. Its partner, Beta Corporation, may serve its technical database application with a "MyBeta.com" type of portal. In that case, each company would operate its own portal server.

The processing flow is as follows:

1. The user logs into the Alpha portal whose access is being managed by a web access management (WAM) system such as Oracle Access Manager or Oracle Single Sign-On.
2. The user initiates a request by clicking on a resource that is actually hosted by Beta Corporation.
3. The Oracle Identity Federation instance at the Alpha portal, acting as the identity provider (IdP), will send the user information to the WAM system.

4. The WAM system will create a session after mapping a user to an identity in its local identity store.
5. The WAM system will return the successful response and the session token to the Oracle Identity Federation IdP server at the Alpha portal.
6. Using the above information, the IdP at the Alpha portal creates a SAML identity assertion and signs it using its private signing key. This response is sent to the Oracle Identity Federation instance acting as a Service Provider (SP) at Beta Corporation.
7. The Oracle Identity Federation server acting as SP at Beta Corporation verifies the signed response using the IdP's public certificate associated with its signing key.
8. The Oracle Identity Federation service provider at Beta Corporation extracts the assertion, and creates a user session for the assertion after mapping the user session to its local authorization system.
9. The Oracle Identity Federation service provider sends the user's browser a redirect to the requested resource.
10. The user's browser sends a request to the target resource over the user session created by the service provider.

**8.13.1.1.4 Configuration Artifacts** The configuration of the Oracle Identity Federation server is managed by MBeans. The configuration data for Oracle Identity Federation is stored in three main files:

- `config.xml`: contains server-wide configuration
- `cot.xml`: stores provider-specific configuration
- `datastore.xml`: stores back-end data store configuration

The Oracle Identity Federation application is deployed as an EAR file by the Oracle Identity Management 11g Installer. The Oracle Identity Federation application is located under the `ORACLE_HOME/fed/install/oif.ear` directory.

The Oracle Identity Federation configuration artifacts are located under the `DOMAIN_HOME/config/fmwconfig/servers/serverName/applications/OIF_11.1.1.2.0/configuration` directory on the managed servers where the application is deployed.

Using Oracle Enterprise Manager Fusion Middleware Control to make configuration changes can help prevent inconsistent configuration across different nodes. This is especially true when the application is deployed in a high availability configuration. However, WLST scripts and JMX MBeans can also be used to configure the Oracle Identity Federation server.

**8.13.1.1.5 External Dependencies** The Oracle Identity Federation server is a J2EE application that is deployed on an Oracle WebLogic Managed Server. The Oracle Identity Federation server interacts with external repositories to store configuration data, runtime transient data (message, user session) and federation data and to authenticate users. The Oracle Identity Federation server requires these repositories to be available during initialization or during runtime. These are the external components required for the Oracle Identity Federation server to function:

- Oracle WebLogic Server
  - Administration Server
  - Managed Server

- Data Repositories
  - Message data store and user session data store (transient data store)
  - Configuration data store
  - User data store
  - Federation data store
- Authentication Engines
- Service Provider Engines

See the following subsections for more information about how these external components function with Oracle Identity Federation.

### Oracle WebLogic Server

The Oracle Identity Federation server is a J2EE application that is deployed on an Oracle WebLogic Managed Server. The server is administered and managed using the Oracle Enterprise Manager Fusion Middleware Control.

The application cannot start if the Managed Server is not running. If the Administration Server is not available, the server continues to run in its current state but changes cannot be made to its configuration.

### Data Repositories

The Oracle Identity Federation server uses the various external data repositories to store configuration, user session, message and federation data. The server requires these data stores to be available during initialization, runtime or both. Details of the repositories are shown below:

- Message data store and user session data store (transient data store):

The Oracle Identity Federation server uses the message data store and the user session data store for storing transient data like federation protocol/session state. The message data store together with the user session data store is also referred to as a the transient data store.

Depending on the deployment option for the Oracle Identity Federation server, the transient data can either be stored in memory or in a relational database.

[Table 8–11](#) shows the relationship between the deployment option and the required transient data store type:

**Table 8–11 Transient Data Store Types for Oracle identity Federation Deployment Options**

| Deployment Option                    | Store Type                |
|--------------------------------------|---------------------------|
| Non-high Availability/<br>Standalone | In Memory Store           |
|                                      | Relational Database Store |
| High Availability                    | Relational Database Store |

- Configuration data store:

The Oracle Identity Federation application uses the configuration data store to store its configuration artifacts. Depending on the deployment option of the Oracle Identity Federation server, the configuration store can either be stored in an XML file or in a relational database.

Table 8–12 shows the relationship between the deployment option and the required configuration data store type:

**Table 8–12 Configuration Data Store Types for Oracle Identity Federation Deployment Option**

| Deployment Option                    | Store Type                       |
|--------------------------------------|----------------------------------|
| Non-high Availability/<br>Standalone | XML<br>Relational Database Store |
| High Availability                    | Relational Database Store        |

The Oracle Identity Federation application connects to the configuration data store to retrieve its configuration data during the start up process. The application fails to start up if the configuration data store is unavailable.

- User data store:

The Oracle Identity Federation server uses the user data store to locate users after local authentication or after processing an incoming SAML Assertion and to retrieve user specific attributes. The user data repository can be either a relational database or a LDAP repository.

The Oracle Identity Federation server is certified to work with LDAP repositories such as Oracle Internet Directory, Sun Java System Directory Server, and Microsoft Active Directory, as well as with a relational database.

The role played by the user data repository depends on whether Oracle Identity Federation will be configured as an identity provider (IdP) or a service provider (SP).

- When configured as an IdP:
  - \* Oracle Identity Federation uses the repository to verify user identities and to build protocol assertions.
- When configured as an SP:
  - \* Oracle Identity Federation uses the repository to map information in received assertions to user identities at the destination, and subsequently to authorize users for access to protected resource.
  - \* When creating a new federation, Oracle Identity Federation uses the repository to identify the user and link the new federation to that user's account.

The Oracle Identity Federation application does not require the user data store to be available while starting up. The user data store is required at runtime.

---

**Note:** Oracle Identity Federation is only certified to work with Oracle Database versions 10.2.0.4 and higher or Oracle Database 11.1.0.7 and higher.

---

- Federation Data Store:

The federation data store can be XML, RDBMS or LDAP. In high availability mode, use only RDBMS or LDAP so that the federation data store is available to all members of the cluster.

The Oracle Identity Federation server uses the federation data store to persist user federation records. Identity federation is the linking of two or more accounts that a principal may hold with one or more identity providers or service providers within a given circle of trust. When users federate the otherwise isolated accounts they have with businesses (known as their local identities), they are creating a relationship between two entities, an association comprising any number of service providers and identity providers.

The Oracle Identity Federation server is certified with industry-standard LDAP repositories including Oracle Internet Directory, Sun Java System Directory Server, and Microsoft Active Directory or a relational database.

Table 8–13 shows the relationship between the deployment option and the required federation data store type:

**Table 8–13 Federation Data Store Types for Oracle Identity Federation Deployment Option**

| Deployment Option                    | Store Type                               |
|--------------------------------------|------------------------------------------|
| Non-high Availability/<br>Standalone | None or XML                              |
| High Availability                    | None, LDAP, or Relational Database Store |

The federation data store is required at runtime only if persistent name identifier formats are used during protocol exchanges. The federation data store is optional at runtime if non-opaque name identifier formats are used, such as email address.

---

**Note:** Oracle Identity Federation is only certified to work with Oracle Database versions 10.2.0.4 and higher or Oracle Database 11.1.0.7 and higher.

---

For more information about the federation data store for Oracle Identity Federation, see the "Federation Data Store" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

- Authentication engines:

Oracle Identity Federation IdP and SP protocol operations, such as single sign-on, federation creation, federation termination, and NameID registration, require users to be authenticated. The Oracle Identity Federation authentication module handles user authentication. The authentication module can perform two distinct roles in user authentication:

- When Oracle Identity Federation is deployed as an Identity Provider, the authentication module acts as a local authentication mechanism. The authentication module can delegate authentication or directly interact with a number of repositories and identity management solutions.
- When Oracle Identity Federation is deployed as a service provider, it uses federation protocols to have the user authenticated at a peer identity provider. Oracle Identity Federation then forwards the user to the authentication module, which propagates and creates an authenticated user session in the deployed identity management solution at the SP. In turn, this enables access to the requested protected resource.



The Oracle database, Oracle Internet Directory, Sun Java System Directory Server, and Microsoft Active Directory are some examples of supported RDBMS and LDAP repositories.

Oracle Access Manager and Oracle Single Sign On are examples of supported repositories and Identity Management solutions.

The authentication module only requires the external authentication engines to be available at runtime.

- Service provider engines:

A Service Provider Integration Engine (SP Integration Engine) is used to create a user authenticated session at the Identity and Access Management (IAM) server. Oracle Identity Federation is shipped with an SP engine includes several internal plug-ins that allow it to interact with different IAM servers, such as:

- Oracle Single Sign-On
- Oracle Access Manager
- Oracle Identity Federation Test Application

Oracle Identity Federation also provides a framework to integrate with third party IAM frameworks; the customized SP Integration Module will interact with Oracle Identity Federation using internal J2EE Servlet forwards, and it will communicate with the third party IAM system to create the user authenticated session.

Any custom SP Engine modules should be deployed on each Managed Server in the cluster.

**8.13.1.1.6 Oracle Identity Federation Log File Location** Oracle Identity Federation is a J2EE application deployed on Oracle WebLogic Server. All server related logs messages are logged in the server log file and all Oracle Identity Federation specific messages are logged into the diagnostic log file of the Oracle WebLogic Server where the application is deployed.

The default location of the Oracle WebLogic Server log file is:

```
WEBLOGIC_SERVER_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

The Oracle Identity Federation log file is named `serverName-diagnostic.log`. For example, if the Oracle WebLogic Server name is `wls_oif1`, then the log file name is `wls_oif1-diagnostic.log`.

Use the Oracle Enterprise Manager Fusion Middleware Control to view the log files.

## 8.13.2 Oracle Identity Federation High Availability Concepts

This section provides an introduction to Oracle Identity Federation concepts and describes how to design and deploy a high availability environment for Oracle Identity Federation.

Listed below are a few guidelines to follow when deploying Oracle Identity Federation in a highly availability configuration:

- The transient and configuration data should always be stored in a shared relational database. This is required when:
  - Oracle Identity Federation acts as an IdP and the Artifact SSO profile is used. In this case, the user might interact with IdP #1 where the assertion will be created, and later when the artifact will de-referenced, the SP will connect to

IdP #2. If the two Oracle Identity Federation servers are not sharing the same message store, then IdP #2 will not be able to locate the assertion for the artifact created by IdP #1.

- Oracle Identity Federation acts as an SP and the POST SSO profile is used. In the POST profile, the assertion is being carried by the user's browser, thus opening the possibility of that assertion being re-submitted to the Oracle Identity Federation SP. Because of potential replay attacks, the Oracle Identity Federation SP Server keeps a trace of the assertion it received, so that no assertion can be used twice.
  - Oracle Identity Federation acts as an Authentication Query Authority, where the Oracle Identity Federation Server will answer queries from remote providers by returning assertions representing the existing sessions for a specific user at the Oracle Identity Federation Server. For this reason, the Oracle Identity Federation servers will need to share the transient session stores.
  - Oracle Identity Federation acts as an IdP/Attribute Authority and has the Assertion ID Responder functionalities enabled. This feature allows remote providers to query the IdP to retrieve assertions that were already created during previous SAML flows. For that reason, the Responder needs to have access to a store containing all the assertions created.
- In high availability environments, when the Oracle Identity Federation Application is not front-ended by an Oracle HTTP Serve Instance, oracle recommends:
    - Enabling sticky sessions on the hardware load balancer that front-ends the Oracle Identity Federation Application.
    - Sticky sessions ensue that user requests to go to the same Oracle Identity Federation server for that session

---

**Note:** HTTP session state is used at runtime by Oracle Identity Federation when processing HTTP requests. The information stored in the HTTP session state is short-lived. The life of the information corresponds to the duration of the federation operation, such as single sign-on.

HTTP session replication replicates session information across nodes, is memory intensive, and is not recommended.

HTTP session replication is not enabled by default.

See [Section 8.2.5.4.1, "Load Balancers"](#) for a description of the features to enable in the load balancer that is required when you deploy Oracle Identity Management software in a high availability configuration.


---

- The Oracle Identity Federation server supports two types of HTTP Session replication:
  - In-memory session replication
  - JDBC session replication
- These are some characteristics of in-memory session replication:
  - All Managed Servers running the Oracle Identity Federation application in a cluster have to be up when the session is created.

- It can slow down performance.
- If a server receives a request before session replication is completed, the server throws an error.
- To avoid this, update the parameters listed below in the `config.xml` file located under the `DOMAIN_HOME/config/fmwconfig/servers/serverName/applications/OIF_11.1.1.2.0/configuration` directory. This must be done on all nodes running the Oracle Identity Federation application.
  - \* `sessionreplicationenabled`: Set this to **true**.
  - \* `sessionreplicationtimeout`: This is set to 2000 by default. Increase this if necessary.
- These are some characteristics of JDBC session replication:
  - Robust in nature.
  - Slower in performance than in-memory session replication due to the additional overhead of database calls.
  - The session persistence state is stored in a shared relational database.
  - To configure Oracle WebLogic Server for JDBC session persistence see "Using a Database for Persistent Storage (JDBC persistence)" in *Oracle Fusion Middleware Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server*.

### 8.13.2.1 Oracle Identity Federation High Availability Architecture

[Figure 8–21](#) shows the Oracle Identity Federation high availability architecture in an active-active configuration.

**Figure 8–21 Oracle Identity Federation in a High Availability Architecture**

In Figure 8–21, the application tier includes the IDMHOST1 and IDMHOST2 computers.

On IDMHOST1, the following installations have been performed:

- An Oracle Identity Federation instance has been installed in the WLS\_OIF1 Managed Server. The Oracle RAC database has been configured in a JDBC multi data source to protect the instance from Oracle RAC node failure.
- A WebLogic Server Administration Server has been installed. Under normal operations, this is the active Administration Server.

On IDMHOST2, the following installations have been performed:

- An Oracle Identity Federation instance has been installed in the WLS\_OIF2 Managed Server. The Oracle RAC database has been configured in a JDBC multi data source to protect the instance from Oracle RAC node failure.

The instances in the WLS\_OIF2 Managed Server on IDMHOST2 and the instances in the WLS\_OIF1 Managed Server on IDMHOST1 are configured as the CLUSTER\_OIF cluster.

- A WebLogic Server Administration Server has been installed. Under normal operations, this is the passive Administration Server. You make this Administration Server active if the Administration Server on IDMHOST1 becomes unavailable.

**8.13.2.1.1 Starting and Stopping the Cluster** In a high availability architecture, Oracle Identity Federation server is deployed on an Oracle WebLogic Cluster, which has at least two servers as a part of the cluster.

By default, Oracle WebLogic Server starts, stops, monitors and manages the various lifecycle events for the application. The Oracle Identity Federation application leverages the high availability features of the underlying Oracle WebLogic Clusters. In case of hardware or other failures, session state is available to other cluster nodes that can resume the work of the failed node.

In a high availability environment, WebLogic Node Manager is configured to monitor the Oracle WebLogic Servers. In case of failure, Node Manager restarts the WebLogic Server.

In a high availability environment, a hardware load balancer is used to load balance requests between the multiple Oracle Identity Federation instances. If one of the Oracle Identity Federation instances fails, the load balancer detects the failure and routes requests to the surviving instances.

In high availability environments, the state and configuration information is stored in a database that is shared by all the members of the cluster.

The surviving Oracle Identity Federation instances will continue to seamlessly process any unfinished transactions started on the failed instance since the state information is in the shared database and is available to all the members in the cluster.

When an Oracle Identity Federation instance fails, its database and LDAP connections are released. The surviving instances in the active-active deployment make their own connections to continue processing unfinished transactions on the failed instance.

You can use one of the following command line tools or consoles to manage the lifecycle events for the Oracle Identity Federation application:

- Oracle WebLogic Server Administration Console
- Oracle Enterprise Manager Fusion Middleware Control
- Oracle WebLogic Scripting Tool (WLST)

**8.13.2.1.2 Cluster-Wide Configuration Changes** Configuration changes made through one cluster member propagate automatically to all others because the configuration is stored in the shared database.

---

**Note:** HTTP session replication is disabled by default and not recommended. It replicates session information across nodes and is memory intensive. However, if your environment requires HTTP Session Replication to be enabled, follow the steps below:

---

To turn session replication on or off, make updates in the `weblogic.xml` file on all Managed Servers where Oracle Identity Federation is deployed:

1. Copy the file `ORACLE_HOME/fed/install/oif.ear` to a temporary location.
2. Extract the `META-INF/weblogic.xml` file from the `web.war` file contained in the `oif.ear` file
3. Update the parameter set `persistent-store-type` to `replicated_if_clustered`.
4. Save the `weblogic.xml` file
5. Re-package the Oracle Identity Federation application using the appropriate tools.
6. Copy the updated `oif.ear` file to the `ORACLE_HOME/fed/install` directory on all the nodes running Oracle Identity Federation.

7. Restart the updated Oracle Identity Federation application on all nodes in your environment running the application.
8. **Restart the managed servers.**

To disable HTTP Session Replication, follow the previous steps and update the parameter `set persistent-store-type` to `memory` in Step 3.

---

**Note:** You must perform these manual steps after updating your environment with every patch set, otherwise the session replication changes are lost.

---

### 8.13.2.2 High Availability Considerations for Integration with Oracle Access Manager

This section describes the steps to take when you are integrating Oracle Identity Federation with Oracle Access Manager in a high availability topology:

1. In addition to deploying Oracle Identity Federation in high availability mode, Oracle Access Manager should also be deployed in high availability mode.
2. The Oracle Access Server SDK must be installed on every Oracle Identity Federation server in the cluster. Oracle Identity Federation must be configured to reference the directory where the SDK is installed. If the SDK is installed in the Domain Home directory, then you can reference the SDK folder using a relative path from the Domain Home folder. If the SDK is installed elsewhere, Oracle Identity Federation will need to reference the SDK folder using an absolute path.

When Oracle Identity Federation is used in a high availability environment, it is recommended that the Access Server SDK be installed under the Domain Home folder, using the same directory name on all the computers where Oracle Identity Federation is installed. This is a requirement for Oracle Identity Federation high availability integration with Oracle Access Manager because all the Oracle Identity Federation instances will share the same configuration, specifically the directory where the Access Server SDK is installed. Using a relative path allows the Oracle Identity Federation instances to share the same configuration.

3. Follow the instructions for integrating Oracle Access Manager as an SP integration module in the "Integrate Oracle Access Manager as an SP Integration Module" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation* to integrate Oracle Access Manager with the SDK instance on each Oracle Identity Federation server.
4. Copy over the required files to the domain library and update the WebLogic Server startup script for each Oracle Identity Federation server to add the SDK jar file to the classpath and to set the `LD_LIBRARY_PATH` And `LD_ASSUME_KERNEL` environment variables, if required. See the "Update the Oracle WebLogic Server Environment" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation* for more information.

### 8.13.2.3 Oracle Identity Federation Prerequisites

Oracle Identity Federation requires the following components:

- Oracle JRockit SDK that is bundled with the installation.
- Oracle WebLogic Server that is bundled with the installation
- User data store. This is typically an LDAP directory, but can optionally be a database store.

- Federation data store. This is a standard LDAP directory such as Oracle Internet Directory, Microsoft Active Directory or Sun Java System Directory Server.

---

**Note:** A user federation data store is not absolutely required for Oracle Identity Federation in all cases: it is required for Liberty 1.x and SAML 2.0 opaque persistent identifiers, but is optional for SAML 1.x, WS-Federation, and SAML 2.0 non-opaque identifiers (such as email address, subject DN, and so on).

---

- Oracle Database version 10.2.0.4.0 and higher or 11.1.0.7 and higher for the RDBMS transient data store.
- Oracle HTTP Server for proxy implementation; this is the only proxy server supported by Oracle Identity Federation, and is bundled with the installation.

---

**Note:** Oracle requires that you synchronize the system clocks on the cluster nodes when you are using Oracle Identity Federation in a high availability deployment.

---

**8.13.2.3.1 Using RCU to Create Oracle Identity Federation Schemas in the Repository** Before you can install the Oracle Internet Directory instances on OIFHOST1 and OIFHOST2, you must use the latest version of the Repository Creation Utility (RCU) to create the collection of schemas used by Oracle Identity Federation.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

Follow these steps to run RCU and create the Oracle Identity Federation schemas in an Oracle RAC database repository:

1. Issue this command:

```
prompt> RCU_HOME/bin/rcu &
```

2. On the Welcome screen, click **Next**.
3. On the Create Repository screen, select the **Create** operation to load component schemas into an existing database.

Click **Next**.

4. On the Database Connection Details screen, enter connection information for the existing database as follows:

**Database Type:** Oracle Database

**Host Name:** Name of the computer on which the database is running. For an Oracle RAC database, specify the VIP name or one node name. Example:  
INFRADBHOST1-VIP or INFRADBHOST2-VIP

**Port:** The port number for the database. Example: 1521

**Service Name:** The service name of the database. Example: oif.mycompany.com

**Username:** SYS

**Password:** The SYS user password

**Role:** SYSDBA

Click **Next**.

5. On the Select Components screen, create a new prefix and select the components to be associated with this deployment:  
**Create a New Prefix:** `idm`  
**Components:** Select **Identity Management** (Oracle Identity Federation - OIF). De-select all other schemas.  
Click **Next**.
6. On the Schema Passwords screen, enter the passwords for the mail and additional (auxiliary) schema users.  
Click **Next**.
7. On the Map Tablespaces screen, select the tablespaces for the components.
8. On the Summary screen, click **Create**.
9. On the Completion Summary screen, click **Close**.

### 8.13.3 Oracle Identity Federation High Availability Configuration Steps

In a high availability environment, it is recommended that Oracle WebLogic Server utilities be used for clustering, load balancing, and failover of Oracle Identity Federation instances.

Ensure that the schema database is running then follow these steps to install.

This section describes the steps to install and configure Oracle Identity Federation instances on OIFHOST1 and OIFHOST2:

- [Section 8.13.3.1, "Configuring Oracle Identity Federation on OIFHOST1"](#)
- [Section 8.13.3.2, "Creating boot.properties for the Administration Server on OIFHOST1"](#)
- [Section 8.13.3.3, "Configuring Oracle Identity Federation on OIFHOST2"](#)
- [Section 8.13.3.4, "Post-Installation Steps for Oracle Identity Federation"](#)
- [Section 8.13.3.5, "Configuring the Load Balancer"](#)
- [Section 8.13.3.6, "Validating Oracle Identity Federation High Availability"](#)
- [Section 8.13.3.7, "Enabling Oracle Identity Federation Integration with Highly Available LDAP Servers"](#)

#### 8.13.3.1 Configuring Oracle Identity Federation on OIFHOST1

Follow these steps to configure the first instance of Oracle Identity Federation on OIFHOST1:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OIFHOST1 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. Ensure that port number 7499 is not in use by any service on the computer by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.



On UNIX:

```
netstat -an | grep LISTEN | grep ":7499"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr "7499"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entry for port 7499 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom port (uncomment the line where you specify the port number for Oracle Identity Federation):

```
The port for OIF Server port
OIF Server Port No = 7499
```

7. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select **Create a New Domain** and specify these values:

**HostName:** OIFHOST1.MYCOMPANY.COM

**Port:** 7001

**UserName:** weblogic

**User Password:** <password for weblogic user>

Click **Next**.

10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be updated.  
`/u01/app/oracle/product/fmw`
  - **Oracle Home Directory:** This value is prefilled and cannot be updated.  
`oif`
  - **WebLogic Server Directory:**  
`/u01/app/oracle/product/fmw/wlserver_10.3`
  - **Oracle Instance Location:**

/u01/app/oracle/admin/oif\_inst1

- **Instance Name:** oif\_inst1

---

**Note:** Ensure that the Oracle Home Location directory path for OIFHOST1 is the same as the Oracle Home Location path for OIFHOST2. For example, if the Oracle Home Location directory path for OIFHOST1 is: /u01/app/oracle/product/fmw/oif, then the Oracle Home Location directory path for OIFHOST2 must also be /u01/app/oracle/product/fmw/oif.

---

Click **Next**.

11. On the Specify Oracle Configuration Manager Details screen, specify the values shown in the example below:
  - **Email Address:** Provide the email address for your My Oracle Support account.
  - **Oracle Support Password:** Provide the password for your My Oracle Support account.
  - Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.


12. On the Configure Components screen, de-select all the components except **Oracle Identity Federation components**. The Oracle Identity Federation components include Oracle Identity Federation and Oracle HTTP Server. Select the **Clustered** checkbox.

Click **Next**.

---

**Note:** The default Oracle WebLogic Server clustering mode set by the installer is unicast (not multicast).

---




13. On the Configure Ports screen, select **Specify Ports using Configuration File**. Provide the path to the `staticports.ini` file that you copied to the temporary directory.  
Click **Next**.
14. On the Specify OIF Details screen, specify these values:
  - **PKCS12 Password:** <password>
  - **Confirm Password:** <confirm the password entered above>
  - **Server Id:** oif\_OIFDomain
 Click **Next**.
15. On the Select OIF Advanced Flow Attributes screen, specify these values:
  - **Authentication Type:** LDAP
  - **User Store:** LDAP
  - **Federation Store:** LDAP
  - **User Session Store:** RDBMS (default selection, which cannot be changed for a cluster)
  - **Message Store:** RDBMS (default selection, which cannot be changed for a cluster)
  - **Configuration Store:** RDBMS (default selection, which cannot be changed for a cluster)

---

**Note:** When you choose RDBMS for the session, message, and configuration data stores during an Advanced installation, the installer creates one data source for all three data stores.

If you want to have separate databases for each of these stores, you must configure this after the installation.

---



16. On the Authentication LDAP Details screen, specify the values shown in the example below:

- **LDAP Type:** Select Oracle Internet Directory from the drop down.
- **LDAP URL:** Provide the LDAP URL to connect to your LDAP store in the following format: `ldap://host:port` or `ldaps://host:port`. For example:  
`ldaps://oid.mycompany.com:636`
- **LDAP Bind DN:** `cn=orcladmin`
- **LDAP Password:** `<password for orcladmin>`
- **User Credential ID Attribute:** `uid`
- **User Unique ID Attribute:** `orclguid`
- **Person Object Class:** `inetOrgPerson`

Click Next.

**Specify Authentication LDAP Details**

LDAP Type: Oracle Internet Directory

LDAP URL: ldaps://oid.mycompany.com:636

LDAP Bind DN: cn=orcladmin

LDAP Password: .....

User Credential ID Attribute: uid

User Unique ID Attribute: orclguid

Person Object Class: inetOrgPerson

Base DN: dc=us,dc=oracle,dc=com

Buttons: < Back, Next >, Finish, Cancel

Elapsed Time: 9m 25s

17. On the LDAP Attributes for User Data Store screen, specify the values shown in the example below:

- **LDAP Type:** Select Oracle Internet Directory from the drop down.
- **LDAP URL:** Provide the LDAP URL to connect to your LDAP store in the following format: `ldap://host:port` or `ldaps://host:port`. For example:

`ldaps://oid.mycompany.com:636`

- **LDAP Bind DN:** `cn=orcladmin`
- **LDAP Password:** `<password for orcladmin>`
- **User Description Attribute:** `uid`
- **User ID Attribute:** `orclguid`
- **Person Object Class:** `inetOrgPerson`
- **Base DN:** `dc=us,dc=mycompany,dc=com`

Click Next.



18. On the LDAP Attributes for Federation Data Store screen, specify the values shown in the example below:

- **LDAP Type:** Select Oracle Internet Directory from the drop down.
- **LDAP URL:** Provide the LDAP URL to connect to your LDAP store in the following format: `ldap://host:port` or `ldaps://host:port`. For example:  
`ldaps://oid.mycompany.com:636`
- **LDAP Bind DN:** `cn=orcladmin`
- **LDAP Password:** <password for orcladmin>
- **User Federation Record Context:** `cn=myfed,dc=us,dc=mycompany,dc=com`
- **Container Object Class:** This is the type of User Federation Record Context that Oracle Identity Federation should use when creating the LDAP container, if it does not exist already. If that field is empty, its value will be set to `applicationprocess`. For Microsoft Active Directory this field has to be set to `container`.

Click **Next**.

**Specify LDAP Attributes For Federation Data Store**

LDAP Type: Oracle Internet Directory

LDAP URL: ldaps://oid.mycompany.com:636

LDAP Bind DN: cn=orcladmin

LDAP Password: .....

User Federation Record Context: cn=myfed,dc=us,dc=mycompany,dc=com

LDAP Container Object Class: applicationprocess

Enter the object class of the User Federation Record Context that Oracle Identity Federation should use when creating the LDAP entry, if one does not already exist. If this field is empty, its value will be set to applicationProcess.

< Back Next > Finish Cancel

Elapsed Time: 13m 50s

19. On the Transient Store Database Details screen, specify the values shown in the example below:

- **Connect String:** Provide the connect string to your database. For example:

```
infradbhost1-vip.mycompany.com:1521:oifdb1^infradbhost2-vip.mycompany.com:
1521:oifdb2@oif.mycompany.com
```

---

**Note:** The Oracle RAC database connect string information needs to be provided in the format  
*host1:port1:instance1^host2:port2:instance2@servicename.*

During this installation, it is not required for all the Oracle RAC instances to be up. If one Oracle RAC instance is up, the installation can proceed.


It is required that the information provided above is complete and accurate. Specifically, the correct host, port, and instance name must be provided for each Oracle RAC instance, and the service name provided must be configured for all the specified Oracle RAC instances.

Any incorrect information entered in the Oracle RAC database connect string has to be corrected manually after the installation.

---

- **UserName:** Enter the username for the OIF Schema. For example: ha\_oif
- **Password:** password for the oif user

Click **Next**.



20. On the Installation Summary screen, review the selections to ensure that they are correct. If they are not correct, click **Back** to modify selections on previous screens. Then click **Install**.

21. On the Installation Progress screen, view the progress of the installation.

Once the installation is done, the oracleRoot.sh confirmation dialog box opens. This dialog box advises you that a configuration script needs to be run as root before installation can proceed. Leave this dialog box open, open another shell window, log in as root, and run the following script:

```
/u01/app/oracle/product/fmw/oif/oracleRoot.sh
```

22. On the Configuration Progress screen, view the progress of the configuration.

23. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

### 8.13.3.2 Creating boot.properties for the Administration Server on OIFHOST1

The boot.properties file enables the Administration Server to start without prompting for the administrator username and password.

To create the boot.properties file:

1. On OIFHOST1, go the following directory:

```
MW_HOME/user_projects/domains/domainName/servers/AdminServer/security
```

For example:

```
cd /u01/app/oracle/product/fmw/user_projects/domains/IDMDomain/servers/AdminServer/security
```

2. Use a text editor to create a file called boot.properties under the security directory. Enter the following lines in the file:

```
username=adminUser
password=adminUserPassword
```



---

**Note:** When you start the Administration Server, the username and password entries in the file get encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, you should start the server as soon as possible so that the entries get encrypted.

---

3. Stop the Administration Server if it is running.

See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.

4. Start the Administration Server on OIFHOST1 using the `startWebLogic.sh` script located under the `MW_HOME/user_projects/domains/domainName/bin` directory.
5. Validate that the changes were successful by opening a web browser and accessing the following pages:
  - WebLogic Server Administration Console at:  
`http://oidhost1.mycompany.com:7001/console`
  - Oracle Enterprise Manager Fusion Middleware Control at:  
`http://oidhost1.mycompany.com:7001/em`

Log into these consoles using the `weblogic` user credentials.

### 8.13.3.3 Configuring Oracle Identity Federation on OIFHOST2

Follow these steps to configure the second instance of Oracle Identity Federation on OIFHOST2:

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that Oracle Identity Management software has been installed and upgraded on OIFHOST2 as described in [Section 8.3.3.1, "Installing Oracle Fusion Middleware Components."](#)
3. On OIFHOST1, port 7499 was used for Oracle Identity Federation. The same port should be used for the Oracle Identity Federation instance on OIFHOST2. Therefore, ensure that port 7499 is not in use by any service on OIFHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep LISTEN | grep ":7499"
```

On Windows:

```
netstat -an | findstr "LISTEN" | findstr "7499"
```

4. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entry for ports 7499 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

5. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
6. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom port (uncomment the line where you specify the port number for Oracle Identity Federation):

```
The port for OIF Server port
OIF Server Port No = 7499
```

7. Start the Oracle Identity Management 11g Configuration Assistant located under the `ORACLE_HOME/bin` directory as follows:

On UNIX, issue this command: **`./config.sh`**

On Windows, double-click **`config.exe`**

8. On the Welcome screen, click **Next**.
9. On the Select Domain screen, select the **Expand Cluster** option and specify these values:

**HostName:** `OIFHOST1.MYCOMPANY.COM`

**Port:** `7001`

**UserName:** `weblogic`

**User Password:** `<password for weblogic user>`

Click **Next**.

10. On the Specify Installation Location screen, specify the following values:
  - **Oracle Middleware Home Location:** This value is prefilled and cannot be updated.  
`/u01/app/oracle/product/fmw`
  - **Oracle Home Directory:** This value is prefilled and cannot be updated.  
`oif`
  - **WebLogic Server Directory:**  
`/u01/app/oracle/product/fmw/wlserver_10.3`
  - **Oracle Instance Location:**  
`/u01/app/oracle/admin/oif_inst2`
  - **Instance Name:** `oif_inst2`

---

**Note:** Ensure that the Oracle Home Location directory path for OIFHOST1 is the same as the Oracle Home Location path for OIFHOST2. For example, if the Oracle Home Location directory path for OIFHOST1 is: `/u01/app/oracle/product/fmw/oif`, then the Oracle Home Location directory path for OIFHOST2 must also be `/u01/app/oracle/product/fmw/oif`.

---

Click **Next**.

11. On the Specify Oracle Configuration Manager Details screen, specify the values shown in the example below:

- **Email Address:** Provide the email address for your My Oracle Support account.
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- Check the checkbox next to the **I wish to receive security updates via My Oracle Support** field.

Click **Next**.

12. On the Configure Components screen, de-select all the components except **Oracle Identity Federation components**. The Oracle Identity Federation components include Oracle Identity Federation and Oracle HTTP Server.

Click **Next**.

13. On the Installation Summary screen, review the selections to ensure that they are correct. If they are not correct, click **Back** to modify selections on previous screens. Then click **Install**.

14. On the Installation Progress screen, view the progress of the installation.

Once the installation is done, the `oracleRoot.sh` confirmation dialog box opens. This dialog box advises you that a configuration script needs to be run as root before installation can proceed. Leaving this dialog box open, open another shell window, log in as root, and run the following script:

```
/u01/app/oracle/product/fmw/oif/oracleRoot.sh
```

15. On the Configuration Progress screen, view the progress of the configuration.

16. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

### 8.13.3.4 Post-Installation Steps for Oracle Identity Federation

Follow the post-installation steps in this section to complete the installation and configuration of the Oracle Identity Federation application.

#### 8.13.3.4.1 Copy the Oracle Identity Federation Configuration Directory from OIFHOST1 to OIFHOST2

To copy the Oracle Identity Federation Configuration directory from OIFHOST1 to OIFHOST2:

Copy the following directory on OIFHOST1: `MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_oif1/applications` to the following location on OIFHOST2: `MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_oif2/applications`.

For example, from OIFHOST1, execute the following command:

```
scp -rp MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_oif1/applicationsuser@OIFHOST2:/MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_oif2/applications
```

**8.13.3.4.2 Start the Managed Server on OIFHOST2 in a Cluster** Follow these steps to start the newly created wls\_oif2 Managed Server in a cluster on OIFHOST2:

1. In the left pane of the Oracle WebLogic Server Administration Console, expand **Environment** and select **Clusters**.  
See the "Starting and Stopping Oracle Fusion Middleware" chapter of the *Oracle Fusion Middleware Administrator's Guide* for information on starting and stopping WebLogic Servers.
2. Click on the link for the cluster (cluster\_oif) containing the Managed Server (wls\_oif2) you want to stop.
3. Select **Control**.
4. Under **Managed Server Instances in this Cluster**, select the check box next to the Managed Server (wls\_oif2) you want to start and click **Start**.
5. On the Cluster Life Cycle Assistant page, click **Yes** to confirm.

WebLogic Node Manager starts the server on the target machine. When the Node Manager finishes its start sequence, the server's state is indicated in the **State** column in the Server Status table.

**8.13.3.4.3 Configure Oracle HTTP Server** Oracle HTTP Server is installed on OIFHOST1 and OIFHOST2 along with the Oracle Identity Federation server. Configure the Oracle HTTP Server by following these steps:

1. On OIFHOST1, edit the oif.conf file located under the *INSTANCE\_HOME/config/OHS/ohsName/moduleconf* directory.
2. If the Identity Management installation is in standalone mode, uncomment and set the WebLogicHost and WebLogicPort variables to reference the WebLogic Server Managed Server where Oracle Identity Federation is running (for example: oifhost1.mycompany.com and 7499).
3. If the Identity Management installation is in clustered mode, uncomment and set the WebLogicCluster variable to reference the WebLogic Server Managed Servers where Oracle Identity Federation is running (for example: oifhost1.mycompany.com:7499,oifhost2.mycompany.com:7499).
4. Save and exit the oif.conf file.
5. Restart Oracle HTTP Server.

### 8.13.3.5 Configuring the Load Balancer

In a high availability configuration, Oracle recommends using an external load balancer to front end and load balance requests between the various Oracle Identity Federation instances.

In high availability environments, where the Oracle Identity Federation Application is not front-ended by an Oracle HTTP Server Instance, Oracle recommends enabling sticky sessions on the hardware load balancer.

**8.13.3.5.1 Load Balancer Virtual Server Name Setup** Refer to [Section 8.2.5.4, "Configuring Virtual Server Names and Ports for the Load Balancer"](#) for details.

**8.13.3.5.2 Oracle Identity Federation Configuration** To configure the Oracle Identity Federation application to use the load balancer VIP:

1. In the Oracle Enterprise Manager Fusion Middleware Control, navigate to **Administration**, and then **Server Properties**.
2. Change the host name and port to reflect the load balancer host and port.
3. In the Oracle Enterprise Manager Fusion Middleware Control, navigate to **Administration**, and then **Identity Provider**.and
4. Change the URL to **http://LoadBalancerHost:LoadBalancerPort**.
5. In the Oracle Enterprise Manager Fusion Middleware Control, navigate to **Administration**, and then **Service Provider**.
6. Change the URL to **http://LoadBalancerHost:LoadBalancerPort**.
7. Repeat these steps for each Managed Server where Oracle Identity Federation is deployed.

### 8.13.3.6 Validating Oracle Identity Federation High Availability

To validate Oracle Identity Federation in a high availability configuration.

1. In a web browser, you will be able to access the following URLs if the configuration is correct:

```
http://<LoadBalancerHost>:<LoadBalancerPort>/fed/sp/metadata
```

```
http://<LoadBalancerHost>:<LoadBalancerPort>/fed/idp/metadata
```

2. Follow the instructions in the "Obtain Server Metadata" and "Add Trusted Providers" sections of *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation* to import metadata from the SP into the IdP and the IDP metadata into the SP.
3. Go to the following URL and do a Single Sign-On operation:

```
http://<SP_Host>:<SP_port>/fed/user/testspssso
```

### 8.13.3.7 Enabling Oracle Identity Federation Integration with Highly Available LDAP Servers

By default, Oracle Identity Federation is not configured to be integrated with LDAP Servers deployed in a high availability configuration. To integrate Oracle Identity Federation with highly available LDAP Servers to serve as user data store, federation data store, or authentication engine, Oracle Identity Federation needs to be configured based on the LDAP server's function. Use the WLST script located under the `$ORACLE_HOME/common/bin` directory.

Enter the WLST script environment for Oracle Identity Federation, then set the following properties as needed:

- To integrate the user data store with a highly available LDAP Server, set the `userldaphaenabled` boolean property from the `datastore` group to `true`; otherwise set it to `false`:

```
setConfigProperty('datastore','userldaphaenabled', 'true', 'boolean')
```

- To integrate the federation data store with a highly available LDAP Server, set the `fedldapenabled` boolean property from the `datastore` group to `true`; otherwise set it to `false`:

```
setConfigProperty('datastore', 'fedldapenabled', 'true', 'boolean')
```

- To integrate the LDAP authentication engine with a highly available LDAP Server, set the `ldapenabled` boolean property from the `authnengines` group to `true`; otherwise set it to `false`:

```
setConfigProperty('authnengines', 'ldapenabled', 'true', 'boolean')
```

### 8.13.4 Oracle Identity Federation Failover and Expected Behavior

This section describes steps for performing various failover operations on Oracle Identity Federation instances deployed in a high availability environment and their expected behavior. Follow the steps in this section to perform:

- Oracle Identity Federation instance failover
- Oracle Real Application Clusters failover

#### 8.13.4.1 Performing an Oracle Identity Federation Failover

Follow these steps to perform a test of a failover of an Oracle Identity Federation instance and to check the status of Oracle Identity Federation:

---

**Note:** The `testspso` URL referred to in the steps below is the Test SP SSO service that is bundled with Oracle Identity Federation 11g.

The testing service enabled by default, but can be disabled by the administrator. In a production environment, the Test SP SSO Service may be disabled.

if the Test SP SSO Service is disabled, you can use whatever service you have integrated to start the Federation SSO Flow from the SP.

---

1. Set up Oracle Identity Federation to be able to perform a federation single sign-on operation.
2. Start Single Sign-On operation from Oracle Identity Federation, acting as a Service Provider. One possible way to do this is to use the `http://<SPhost>:<SPport>/fed/user/testspso` URL choosing Artifact profile.
3. On the IdP login page, shut down `wls_oif1` through the Managed Server page and enter the username and password.
4. The Single Sign-On operation should succeed.

#### 8.13.4.2 Performing an Oracle RAC Failover

To perform an Oracle RAC failover:

1. On one of the database hosts (`infradbhost1-vip`) where the OIF schema is installed, use the `srvctl` command to stop a database instance:

```
srvctl stop instance -d db_unique_name -i inst_name_list
```

2. Use the `srvctl` command to check the status of the database:

```
srvctl status database -d db_unique_name -v
```

3. Perform an operation on Oracle Identity Federation:

```
ORACLE_INSTANCE/bin/opmnctl status ias-component=oif1
```

4. Use the `srvctl` command to start the database instance:

```
srvctl start instance -d db_unique_name -i inst_name_list
```

### 8.13.5 Troubleshooting Oracle Identity Federation High Availability

The following information may help you troubleshoot Oracle Identity Federation issues:

- Oracle Identity Federation logs its messages in the Oracle WebLogic Server Managed Server log files, for example:

```
DOMAIN_HOME/servers/wls_oif1/logs/wls_oif1-diagnostic.log
```

It is recommended that you use the Oracle Enterprise Manager Fusion Control to view log files by choosing **Identity And Access > OIF > Logs > View Log Messages**.

- Ensure that the database that you use for Oracle Identity Federation does not have old configuration data. If so, newer data may be overwritten. Delete old data in the Oracle Identity Federation configuration database tables or recreate the schema before you point Oracle Identity Federation to the database.
- The host name and port in the Oracle Identity Federation server configuration should be set to the load balancer host and port - otherwise there will be errors during Single Sign-On operation.
- If the clocks on the computers on which the IDP and SP are running have different times, you will see errors during Single Sign-On. Fix this by setting the system clocks to have the same time or by adjusting the server drift using the Server Properties page in the Oracle Enterprise Manager Fusion Middleware Control.
- The ProviderId is a string that uniquely identifies the IDP/SP. The ProviderId for all servers in a cluster must be the same. The ProviderId is defaulted to: `host:port/fed/<sp|idp>` at installation time. If necessary change or set this value after installation. Do not change it again during operation.

For more information on troubleshooting Oracle Identity Federation, see *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Federation*.

## 8.14 Starting and Stopping Oracle Identity Management Components

This section describes how to start, stop and restart various components described in this chapter.

This section contains the following topics:

- [Section 8.14.1, "Oracle Internet Directory"](#)
- [Section 8.14.2, "Oracle Virtual Directory"](#)
- [Section 8.14.3, "Oracle HTTP Server"](#)
- [Section 8.14.4, "Node Manager"](#)
- [Section 8.14.5, "WebLogic Administration Server"](#)

- [Section 8.14.6, "Oracle Identity Manager"](#)
- [Section 8.14.7, "Oracle Access Manager Managed Servers"](#)
- [Section 8.14.8, "Oracle Adaptive Access Manager Managed Servers"](#)
- [Section 8.14.9, "Oracle Identity Federation"](#)

## 8.14.1 Oracle Internet Directory

### Starting

Start system components such as Oracle Internet Directory by typing

```
ORACLE_INSTANCE/bin/opmnctl startall
```

You can verify that the system components have started by executing:

```
ORACLE_INSTANCE/bin/opmnctl status -l
```

### Stopping

Stop system components such as Oracle Internet Directory by executing the following command:

```
ORACLE_INSTANCE/bin/opmnctl stopall
```

## 8.14.2 Oracle Virtual Directory

### Starting

Start system components such as Oracle Virtual Directory by typing:

```
ORACLE_INSTANCE/bin/opmnctl startall
```

You can verify that the system components have started by executing:

```
ORACLE_INSTANCE/bin/opmnctl status -l
```

### Stopping

Stop system components such as Oracle Virtual Directory by executing the following command:

```
ORACLE_INSTANCE/bin/opmnctl stopall
```

## 8.14.3 Oracle HTTP Server

Prior to starting/stopping the Oracle HTTP server ensure that the environment variables `ORACLE_HOME` and `ORACLE_INSTANCE` are defined and that `ORACLE_HOME/opmn/bin` appears in the `PATH`. For example:

```
export ORACLE_HOME=/u01/app/oracle/product/fmw/web
export ORACLE_INSTANCE=/u01/app/oracle/admin/web[1-2]
export PATH=$ORACLE_HOME/opmn/bin:$PATH
```

### Starting

Start the Oracle web tier by issuing the command:



```
opmnctl startall
```

### Stop

Stop the web tier by issuing the command

```
opmnctl stopall
```

to stop the entire Web tier or

```
opmnctl stopproc process-type=OHS
```

to stop Oracle HTTP Server only.

### Restarting

You can restart the web tier by issuing a **Stop** followed by a **Start** as the previous sections describe.

To restart the Oracle HTTP server only, use the following command.

```
opmnctl restartproc process-type=OHS
```

## 8.14.4 Node Manager

Start and stop the Node Manager as follows:

### Starting

To start Node Manager, issue the commands:

```
IDMHOST> cd ORACLE_BASE/product/fmw/wlserver_10.3/server/bin
IDMHOST> ./startNodeManager.sh
```

### Stopping

To stop node manager, kill the process started in the previous section

## 8.14.5 WebLogic Administration Server

Start and stop the WebLogic Administration Server as follows:

### Starting

The recommended way to start the Administration Server is to use WLST and connect to Node Manager:

```
IDMHOST1> cd ORACLE_BASE/product/fmw/oracle_common/common/bin
IDMHOST1> ./wlst.sh
```

Once in `wlst` shell, execute

```
wls:/offline>nmConnect(Admin_User,'Admin_Password',ADMINHOST1,'5556',
'IDMDomain','/u01/app/oracle/admin/domain_name/aserver/IDMDomain')
wls:/nm/domain_name> nmStart('AdminServer')
```

Alternatively, you can start the Administration Server by using the command:

```
DOMAIN_HOME/bin/startWeblogic.sh
```

**Stopping**

To stop the Administration Server, log in to the WebLogic console using the URL: `http://adminhost.mycompany.com:7001/console`, where `adminhost` is the name of the host where the Administration Server is running.

Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click on the **Control** tab
3. Select **AdminServer(admin)**
4. Click **Shutdown** and select **Force Shutdown now**.
5. Click **Yes** when asked to confirm that you wish to shutdown the administration server.

**Restarting**

Restart the server by following the **Stop** and **Start** procedures in the previous sections.

### 8.14.6 Oracle Identity Manager

Start and stop Oracle Identity Manager and Oracle SOA Suite servers as follows:

**Starting**

To start the Oracle Identity Manager managed server(s), log in to the WebLogic console using the URL: `http://oimhost1.mycompany.com:7001/console`.

Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click on the **Control** tab
3. Select **SOA Servers (WLS\_SOA1 and/or WLS\_SOA2)**

---

---

**Note:** You can start the Oracle Identity Manager and Oracle SOA Suite servers independently of each other. There is no dependency in their start order. However, the SOA server must be up and running for all of the OIM functionality to be available.

---

---

4. Click on the **Start** button.
5. Click **Yes** when asked to confirm that you wish to start the server(s).
6. After **WLS\_SOA1** and/or **WLS\_SOA2** have started, select **WLS\_OIM1** and/or **WLS\_OIM2**
7. Click **Start**.
8. Click **Yes** when asked to confirm that you wish to start the server(s).

**Stopping**

To stop the OIM managed server(s), log in to the WebLogic console using the URL: `http://oimhost1.mycompany.com:7001/console`. Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click the **Control** tab

3. Select **OIM Servers (WLS\_OIM1 and/or WLS\_OIM2)** and **(WLS\_SOA1 and/or WLS\_SOA2)**
4. Click the **Shutdown** button and select **Force Shutdown now**.
5. Click **Yes** when asked to confirm that you wish to shutdown the server(s).

### Restarting

Restart the server by following the **Stop** and **Start** procedures in the previous sections.

## 8.14.7 Oracle Access Manager Managed Servers

Start and stop Oracle Access Manager managed servers as follows:

### Starting

To start the OAM managed server(s), log in to the WebLogic console using the URL: `http://oamhost1.mycompany.com:7001/console`.

Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click on the **Control** tab
3. Select **OAM Servers (WLS\_OAM1 and/or WLS\_OAM2)**
4. Click on the **Start** button.
5. Click **Yes** when asked to confirm that you wish to start the server(s).

### Stopping

To stop the OAM managed server(s), log in to the WebLogic console using the URL: `http://oamhost1.mycompany.com:7001/console`. Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click the **Control** tab
3. Select **OAM Servers (WLS\_OAM1 and/or WLS\_OAM2)**
4. Click the **Shutdown** button and select **Force Shutdown now**.
5. Click **Yes** when asked to confirm that you wish to shutdown the server(s).

### Restarting

Restart the server by following the **Stop** and **Start** procedures in the previous sections.

## 8.14.8 Oracle Adaptive Access Manager Managed Servers

Start and stop Oracle Adaptive Access Manager as follows:

### Starting

To start the OAAM managed server(s), log in to the WebLogic console using the URL: `http://oaamhost1.mycompany.com:7001/console`. Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click the **Control** tab
3. Select **OAAM Servers (WLS\_OAAM1 and/or WLS\_OAAM2)**

4. Click the **Start** button.
5. Click **Yes** when asked to confirm that you wish to start the server(s).

### Stopping

To stop the OAM managed server(s), log in to the WebLogic console using the URL: `http://oaamhost1.mycompany.com:7001/console`. Then proceed as follows:

1. Select **Environment - Servers** from the Domain Structure menu
2. Click the **Control** tab
3. Select **OAAM Servers (WLS\_OAAM1 and/or WLS\_OAAM2)**
4. Click **Shutdown** and select **Force Shutdown now**.
5. Click **Yes** when asked to confirm that you wish to shutdown the server(s).

### Restarting

Restart the server by following the Stop and Start procedures above.

## 8.14.9 Oracle Identity Federation

Start and stop Oracle Identity Federation managed servers as follows:

### Starting

To start the OIF managed server(s), log in to the WebLogic console at: `http://oifhost1.mycompany.com:7001/console`. Then proceed as follows:

1. Select **Environment - Servers** from the **Domain Structure** menu.
2. Click the **Control** tab.
3. Select **OIF Servers (WLS\_OIF1 and/or WLS\_OIF2)**.
4. Click **Start**.
5. Click **Yes** when asked to confirm that you wish to start the server(s).

### Stopping

To stop the OIF managed server(s), log in to the WebLogic console at: `http://oifhost1.mycompany.com:7001/console`. Then proceed as follows:

1. Select **Environment - Servers** from the **Domain Structure** menu.
2. Click the **Control** tab.
3. Select **OIF Servers (WLS\_OIF1 and/or WLS\_OIF2)**.
4. Click **Shutdown** and select **Force Shutdown now**.
5. Click **Yes** when asked to confirm that you wish to shut down the server(s).

### Restarting

Restart the server by following the Stop and Start procedures above.

### Starting the OIF Instances and EMAgent

Start the OIF Instance and EMAgent by executing the following command:

```
ORACLE_INSTANCE/bin/opmnctl startall
```

You can verify that the instance started successfully by executing:

```
ORACLE_INSTANCE/bin/opmnctl status -l
```

**Stopping the OIF Instances and EMAgent**

Stop the OIF Instance and EMAgent by executing the following command:

```
ORACLE_INSTANCE/bin/opmnctl stopall
```



---

# Configuring High Availability for Web Tier Components

The Web tier receives each incoming HTTP request and invokes the requested business logic operation in the application. Based on the results of the operation and state of the model, the next view is selected to display. The selected view is transmitted to the client for presentation.

This chapter describes high availability concepts and configuration procedures for Oracle HTTP Server and Oracle Web Cache. This chapter includes the following topics:

- [Section 9.1, "About the Web Tier"](#)
- [Section 9.2, "Oracle HTTP Server and High Availability Concepts"](#)
- [Section 9.3, "Oracle Web Cache and High Availability Concepts"](#)

## 9.1 About the Web Tier

The Web tier of Java EE application server is responsible for interacting with the end user such as web browsers primarily in the forms of HTTP requests and responses. It is the outermost tier in the application server, closest to the end user. At the highest level, the Web tier does four basic tasks:

- Interprets client requests
- Dispatches those requests to an object (for example, an enterprise Java bean) that encapsulates business logic
- Selects the next view for display
- Generates and delivers the next view

The Web tier receives each incoming HTTP request and invokes the requested business logic operation in the application. Based on the results of the operation and state of the model, the next view is selected to display. The selected view is transmitted to the client for presentation.

Oracle Web Cache is a content-aware server accelerator, or reverse proxy, for the Web tier that improves the performance, scalability, and availability of Web sites that run on Oracle HTTP Server.

Oracle Web Cache is the primary caching mechanism provided with Oracle Fusion Middleware. Caching improves the performance, scalability, and availability of Web sites that run on Oracle Application Server by storing frequently accessed URLs in memory.

By storing frequently accessed URLs in memory, Oracle Web Cache eliminates the need to repeatedly process requests for those URLs on the application Web server and database tiers. Unlike legacy proxies that handle only static objects, Oracle Web Cache caches both static and dynamically generated content from one or more application Web servers. Because Oracle Web Cache is able to cache more content than legacy proxies, it provides optimal performance by greatly reducing the load on application Web server and database tiers. As an external cache, Oracle Web Cache is also an order of magnitude faster than object caches that run within the application tier.

Because Web Cache is fully compliant with HTTP 1.0 and 1.1 specifications, it can accelerate Web sites that are hosted by any standard web servers, such as Apache and IIS. In Oracle Fusion Middleware, Oracle Web Cache resides in front of one or more instances of Oracle HTTP Server. Responses to browser based HTTP requests are directed to the Oracle HTTP Server instance and transmitted through Oracle Web Cache. The Oracle Web Cache instance can handle any Web content transmitted with standard HTTP protocol.

## 9.2 Oracle HTTP Server and High Availability Concepts

Oracle HTTP Server is the Web server component for Oracle Fusion Middleware. It provides a listener for Oracle WebLogic Server and the framework for hosting static pages, dynamic pages, and applications over the Web.

### 9.2.1 Oracle HTTP Server Single-Instance Characteristics

Oracle HTTP Server is based on Apache 2.2.10 infrastructure, and includes modules developed specifically by Oracle. The features of single sign-on, clustered deployment, and high availability enhance the operation of the Oracle HTTP Server. Oracle HTTP Server has the following components to handle client requests:

- HTTP listener, to handle incoming requests and route them to the appropriate processing utility.
- Modules (mods), to implement and extend the basic functionality of Oracle HTTP Server. Many of the standard Apache modules are included with Oracle HTTP Server. Oracle also includes several modules that are specific to Oracle Fusion Middleware to support integration between Oracle HTTP Server and other Oracle Fusion Middleware components.
- Perl interpreter, a persistent Perl runtime environment embedded in Oracle HTTP Server through mod\_perl.

Oracle HTTP Server enables developers to program their site in a variety of languages and technologies, such as the following:

- Perl (through mod\_perl and CGI)
- C (through CGI and FastCGI)
- C++ (through FastCGI)
- PHP (through mod\_php)
- Oracle PL/SQL

Oracle HTTP Server can also be a proxy server, both forward and reverse. A reverse proxy enables content served by different servers to appear as if coming from one server.

The Oracle HTTP Server is essentially a stateless application. Session State can be maintained using browser-based cookies.




**Figure 9–1 Oracle HTTP Server Architecture**

Figure 9–1 illustrates the following Oracle HTTP Server components:

- The Oracle Process Manager and Notification Service (OPMN) is used to start, stop, and monitor Oracle HTTP Server. OPMN periodically polls Oracle HTTP server to ensure that it is still functioning. If Oracle HTTP Server is not functioning, OPMN takes appropriate action to restart the failed component to ensure that service is maintained.
- Oracle HTTP server is based on the industry leading Apache Web Server. You can extend Oracle HTTP Server’s core functionality using the following *modules*:
  - The **mod\_dms** module enables performance monitoring of site components using Oracle Dynamic Monitoring Service (DMS).
  - The **mod\_onsint** module provides integration support with Oracle Notification Service (ONS) and OPMN.
  - The **mod\_oradav** module is an Oracle Call Interface (OCI) application written in C, that extends the implementation of **mod\_dav**. The **mod\_oradav** directive can read and write to local files, or to an Oracle database.
  - The **mod\_oss1** module enables strong cryptography for Oracle HTTP Server. This Oracle module is a plug-in to Oracle HTTP Server that enables the server to use SSL.
  - The **mod\_osso** module enables Oracle Single Sign-on.
  - The **mod\_perl** module embeds the Perl interpreter into Oracle HTTP Server. This eliminates start-up overhead and enables Perl applications to be accessed through Oracle HTTP Server.
  - The **mod\_plsql** module connects Oracle HTTP Server to an Oracle database, enabling Web application creation using Oracle stored procedures.

- The **mod\_wl\_ohs** module allows requests to be proxied from Oracle HTTP Server to Oracle WebLogic Server.

### 9.2.1.1 Oracle HTTP Server and Oracle WebLogic Server

Oracle HTTP Server does not require an Oracle WebLogic domain. However, Oracle HTTP Server is typically used in conjunction with an Oracle WebLogic Domain.

The link to Oracle WebLogic managed servers is handled through the module `mod_wl_ohs`. This module is configured by routing requests of a particular type, for example, JSPs, or by routing requests destined to a URL, to specific WebLogic Managed Servers. In a non high availability deployment, this destination is defined by specifying the host name and port on which the WebLogic managed server resides.

In a high availability deployment, the WebLogic managed servers are typically clustered together. In such a deployment, a special `mod_wl_ohs` directive, `WebLogicCluster`, specifies a comma-separated list of cluster members. This list is not necessarily a complete list of cluster members. When a request requiring a WebLogic managed server is received, `mod_wl_ohs` sends that request to one of the WebLogic cluster members listed in the directive. When the WebLogic managed server receives the request, it processes it and sends a complete list of cluster members back to `mod_wl_ohs`. When `mod_wl_ohs` receives the updated list, it dynamically adds any previously unknown servers to the list of known servers, allowing all future requests to be load balanced across the full cluster member list. This process has the advantage of allowing new managed servers to be added to the cluster without updating `mod_wl_ohs`, or adding the Oracle HTTP Server.

In an example scenario, assume that the WebLogic Cluster consists of WLS1, WLS2, and WLS3. Also assume that WLS1 and WLS2 are known to `mod_wl_ohs`.

When `mod_wl_ohs` first receives a request, it attempts to send that request to either WLS1 or WLS2, if it is successful then the list is updated to include WLS3 for future requests. If, however, WLS1 and WLS2 are unavailable, but WLS3 is available, the request is rejected because `mod_wl_ohs` has no way of knowing WLS3 exists.

When using the Oracle HTTP server with an Oracle WebLogic domain, Oracle recommends associating Oracle HTTP Server with the WebLogic domain. This allows the Oracle HTTP Server to be incorporated into the Oracle Fusion Middleware Console for centralized management and monitoring.

For more information Oracle WebLogic clusters, see *Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server*.

### 9.2.1.2 Oracle HTTP Server External Dependencies

Oracle HTTP Server is dependent on the Oracle Process Management and Notification Server to start the HTTP server itself.

In order to access other resources, such as Oracle WebLogic, or Oracle Single Sign-On, Oracle HTTP Server is also dependent on the appropriate loading of Apache module.

## 9.2.2 Oracle HTTP Server Startup and Shutdown Lifecycle

The Oracle HTTP Server process is invoked indirectly through OPMN. When OPMN receives a request to start the Oracle HTTP Server, it starts the Oracle HTTP Server process (`httpd`).

After Oracle HTTP Server is started, it is ready to listen for and respond to HTTP(S) requests. The request processing model on Microsoft Windows systems differs from that on UNIX systems.

- For Microsoft Windows, there is a single parent process and a single child process. The child process creates threads that are responsible for handling client requests. The number of created threads is static and can be configured for performance.
- For UNIX, there is a single parent process that manages multiple child processes. The child processes are responsible for handling requests. The parent process brings up additional child processes as necessary, based on configuration. Although the server has the ability to dynamically bring up additional child processes, it is best to configure the server to start enough child processes initially so that requests can be handled without a need for more child processes.

## 9.2.3 Starting and Stopping Oracle HTTP Server

You can use Fusion Middleware Control or the `opmnctl` command to start and stop Oracle HTTP Server.

---

**Note:** Do not use the `apachectl` utility to manage Oracle HTTP Server.

---

You can determine the status of Oracle HTTP Server using `opmnctl`:

```
opmnctl status
```

Processes in Instance: instance1

| ias-component | process-type   | pid   | status |
|---------------|----------------|-------|--------|
| webcache1     | WebCache-admin | 19556 | Alive  |
| webcache1     | WebCache       | 19555 | Alive  |
| ohs1          | OHS            | 7249  | Alive  |

### 9.2.3.1 Understanding the PID File

When Oracle HTTP Server starts up, it writes the process ID (PID) of the parent `httpd` process to the `httpd.pid` file located, by default, in the following directory:

```
ORACLE_INSTANCE/diagnostics/logs/OHSComponent/ohs1/
```

Administrators can use the process ID when restarting and terminating the daemon. If a process stops abnormally, you must stop the `httpd` child processes using the `kill` command.

The `PidFile` directive in `httpd.conf` specifies the location of the PID file.

---

**Note:** For UNIX platforms, if you edit the `PidFile` directive, you must also edit the `ORACLE_HOME/ohs/bin/apachectl` file to specify the new location of the PID file.

---

**See Also:** `PidFile` directive in the Apache Server documentation at:

<http://httpd.apache.org/docs/>

### 9.2.3.2 Starting and Stopping Oracle HTTP Server Using Oracle Fusion Middleware Control

To start Oracle HTTP Server using Fusion Middleware Control:

1. Navigate to the Oracle HTTP Server home page.
2. Select **Control** from the Oracle HTTP Server menu.
3. Select **Start Up** from the Control menu.

To stop Oracle HTTP Server using Fusion Middleware Control:

1. Navigate to the Oracle HTTP Server home page.
2. Select **Control** from the Oracle HTTP Server menu.
3. Select **Shut Down** from the Control menu.

### 9.2.3.3 Starting and Stopping Oracle HTTP Server Using `opmnctl`

To start all Oracle HTTP Server components in an Oracle instance using `opmnctl`:

```
opmnctl startproc process-type=OHS
```

To start a specific Oracle HTTP Server component, such as `ohs1`, using `opmnctl`:

```
opmnctl startproc ias-component=ohs1
```

To stop all Oracle HTTP Server components in an Oracle instance using `opmnctl`:

```
opmnctl stopproc process-type=OHS
```

To stop a specific Oracle HTTP Server component, such as `ohs1`, using `opmnctl`:

```
opmnctl stopproc ias-component=ohs1
```

## 9.2.4 Oracle HTTP Server Configuration Artifacts

Oracle HTTP Server is configured using several operating system files. When Oracle HTTP Server is deployed, its configuration information is placed into the following directory structure:

- `ORACLE_INSTANCE/config/OHS/OHS_NAME/`
- `ORACLE_INSTANCE/config/OHS/OHS_NAME/moduleconf`

The primary configuration file is `httpd.conf`, located in the `ORACLE_INSTANCE/config/OHS/OHS_NAME`. All other configuration files are invoked through this main control file.

---

---

**Note:** The `http.conf` file has a general include directive. As a result, all files with a `.conf` extension in the directory `moduleconf` are automatically included in the configuration. Many Oracle applications, such as Oracle Portal and Oracle Forms, place their Oracle HTTP directives into files located in this directory.

---

---

## 9.2.5 Oracle HTTP Server Log File Locations

There are two types of log files for Oracle HTTP Server:

- Error logs, which record server problems.

- Access logs, which record which components and applications are being accessed, and by whom.

You can view Oracle Fusion Middleware log files using either Oracle Fusion Middleware Control or a text editor. The log files for Oracle HTTP Server are located in the following directory:

`ORACLE_INSTANCE/diagnostics/logs/OHS/OHS_NAME.`

The default name of a log file is `ohs_name.log`.

For more information about Oracle HTTP Server log files, see the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

## 9.2.6 Oracle HTTP Server High Availability Architecture and Failover Considerations

Figure 9–2 shows two Oracle HTTP Servers, which have been placed behind a load balancer. The load balancer receives requests from users and forwards them on to the connected Oracle HTTP Servers. In the example, the Load Balancer receives the requests on the standard HTTP/HTTPS ports (80/443), however, it then passes them on to the Oracle HTTP Servers using completely different ports. This has the following advantages:

- Actual ports are hidden from users.
- Users do not have to add the port numbers to the URL.

On Unix-based systems, it is not mandatory to start Oracle HTTP Server with root privileges. Only root can start a process which uses a port less than 1024.

The load balancer routes requests to the functioning Oracle HTTP Servers.

**Figure 9–2 Oracle HTTP Server High Availability Architecture**




Figure 9–2 also shows how Oracle HTTP Server distribute requests to Web Logic Managed Servers. For High Availability, it is assumed that each pair of components (Oracle HTTP Server and Web Logic Managed Servers) exist on different hosts.

This architecture could simply be separated across two servers. Alternatively, in more complex implementations, each component could reside on a completely separate server.

## 9.2.7 Oracle HTTP Server Protection from Failures and Expected Behaviors

There are two categories of Oracle HTTP Server failures: process failures and node failures. An individual operating system process may fail, or a node failure could involve the entire host upon which the Oracle HTTP Server is failing.

### Process Failure

The Oracle HTTP Server processes are protected by the Oracle Process Manager and Notification system (OPMN). If an Oracle HTTP Server process fails, OPMN automatically restarts the process.

### Node Failure

If an entire node fails, the load balancer or Oracle Web Cache, in front of Oracle HTTP Server, sends a request to another Oracle HTTP Server if the first one does not respond, or is determined to be failed through URL pings.

### WebLogic Managed Server Failure

In a high availability deployment, Oracle WebLogic managed servers are part of a cluster. If one of the managed servers fails, `mod_wl_ohs` automatically redirects requests to one of the active cluster members. If the application stores state, state replication is enabled within the cluster, which allows redirected requests access to the same state information.

### Database Failure

Database failures are only likely to be an issue where `mod_oradav` or `mod_plsql` is used. If this is an Oracle Real Application Clusters (Oracle RAC) database, the failure characteristics are determined by the defined Oracle RAC connection.

If client connection failover is configured, any in-flight transactions are rolled back, and a database reconnection is required.

If Transparent Application Failover (TAF) is configured, any in-flight database write is rolled back, but an automatic database reconnection takes place, and select statements are automatically recovered. In this scenario, TAF only fails over select statements, and package variables are lost.

## 9.2.8 Oracle HTTP Server Cluster-Wide Configuration Changes

Oracle HTTP Servers are not deployed in a clustered framework. The Oracle HTTP server configuration is file-based, so changes made to one Oracle HTTP Server must be manually copied to other Oracle HTTP Servers in the configuration. This also applies to static HTML files stored in the `htdocs` directory.

## 9.2.9 Configuring Oracle HTTP Server for High Availability

This section describes the prerequisites and procedures for configuring an example high availability deployment of Oracle HTTP Server.

### 9.2.9.1 Prerequisites

Consider the following prerequisites before configuring a high availability Oracle HTTP Server deployment.

**9.2.9.1.1 Load Balancer** To distribute requests against Oracle HTTP Servers, a load balancer is required in the form of an external load balancer or Oracle Web Cache. If using an external load balancer, it should have the following features:

- Virtual server name and port configuration
- Process failure detection
- Monitoring of ports (HTTP, HTTPS) for Oracle HTTP and HTTPS
- SSL Protocol Conversion (if required)

### Configuring Virtual Server Names and Ports for the Load Balancer

For each application, such as `myapp.mycompany.com`, configure the load balancer with a virtual server name and associated ports. In an Oracle HTTP Server installation, these virtual servers are configured for HTTP connections, which are distributed across the HTTP servers.

If your site is serving requests for HTTP and HTTPS connections, Oracle recommends that HTTPS requests are terminated at the load balancer and passed through as HTTP requests. To do this the load balancer should be able to perform the protocol conversion, and must be configured for persistent HTTP sessions.

For this example configuration, it is assumed that a the load balancer has been configured as:

- Virtual Host: `Myapp.mycompany.com`
- Virtual Port: `7777`
- Server Pool: `Map`
- Server: `WEBHOST1, Port 7777, WEBHOST2, Port 7777`

### Managing Port Numbers

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

Most port numbers are assigned during installation. It is important that any traffic going from the Oracle HTTP Servers to the Oracle WebLogic Servers has access through any firewalls.

**9.2.9.1.2 Associating Oracle HTTP Server with a WebLogic Domain** If registering Oracle HTTP Server with a WebLogic Administration Server, that server must be up and running and configured to accept JMX requests.

### 9.2.9.2 Install Oracle HTTP Server on WEBHOST1

Start the Oracle Universal Installer for Oracle Fusion Middleware Web Tier and Plug-ins 11g (11.1.1.1.0) DVD installation as follows:

For UNIX, run the following command: `runInstaller`

For Windows, double-click `setup.exe`

1. In the Welcome screen, click **Next**.
2. In the Select Install Type screen, select **Install and Configure**, and click **Next**.
3. In the Specify Installation Location screen, enter the following installation location:  
`/u01/app/oracle/product/FMW/Web1`
4. In the Prerequisite Checks screen, click **Next**.
5. In the Configure Components screen, select **Oracle HTTP Server** and click **Next**

---

**Note:** If this installation is to be associated with a WebLogic Domain, select the **Associate Selected Components with WebLogic Domain** check box.

If this installation is not to be associated with a WebLogic domain, deselect the **Associate Selected Components with WebLogic Domain** check box. If required, this association can be created after the installation.

---

6. In the Specify WebLogic Domain Details screen (optional), enter the following values:
  - Domain Host Name: The name of the server hosting the WebLogic administration server, for example, **wladmin.mycompany.com**.
  - Domain Port Number: The WebLogic Administration server Port, for example, **7001**.
  - Username: The WebLogic Administration Server user, for example **WebLogic**.
  - Password: Administration Server password
7. Click **Next**.
8. In the Specify Component Details screen, enter the following values:
  - Instance Home Location: **/u01/app/oracle/admin/web1**
  - AS Instance Name: **web1**
  - OHS Component Name: **ohs1**
  - WebCache Component Name: **webcache1**
9. Click **Next**.
10. In the Web Cache Administrator Password screen, specify the password for your Web Cache administrator.

Valid passwords are 5 to 30 characters long, must begin with an alphabetic character, use only alphanumeric, underscore (\_), dollar (\$) or pound (#) characters and include at least one number.
11. Click **Next**.
12. In the Specify Web Tier Port Details screen, create a file with the following entries:
  - [OHS]#: The http\_main port for the OHS component.
  - OHS Port = 7777

Save the file and specify the name of the file in the **File name** field after highlighting **Specify Ports** using the **Configuration** file option.

---

**Note:** It is not mandatory to have the same port numbers on all Web tier instances, however, it is recommended. For this configuration example, it is assumed.

---

13. Click **Next**.
14. In the Configuration Summary screen, review the selections to ensure that they are correct. If they are not correct, click **Back** to modify selections on previous screens.



**15. Click Install.**

- 16.** In the Configuration screen, multiple configuration assistants are launched in succession. When this process completes, click **Next**.

The Installation Complete screen appears.

**17. Click Finish.****Validate the Installation**

Validate the installation using the following URL to access the Oracle HTTP Server home page:

`http://webhost1:7777/`

**9.2.9.2.1 Configure Virtual Host(s)** For each virtual host or site name used add an entry to the Oracle HTTP server configuration. Create a file called `virtual_hosts.conf` file located in the `ORACLE_INSTANCE/config/OHS/ohs_component_name/moduleconf` directory as follows:

```
NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName http://myapp.mycompany.com:80
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>
```

If you are using SSL/SSL Termination (\*):

```
NameVirtualHost *:7777
<VirtualHost *:7777>
 ServerName https://myapp.mycompany.com:443
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>
```

**9.2.9.2.2 Configure mod\_wl\_ohs** After installing and configuring Oracle HTTP Server, link it to any defined WebLogic managed servers by editing the `mod_wl_ohs.conf` file located in `ORACLE_INSTANCE/config/OHS/ohs_name` directory.

The following is an example of `mod_wl_ohs.conf` entries:

```
LoadModule weblogic_module ORACLE_HOME/ohs/modules/mod_wl_22.so

<IfModule mod_weblogic.c>
 WebLogicCluster apphost1.mycompany.com:7050, apphost2.mycompany.com:7050
 Debug ON
 WLLogFile /u01/app/oracle/admin/WL1/logs/mod_wl_ohs.log
 MatchExpression *.jsp
</IfModule>

<Location /weblogic>
 SetHandler weblogic-handler
 WebLogicCluster apphost1.mycompany.com:7050, apphost2.com:7050
 Debug ERR
 WLLogFile /u01/app/oracle/admin/WL1/logs/mod_wl_ohs.log
 DefaultFileName index.jsp
</Location>
```

---

**Note:** If using SSL termination AND routing requests to WebLogic then the following additional configuration is required.

In the WebLogic console, WebLogic Plugin Enabled must be set to true, either at the domain, cluster or managed server level.

In the Location block which directs requests to the WebLogic managed servers, the following lines also need to be added.

```
WLProxySSL ON
WLProxySSLPassThrough ON
```

For example:

```
<Location /weblogic>
 SetHandler weblogic-handler
 WebLogicCluster apphost1.mycompany.com:7050,apphost2.com:7050
 WLProxySSL On
 WLProxySSLPassThrough ON
 Debug ERR
 WLLogFile /u01/app/oracle/admin/WL1/logs/mod_wl_ohs.log
 DefaultFileName index.jsp
</Location>
```

After enabling the WebLogic plugin, restart the Administration Server.

---

These examples show two different ways of routing requests to Oracle WebLogic managed servers:

- The `<ifModule>` block sends any requests ending in \*.jsp to the WebLogic managed server cluster located on Apphost1 and Apphost2.
- The `<Location>` block sends any requests with URLs prefixed by /weblogic to the WebLogic managed server cluster located on Apphost1 and Apphost2.

**9.2.9.2.3 Restart Oracle HTTP Server** Restart the Oracle HTTP Server using the following commands:

```
opmnctl restartproc process-type=OHS
```

**9.2.9.2.4 Validate the Oracle HTTP Server Configuration** Validate the configuration using the following URLs:

```
http://myapp.mycompany.com:7777/weblogic
```

```
http://myapp.mycompany.com:7777/weblogic
```

### 9.2.9.3 Install Oracle HTTP Server on WEBHOST2

On WEBHOST2, perform the steps from [Section 9.2.9.2, "Install Oracle HTTP Server on WEBHOST1."](#)

#### Validate the Installation

Validate the installation using the following URL to access the Oracle HTTP Server home page:

```
http://webhost2:7777/
```

**9.2.9.3.1 Configure Virtual Host(s)** Add an entry for each virtual host or site name to the Oracle HTTP Server configuration. Copy the file `virtual_hosts.conf` from `WEBHOST1` to `WEBHOST2`. This is located in the `ORACLE_INSTANCE/config/OHS/ohs_component_name/moduleconf` directory.

**9.2.9.3.2 Configure mod\_wl\_ohs** After installing and configuring the Oracle HTTP Server, link it to any defined WebLogic Server instances by copying the file `mod_wl_ohs.conf` located in the `ORACLE_INSTANCE/config/OHS/ohs_component_name` directory from `WEBHOST1`.

**9.2.9.3.3 Restart Oracle HTTP Server** Restart the HTTP Server using the following commands:

```
opmnctl restartproc process-type=OHS
```

**9.2.9.3.4 Validate the Oracle HTTP Server Configuration** Validate the configuration using the following URLs:

`http://myapp.mycompany.com/`

`https://myapp.mycompany.com` (if using SSL/SSL termination)

`http://myapp.mycompany.com:7777/weblogic`

## 9.3 Oracle Web Cache and High Availability Concepts

Oracle Web Cache is a content-aware server accelerator, or *reverse proxy*, for the Web tier that improves the performance, scalability, and availability of Web sites that run on Oracle HTTP Server.

Oracle Web Cache is the primary caching mechanism provided with Oracle Fusion Middleware. Caching improves the performance, scalability, and availability of Web sites that run on Oracle Application Server by storing frequently accessed URLs in memory.

By storing frequently accessed URLs in memory, Oracle Web Cache eliminates the need to repeatedly process requests for those URLs on the application Web server and database tiers. Unlike legacy proxies that handle only static objects, Oracle Web Cache caches both static and dynamically generated content from one or more application Web servers. Because Oracle Web Cache is able to cache more content than legacy proxies, it provides optimal performance by greatly reducing the load on application Web server and database tiers. As an external cache, Oracle Web Cache is also an order of magnitude faster than object caches that run within the application tier.

### 9.3.1 Oracle Web Cache Single-Node Characteristics

Oracle Web Cache is fully compliant with HTTP 1.0 and 1.1 specifications. Therefore, it can accelerate Web sites that are hosted by any standard web servers, such as Apache and IIS. In Oracle Fusion Middleware, Oracle Web Cache resides in front of one or more instances of Oracle HTTP Server. Responses to browser based HTTP requests are directed to the Oracle HTTP Server instance and transmitted through Oracle Web Cache. The Oracle Web Cache instance can handle any Web content transmitted with standard HTTP protocol.

A reverse proxy appears to be the content server to clients but internally retrieves its objects from other backend origin servers as a proxy. A reverse proxy acts as a gateway to the origin servers. It relays requests from outside the firewall to origin servers behind the firewall, and delivers retrieved content back to the client.


Figure 9–3 shows an overview of how reverse proxy Web caching works. Oracle Web Cache has an IP address of 144.25.190.241 and the application Web server has an IP address of 144.25.190.242.

The steps for browser interaction with Oracle Web Cache are as follows:

1. A browser sends a request to a Web site named `www.company.com:80`.  
This request in turn generates a request to Domain Name System (DNS) for the IP address of the Web site.
2. DNS returns the IP address of the load balancer for the site, that is, 144.25.190.240.
3. The browser sends the request for a Web page to the load balancer. In turn, the load balancer sends the request to Oracle Web Cache server 144.25.190.241.
4. If the requested content is in its cache, then Oracle Web Cache sends the content directly to the browser. This is called a cache hit.
5. If Oracle Web Cache does not have the requested content or the content is stale or invalid, it hands the request off to application Web server 144.25.190.242. This is called a cache miss.
6. The application Web server sends the content to Oracle Web Cache.
7. Oracle Web Cache sends the content to the client and stores a copy of the page in cache.

A page stored in the cache is removed when it becomes invalid or outdated.

**Figure 9–3 Web Server Acceleration**



### 9.3.1.1 Oracle Web Cache Component Characteristics

An Oracle Web Cache system component consists of two performance-oriented native processes, the cache server process and the admin server process. The cache server

process handles client requests serving content back to the client. The `admin` server process provides administration, configuration, and monitoring capabilities.

### 9.3.1.2 Oracle Web Cache Process Monitoring

You manage the `cache` and `admin` server processes with Oracle Process Manager and Notification Server (OPMN) using the Fusion Middleware Control, Oracle Web Cache Manager, or the `opmnctl` utility, as described in the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

### 9.3.1.3 Oracle Web Cache Startup and Shutdown Lifecycle

OPMN is responsible for the direct start, stop, restart, and monitoring of the `cache` server and `admin` server processes. Anytime the Oracle Web Cache configuration is statically modified, you must stop and restart Oracle Web Cache processes:

The executable for the `cache` process is `webcached`, and the executable for the `admin` server process is `webcachea`. These executables reside in the following directories:

```
(UNIX) ORACLE_HOME/webcache/bin
(Windows) ORACLE_HOME\bin
```

When you stop Oracle Web Cache, all objects are cleared from the cache. In addition, all statistics are cleared.

After you configure Oracle Web Cache, restart Oracle Web Cache. In addition, if you change the administrator password, restart the `admin` server.

To restart Oracle Web Cache, use one of the following tools:

- Use Fusion Middleware Control or the `opmnctl` command-line utility to restart the `cache` or `admin` server processes.
- Use Oracle Web Cache Manager to restart the `cache` server process.


You must restart *both* the `cache` server and `admin` server processes if you modified one of the following configuration settings:

- Administration port properties
- Trusted subnets
- User and group ID information

See the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache* for further information about starting and stopping Oracle Web Cache.

### 9.3.1.4 Oracle Web Cache Request Flow

Figure 9–4 shows further details of the request flow within the Oracle Web Cache tier.

**Figure 9–4 Request Flow to Oracle Web Cache within the Web Tier**

As shown in [Figure 9–4](#), the following occurs within the Oracle Web Cache tier:

1. The incoming browser request is analyzed for the correct HTTP format.
2. The browser request is then further analyzed to determine if it is in HTTPS format:
  - a. If the browser request is in HTTPS format, Oracle Web Cache decrypts SSL.
  - b. If the browser request is not in HTTPS format, Oracle Web Cache parses the request.
3. After the request is understood, it is filtered by a set of prescribed filtering rules.
4. A cache lookup is performed to see if the HTTP request was sent previously and is present in the cache.

If the request is present in the cache, a cache hit, the request is compressed and the content is sent directly to the browser.

If the request is not present in the cache, a cache miss, then either:

- a. The request is sent directly to a single origin server.
- b. The request is sent to load-balanced origin servers.

Each load balanced origin server pings each Oracle Web Cache server on a periodic basis to check the status of the cache. The load balancer distributes any incoming requests among cache cluster members. If Oracle Web Cache does not have the requested content or the content is stale or invalid, it hands the request off to the application Web server. The application Web server sends the content to Oracle Web Cache. Oracle Web Cache sends the content to the client and stores a copy of the page in cache.

The proxy server is placed in a less secure zone, the Demilitarized Zone (DMZ), instead of the origin server.

Caching rules determine which objects are cached. When you establish a caching rule for a particular URL, those objects contained within the URL are not cached until there is a client request for them. When a client first requests an object, Oracle Web Cache sends the request to the origin server. This request is a cache miss. Because this URL has an associated caching rule, Oracle Web Cache caches the object for subsequent requests. When Oracle Web Cache receives a second request for the same object, Oracle Web Cache serves the object from its cache to the client. This request is a cache hit.

When you stop Oracle Web Cache, the cache clears all objects. In addition, Oracle Web Cache clears and resets statistics.

### 9.3.1.5 Oracle Web Cache Configuration Artifacts

Oracle stores configuration for Oracle Web Cache in the `webcache.xml` file, located in the following directories:

```
(UNIX) ORACLE_INSTANCE/instance_name/config/WebCache/webcache_name
(Windows) ORACLE_INSTANCE\instance_name\config\WebCache\webcache_name
```

Oracle offers two tools for managing the configuration files.

- Fusion Middleware Control.
- Oracle Web Cache Manager.

Use these tools rather than edit the `webcache.xml` configuration file, to perform all administrative tasks unless a specific procedure requires you to edit a file. Editing a file may cause the settings to be inconsistent and generate problems.

For further information about these tools, see the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

### 9.3.1.6 Log File Locations

Oracle Web Cache records event and error information in event logs. An event log entry can help you determine what objects have been inserted in the cache and alert you to any cache-related issues. Oracle Web Cache stores every request internally and then writes them in bulk at the end of the request. Request-based logging groups all the requests together.

By default, the event log has a file name of `event_log` for the Oracle Web Cache and Oracle Diagnostic Logging (ODL) text formats and `log.xml` for the ODL XML format.

Oracle Web Cache records information about the received HTTP requests in access logs. By default, the access log has a file name of `access_log`.

By default, Oracle Web Cache stores logs files in the following directories:

```
(UNIX) ORACLE_INSTANCE/diagnostics/logs/WebCache/<webcache_name>
(Windows) ORACLE_INSTANCE\diagnostics\logs\WebCache\<webcache_name>
```

For further information about these tools, see the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

## 9.3.2 Oracle Web Cache High Availability Considerations

The following sections describe the high availability solutions available with Oracle Web Cache:

- [Section 9.3.2.1, "Oracle Web Cache Stateless Load Balancing"](#)
- [Section 9.3.2.2, "Oracle Web Cache Backend Failover"](#)
- [Section 9.3.2.3, "Oracle Web Cache Session Binding"](#)
- [Section 9.3.2.4, "Oracle Web Cache Cluster-Wide Configuration Changes"](#)
- [Section 9.3.2.5, "Oracle Web Cache as a Software Load Balancer"](#)

### 9.3.2.1 Oracle Web Cache Stateless Load Balancing

Most Web sites are served by multiple origin servers running on multiple computers that share the load of HTTP and HTTPS requests. All requests that Oracle Web Cache cannot serve are passed to the origin servers. Oracle Web Cache balances the load among origin servers by determining the percentage of the available capacity, the weighted available capacity of each origin server. Oracle Web Cache sends a request to the origin server with the most weighted available capacity. The weighted available capacity is determined by the following formula:

$$(\text{Capacity} - \text{Load}) / \text{Capacity}$$

where:

- `Capacity` is the maximum number of concurrent connections that the origin server can accept
- `Load` is the number of connections currently in use


If the weighted available capacity is equal for multiple origin servers, Oracle Web Cache sends requests to the origin servers using *round robin*. With round robin, the first origin server in the list of configured servers receives the request, then the second origin server receives the second request. If the weighted available capacity is not equal, Oracle Web Cache sends the request to the origin server with the most available capacity.

If the load of origin servers is equivalent, Oracle Web Cache continues to use round robin, even when capacity is not equal for origin servers. Therefore, it is possible to see an even distribution of requests to origin server when the capacities are not configured to be the same.

To configure load balancing for a site, set the capacity of each origin server, and create *one* site-to-server mapping that maps *all* the applicable origin servers to the site.

[Figure 9–5](#) shows two sites, `www.company.com:80` and `www.server.com:80`. The site `www.company.com:80` is supported by application Web servers `company-host1` and `company-host2` with capacities of 50 each. The site `www.server.com:80` is supported by application Web servers `server-host1`, `server-host2`, and `server-host3` with capacities of 150, 50, and 50, respectively.



**Figure 9–5 Load Balancing**

Assuming all application Web servers have an initial load of 0, the requests to `www.company.com:80` and `www.server.com:80` will be distributed in the following manner:

- The requests to `www.company.com:80` are distributed between the two origin servers using round robin.

The requests to `company-host1` and `company-host2` will be distributed between the two origin servers so that they maintain an equal load. The first request is sent to `company-host1`. The second request is sent to `company-host2` if `company-host1` is still processing the first request. The third and subsequent requests are sent to the origin server that has the highest weighted available capacity.

When the capacities are equal, Oracle Web Cache uses round robin to distribute requests.

- The requests to `www.server.com:80` are distributed between three origin servers using the weighted available capacity percentage.

The first request to `www.server.com:80` is sent to `server-host1`, because it is the first in the configured list. The second request is sent to `server2-host`, because `server-host1` is still processing the first request and has a weighted available capacity of 99.3 percent and `server-host2` has a weighted available capacity of 100 percent. The third request is sent to `server-host3` because `server2-host` is still processing a request and has a weighted available capacity of 98 percent and `server3-host` has a weighted available capacity of 100 percent. The fourth request is sent to `server-host1` because `server-host2` and `server3-host` are still processing requests and have weighted available capacities of 98 percent. The fifth request is sent to `server-host1` because its weighted available capacity is 98.6 percent, which is still greater than `server-host2` and `server-host3`, respectively.


When the capacities and loads are not equal, Oracle Web Cache uses the weighted available capacity to distribute requests. If requests were processed before new requests came in, then it is possible for all three origin servers to have loads of 0. In this case, Oracle Web Cache uses round robin.

See the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache* for instructions on specifying capacity and creating site-to-server mappings.

### 9.3.2.2 Oracle Web Cache Backend Failover

After a specified number of continuous request failures, Oracle Web Cache considers an origin server as failed. When an origin server fails, Oracle Web Cache automatically distributes the load over the remaining origin servers and polls the failed origin server for its current up or down status until it is back online. Existing requests to the failed origin server result in errors. However, new requests are directed to the other origin servers. When the failed server returns to operation, Oracle Web Cache includes it in its weighted available capacity to load balance requests.

The failure feature is shown in [Figure 9-6](#). An outage of `server-host3`, which had a capacity of 50, results in 75 percent of requests being distributed to `server-host1` and 25 percent request being distributed to `server-host2`.

**Figure 9–6 Failover**

See the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache* for instructions on specifying the failover threshold.

### 9.3.2.3 Oracle Web Cache Session Binding


You can configure Oracle Web Cache to support *session binding*, whereby a user session for a particular site is bound to an origin server in order to maintain state for a period of time. To use this feature, the origin server itself must maintain state; that is, it must be stateful.

If a request is forwarded to an origin server for an object requiring session binding, the origin server creates the user session by including the session information to clients through Oracle Web Cache in the form of a *session cookie*, or an embedded URL parameter. Oracle Web Cache does not process the value of the parameter or cookie; it simply passes the information back to the client browser. When a client includes the session information in a subsequent request, Oracle Web Cache forwards the request

to the origin server that originally created the user session. Oracle Web Cache binds the user session to that particular origin server.

Figure 9–7 shows how Oracle Web Cache supports objects that use session binding.

**Figure 9–7 Session Binding**



The steps for how session binding works for requests are as follows:

1. When a request first comes in, Oracle Web Cache uses load balancing to determine to which origin server the request is forwarded. In this example, application Web server `www.server2.com` is selected.
2. If the requested object requires session binding, the origin server sends the session information back to the client through Oracle Web Cache in the form of a cookie or an embedded URL parameter.
3. Oracle Web Cache sends subsequent requests for the session to the origin server that established the session, bypassing load balancing. In this example, application Web server `www.server2.com` handles the subsequent requests.

If you configure a cache cluster, when you configure session binding, do not select the **Internal-Tracking** mechanism option, as it does not work for cache clusters. The other mechanisms work for cache clusters.


See [Section 9.3.2.3, "Oracle Web Cache Session Binding"](#) for instructions on configuring session binding.

#### 9.3.2.4 Oracle Web Cache Cluster-Wide Configuration Changes

Oracle Web Cache provides cluster-wide capabilities through *cache clusters*. In a cache cluster, multiple system components of Oracle Web Cache operate as one logical cache. This one logical cache is referred to as the *cache cluster member*. The cache cluster members communicate with one another to request cacheable content that is cached by another cache cluster member and to detect when a cache cluster member fails.

Figure 9–8 shows an Oracle Web Cache cluster that contains three cache cluster members. As the figure shows, the cluster members communicate with one another as well as with the application Web servers and with the clients.

**Figure 9–8 Oracle Web Cache Cluster Architecture**



Oracle Web Cache uses the relative capacity of each cache instance to distribute the cached content among the cache cluster members. In effect, it assigns a cache cluster member to be the owner of a particular object. This content is called *owned content*.

In addition to the owned content, Oracle Web Cache stores popular objects in the cache of each cluster member. These objects are known as *on-demand content*. By storing the on-demand content, Oracle Web Cache responds to requests for those objects quickly and decreases the number of cache misses. Fewer requests are sent to the application Web server. The result is improved performance.

A cache cluster uses one configuration that is synchronize to all cluster members. The configuration contains general information, such as security, session information, and caching rules, which is the same for all cluster members. It also contains cache-specific information, such as capacity, administration and other ports, resource limits, and log files, for each cluster member.

Each member must be authenticated before it is added to the cache cluster. The authentication requires that the administration username and password of the Oracle Web Cache instance to be added be the same as the administration username and password of the cluster.

When you add a cache to the cluster, the cache-specific information of the new cluster member is added to the configuration of the cache cluster. Then, Oracle Web Cache synchronizes the configuration to all members of the cluster. Because adding a new member changes the relative capacity of each Web cache, Oracle Web Cache uses the information about capacity to recalculate which cluster member owns which content.

When cache cluster members detect the failure of another cluster member, the remaining cache cluster members automatically take over ownership of the content of the failing member. When the cache cluster member is reachable again, Oracle Web Cache again reassigns the ownership of the content.

When you remove a Web cache from a cache cluster, the remaining cache cluster members take over ownership of the content of the removed member. In addition, the configuration information about the removed member is deleted from the configuration and the revised configuration is synchronized with the remaining cache cluster members.

See [Section 9.3.3.2, "Configuring a Cache Cluster"](#) for instructions on configuring a cache cluster.

### 9.3.2.5 Oracle Web Cache as a Software Load Balancer

You can configure a special mode of Oracle Web Cache that enables you to use Oracle Web Cache solely as a software load balancer of HTTP traffic. This configuration mode is useful to:

- Manage low-volume, departmental, or test Web sites
- Manage traffic in the DMZ between a hardware load balancer and an application using Oracle Portal or Oracle Single Sign-On

This mode does not cache *any* content or provide support for the following features:

- Compression: Oracle Web Cache ignores all compression settings.
- Request filtering: Oracle Web Cache ignores any configure request filters and rules.
- ESI: Oracle Web Cache does not assemble ESI content.
- Cache hierarchies: If you plan to configure two caches in a cache hierarchy, then you should not configure one of the caches as a load balancer.

You can deploy a single Oracle Web Cache server as a load balancer. However, this deployment makes the Oracle Web Cache server a single point of failure for your application. You can instead configure a cache cluster with multiple Oracle Web Cache servers in conjunction with operating system load balancing capabilities. Take note of the capacity changes mentioned earlier in this section.

In this mode, you can configure Oracle Web Cache to load balance HTTP traffic in front of an application using ESI or in front of another Oracle Web Cache. The Oracle Web Cache load balancer does not process ESI content or participate in hierarchical caching. For example, a typical Oracle Portal deployment has a built-in Oracle Web Cache used for ESI assembly. For these configurations, do not configure the Oracle Web Cache used for ESI assembly as a load balancer.

If you require other Oracle Web Cache features, such as caching or compression support, do not configure this mode. Instead, configure a hardware load balancer or operating system load balancing support, and use the load balancing feature to manage requests to origin servers.

See [Section 9.3.3.3, "Configure Oracle Web Cache as a Software Load Balancer"](#) for instructions on configuring this mode.

### 9.3.3 Configuring Oracle Web Cache High Availability Solutions

The following sections describe how to configure the following high availability solutions:

- [Section 9.3.3.1, "Configure Oracle Web Cache Session Binding"](#)
- [Section 9.3.3.2, "Configuring a Cache Cluster"](#)
- [Section 9.3.3.3, "Configure Oracle Web Cache as a Software Load Balancer"](#)

#### 9.3.3.1 Configure Oracle Web Cache Session Binding

To configure session binding, you specify a set of session binding rules, and then apply them to the sites. By default, sites are configured to use a default rule. You can use the default rule or select another rule customized for the site.

If you decide to change the value of the default session binding rule, ensure all named sites currently configured with this rule require session binding. If some sites do not require session binding, leave the value of default rule, and instead specify session binding settings for the site.

To configure session binding:

1. Navigate to the Web Cache Home page in the Fusion Middleware Control.
2. From the Web Cache menu, select **Administration** and then select **Session Configuration**.
3. From the **Site** list, select the site to create customized session-bindings.

Select **Global** to specify default settings for sites and for requests which do not match any defined site. If you do not specify customized session-binding settings for a site, then you can click the **Use global settings** option to apply the settings you specify for **Global**. The default selection for the Global selection is the **Disable session binding**. You change the default setting by selecting **Global** from the **Site** list and selecting on of the other session-binding selections.

4. Create a session definition in the **Session Definition** table.

- **Use global settings:** Select this option if you want to apply the session-binding settings you configured for the **Global** selection from the **Site** list.

By default, Oracle Web Cache disables session binding for all requests. The default selection for **Global** is the **Disable session binding** option. When you first create a site, it is set by default to use the global session binding settings

- **Disable session binding:** Select this option to disable session binding. This selection is the default for the Global site. You change the default setting by selecting **Global** from the **Site** list and selecting on of the other session-binding selections.

- **Cookie based session binding with any Set-Cookie:** Select this option if the client supports cookies and your application server uses one or more cookies for session state. Oracle Web Cache uses its own cookie to track a session. This cookie, which contains routing information, is maintained between Oracle Web Cache and the client browser. The client to application server binding will be initiated by the first cookie that the application server sends to the client.

- **Bind using a session:** Select this option to enable binding for a specific session, and then perform the following steps:

- a. From the **Session Name** list, select a session to enable binding for a specific session.

- b. From the Session Binding Mechanism list, select one of the following binding mechanisms for the selected session definition:
  - c. Select one of the following binding mechanisms for the selected session definition:
    - **Cookie Based:** Select if the client supports cookies. Oracle Web Cache uses its own cookie to track a session. This cookie, which contains routing information, is maintained between Oracle Web Cache and the client browser.
    - **Session Binding IAS:** Select if the application is based on OC4J. Oracle Web Cache forwards routing information with each request to OC4J through Oracle HTTP Server.
    - **Internal-Tracking:** Select if the client does not support cookies and the application is not based on Oracle HTTP Server and OC4J. This option is intended for backward compatibility with earlier releases of Oracle Web Cache. Oracle Web Cache maintains an in-memory routing table, of which each entry maps a session ID to an origin server. The routing table is not shared among cluster nodes. If you select this option and you have a cache cluster configuration, then you must also bind at the load balancer layer.
- 5. Click **Apply** to submit changes.
  - 6. Restart Oracle Web Cache.

### 9.3.3.2 Configuring a Cache Cluster

To configure a cache cluster, you configure two or more Oracle Web Cache instances as cache cluster members, and specify properties for the cluster.

A cache cluster uses one configuration that is synchronized from the current cache (the cache to which your client browser is connected) to all cluster members. The configuration contains settings that are the same for all cluster members as well as cache-specific settings for each cluster member.

This section contains the following topics to aid you in configuring a cache cluster in a environment in which all the caches are associated with the same Oracle WebLogic Server. These instruction explain how to configure a cluster using Fusion Middleware Control, which requires all the cache members to use the same Oracle WebLogic Server:

- [Section 9.3.3.2.1, "Configuration Prerequisites"](#)
- [Section 9.3.3.2.2, "Understanding Failover Threshold and Capacity Settings"](#)
- [Section 9.3.3.2.3, "Task 1: Add Caches to the Cluster and Configure Properties"](#)
- [Section 9.3.3.2.4, "Task 2: Enable Tracking of Session Binding"](#)
- [Section 9.3.3.2.5, "Task 3: Synchronize Configuration to Cluster Members"](#)

In addition, see the following information about configuring clusters:

- [Section 9.3.3.2.6, "Removing a Cache Member from a Cluster"](#)
- [Section 9.3.3.2.7, "Configuring Administration and Invalidation-Only Clusters"](#)

If you have want to configure a cache cluster for a configuration in which the caches are associated with different Oracle WebLogic Servers, follow the instructions in *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.

**9.3.3.2.1 Configuration Prerequisites** Because a cache cluster contains two or more instances of Oracle Web Cache, you must have two or more instances of Oracle Web



Cache installed on one or more nodes before you configure a cache cluster. The instances must be the same version of Oracle Web Cache. In addition, the respective passwords for the Oracle Web Cache administrator, and the invalidator user, `invalidator`, must be the same across the cluster members. If they are different, you must connect to the cache's admin server and modify the administration password.

**9.3.3.2.2 Understanding Failover Threshold and Capacity Settings** To ease with configuration, take the time to understand the following key configuration settings for a cache cluster and its members:

- [Failover Threshold for the Cache Cluster](#)
- [Capacity for Cache Cluster Members](#)

#### **Failover Threshold for the Cache Cluster**

You set the failover threshold when you configure cache cluster properties. This setting reflects the number of allowed consecutive request failures before Oracle Web Cache considers another cache cluster member to have failed. In other words, Oracle Web Cache consecutively retries a failed member for certain number of times, before considering the cache-member as down. The number of times Oracle Web Cache is allowed to retry is termed as failover threshold.

Oracle Web Cache considers a request to another cache cluster member to have failed if:

- There are any network errors
- The HTTP response status code is either less than 100, or is one of the following: 500 Internal Server Error, 502 Bad Gateway, 503 Service Unavailable, or 504 Gateway Timeout.

For each failed request, Oracle Web Cache increments the failure counter for that cluster member. This counter is kept separately by each cluster member. When a request is successfully processed by a cluster member, Oracle Web Cache resets the failure counter.

When the failover threshold is met, Oracle Web Cache considers the cache cluster member to have failed. Oracle Web Cache recalculates the relative capacity of the remaining cache cluster members. It then reassigns ownership of cache content.

When a cache cluster member is down, Oracle Web Cache starts polling the cache cluster member. It does this by sending requests to the ping URL you specify. When Oracle Web Cache receives a success response from the cache cluster member, it considers that cache cluster member to be up again. It recalculates the relative capacity of the cache cluster members and it reassigns ownership of cache content.

#### **Capacity for Cache Cluster Members**

When you configure a cache cluster member, you specify capacity for that member.

Oracle Web Cache uses capacity in two different ways:

- As the absolute capacity for the number of concurrent incoming connections to this cache cluster member from all other cache cluster members.

The connections are used to receive requests for owned content from other cache cluster members. The number of connections are divided among the other cluster members. For example, in a three-cache cluster, if the capacity of Cache\_A is 50, Cache\_B can open 25 connections to Cache\_A and Cache\_C can open 25 connections to Cache\_A.

More connections are used when another cache cluster member contains little or no data in its cache, such as when it is initially started, when it recovers from a failure, or after invalidation. During this time, the cluster member sends many of the requests to its peers, the owners of the content. In most cases, these requests are served more quickly than requests to the origin server. Having a higher number of connections increases performance during this time and shortens the time it takes to fully load the cache. After a cache is fully loaded, fewer of the connections are used. There is no overhead for unused connections.

- As the relative capacity of the cache cluster member.

The capacity of a cache cluster member is weighted against the total capacity of all active cache cluster members. When you set the capacity, Oracle Web Cache assigns a percentage of the ownership array to the cluster member, indicating how much of the cached content will be owned by the cluster member. The percentage is calculated using the following formula:

$$\text{cluster\_member\_capacity} / \text{total\_capacity\_of\_all\_active\_cluster\_members}$$

For example, if cache cluster member Cache\_A has a capacity of 100 and cache cluster member Cache\_B has a capacity of 300, for a total capacity of 400, Cache\_A is assigned 25 percent of the ownership array and Cache\_B is assigned 75 percent of the ownership array. That means that Cache\_A owns 25 percent of the cached content.

Note that in calculating the relative capacity, Oracle Web Cache considers the capacity of active cluster members; it does not consider the capacity of cluster members that it has determined to have failed.

Set the initial capacity for each cache cluster member to 10 percent of the **Maximum Incoming Connections** setting.

Once you have a better idea of an application's capacity needs and hit rates, fine tune the capacity. If these two assumptions apply to your cache cluster, then apply the following formula to determine the capacity for each cluster member:

1. Incoming traffic will be distributed equally to all the cache cluster members.
2. Ownership of content will be distributed equally among all the cache cluster members.

In the following formula, pick the highest value between the default value or the *max\_incoming\_connections* formula:

$$\text{max}(\text{default\_value}, (\text{max\_incoming\_connections} * (\text{cacheable\_misses}\%/100) * (\text{number\_of\_caches} - 1) / \text{number\_of\_caches}))$$

In the formula:

- *default\_value* is one of the following:
  - 100 for production environments
  - 30 for test environments
  - 0 for invalidation-only clusters

When the capacity increases, the number of file descriptors needed by Oracle Web Cache also increases.

- *max\_incoming\_connections* is the **Maximum Incoming Connections** setting from the Resource Limits page of Fusion Middleware Control.

- *cacheable\_misses%* is the percentage of requests for cacheable objects that were not served directly by Oracle Web Cache, but were served by Oracle Web Cache after it fetched the content from the origin server.

You can find the **Cacheable Misses** setting in the Web Cache Statistics page of Fusion Middleware Control.

For example, assume a cache cluster with four members. If Oracle Web Cache is operating at 1500 maximum incoming connections, with a 30 percent cacheable miss rate, then the equation to calculate capacity for this configuration looks like the following:

$$(1500 * (30/100) * (4 - 1) / 4$$

The equation calculates to 337.5. You would round up to 338, which is the capacity you would then enter for each cache cluster member.

$$1500 * .3 * 3 / 4 = 337.5$$

If you assign a capacity of 0 to a cluster member, that cluster member will not receive requests from other cluster members. However, that cluster member will forward requests to other cluster members, the owners of the content. If you assign a capacity of 0 to *all* cluster members, no requests will be forwarded between cluster members. Even when capacity is set to 0, you can still synchronize the configuration and Oracle Web Cache can automatically propagate invalidation requests to cluster members.

**9.3.3.2.3 Task 1: Add Caches to the Cluster and Configure Properties** Before you add a cache to the cluster, ensure the conditions described in [Section 9.3.3.2.1, "Configuration Prerequisites"](#) are met.

To add cache members to a cluster:

1. Navigate to the Web Cache Home page in the Fusion Middleware Control for one of the Oracle Web Cache instances.
2. From the Web Cache menu, select **Administration** and then select **Cluster**.
3. For each cache member you want to add:
  - a. Click **Add**.
  - b. In the **Component** field, enter the name of the cache member.

The **Capacity** field is auto-filled with a default value. You can modify this value. See [Section 9.3.3.2.2, "Understanding Failover Threshold and Capacity Settings"](#) for more information about capacity.

4. In the **Failover Threshold** field, enter the number of allowed consecutive request failures before Oracle Web Cache considers another cache cluster member to have failed.

The default is five failures.

See [Section 9.3.3.2.2, "Understanding Failover Threshold and Capacity Settings"](#) for further information about this field.

5. In the **Ping URL** field, enter the URL that cache cluster members will use to attempt to contact a cache cluster member that has reached its failover threshold.

Use a URL that is cacheable and that you can guarantee is stored in each cache. The default is `_oracle_http_server_webcache_static_.html`, which is stored in the cache.

6. In the **Ping Frequency** field, enter the time, in seconds, between attempts by a cluster member to reach the failed cluster member.

The default, 10 seconds, is a reasonable interval for most situations.

7. Click **Apply**.

**9.3.3.2.4 Task 2: Enable Tracking of Session Binding** In a cache cluster, all cache cluster members must be able to determine which origin server established the session, although the request was routed originally through only one cache cluster member.

For the Oracle Web Cache you established properties for in [Section 9.3.3.2.3, "Task 1: Add Caches to the Cluster and Configure Properties,"](#) configure session binding with a session binding mechanism. Do not use the **Internal-Tracking** option, as it does not work for cache clusters.

To configure session binding, see [Section 9.3.3.1, "Configure Oracle Web Cache Session Binding."](#)

**9.3.3.2.5 Task 3: Synchronize Configuration to Cluster Members** When you modify the cluster and apply changes, Oracle Web Cache adds the cache-specific information from the new cache cluster members to the configuration. For those changes to take affect in all cluster members, you must synchronize the configuration and restart the cluster members.

To synchronize configuration from a newly-added cache member to the other caches in the cluster:

1. Navigate to the Web Cache Home page in the Fusion Middleware Control for the Oracle Web Cache you established properties for in [Section 9.3.3.2.3, "Task 1: Add Caches to the Cluster and Configure Properties."](#)
2. From the Web Cache menu, select **Administration** and then select **Cluster**.
3. Select the other cache members in the cluster, click **Synchronize**.
4. Select all the cache members, select an interval for staggering the time that the operation begins on the cache sand click **Start Up**.

The cache cluster is ready to use.

**9.3.3.2.6 Removing a Cache Member from a Cluster** To remove a cache-member from a cluster, you must ensure that remaining cluster members no longer include that cache in cluster, and that the removed cache no longer considers itself to be part of the cluster.

To remove a cache from a cluster in Oracle Web Cache Manager:

1. Navigate to the Web Cache Home page in the Fusion Middleware Control for one of the Oracle Web Cache instances, but *not* the cache that you want to remove from the cluster.
2. From the Web Cache menu, select **Administration** and then select **Cluster**.
3. Select the cache you want to remove and click **Delete**.
4. Select the other cache members in the cluster, click **Synchronize**.
5. With the other caches still selected, click **Restart**.

All remaining caches in the cluster no longer consider the removed cache to be part of the cluster. However, the removed cache still considers itself to be part of the cluster. You will remedy this situation in the next steps.

6. Navigate to the Web Cache Home page in the Fusion Middleware Control of the cache you removed from the cluster.
7. From the Web Cache menu, select **Administration** and then select **Cluster**.
8. Select any cache except the current one and click **Delete**. Repeat until only the current cache remains in the **Cluster Members** list.
9. Click **Restart**.

**9.3.3.2.7 Configuring Administration and Invalidation-Only Clusters** You can configure a cluster that supports synchronizing the configuration and invalidation requests across all cache cluster members, but that does not forward requests between cache cluster members. That is, in processing requests, each cluster member acts as an individual cache and does not request objects from its peer cluster members. However, configuration changes and invalidation requests can be synchronized to cluster members.

You can use this configuration to simplify administration of many caches. It may be needed in a cluster where members are separated by a firewall. For example, you can have a cluster where two caches are located on either side of a firewall that separates the intranet from Internet. (The network settings of such a setup—of sending Internet traffic to one cache and intranet traffic to another—is beyond the scope of this document.)

To manage these caches as a cluster and avoid sharing contents between the caches, take the following steps:

1. Create a cluster and add members to it as discussed in [Section 9.3.3.2.3, "Task 1: Add Caches to the Cluster and Configure Properties"](#) and [Section 9.3.3.2.4, "Task 2: Enable Tracking of Session Binding,"](#) with the following exceptions:
  - For each cluster member, set the capacity to 0.
  - In the **Cluster Properties** section, select option **Invalidation requests sent to any cluster member will be propagated to all cluster members**.
2. Synchronize the configuration to all cluster members, as described in [Section 9.3.3.2.5, "Task 3: Synchronize Configuration to Cluster Members."](#)

### 9.3.3.3 Configure Oracle Web Cache as a Software Load Balancer

To configure a single Oracle Web Cache server as a software load balancer:

1. Use a text editor to open `webcache.xml`, located in:

```
(UNIX) ORACLE_INSTANCE/<instance_name>/config/WebCache/<webcache_name>
(Windows) ORACLE_INSTANCE\<instance_name>\config\WebCache\<webcache_name>
```

2. Locate the **CACHE** element.
3. Add the **ROUTINGONLY** attribute to the **CACHE** element. For example:

```
...
<CACHE WCDEBUGON="NO" CHRONOSONPERNODE="NO" CAPACITY="301" VOTES="1"
INSTANCENAME="instance_name" COMPONENTNAME="component_name"
ORACLEINSTANCE="instance" HOSTNAME="host"
ORACLEHOME="directory" NAME="web_cache_name"
ROUTINGONLY="YES">
...

```

4. Save `webcache.xml`.
5. Restart Oracle Web Cache with the following command:

```
opmnctl restartproc ias-component=component_name
```

This executable is found in the following directory:

```
(UNIX) ORACLE_INSTANCE/bin
(Windows) ORACLE_INSTANCE\bin
```

6. Verify Oracle Web Cache is running in the load balancer mode from the Oracle Web Cache Manager by verifying the following status message appears beneath the **Apply Changes** and **Cancel Changes** buttons:

Web Cache running in Routing Only Mode with current configuration

Fusion Middleware Control Console does not provide an equivalent verification status.

7. Configure origin servers, as described in the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.
8. Create site definitions and map them to the origin servers, as described in the *Oracle Fusion Middleware Administrator's Guide for Oracle Web Cache*.
9. If your application deployment requires session stickiness, enable session binding. See [Section 9.3.2.3, "Oracle Web Cache Session Binding."](#)

---

## Configuring Identity Management for Maximum High Availability

This chapter provides high-level instructions for setting up a maximum high availability deployment for Oracle Identity Management. This deployment includes two sites in different geographic locations. This is an active-active deployment where both sites are active at the same time when the deployment is functioning normally. If one site fails, the surviving site continues to function.

Each site includes a two-node Oracle Internet Directory cluster configuration, which provides high availability for Oracle Internet Directory. The Oracle Internet Directory cluster configuration at each site uses an Oracle Real Applications Cluster (Oracle RAC) database as the security store, which provides high availability for the database. [Chapter 8, "Configuring High Availability for Identity Management Components"](#) provides an introduction to the high availability Oracle Internet Directory cluster configurations.


Multimaster replication is used to replicate data from the master site to the replica site.

This chapter includes the following topics:

- [Section 10.1, "Introduction to the Maximum High Availability Identity Management Deployment"](#)
- [Section 10.2, "Overview of Replication"](#)
- [Section 10.3, "Setting up Multimaster Replication"](#)

### 10.1 Introduction to the Maximum High Availability Identity Management Deployment

[Figure 10–1](#) shows the maximum high availability deployment for Oracle Identity Management.

**Figure 10–1 Maximum High Availability Multimaster Replication Deployment**

The master site is located in New York and the replica site is located in Los Angeles.

Each site includes a highly available two-node Oracle Internet Directory cluster configuration that uses an Oracle RAC database as a highly available security store. Each two-node cluster has a load balancer. See [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps"](#) for information on setting up a two-node Oracle Internet Directory cluster.

The master site in New York consists of:

- OIDHOST1 and OIDHOST2

These are the two clustered hosts on which Oracle Internet Directory is installed.

- RAC\_DB1

This is the Oracle RAC database which serves as the security store for the Oracle Internet Directory instances on OIDHOST1 and OIDHOST2. Multimaster replication is used to replicate data between RAC\_DB1 in New York and RAC\_DB2 in Los Angeles.

The replica site in Los Angeles consists of:

- OIDHOST3 and OIDHOST4

These are the two clustered hosts on which Oracle Internet Directory is installed.

- RAC\_DB2

This is the Oracle RAC database which serves as the security store for the Oracle Internet Directory instances on OIDHOST3 and OIDHOST4. Multimaster replication is used to replicate data between RAC\_DB1 in New York and RAC\_DB2 in Los Angeles.



## 10.2 Overview of Replication

The following types of replication are available for Oracle Internet Directory:

- LDAP multimaster replication  
Uses the industry-standard Lightweight Directory Access Protocol Version 3 as the replication transport mechanism. This is the recommended protocol to use for replication.
- Oracle Advanced Database multimaster replication  
Uses the replication feature of Oracle Database. This is also called Advanced Replication.
- Two-way fan-out replication  
With this replication method, the replicated data is read/write at both the master site and replica site. Fan-out uses LDAP as its transport mechanism.
- One-way fan-out replication  
With this replication method, the replicated data is read-only at the replica site. Fan-out uses LDAP as its transport mechanism.

For more information about the replication types for Oracle Internet Directory, refer to *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

For the maximum availability deployment shown in [Figure 10–1](#), either LDAP or Oracle Advanced Database multimaster replication can be set up.

## 10.3 Setting up Multimaster Replication

This section describes how to set up LDAP multimaster replication or Oracle Advanced Database multimaster replication for the maximum high availability Oracle Internet Directory deployment shown in [Figure 10–1](#).

---

**Note:** See [Section 8.3.3, "Oracle Internet Directory High Availability Configuration Steps"](#) for information on installing the Oracle Internet Directory two-node clusters for the New York and Los Angeles multimaster topology shown in [Figure 10–1](#).

---

It is recommended that you use LDAP multimaster replication for the maximum availability Oracle Internet Directory deployment.

---

**Note:** New Oracle Fusion Middleware 11g customers who want to install and configure 10.1.4.3 or later Oracle Single Sign-On and Oracle Delegated Administration Services against 11g Oracle Internet Directory and to set up multimaster replication should refer to these steps:

1. To configure 10.1.4.3 Oracle Single Sign-On and Oracle Delegated Administration Services to run against 11g Oracle Internet Directory, see the "Special Instructions for Oracle Single Sign-On and Oracle Delegated Administration Services Users" section of *Oracle Fusion Middleware Installation Guide for Oracle Identity Management*.
  2. Perform the steps in the following sections of the "Deploying Identity Management with Multimaster Replication" chapter in the *Oracle Fusion Middleware High Availability Guide* for release 10.1.4.0.1 (part number B28186-01) to install and configure 10.1.4.3 Oracle Single Sign-On and Oracle Delegated Administration Services for multimaster replication:
    - Section 10.1.4 "Installing OracleAS Single Sign-On and Oracle Delegated Administration Services on the Master Node"
    - Section 10.1.5 "Synchronizing the OracleAS Single Sign-On Schema Password"
    - Section 10.1.6 "Installing OracleAS Single Sign-On and Oracle Delegated Administration Services on the Replica Node"
    - Section 10.1.7 "If You Are Running in SSL Mode"
- 

### 10.3.1 Setting Up LDAP Multimaster Replication

Follow these steps in the Oracle Enterprise Manager Fusion Middleware Control to set up LDAP multimaster replication:

1. From the Oracle Internet Directory menu on the Oracle Internet Directory instance home page, choose **Administration**, and then **Replication Management**.
2. You are prompted to log into the replication DN account. Provide the host, port of one of the Oracle Internet Directory servers at the master site (the New York cluster in [Figure 10–1](#)), replication DN, and replication DN password. If anonymous binds are enabled on this Oracle Internet Directory component, the replication DN field will fill in automatically when you enter the host and port.
3. Click the Create icon.
4. On the Type screen, select the replication type: **Multimaster Replication**.
5. Click **Next**. The Replicas screen shows the replication type you selected.
6. Provide an agreement name that is unique across all nodes.
7. For multimaster replication, enter the host, port, user name (replication DN), and replication password for the primary node and all the secondary nodes.

---

**Note:** Enter the host/port of any of the Oracle Internet Directory instances in the cluster.

---

8. Click **Next** to go the Settings page.
9. In the **LDAP Connection** field, select **Keep Alive** if you want the replication server to use same connection for performing multiple LDAP operations. Select **Always Use New Connection** if you want the server to open a new connection for each LDAP operation.
10. Enter the **Replication Frequency**.
11. Enter the **Human Intervention Queue Schedule**. This is the interval, in minutes, at which the directory replication server repeats the change application process.
12. The Replication Server Start Details section has options to start the replication server and enable bootstrap. Choose **Start Server** to start the appropriate server instance. You can also enable bootstrap by choosing **Enable Bootstrap**. You must select the **Instance Name** and **Component Name** from the dropdown lists to start the server.
13. Click **Next** to go to the Scope page. The default primary naming context will be filled in. Keep the default settings.
14. Click **Next**. The Summary page shows a summary of the replication agreement you are about to create. To make any changes to information on the Summary page, click **Back**.
15. Click **Finish** to create the replication agreement.

For detailed instructions on setting up LDAP multimaster replication, read the Oracle Enterprise Manager tool tips or refer to *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

#### 10.3.1.1 Adding a Node in LDAP Multimaster Replication

Follow these steps in the Oracle Enterprise Manager Fusion Middleware Control to add a node in an LDAP multimaster replication deployment:

1. From the Oracle Internet Directory menu on the Oracle Internet Directory instance home page, select **Administration** then **Replication Management**.
2. You will be prompted to log into the replication DN account. Provide the host, port and replication DN password of any of the replicas in the multimaster replication deployment.
3. In the upper half of the screen, click on the appropriate multimaster replication agreement row to enable editing.
4. Click **Edit** on the Replication Agreements page.
5. In the lower half of the screen, click the Replicas tab.
6. To add a new replica to the multimaster replication deployment, click the **Create** icon.
7. Enter the host, port and replication DN password details for the new node. Click **Add**.
8. Click **Apply**. The replica is added to the existing multimaster directory replication group (DRG).

For more information, see *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

### 10.3.1.2 Deleting a Node in LDAP Multimaster Replication

Follow these steps in the Oracle Enterprise Manager Fusion Middleware Control to delete a node in an LDAP multimaster replication deployment:

1. From the Oracle Internet Directory menu on the Oracle Internet Directory instance home page, select **Administration**, and then **Replication Management**.
2. You will be prompted to log into the replication DN account. Provide the host, port and replication DN password of any of the replicas in the multimaster replication deployment.
3. In the upper half of the screen, click on the appropriate multimaster replication agreement row to enable editing.
4. Click **Edit** on Replication Agreement screen.
5. In the lower half of the screen, click the Replicas tab.
6. Click the replica you want to delete from the multimaster replication deployment. The **Delete** icon becomes enabled.
7. Click the **Delete** icon.
8. Click the **Apply** button to remove the replica from the LDAP multimaster directory replication group (DRG).

For more detailed information on deleting a node in an LDAP multimaster replication deployment, see *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

## 10.3.2 Setting Up Oracle Advanced Database Multimaster Replication

The detailed steps for setting up Oracle Advanced Database multimaster replication are available in the "Installing and Setting Up an Oracle Database Advanced Replication-Based Multimaster Replication Group" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

[Table 10-1](#) shows the subsections of the "Installing and Setting Up an Oracle Database Advanced Replication-Based Multimaster Replication Group" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*. It also describes the instructions to perform in each subsection to set up Oracle Advanced Database multimaster replication for the maximum high availability Oracle Internet Directory deployment shown in [Figure 10-1](#).

**Table 10–1 Steps for Setting Up Oracle Database Advanced Multimaster Replication**

| Subsection                                                                         | Instructions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Task 1: Install Oracle Internet Directory on the Master Definition Site (MDS)      | This task should already have been performed based on the Note in <a href="#">Section 10.3, "Setting up Multimaster Replication."</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Task 2: Install the Oracle Internet Directory on the Remote Master Sites (RMS)     | This task should already have been performed based on the Note in <a href="#">Section 10.3, "Setting up Multimaster Replication."</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| If an Existing Master is Used as a Remote Master Site                              | <p>Set up one site as the master definition site (MDS), for example, the New York site. Then set up one Oracle Internet Directory node in the cluster at the master site (for example, OIDHOST1) to be the master host. OIDHOST1 will be the host in the master site cluster where the replication server will be configured and will run. When the setup steps require a reference to a replication server, process, or port for the MDS, specify the correct value for OIDHOST1.</p> <p>Set up another site as the remote master site (RMS), for example, the Los Angeles site. Then set up one Oracle Internet Directory node in the cluster at the remote site (for example, OIDHOST3) to be the replica host. The replica host is referred to as the "new node" in the "If an Existing Master is Used as a Remote Master Site" section). OIDHOST3 will be the host in the Los Angeles cluster where the replication server will be configured and will run. When the setup steps require a reference to a replication server, process, or port for the RMS, specify the correct value for OIDHOST3.</p> |
| Task 3: Set Up Advanced Replication for a Directory Replication Group              | Follow the instructions for this task.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| On All Nodes, Prepare the Oracle Net Services Environment for Replication          | Perform the steps in this subsection in all the database Oracle homes and in all the Oracle Internet Directory Oracle homes in the New York site and the Los Angeles site.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| From the MDS, Configure Advanced Replication For Directory Replication             | <p>Perform the steps in this subsection in all of the Oracle Internet Directory Oracle homes in New York and Los Angeles, with one exception:</p> <p>When you configure Advanced Replication using the Replication Environment Management Tool, execute the command on only the master host at the MDS site (for example, on OIDHOST1 in New York). The replication must be started on only one Oracle Internet Directory host.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Task 4 (Optional): Load Data into the Directory                                    | If you choose to use the bulkload utility, stop all the Oracle Internet Directory instances in all the Oracle Internet Directory homes, and use only one of the Oracle Internet Directory instances to perform the bulkload operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Task 5: Ensure that Oracle Directory Server Instances are Started on All the Nodes | Perform this task in all the Oracle Internet Directory Oracle home directories.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Task 6: Start the Replication Servers on All Nodes in the DRG                      | Start the replication server on only OIDHOST1 in New York and OIDHOST3 in Los Angeles.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

### 10.3.2.1 Adding a Node in Oracle Advanced Database Multimaster Replication

The detailed steps for adding a node in an Oracle Advanced Database multimaster replication deployment are in the "Adding a Node for Oracle Database Advanced Replication-Based Multimaster Replication Group" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

[Table 10–2](#) shows the subsections of the "Adding a Node for Oracle Database Advanced Replication-Based Multimaster Replication Group" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*. It also describes the instructions to perform in each subsection to add a node to the maximum high availability Oracle Internet Directory deployment shown in [Figure 10–1](#).

**Table 10–2 Steps for Adding a Node in Oracle Advanced Database Replication**

| Subsection                                                                               | Instructions                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prepare the Oracle Net Services Environment                                              | Perform the steps in this subsection in all the database Oracle homes and in all the Oracle Internet Directory Oracle homes in the master definition site (New York site) and the remote master definition site (Los Angeles site).<br><br>Also, perform these steps in the database Oracle homes and in all the Oracle Internet Directory Oracle homes for the new cluster that is being added. |
| Task 1: Stop the Directory Replication Server on All Nodes                               | Stop the replication server on OIDHOST1 in New York and on OIDHOST3 in Los Angeles.                                                                                                                                                                                                                                                                                                              |
| Task 2: Identify a Sponsor Node and Install Oracle Internet Directory on the Remote Site | OIDHOST1 in the New York cluster will be the sponsor node and the MDS.                                                                                                                                                                                                                                                                                                                           |
| Task 3: Switch the Sponsor Node to Read-Only Mode                                        | Perform this task in the Oracle home for Oracle Internet Directory on OIDHOST1 and OIDHOST2.                                                                                                                                                                                                                                                                                                     |
| Task 4: Back up the Sponsor Node by Using Idifwwrite                                     | Perform this task in the Oracle home for Oracle Internet Directory on OIDHOST1.                                                                                                                                                                                                                                                                                                                  |
| Task 5: Perform Advanced Replication Add Node Setup                                      | Perform this task in the Oracle home for Oracle Internet Directory on OIDHOST1.                                                                                                                                                                                                                                                                                                                  |
| Task 6: Switch the Sponsor Node to Updatable Mode                                        | Perform this task in the Oracle home for Oracle Internet Directory on OIDHOST1 and OIDHOST2.                                                                                                                                                                                                                                                                                                     |
| Task 7: Start the Directory Replication Server on All Nodes Except the New Node          | Perform this task on OIDHOST1 in the New York cluster and on OIDHOST3 in the Los Angeles cluster.                                                                                                                                                                                                                                                                                                |
| Task 8: Load Data into the New Node by Using bulkload                                    | Perform this task on one of the Oracle Internet Directory Oracle homes in the new cluster that is being added.<br><br>Stop all the Oracle Internet Directory processes on the new node before using bulkload.                                                                                                                                                                                    |
| Task 9: Start the Directory Server on the New Node                                       | Perform this task on all of the Oracle Internet Directory nodes in the new cluster that is being added.                                                                                                                                                                                                                                                                                          |
| Task 10: Start the Directory Replication Server on the New Node                          | Perform this task on one of the Oracle Internet Directory Oracle homes in the new cluster that is being added.                                                                                                                                                                                                                                                                                   |

### 10.3.2.2 Deleting a Node in Oracle Advanced Database Multimaster Replication

The detailed steps for deleting a node in an Oracle Advanced Database multimaster replication deployment are in the "Deleting a Node from a Multimaster Replication Group" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

[Table 10–3](#) shows the subsections of the "Deleting a Node from a Multimaster Replication Group" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*. It also describes the instructions to perform in each subsection to delete a node in the maximum high availability Oracle Internet Directory deployment shown in [Figure 10–1](#). In the instructions below, the MDS is assumed to be the New York site.

**Table 10–3 Steps for Deleting a Node in Oracle Advanced Database Replication**

| Subsection                                                                     | Instructions                                                                    |
|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Task 1: Stop the Directory Replication Server on All Nodes                     | Perform this task on each node in the Directory Replication Group (DRG).        |
| Task 2: Stop All Oracle Internet Directory Processes in the Node to be Deleted | Perform this task in the node to be deleted.                                    |
| Task 3: Delete the Node from the Master Definition Site                        | Perform this task in the Oracle home for Oracle Internet Directory on OIDHOST1. |
| Task 4: Start the Directory Replication Server on All Nodes                    | Perform this task on all the remaining nodes in the DRG.                        |





---

# Configuring High Availability for Enterprise Content Management

Oracle Enterprise Content Management (ECM) Suite, an Oracle Fusion Middleware component, is an integrated suite of products designed for managing content. It is the industry's most unified enterprise content management platform, which enables you to leverage industry-leading document management, Web content management, digital asset management, and records management functionality to build your business applications. Building a strategic enterprise content management infrastructure for content and applications helps you to reduce costs, easily share content across the enterprise, minimize risk, automate expensive, time-intensive and manual processes, and consolidate multiple Web sites onto a single platform.

Oracle ECM offers the following benefits:

- Superior usability: Built-in support for end users, workgroups, content experts, process owners, administrators, and webmasters
- Optimized management: Unified architecture for securely managing documents, files, web content and digital assets
- Hot-pluggable: Out-of-the-box support for Oracle and third-party repositories, security systems, and enterprise applications

This chapter describes some of the Oracle Enterprise Content Management components from a high availability perspective. This chapter outlines single-instance concepts that are important for designing a high availability deployment.

The chapter includes the following topics:

- [Section 11.1, "Oracle Imaging and Process Management High Availability"](#)
- [Section 11.2, "Oracle Universal Content Management High Availability"](#)
- [Section 11.3, "Oracle ECM High Availability Configuration Steps"](#)

## 11.1 Oracle Imaging and Process Management High Availability

This section provides an introduction to Oracle Imaging and Process Management (Oracle I/PM) and describes how to design and deploy a high availability environment for Oracle I/PM.

### 11.1.1 Oracle I/PM Component Architecture


Oracle I/PM provides organizations with a scalable solution focused on process-oriented imaging applications and image-enabling business applications, such as Oracle E-Business Suite and PeopleSoft Enterprise. It enables annotation and

markup of images, automates routing and approvals, and supports high-volume applications for trillions of items.

Oracle I/PM is a transactional content management tool; it manages each document as part of a business transaction or process. A transactional document is important to you mostly during the time when the transaction is occurring, typically a relatively short period of time.

Figure 11–1 illustrates the Oracle I/PM internal components (components inside the Oracle I/PM package).

**Figure 11–1 Oracle I/PM Components for a Single Instance**



Note the following before you begin setting up an Oracle I/PM configuration:

- Oracle I/PM is a relatively thin layer of functionality on top of the Oracle Universal Content Manager (Oracle UCM) repository. Oracle I/PM provides an imaging facade on Oracle UCM. The repository performs a major portion of resource utilization. For this reason, Oracle recommends configuring Oracle UCM by putting an Oracle UCM instance with every Oracle I/PM instance. This configuration provides good utilization of hardware.
- Oracle I/PM's largest performance factor is the movement of document content. Oracle I/PM is usually used in high-volume scenarios in which hundreds of thousands of documents are ingested per day. These documents are searched for, retrieved, and viewed frequently during the day. It is the movement of document

content in and out of the system that has the largest impact on performance. Oracle I/PM stores TIF documents that average about 25K bytes in size. However, TIF documents can have multiple pages, which increases their size. Oracle I/PM also supports 400 other file formats of varying sizes, including larger sizes.

#### 11.1.1.1 Oracle I/PM Component Characteristics

Oracle I/PM is a JEE application that runs in the WebLogic Server Managed Server environment. It includes the following components and subcomponents:

- Servlets: Oracle I/PM has a servlet for internal use that provides a mechanism for the REST API and I/PM Applet Viewer to access the public API functions of the system.
- EJBs: Oracle I/PM uses some EJB concepts internally but has no external facing EJBs.
- Oracle I/PM leverages JMS to drive both its Input Agent and its BPEL Agent message-driven beans (MDBs). In [Figure 11-1](#), the BPEL Agent is the collection of BPEL MDBs. These queues should be persistent and provide guaranteed delivery. They should also distribute work across clusters.
- Oracle I/PM leverages JAX-WS and JAX-B to implement its Web services and serialization mechanisms.
- Oracle I/PM leverages TopLink for its definition persistence as well as configuration. The objects are generally the definitions for applications and searches. In addition, TopLink is used to get user and system preferences.
- Oracle I/PM leverages Java Object Cache (JOC) for an internal cache of documents retrieved from the repository and being processed for viewing. These caches exist only locally, and are not shared across the servers.
- The interaction between AXF and BPEL is done through the BPEL remote API, which uses RMI via EJB.

---

**Note:** Oracle I/PM does not use Java Transaction API (JTA) because Oracle I/PM does not support multi-component transactions.

---

**11.1.1.1.1 Oracle I/PM State Information** Because Oracle I/PM has a stateful UI but does not support session replication, it requires a sticky load balancing router.

Invocations to the Oracle I/PM services using the Oracle I/PM API are stateless.

**11.1.1.1.2 Oracle I/PM Runtime Processes** Oracle I/PM's core uses a service-based architecture with a central API, UI and Web Service interface. The UI, which is ADF based, uses Oracle I/PM's public API to access the core services. End users can access this API via Web Services. Oracle I/PM also provides a thin Java client library that exposes those Web services in a more convenient packaging for Java developers.

Oracle I/PM has these agents to perform background processing:

- The Input Agent watches an input directory for incoming lists of documents to ingest.
- The BPEL Agent responds to document create events and creates new BPEL process instances with metadata from the document and a URL to the document attached. In [Figure 11-1](#), the BPEL Agent is the collection of BPEL MDBs.

- The Ticket Agent is a stateless EJB that is invoked off a timer. It is not shown in [Figure 11-1](#). The Ticket Agent is responsible for identifying and removing expired URL tickets. These tickets are used to create clean URLs without parameters.

The Oracle I/PM service uses the Oracle UCM document repository for the storage of its scan documents. The Oracle UCM repository is a standalone product to which Oracle I/PM communicates via a TCP/IP socket-based connection mechanism provided by Oracle UCM called RIDC.

**11.1.1.1.3 Oracle I/PM Process Lifecycle** Oracle I/PM does not provide any specific death detection, but WebLogic Server Node Manager can be used to keep the server running.

The steps in the Oracle I/PM process lifecycle are:

1. Oracle I/PM is mostly initialized after the domain configuration is performed. The remaining initialization happens when Oracle I/PM starts and the first user logs on.
2. When the first user logs on, full security is granted to the user under the assumption that he can subsequently grant permissions to others. Oracle I/PM provides a self seeding database security mechanism. If Oracle I/PM starts and there is no security defined, the first user logging into the system is granted access to all things.
3. After Oracle I/PM's first startup Oracle I/PM is "empty." It is the administrative user's responsibility to create connections to the Oracle UCM document repository and BPEL servers.
4. Once the first UCM connection is created, the Oracle UCM repository initializes for Oracle I/PM's use.
5. Then, appropriate business object must be created including Applications, Searches, and Inputs. The administrative user may choose to import these business objects rather than create them from scratch.
6. If the Agent User is specified in step 1, the agents will be waiting for JMS queue events to begin their work after the servers has been restarted
7. After the server is running and the JAX-WS mechanics are ready, Oracle I/PM can receive client requests.

**11.1.1.1.4 Oracle I/PM Configuration Artifacts** Oracle I/PM has the following configuration artifacts:

- All Oracle I/PM configuration data is stored within the Oracle I/PM database. Configuration does not reside in local files. The configuration parameters are exposed either through the appropriate MBeans (available from Enterprise Manager or WLST) or through Oracle I/PM's UI.
- The other Oracle I/PM configuration artifacts are the definition objects, which are described in [Section 11.1.1.1.5, "Oracle I/PM External Dependencies."](#)
- Oracle I/PM uses Oracle UCM mechanisms for customizing Oracle UCM for use with Oracle I/PM, including management of UCM Profiles and Information Fields.

**11.1.1.1.5 Oracle I/PM External Dependencies** Oracle I/PM has the following external dependencies:

- Oracle I/PM requires a database to store its configuration. The database is initially create via RCU and is managed through WebLogic Server JDBC data sources. The database is accessed via TopLink.

- Oracle I/PM leverages a variety of Oracle and Java technologies, including JAX-WS, JAX-B, ADF, TopLink, and JMS. These are included with the installation and do not require external configuration.
- Oracle I/PM provides Mbeans for configuration. These are available through WLST and the EM Mbean browser. Oracle I/PM provides a few custom WLST commands.
- Oracle I/PM depends on the existence of an Oracle UCM repository.

The following clients depend on Oracle I/PM:

- The Oracle I/PM UI is built on its public toolkit (the Oracle I/PM API). The UI and core API are distributed in the same EAR file.
- Oracle I/PM provides Web services for integrations with other products, and provides a set of Java classes that wrap those Web services as a Java convenience for integrating with other products.
- Oracle I/PM provides a URL toolkit to facilitate other applications interacting with Oracle I/PM searches and content.
- Oracle I/PM provides a REST toolkit for referencing individual pages of documents for web presentation.

Clients connect to Oracle I/PM through either:

- JAX-WS for the Web services and/or Java client. These connections are stateless and perform a single function.
- HTTP for access to URL and REST tools:
  - REST requests are stateless and perform a single function.
  - URL tools log into the Oracle I/PM UI and provide a more stateful experience with the relevant UI components.

**11.1.1.1.6 Oracle I/PM Log File Location** All log messages log to server log files of the WebLogic Server that the application is deployed on.

The default location of the diagnostics log file is:

```
WL_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```


You can use Oracle Enterprise Manager to easily search and parse these log files.

## 11.1.2 Oracle I/PM High Availability Concepts

This section provides conceptual information about using Oracle I/PM in a high availability two-node cluster.

### 11.1.2.1 Oracle I/PM High Availability Architecture

[Figure 11–2](#) shows an Oracle I/PM high availability architecture:

**Figure 11–2 Oracle I/PM High Availability Architecture Diagram**

Oracle I/PM can be configured in a standard two node active-active high availability configuration. Although Oracle I/PM's UI layer will not fail over between active servers, its other background processes will.

In the Oracle I/PM high availability configuration shown in [Figure 11–2](#):

- The Oracle I/PM nodes are exact replicas of each other. All the nodes in a high availability configuration perform the exact same services and are configured through the centralized configuration database, from which all servers pull their configuration.

- A load balancing router with sticky session routing must front-end the Oracle I/PM nodes. Oracle HTTP Server can be used for this purpose.
- Oracle I/PM can run both within a WebLogic Server cluster or without.
- All the Oracle I/PM instances must share the same database connectivity. This can be Oracle RAC for increased high availability potential. Oracle I/PM uses TopLink for database actions.
- Oracle I/PM requires one common shared directory from which its Input Agent draws inbound data files staged for ingestion. Once an Input Agent pulls input definition files, processing of the input file and images associated will stay isolated to that WebLogic Server instance.

**11.1.2.1.1 Starting and Stopping the Cluster** Oracle I/PM's agents (Input Agent, BPEL Agent) start as integral part of the Oracle I/PM application residing in the WebLogic Server in the cluster. Based on outstanding work sitting in the JMS queues (which are persisted and preserved across failures), they begin processing immediately. The Input Agent also looks for work in its inbound file directory. If there are files to be processed, they are pushed to the corresponding JMS queues and a server in the cluster consumes it and processes the associated images. The numbers of threads dedicated to the BPEL Agent and Input Agent are controllable through the corresponding WebLogic Server workload managers.

When servers in a cluster stop, the agents finish their activity. The BPEL Agent has a short run cycle and all work is likely completed before shutdown. In flight BPEL invocations are retried three times and JMS logs preserve pending operations. The Input Agent can process files of considerable size; a large input file can take hours to process. The amount of work pending after a stop is preserved in the associated JMS persistence stores. When you restart the server that hosts the Input Agent, the agent starts processing where it left off. Processing image files continues in the server that initially consumed the corresponding input file after restart or server migration.

**11.1.2.1.2 Cluster-Wide Configuration Changes** Oracle I/PM configuration is stored in the Oracle I/PM database and shared by all the Oracle I/PM instances in the cluster.

The number of threads for the BPEL Agent and Input Agent for a particular Oracle I/PM instance is a WebLogic Server property (stored as configuration of the Oracle I/PM application) and is controlled through the Administration Console or with WLST commands.

### 11.1.2.2 Protection from Failures and Expected Behaviors

Follow these recommendations to protect the Oracle I/PM high availability configuration from failure:

- Use a load balancing router with sticky session routing to front-end the Oracle I/PM nodes. Oracle HTTP Server can be used for this purpose.
- Configure timeouts for the load balancing router based on the size of the documents being uploaded or retrieved by Oracle I/PM for viewing.

Also, the following features help protect the Oracle I/PM high availability configuration from failure:

- JMS failover is provided by WebLogic Servers's JMS implementation.
- TopLink provides database retry logic.
- For connections to Oracle UCM, retry logic exists for immediate retry attempts to resolve network glitches. Oracle I/PM provides a failover mechanism to find other

Oracle UCM servers in a cluster. The customer can use a load balancer between Oracle I/PM and Oracle UCM to also achieve load balancing/failover goals.

- For connections to BPEL, Oracle I/PM performs an immediate retry attempt and then push failed items back into the associated JMS queue for retry after a longer timeout (5 minutes). This JMS retry mechanism is performed three times. Any objects continuing to fail go into a holding queue for user intervention.
- API level integrations (such as JPS) are not retried and are assumed to have internal retry logic.

### **Node Failure**

An Oracle/IPM node failure affects the Agent processes for that node. Any unfinished transactions are abandoned and lost. During a failover, the user sessions in the Oracle I/PM UI for that node are lost. Users must log in again and restart. Since each Oracle I/PM cluster node is independent, failover is managed by the load balancing router and not by Oracle I/PM directly.

Because all the Oracle I/PM configuration information is stored in the database and all the nodes are replicas of each other, the configuration information is immediately available to the new node after failover.

These are the implications of an Oracle I/PM node failure on external resources:

- For JMS transactions, the assumption is that atomic actions done to JMS are internally transactional, and that any messages received from the queue that had not been completed will be resubmitted. BPEL Agent and Input Agent failure recovery relies on persistence of the local messages to a file system. In the event that messages are in the local node, they will be continued upon start of that node. Other messages will be picked up by nodes still running.
- For database transactions, the database is not responsible for a client failure except that it forces an implicit rollback operation when the disconnect occurs. Database transactions from a database outage will be managed by the necessary database facilities (for example, a failover when an Oracle RAC database is being used, and no failover when a single node database is being used). If there is a node failure in the middle of a database access, the user will most likely need to log in again to retry the UI operation.
- User sessions within Oracle I/PM are lost during an Oracle I/PM server failure. If you are uploading documents through the Oracle I/PM UI and an Oracle I/PM server failure occurs, it is possible that you will not receive confirmation of successful manual document uploads. In this case, it is recommended that you search for the existence of the documents that you were uploading when the server failure occurred before you attempt to upload them again; this will help avoid the creation of duplicate documents.

If an external application accesses an Oracle I/PM cluster by means of a WebService invocation, it is expected that the front end load balancer will provide the appropriate failover logic when an Oracle I/PM server fails. If the application interacts with the Oracle I/PM system by directly pushing files to the input directory, it is the application's responsibility to retry should a failure in the file system occur.

If an Input definition file is being processed during the time of the outage, Oracle I/PM will maintain the file details and upon restart of that specific node, the local JMS queue will respawn that JMS message for retry. Upon retry, the Input Ingestor will attempt to determine where specifically in the file the shutdown occurred, to pick up where it left off so that no documents from the batch are lost. The definition file is



maintained in the Stage directory until the specific node that was processing the definition file is able to fully restart and finish the file.

The input file is the smallest unit of work that the Input Agent can schedule and process. There are multiple elements to be taken into consideration to achieve the highest performance, scalability, and high availability:

- All of the machines in an Oracle I/PM cluster share a common input directory.
- Input files from this directory are distributed to each machine via a JMS queue.
- The frequency with which the directory is polled for new files is configurable via the CheckInterval MBeans configuration setting.
- Each machine has a specified number of Input Ingestors, responsible for parsing of the Input Definition files. The number of ingestors is controlled via the Workload manager for the InputAgent and is managed by WebLogic Server.

The optimum performance will be achieved when:

- Each Oracle I/PM cluster node has the maximum affordable number of parsing agents configured via the Work Manager without compromising the performance of the other I/PM activities, such as the user interface and Web services.
- The inbound flow of documents is partitioned into input files containing the appropriate number of documents. On average there should be two input files queued for every parsing agent within the cluster.
- If one or more machines within a cluster fails, active machines will continue processing the input files. Input files from a failed machine will remain in limbo until the server is restarted. Smaller input files ensure that machine failures do not place large numbers of documents into this limbo state.

For example: Consider 10,000 inbound documents per hour being processed by two servers. A configuration of 2 parsing agents per server produces acceptable overall performance and ingests 2 documents per second per agent. The 4 parsing agents at 2 documents per second is 8 documents per second, or 28,800 documents per hour. Note that a single input file of 10,000 documents will not be processed in an hour since a single parsing agent working at 7,200 documents per hour will be unable to complete it. However, if you divide the single input file up into 8 input files of 1,250 documents, this ensures that all 4 parsing agents are fully used, and the 10,000 documents are completed in the one hour period. Also, if a failure should occur in one of the servers, the other can continue processing the work remaining on its parsing agents until the work is successfully completed.

### **Outages in Components on Which Oracle I/PM Depends**

Oracle I/PM integrates with Oracle SOA Suite. If the SOA server is taken down, Oracle I/PM will experience failures. Oracle I/PM will retry actions three times. Each attempt is followed by a configurable delay time before it is repeated. If after the third attempt the action cannot be completed, then these requests are queued in a database table to be resubmitted at a later time. The resubmission is accomplished through WLST commands.

Oracle I/PM is not dependent on Oracle WebCenter Portal.

Other Oracle ADF applications do not affect the operation of Oracle I/PM.

The Oracle I/PM UI allows specifying a list of Oracle UCM servers to which it connects. It uses a primary server to connect to and allows specifying multiple secondary servers. In case of a failure in the primary UCM server, Oracle I/PM retries

the next UCM server in the list. If all the UCM servers are down, an exception is returned and traced in the Oracle I/PM server's log.

### 11.1.2.3 Creation of Oracle I/PM Artifacts in a Cluster

Searches, inputs, applications (all Oracle I/PM structural configurations) are stored within the Oracle I/PM database. They become immediately available to all machines in the cluster, including the Web UI layers. The Web UI does not automatically refresh when a new entity is created in another Oracle I/PM server in the cluster or in the same server by a different user session; however, you can refresh your web browser to cause a navigation bar refresh for the imaging application.

### 11.1.2.4 Troubleshooting Oracle I/PM

This section describes how to trouble shoot the Oracle I/PM server, Advanced Viewer, and Input Agent in a high availability environment.

Additional Oracle I/PM troubleshooting information can be found in *Oracle Fusion Middleware Administrator's Guide for Oracle Imaging and Process Management*.

#### Oracle I/PM Server Troubleshooting

Oracle/PM is a JEE application deployed on WebLogic Server. All log messages log to server log files of the WebLogic Server that the application is deployed on.

The default location of the diagnostics log file is:

```
WL_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

You can use Oracle Enterprise Manager to easily search and parse these log files.

#### Advanced Viewer Troubleshooting

Oracle I/PM primarily provides server side logging, although it does have a client-side component called the Advanced Viewer. This Java Applet uses the standard logging utilities; however, these utilities are disabled in client situations. These utilities can be configured to direct the logging output to either a file or the Java Console via appropriate changes to the logging.properties file for the Java plug-in.

#### Input Agent Troubleshooting

When the Input Agent for Oracle I/PM is uploading documents, if Oracle Universal Content Management fails, messages such as these are generated:

```
Filing Invoices Input completed successfully with 48 documents processed
successfully out of 50 documents.
```

In addition, an exception similar to this is generated:

```
[2009-07-01T16:43:01.549-07:00] [IPM_server2] [ERROR] [TCM-00787]
[oracle.imaging.repository.ucm] [tid: [ACTIVE] .ExecuteThread: '10' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: <anonymous>] [ecid:
0000I8sgmWtCgon54nM6UilAIvEo00005N,0] [APP: imaging#11.1.1.1.0]
A repository error has occurred.
Contact your system administrator for assistance.
[[oracle.imaging.ImagingException: TCM-00787: A repository error has occurred.
```

```
Contact your system administrator for assistance.
stackTraceId: 5482-1246491781549
faultType: SYSTEM
faultDetails:
```

```

ErrorCode = oracle.stellent.ridc.protocol.ServiceException,
ErrorMessage = Network message format error.
 at
oracle.imaging.repository.ucm.UcmErrors.convertRepositoryError(UcmErrors.
java:109
.
.
.

```

The Input Agent handles errors by placing the failed documents in an error file, so this behavior and the error messages that are generated under these circumstances are expected.


## 11.2 Oracle Universal Content Management High Availability

This section provides an introduction to Oracle Universal Content Management (Oracle UCM) and describes how to design and deploy a high availability environment for Oracle UCM.

### 11.2.1 Oracle Universal Content Management Component Architecture

Figure 11–3 shows the Oracle UCM single instance architecture.

**Figure 11–3 Oracle Universal Content Management Single Instance Architecture**



An instance of Oracle UCM runs within an Oracle WebLogic Managed Server environment. WebLogic Server is responsible for the lifecycle of the Oracle UCM application, such as starting and stopping it.

There are two entry points into the Oracle UCM application. One is a listening TCP socket, which can accept connections from clients such as Oracle I/PM or Oracle WebCenter Portal. The second is an HTTP listening port, which can accept Web

services invocations. The HTTP port is the same port as the Managed Server listen port.

The UCM Server also contains several embedded applications. These can run as applets within the client browser. They include a Repository Manager (RM) for file management and indexing functions and a Workflow Admin for setting up user workflows. For more information on any of these applications, see the *Oracle Fusion Middleware Application Administrator's Guide for Content Server*.

The UCM Server has three types of persistent stores: the repository (which can be stored in an Oracle database), the search index (which can be stored in an Oracle database), and metadata and web layout information (which can be stored on the file system).

### 11.2.1.1 Oracle UCM Component Characteristics

Oracle UCM is a servlet that resides in WebLogic Server. A servlet request is a wrapper of a UCM request.

Oracle UCM includes various standalone administration utilities, including BatchLoader and IdcAnalyzer.

JMS and JTA are not used with Oracle UCM.

**11.2.1.1.1 Oracle UCM State Information** A request in Oracle UCM is stateless. The session can be persisted in the database to retain authentication information.

**11.2.1.1.2 Oracle UCM Runtime Processes** Oracle UCM has a service-oriented architecture. Each request comes into the system as a service. The Shared Services layer parses out the requests and hands the request to the proper handler. Each handler typically accesses the underlying repository to process the request. The three types of repositories and the data that they store are shown in [Figure 11-3](#).

**11.2.1.1.3 Oracle UCM Process Lifecycle** You can start and stop Oracle UCM as a WebLogic Server Managed Server. The UCM Administration Server can also be used to start, stop, and configure the UCM instance; the UCM Administration Server still works in the WebLogic Server environment, although it has been deprecated in favor of the WebLogic Administration Console. The load balancer can use the PING\_SERVER service to monitor the status of the UCM instance.

During startup, Oracle UCM loads initialization files and data definitions, initializes database connections, and loads localization strings. Through the Oracle UCM internal component architecture, the startup sequence also discovers internal components that are installed and enabled on the system, and initializes those components. The Search/Indexing engine and file storage infrastructure are also initialized.

A client request is typically serviced entirely by one Oracle UCM instance. Existence of other instance in a cluster does not impact client request.

**11.2.1.1.4 Oracle UCM Configuration Artifacts** Initialization files are stored in the file system and stored in Oracle UCM system directories.

**11.2.1.1.5 Oracle UCM Deployment Artifacts** Oracle UCM uses nostage deployment. That is, all deployment files are local.

**11.2.1.1.6 Oracle UCM External Dependencies** Oracle UCM requires WebLogic Server and an external Oracle database.

The following clients depend on Oracle UCM:

- Oracle Universal Records Management (Oracle URM)
- Oracle Imaging and Process Management (Oracle I/PM)
- Oracle WebCenter Portal

The connection from clients is short-lived, and is only needed for the duration of sessionless service.

Clients can connect to Oracle UCM using the HTTP, SOAP/Web Services, JCR, and VCR protocols.

**11.2.1.1.7 Oracle UCM Log File Locations** Oracle UCM is a J2EE application deployed on WebLogic Server. Log messages are logged in the server log file of the WebLogic Server that the application is deployed on. The default location of the server log is:

```
WL_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```

Oracle UCM can also keep logs in:

```
WebLayoutDir/groups/secure/logs
```

Oracle UCM trace files can be configured to be captured and stored in:

```
IntraDocDir/data/trace
```

To view log files using the Oracle UCM GUI, choose the **UCM** menu and then choose **Administration > Logs**.


To view trace files using the Oracle UCM GUI, choose the **UCM** menu and then choose **System Audit Information**. Click the **View Server Output** and **View Trapped Output** links to view current tracing and captured tracing.

## 11.2.2 Oracle UCM High Availability Concepts

This section provides conceptual information about using Oracle UCM in a high availability two-node cluster.

### 11.2.2.1 Oracle UCM High Availability Architecture

[Figure 11–4](#) shows a two-node active-active Oracle Universal Content Management cluster.

**Figure 11–4 Oracle Universal Content Management Two-Node Cluster**

In the Oracle UCM high availability configuration shown in [Figure 11–4](#):

- Each node runs independently against the shared file system, the same database schema, and the same search indexes. Each client request is served completely by a single node.
- Oracle UCM can run with a WebLogic Server cluster or an external load balancer. Oracle UCM is also transparent to Oracle RAC and multi data source configuration.
- Oracle UCM nodes can be scaled independently, and there is limited inter-node communication. Nodes communicate via writing and reading to a shared file system. This shared file system must support synchronized write operations.

**11.2.2.1.1 Starting and Stopping the Cluster** When the cluster starts, each Oracle UCM node goes through its normal initialization sequence, which parses, prepares, and

caches its resources, prepares its connections, and so on. If a node is part of a cluster then in-memory replication is initiated if other cluster members are available. One or all members of the cluster can be started at any one time.

Shutting down a cluster member will only involve that cluster member being unavailable to service requests. When a server is shut down gracefully, it will finish processing current requests, signal its unavailability, and then release all shared resources and close its file and database connections. All session state will be replicated as well, allowing users originally connected to this cluster member to fail over to another member of the cluster.

**11.2.2.1.2 Cluster-Wide Configuration Changes** At the cluster level, new Oracle UCM features or customization of behaviors can be introduced through Oracle UCM internal components. Nodes need to be restarted to pick up these new changes.

For example, the metadata model can change system wide. For instance, a metadata field can be added, modified, or deleted. The system behavior that was driven by these metadata fields can be changed. This change would automatically be picked up by cluster nodes through notification between nodes.

For changes to occur on each cluster member, the first node needs to have shared folders configured. As long as all the shared folders have the same mount point on each cluster node, the other nodes do not need to make manual changes using WebLogic Server pack/unpack.

### 11.2.2.2 Oracle UCM and Inbound Refinery High Availability Architecture

Oracle UCM can be configured with one or more Inbound Refinery instances to provide document conversions.

Inbound Refinery is a conversion server that manages file conversions for electronic documents such as documents, digital images, and motion video. It also provides thumbnailing functionality for documents and images, the ability to extract and use EXIF data from digital images and XMP data from electronic files generated from programs such as Adobe Photoshop and Adobe Illustrator, and storyboarding for video. For more information about Inbound Refinery, refer to *Oracle Fusion Middleware Administrator's Guide for Conversion*.

The number of Inbound Refinery instances which need to be installed will depend on the amount of conversions which need to be supported. For availability, it is recommended that at least two Inbound Refinery instances be installed. All of the Inbound Refinery instances should be made available as providers to all members of a Content Server cluster.

**11.2.2.2.1 Content Server and Inbound Refinery Communication** All communication between a Content Server and an Inbound Refinery is initiated by the Content Server. This includes:

- Regular checks of all available Inbound Refinery instances to determine their status.
- The initiation of a job request.
- Retrieval of completed jobs.

The status includes information on whether the Inbound Refinery is available to accept requests. This will depend on how busy the Inbound Refinery is. If no Inbound Refinery instances are available, the request is retried.

When jobs are completed, this information will appear as part of the Content Server's status request. If so, the Content Server will initiate download of the completed job.

**11.2.2.2.2 Content Server Clusters and Inbound Refinery Instances** In a Content Server cluster, the list of all available Inbound Refinery instances is shared among members of the cluster.

That is, when a new Inbound Refinery instance is added, it becomes available for any Content Server cluster member to use.

However, each individual Content Server communicates independently with all the available Inbound Refinery instances. There is no shared status in a Content Server cluster.

Likewise, all Inbound Refinery instances operate completely independently, responding to requests from Content Servers. Inbound Refinery instances do not share any information.

**11.2.2.2.3 Inbound Refinery Instances and Load Balancers** Load balancers cannot be used in front of a collection of Inbound Refinery instances for the following reasons:

- A Content Server expects to be able to contact an Inbound Refinery directly to get status and availability information.
- When a job is finished, a Content Server must be able to access the one Inbound Refinery that has the specific job and initiate a download.

**11.2.2.2.4 Inbound Refinery Availability** Requests to Inbound Refinery do not occur in real-time since all requests are asynchronous requests mediated by Content Server. Users access Content Servers, and Content Servers then delegate requests to available Inbound Refinery instances. Enough Inbound Refinery instances should be configured so that if one or more fail, there are always other Inbound Refinery instances available to service requests.

A Content Server will never mark an Inbound Refinery as "failed." If unable to contact an Inbound Refinery, then Content Server will choose another available Inbound Refinery and will continue to check the availability of all Inbound Refinery instances. It is the responsibility of a system administrator to manually remove failed Inbound Refinery instances from the Content Server lists.

### **11.2.2.3 Oracle URM High Availability**

Oracle URM and Oracle UCM share the same architecture. As such, all high availability considerations which apply to Oracle UCM also apply to Oracle URM.

When configuring high availability, there are only two differences to note between Oracle UCM and Oracle URM:

1. Oracle URM will store its file system metadata in a different directory than Oracle UCM.
2. When configuring an HTTP Server in front of an Oracle URM cluster, a different session cookie name must be used than that of Oracle UCM.

More details on these two configuration differences are provided in the section below on the Oracle ECM High Availability Configuration Steps.

### **11.2.2.4 Protection from Failure and Expected Behaviors**

The following features help protect an Oracle UCM node from failure:

- Load balancing can be achieved by using an external load balancer. The session does not need additional sticky state data other than its authentication information.



- The load balancer can be configured in front of Oracle UCM. Multi data sources can be configured with Oracle UCM to support Oracle RAC. Timeouts and retries are also supported.
- Oracle UCM uses standard WebLogic Server persistent session.
- Oracle UCM does not use EJBs.
- Death can be detected using the Oracle UCM PING\_SERVER service. When a node is restarted, it does not affect other nodes in the cluster.
- Failover is handled either by the load balancer in front of the Oracle UCM nodes, or by the WLS module in Apache. Ongoing requests from the client would fail, but subsequent requests are rerouted to active nodes. Unfinished transactions would not complete.
- Configuration information from the node is not lost, because it is stored on the shared file system and in the database.
- When a node fails, a database connection between the node and the database should be reclaimed by the database.
- Although instances normally do not get hung, it can occur because of unknown deadlock conditions in the Oracle UCM product or customized versions of Oracle UCM.
- Users with active sessions should not experience any disruption as their sessions fail over to an active server. This includes users with UCM administration applets open.
- After the installation of Oracle UCM is complete, the following value should be set in the `config.cfg` file. This will allow retry logic during an Oracle RAC failover:

```
ServiceAllowRetry=true
```

If this value is not set, user will need to manually retry any operation that was in progress when the failover began.

The value can be set using the WebLogic Server Administration Console for Oracle UCM at the following URL:

```
http://<hostname>:<port>/cs
```

Go to **Administration > Admin Server > General Configuration** and add `ServiceAllowRetry=true` to Additional Configuration Variables. Save and restart all the UCM managed servers. The new value should appear in `config.cfg` at the following location:

```
SHARED_DISK_LOCATION_FOR_UCM/ucm/cs/config/config.cfg
```

#### 11.2.2.5 Troubleshooting Oracle UCM High Availability

To troubleshoot Oracle UCM, look in the log files and then in the trace files.

See [Section 11.2.1.1.7, "Oracle UCM Log File Locations"](#) for information on the location of Oracle UCM log and trace files and information on how to view log and trace files using the Oracle UCM GUI.

## 11.3 Oracle ECM High Availability Configuration Steps

In a high availability environment such as the one shown in [Figure 11–5](#), it is recommended that Oracle ECM products be set up in a clustered deployment, where

the clustered instances access an Oracle Real Applications Cluster (Oracle RAC) database content repository and a shared disk is used to store the common configuration.

A hardware load balancer routes requests to multiple Oracle HTTP Server instances which, in turn, route requests to the clustered Oracle ECM servers.

This section provides the installation and configuration steps for the Oracle ECM high availability configuration shown in [Figure 11–5](#).


It includes the following topics:

- [Section 11.3.1, "Shared Storage"](#)
- [Section 11.3.2, "Configuring the Oracle Database"](#)
- [Section 11.3.3, "Installing and Configuring ECMHOST1"](#)
- [Section 11.3.4, "Installing and Configuring WEBHOST1"](#)
- [Section 11.3.5, "Configuring the Load Balancer"](#)
- [Section 11.3.6, "Installing and Configuring ECMHOST2"](#)
- [Section 11.3.7, "Installing and Configuring WEBHOST2"](#)
- [Section 11.3.8, "Configuring the I/PM Managed Servers"](#)
- [Section 11.3.9, "Configuration of Inbound Refinery Instances"](#)

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

**Figure 11–5 Oracle ECM High Availability Architecture**

The configuration will include the following products:

- Oracle I/PM
- Oracle UCM
- Oracle Universal Records Manager (URM)

- Oracle Inbound Refinery (IBR)

All products will be installed on the same machines, with the exception of Inbound Refinery, which is assumed to be on remote machines.

If you do not want to install Oracle I/PM or Oracle URM, ignore the references to these products in the configuration steps.

### 11.3.1 Shared Storage

Cluster members require shared storage to access the same configuration files. This should be set up to be accessible from all nodes in the cluster.

When configuring Oracle UCM/URM/IBR clusters, follow these requirements:

- For Oracle UCM clusters, all members of the cluster must point to the same configuration directory. This directory must be on a shared disk accessible to all members of the cluster.
- For Oracle URM clusters, all members of the cluster must point to the same configuration directory. This directory must be on a shared disk accessible to all members of the cluster.
- For Oracle IBR, each member of the IBR cluster must have its own separate configuration directory. This directory can be on local disk or shared disk. If the latter, care must be taken that the same directory is not used by multiple IBRs.

Take care if you install UCM and IBR on the same machines; configuration directories for all products are in the same subdirectory. By default, this directory is under *DOMAIN\_HOME* but you can configure this for any other directory.

If IBR servers share the */ibr* directory, an error appears when you try to start the second server:

```
<@csRefineryFileSystemLocked=Inbound Refinery {1} on machine {2} cannot start
because the file {3} exists.
```

The most likely cause of this failure to start is that this IBR is configured as a cluster node; the IBR does not support a clustered environment.

To disable this check set *EnableReserveDirectoryCheck=false* but do not use it to run IBR nodes on the same file system.

### 11.3.2 Configuring the Oracle Database

Configure the Oracle RAC database before completing this procedure.

The Repository Creation Utility (RCU) ships on its own CD as part of the Oracle Fusion Middleware 11g kit.

To run RCU and create the Oracle UCM schemas in an Oracle RAC database repository, follow these steps on any available host:

1. Issue this command:  

```
prompt> RCU_HOME/bin/rcu &
```
2. On the Welcome screen, click **Next**.
3. On the Create Repository screen, select the **Create** operation to load component schemas into an existing database.  
Click **Next**.

4. On the Database Connection Details screen, enter connection information for the existing database as follows:

**Database Type:** Oracle Database

**Host Name:** Name of the computer on which the database is running. For an Oracle RAC database, specify the VIP name or one node name. Example:  
ECMDBHOST1-VIP or ECMDBHOST2-VIP

**Port:** The port number for the database. Example: 1521

**Service Name:** The service name of the database. Example:  
ecmdb.mycompany.com

**Username:** SYS

**Password:** The SYS user password

**Role:** SYSDBA

Click **Next**.

If you entered everything correctly, a dialog box will display with the words "Operation Completed." If so, click **OK**. If not, you will see an error under **Messages**, such as "Unable to connect to the database" or "Invalid username/password."

5. On the Select Components screen, create a new prefix and select the components to be associated with this deployment:

**Create a New Prefix:** for example, DEV

**Components:** Select only the components for which you intend to create schemas. Under **Oracle AS Repository Components > Enterprise Content Management**, choose these components:

- Oracle Content Server 11g - Complete
- Oracle Universal Records Management 11g

Click **Next**.

If you entered everything correctly, a dialog box will display with the words "Operation Completed." If so, click **OK**. If not, you will see an error under **Messages** that you will need to resolve before continuing.

6. On the Schema Passwords screen, ensure that **Use same passwords for all schemas** is selected. This password will be used in a later step for Content Server to connect to this database schema.

Click **Next**.

7. On the Map Tablespaces screen, click **Next**.

When the warning dialog box appears, click **OK**.

The tablespace creation results should appear in a dialog with the text **Operation completed**. Click **OK**.

8. On the Summary screen, click **Create**.
9. On the Completion Summary screen, click **Close**.

### 11.3.3 Installing and Configuring ECMHOST1

This section describes the installation and configuration steps for ECMHOST1.

### 11.3.3.1 Installing Oracle WebLogic Server on ECMHOST1

On ECMHOST1, start the Oracle WebLogic Server installation by running the installer executable file.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

Start the Oracle WebLogic Server installer.

Then follow these steps in the installer to install Oracle WebLogic Server:

1. On the Welcome screen, click **Next**.
2. On the Choose Middleware Home Directory screen, choose a directory on your computer into which the Oracle WebLogic software is to be installed.

For the **Middleware Home Directory**, specify this value:

```
/u01/app/oracle/product/fmw
```

Click **Next**.

3. On the Register for Security Updates screen, enter your "My Oracle Support" User Name and Password.
4. On the Choose Install Type screen, a window prompts you to indicate whether you want to perform a complete or a custom installation.

Choose **Typical** or **Custom**.

Click **Next**.

5. On the Choose Product Installation Directories screen, specify the following values:

WebLogic Server:

```
/u01/app/oracle/product/fmw/wlserver_10.3
```

Oracle Coherence:

```
/u01/app/oracle/product/fmw/coherence_3.5.2
```

6. On the Installation Summary screen, the window contains a list of the components you selected for installation, along with the approximate amount of disk space to be used by the selected components once installation is complete.

Click **Next**.

7. On the Installation Complete screen, deselect the **Run Quickstart** checkbox.

Click **Done**.

### 11.3.3.2 Installing Oracle ECM on ECMHOST1

To start the Oracle Enterprise Content Management (Oracle ECM) installer:

- On UNIX (Linux in the following example):

```
./runInstaller -jreLoc file-specification-for-jdk6
```

Ensure that your DISPLAY environment variable is set correctly before issuing the command above.

On Windows:

setup.exe

Specify the JRE location, for example: C:\Oracle\Middleware\jdk160\_11

Follow these steps in the ECM installer to install Oracle Enterprise Content Management on the system:

1. If this is the first installation, the Specify Oracle Inventory screen opens.  
Accept the defaults and click **OK**.  
If you are installing as a normal (non-root) user, the Inventory Location Confirmation dialog box opens. Follow its directions, or select **Continue Installation with local inventory** and click **OK**.
2. On the Welcome screen, click **Next**.
3. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.  
Click **Next**.
4. On the Specify Installation Location screen, enter these values:  
For the **Middleware Home Directory**, specify this value:  
`/u01/app/oracle/product/fmw`  
  
For the **Oracle Home Directory**, enter the name of the subdirectory where Oracle ECM will be installed (or use the default). This will be referred to as the ECM\_HOME:  
`/u01/app/oracle/product/fmw/ecm`  
  
Click **Next**.
5. On the Installation Summary screen, review the selections to ensure that they are correct (if they are not, click **Back** to modify selections on previous screens), and click **Install**.
6. On the Installation Progress screen on UNIX systems, a dialog box appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.  
Click **Next**.
7. On the Installation Complete screen, click **Finish** to confirm your choice to exit.

---

**Note:** Before you run the Configuration Wizard by following the instructions in [Section 11.3.3.3, "Create a High Availability Domain,"](#) ensure that you applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home so that you have the latest version of Oracle Fusion Middleware.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

---

### 11.3.3.3 Create a High Availability Domain

This section describes how to run the Oracle Fusion Middleware Configuration Wizard to create and configure a high availability domain into which you can deploy applications and libraries.

To start the Oracle Fusion Middleware Configuration Wizard:

- On UNIX (Linux in the following example):

`ECM_HOME/common/bin/config.sh`

Ensure that your `DISPLAY` environment variable is set correctly before issuing the command above.

On Windows:

`ECM_HOME/config.cmd`

To create and configure the Oracle ECM domain on the computer:

1. On the Welcome screen, select **Create a new WebLogic domain**.

Click **Next**.

2. On the Select Domain Source screen, select the applications that you want to install, for example:

- Oracle WebCenter Content Imaging - 11.1.1.0
- Oracle Universal Records Management - 11.1.1.0
- Oracle Universal Content Management - Inbound Refinery - 11.1.1.0
- Oracle Imaging and Process Management - 11.1.1.0

Other products such as Oracle JRF and Oracle Enterprise Manager are also automatically selected.

3. On the Specify Domain Name and Location screen, make these entries:

- Change the domain name to `ecmdomain` or use the default, `base_domain`. A directory will be created for this domain, for example:

`MW_HOME/user_projects/domains/ecmdomain`

- Domain Location: accept the default
- Application Location: accept the default

Click **Next**.

4. On the Configure Administrator Username and Password screen:

- In the **Username** field, enter a user name for the domain administrator or use the default username, `weblogic`.
- In the **Password** field, enter a user password that is at least 8 characters, for example: `welcome1`
- In the **Configure User Password** field, enter the user password again.

Click **Next**.

5. On the Configure Server Start Mode and JDK screen:

- **WebLogic Domain Startup Mode:** Select **Production Mode**.
- **JDK Selection:** Choose a JDK from the list of available JDKs or accept the default.

Click **Next**.

6. On the Configure JDBC Component Schema screen:



- a. Select all component schemas that appear in the table at the bottom: **UCM Schema**, **IPM Schema**, and **URM Schema**.
- b. For the RAC configuration for component schemas, select **Convert to GridLink** or **Convert to RAC multi data source**. For single database configuration, select **Don't convert**.
- c. Ensure that the following data source appears on the screen. The user names in [Table 11-1](#) assume that you used **DEV** as the prefix for schema creation from RCU.

For information about GridLink data source configuration with Oracle RAC, see [Section 4.1.2, "Using GridLink Data Sources with Oracle RAC."](#)

For information about multi data source configuration with Oracle RAC and the MDS repository, see [Section 4.1.3, "Using Multi Data Sources with Oracle RAC."](#)

**Table 11-1 Configuring Values for Data Sources**

| Component Schema | Schema Owner  |
|------------------|---------------|
| UCM Schema       | DEV_OCS       |
| IPM Schema       | DEV_IPM       |
| URM Schema       | DEV_URMSERVER |

- d. Click **Next**.

7. Enter values for the following fields, specifying connect information for the Oracle RAC database that was seeded with RCU.

For multi data sources, enter values for these fields:

**Figure 11-6 Configure RAC Multi Data Source Component Schema Screen**

**Fusion Middleware Configuration Wizard**

**Configure RAC Multi Data Source Component Schema**

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for RAC Service-Instance co

Service Name: ecmdb.my.company.com

Username: Varies among component schemas

Password: \*\*\*\*\*

| Host Name    | Instance Name | Port |
|--------------|---------------|------|
| ecmdbhost1.m | ecmdb1        | 1521 |
| ecmdbhost2.m | ecmdb2        | 1521 |

[Add](#) [Delete](#)

| RAC Component Schema                           | Service Name         | Schema Owner  | Schema Password |
|------------------------------------------------|----------------------|---------------|-----------------|
| <input checked="" type="checkbox"/> UCM Schema | ecmdb.my.company.com | DEV_OCS       | *****           |
| <input checked="" type="checkbox"/> IPM Schema | ecmdb.my.company.com | DEV_IPM       | *****           |
| <input checked="" type="checkbox"/> URM Schema | ecmdb.my.company.com | DEV_URMSERVER | *****           |

[Exit](#) [Help](#) [Previous](#) [Next](#)

For **GridLink data sources**, enter values for these fields:

**Figure 11–7 Configure GridLink RAC Component Schema**

**Fusion Middleware Configuration Wizard**  
**Configure GridLink RAC Component Schema**

**Note:** Change only the input fields below that you wish to modify and values will be applied to all selected rows.

Driver: \*Oracle's Driver (Thin) for GridLink Connections; Versior

Service Name: ecmdb.my.company.com

Username: Varies among component schemas

Password: \*\*\*\*\*

Enable FAN: ☒

Enable SSL: ☐

Wallet File: Enter a value

Wallet Password: Enter a value

| Service Listener       | Port | Protocol |
|------------------------|------|----------|
| ecmdb-r-my.company.com | 1521 | TCP      |

| ONS Host                  | Port |
|---------------------------|------|
| ecmdbhost1.my.company.com | 6200 |
| ecmdbhost2.my.company.com | 6200 |

| RAC Component Schema                           | Service Name         | Schema Owner  | Schema Password |
|------------------------------------------------|----------------------|---------------|-----------------|
| <input checked="" type="checkbox"/> UCM Schema | ecmdb.my.company.com | DEV_OCS       | *****           |
| <input checked="" type="checkbox"/> IPM Schema | ecmdb.my.company.com | DEV_IPM       | *****           |
| <input checked="" type="checkbox"/> URM Schema | ecmdb.my.company.com | DEV_URMSERVER | *****           |

- a. Driver: For RAC: Select **Oracle driver (Thin) for RAC Service-Instance connections, Versions:10, 11**. For **GridLink**: **Oracle Driver (Thin) for GridLink Connections, Versions:10 and later**
- b. Service Name: Enter the service name of the database. **For a GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example,  
`<mydbservice>.us.oracle.com`

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example,  
`us.oracle.com`

---

- c. Username prefix: Enter the prefix for the schemas. The user names in [Table 11–1](#) assume that **DEV** is the prefix used in schema creation from RCU.
- d. Password and Confirm Password: Enter the password for access to the schemas.
- e. **For a GridLink data source**, enter the SCAN address in the Listener Address field and SCAN port in the Port field. Enter the ONS host and port information in the ONS Host and Port fields, respectively.

---

**Note:** Oracle recommends that you use SCAN addresses to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. You do not need to update a GridLink data source containing SCAN addresses if you add or remove Oracle RAC nodes. Contact your network administrator for appropriately configured SCAN urls for your environment. See SCAN Addresses in the *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* guide.

---

**For a multi data source**, click **Add**. Enter details for Oracle RAC instances.

- f. Update each schema: select one data source at a time and add the appropriate details.

Ensure that the information is entered for all schemas: **UCM Schema**, **IPM Schema**, and **URM Schema**.

- g. Click **Next**.

8. In the Test JDBC Data Sources screen, the connections are tested automatically. The **Status** column shows the results. Ensure that all connections were successful. If not, click **Previous** to correct your entries.

After every schema is specified correctly, click **Next**.

9. Select **JMS Distributed Destination**, **Administration Server** and **Managed Servers, Clusters and Machines**.
10. In the Select JMS Distributed Destination Type screen, select **UDD** from the drop-down list for the JMS modules of all Oracle Fusion Middleware components. Click **Next**.

11. In the Configure the Administration Server screen, enter the following values:

- **Name:** AdminServer
- **Listen address:** Enter a host or virtual hostname.
- **Listen port:** 7001
- **SSL listen port:** Not applicable.
- **SSL enabled:** Leave this checkbox unselected.

Click **Next**.

12. On the Configure Managed Servers screen, add an additional Managed Server for each of the existing Servers. For example, add an WLS\_UCM2 to match WLS\_UCM1. The Listen Port for the second Managed Server should be the same as the first. Set the Listen Address for all the servers to the host name that the Managed Server runs on.

For example:

WLS\_UCM1

Listen Address: ECMHOST1

Listen Port: 16200

WLS\_UCM2

Listen Address: ECMHOST2

Listen Port: 16200

Do this for all the Managed Servers.

The Oracle I/PM servers require a virtual host name as the listen addresses for the managed server. These virtual host names and corresponding virtual IPs are required to enable server migration for the I/PM component. You must enable a VIP mapping to ECMHOST1VHN1 on ECMHOST1 and ECMHOST2VHN1 on ECMHOST2, and must also correctly resolve the host names in the network system used by the topology (either by DNS Server or hosts resolution).

The Inbound Refineries should be configured to reside on separate remote servers, for example: IBR\_Server1 on ECMHOST3 and IBR\_Server2 on ECMHOST4.

13. On the Configure Clusters screen, create a cluster for each pair of servers. For example, for UCM:

- **Name:** UCM\_Cluster
- **Cluster Messaging Mode:** unicast
- **Multicast Address:** Not applicable
- **Multicast Port:** Not applicable
- **Cluster Address:** Not applicable

Add similarly for other products: IPM\_Cluster, URM\_Cluster, and so on.

Click **Next**.

14. On the Assign Servers to Clusters screen, assign each pair of Managed Servers to the newly created cluster:

UCM Cluster:

- WLS\_UCM1
- WLS\_UCM2

Do this for all the clusters you created in the previous step in this section, such as IPM\_Cluster, URM\_Cluster, and IBR\_Cluster.

Click **Next**.

15. On the Configure Machines screen, click the **Unix Machine** tab, and add the following machines:

- ECMHOST1
- ECMHOST2
- ECMHOST3 (if Inbound Refinery is being configured)
- ECMHOST4 (if Inbound Refinery is being configured)

16. On the Assign Servers to Machines screen:

- To the ECMHOST1 machine, assign the **AdminServer** and all \*1 Servers (WLS\_UCM1, WLS\_IPM1, WLS\_URM1).
- To the ECMHOST2 machine, assign all the \*2 Servers (WLS\_UCM2, WLS\_IPM2, WLS\_URM2).

- If configuring Inbound Refinery, add WLS\_IBR1 to ECMHOST3 and WLS\_IBR2 to ECMHOST4.

17. On the Configuration Summary screen, click **Create** to create the domain.

### 11.3.3.4 Start the Administration Server and the Managed Servers on ECMHOST1

Start the Administration Server:

```
cd DOMAIN_HOME/bin
startWebLogic.sh &
```

Configure and start the Node Manager:

```
ECMHOST1> MW_HOME/oracle_common/common/bin/setNMProps.sh
ECMHOST1> MW_HOME/wlserver 10.3/server/bin/startNodeManager.sh
MW_HOME/oracle_common/common/bin/setNMProps.sh
MW_HOME/wlserver 10.3/server/bin/startNodeManager.sh
```

Access the Administration Console at <http://ECMHOST1:7001/console>. In the Administration Console, start the WLS\_UCM1, WLS\_URM1, and WLS\_IPM1 managed servers.

### 11.3.3.5 Disabling Host Name Verification for the Administration Server and the Managed Servers for ECMHOST1 and ECMHOST2

This step is required if you have not set up SSL communication between the Administration Server and the Node Manager. If SSL is not set up, you receive an error message unless you disable host name verification.

You can re-enable host name verification when you have set up SSL communication between the Administration Server and the Node Manager.

To disable host name verification on ECMHOST1:

1. In Oracle WebLogic Server Administration Console, select **Servers**, and then **AdminServer**.
2. Select **SSL**, and then **Advanced**.
3. In the Change Center, click **Lock & Edit**.
4. When prompted, save the changes and activate them.
5. Set **Hostname Verification** to **None**.
6. Select **WLS\_UCM1**, **SSL**, and then **Advanced**.
7. Set **Hostname Verification** to **None**.
8. Repeat Steps 6 and 7 for WLS\_URM1 and WLS\_IPM1.
9. Restart the AdminServers and all the Managed Servers (for example, WLS\_UCM1, WLS\_URM1, and WLS\_IPM1).

To disable host name verification on ECMHOST2:

1. In Oracle WebLogic Server Administration Console, select **WLS\_UCM2**, **SSL**, and then **Advanced**.
2. Set **Hostname Verification** to **None**.
3. Repeat Steps 1 and 2 for WLS\_URM2 and WLS\_IPM2.
4. Restart the AdminServers and all the Managed Servers (for example: WLS\_UCM2, WLS\_URM2, and WLS\_IPM2).

### 11.3.3.6 Configure the WLS\_UCM1 Managed Server

Access the WLS\_UCM1 configuration page at `http://ECMHOST1:16200/cs`.

After you log in, the configuration page opens. Oracle UCM configuration files must be on a shared disk. This example assumes that the shared disk is at `/u01/app/oracle/admin/domainName/ucm_cluster`.

Set the following values on the configuration page to the values below:

- **Content Server Instance Folder:** `/u01/app/oracle/admin/domainName/ucm_cluster/cs`
- **Native File Repository Location:** `/u01/app/oracle/admin/domainName/ucm_cluster/cs/vault`
- **Web Layout Folder:** `/u01/app/oracle/admin/domainName/ucm_cluster/cs/weblayout`
- **Server Socket:** Port 4444
- **Socket Connection Address Security Filter:** Set to a pipe-delimited list of localhost and HTTP Server Hosts: `127.0.0.1|WEBHOST1|WEBHOST2`
- **WebServer HTTP Address:** Set to the host and port of the load balancer HTTP: `ecm.mycompany.com:80`

To apply these updates, click **Submit** and restart the WLS\_UCM1 managed server.

### 11.3.3.7 Configure the WLS\_URM1 Managed Server

Access the WLS\_URM1 configuration page at `http://ECMHOST1:16250/urm`.

After you log in, the configuration page opens. The Oracle URM configuration files must be on a shared disk. This example assumes that the shared disk is at `/u01/app/oracle/admin/domainName/ucm_cluster`.

Set the following values on the configuration page to the values below:

- **Content Server Instance Folder:** `/u01/app/oracle/admin/domainName/ucm_cluster/urm`
- **Native File Repository Location:** `/u01/app/oracle/admin/domainName/ucm_cluster/urm/vault`
- **Web Layout Folder:** `/u01/app/oracle/admin/domainName/ucm_cluster/urm/weblayout`

Verify that the **Is new Content Server Instance?** check box is not selected.

After making these updates, click **Submit** and then restart the WLS\_URM1 managed server.

## 11.3.4 Installing and Configuring WEBHOST1

This section describes the installation and configuration steps to perform on WEBHOST1.

### 11.3.4.1 Installing Oracle HTTP Server on WEBHOST1

This section describes how to install Oracle HTTP Server on WEBHOST1.

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the

Oracle Fusion Middleware documentation library for the platform and version you are using.

2. Ensure that port 7777 is not in use by any service on WEBHOST1 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep "7777"
```

On Windows:

```
netstat -an | findstr :7777
```

3. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for port 7777 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

4. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
5. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port number for Oracle HTTP Server):

```
The port for Oracle HTTP server
Oracle HTTP Server port = 7777
```

6. Start the Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:

On UNIX, issue this command: **runInstaller.**

On Windows, double-click **setup.exe.**

The `runInstaller` and `setup.exe` files are in the `../install/platform` directory, where *platform* is a platform such as Linux or Win32.

7. On the Specify Inventory Directory screen, enter values for the Oracle Inventory Directory and the Operating System Group Name. For example:

**Specify the Inventory Directory:** `/u01/app/oraInventory`

**Operating System Group Name:** `oinstall`

A dialog box opens with the following message:

"Certain actions need to be performed with root privileges before the install can continue. Please execute the script  
`/u01/app/oraInventory/createCentralInventory.sh` now from another window and then press "Ok" to continue the install. If you do not have the root privileges and wish to continue the install select the "Continue installation with local inventory" option"

Log in as root and run the `"/u01/app/oraInventory/createCentralInventory.sh"`

This sets the required permissions for the Oracle Inventory Directory and then brings up the Welcome screen.

---

**Note:** The Oracle Inventory screen is not shown if an Oracle product was previously installed on the host. If the Oracle Inventory screen is not displayed for this installation, make sure to check and see:

1. If the `/etc/oraInst.loc` file exists
  2. If the file exists, the Inventory directory listed is valid
  3. The user performing the installation has write permissions for the Inventory directory
- 

8. On the Welcome screen, click **Next**.
9. On the Select Installation Type screen, select **Install and Configure** and click **Next**.
10. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.

Click **Next**.

11. On the Specify Installation Location screen:

- On WEBHOST1, set the **Location** to:

`/u01/app/oracle/admin`

Click **Next**.

12. On the Configure Components screen:

- Select **Oracle HTTP Server**.
- Select **Associate Selected Components with Weblogic Domain**.

Click **Next**.

13. On the Specify WebLogic Domain screen:

Enter the location where you installed Oracle WebLogic Server. Note that the Administration Server must be running.

- **Domain Host Name:** ECMHOST1
- **Domain Port No:** 7001
- **User Name:** weblogic
- **Password:** \*\*\*\*\*

Click **Next**.

14. On the Specify Component Details screen:

- Enter the following values for WEBHOST1:
  - **Instance Home Location:** `/u01/app/oracle/admin/ohs_inst1`
  - **Instance Name:** `ohs_inst1`
  - **OHS Component Name:** `ohs1`

Click **Next**.

15. On the Specify Webtier Port Details screen:



- Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
- Enter the Oracle HTTP Server port, for example, 7777.

Click **Next**.

16. On the Oracle Configuration Manager screen, enter the following:

- **Email Address:** Provide the email address for your My Oracle Support account
- **Oracle Support Password:** Provide the password for your My Oracle Support account.
- **I wish to receive security updates via My Oracle Support:** Click this checkbox.

17. On the Installation Summary screen, ensure that the selections are correct. If they are not, click **Back** and make the necessary fixes. After ensuring that the selections are correct, click **Next**.

18. On the Installation Progress screen on UNIX systems, a dialog appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

19. On the Configuration Progress screen, multiple configuration assistants are launched in succession; this process can be lengthy.

20. On the Configuration Completed screen, click **Finish** to exit.

#### 11.3.4.2 Configuring Oracle HTTP Server on WEBHOST1

After installing Oracle HTTP Server on WEBHOST1, add the following lines to the `OHS_HOME/instances/ohs_instance1/config/OHS/ohs1/mod_wl_ohs.conf` file:

```
UCM
<Location /cs>
WebLogicCluster ecmhost1:16200,ecmhost2:16200
SetHandler weblogic-handler
WLCookieName JSESSIONID

</Location>

<Location /adfAuthentication>
WebLogicCluster ecmhost1:16200,ecmhost2:16200
SetHandler weblogic-handler
WLCookieName JSESSIONID

</Location>

URM
<Location /urm>
WebLogicCluster ecmhost1:16250,ecmhost2:16250
SetHandler weblogic-handler
WLCookieName JSESSIONID

</Location>

<Location /urm/adfAuthentication>
WebLogicCluster ecmhost1:16250,ecmhost2:16250
SetHandler weblogic-handler
```

```
WLCookieName JSESSIONID

</Location>

IBR
<Location /ibr>
WebLogicCluster ecmhost1:16300,ecmhost2:16300
SetHandler weblogic-handler
WLCookieName JSESSIONID

</Location>

<Location /ibr/adfAuthentication>
WebLogicCluster ecmhost1:16300,ecmhost2:16300
SetHandler weblogic-handler
WLCookieName JSESSIONID

</Location>

I/PM Application
<Location /imaging>
 SetHandler weblogic-handler
 WebLogicCluster ECMHOST1VHN1:16000,ECMHOST2VHN1:16000
</Location>

AXF WS Invocation
<Location /axf-ws>
 SetHandler weblogic-handler
 WebLogicCluster ECMHOST1VHN1:16000,ECMHOST2VHN1:16000
</Location>
```

These lines configure Oracle HTTP Server on WEBHOST1 to route requests to the clustered applications on ECMHOST1 and ECMHOST2.

---

---

**Note:** The value of WLCookieName is different in earlier software releases. The value was IDCCS\_SESSIONID in earlier releases but is JSESSIONID in this release.

---

---

After adding these lines, restart Oracle HTTP Server, and then ensure that the applications can be accessed at:

```
http://WEBHOST1:7777/cs
http://WEBHOST1:7777/urm
http://WEBHOST1:7777/imaging
```

---

---

**Note:** More Location tags may need to be added depending on how the applications will be used. For example, Oracle UCM also requires access at /\_dav for WebDav access and /idcws for web services access. Consult the product documentation for more details.

---

---

### 11.3.5 Configuring the Load Balancer

Configure the load balancer so that the virtual host (for example, ucm.mycompany.com) routes to the available Oracle HTTP Servers using round-robin load balancing.

You should also configure the load balancer to monitor the HTTP and HTTPS listen ports for each Oracle HTTP Server.

Verify that Oracle Content Server can be accessed at `http://WEBHOST1:7777/cs`.

## 11.3.6 Installing and Configuring ECMHOST2

This section describes the installation and configuration steps to perform on ECMHOST2.

### 11.3.6.1 Installing Oracle WebLogic Server on ECMHOST2

To install Oracle WebLogic Server on ECMHOST2, perform the steps in [Section 11.3.3.1, "Installing Oracle WebLogic Server on ECMHOST1"](#) on ECMHOST2.

### 11.3.6.2 Installing ECM on ECMHOST2

To install Oracle ECM on ECMHOST2, perform the steps in [Section 11.3.3.2, "Installing Oracle ECM on ECMHOST1"](#) on ECMHOST2.

### 11.3.6.3 Using pack and unpack to Join the Domain on ECMHOST1

The pack and unpack commands are used to enable the WLS\_UCM2 managed server on ECMHOST2 to join the ECM domain on ECMHOST1.

First, execute the pack command on ECMHOST1:

```
pack -managed=true -domain=/u01/app/oracle/product/WLS/11G/user_projects/domains/ecmdomain -template=ecm_template.jar -template_name="my ecm domain"
```

After copying the `ecm_template.jar` file to ECMHOST2, execute the unpack command on ECMHOST2:

```
unpack -domain=/u01/app/oracle/product/WLS/11G/user_projects/domains/ecmdomain -template=ecm_template.jar
```

### 11.3.6.4 Start Node Manager and the WLS\_UCM2 Server on ECMHOST2

To start Node Manager on ECMHOST2, use this command:

```
ECMHOST2> MW_HOME/oracle_common/common/bin/setNMProps.sh
ECMHOST2> WL_HOME/wlserver_10.3/bin/startNodeManager.sh
```

Access the Administration Console at `http://ECMHOST1:7001/console`. In the Administration Console, start the WLS\_UCM2 managed server.

### 11.3.6.5 Start the Managed Servers on ECMHOST2

Access the Administration Console at `http://ECMHOST2:7001/console`. In the Administration Console, start the WLS\_UCM2, WLS\_URM2, and WLS\_IPM2 managed servers.

### 11.3.6.6 Configure the WLS\_UCM2 Managed Server

Access the WLS\_UCM2 configuration page at `http://ECMHOST2:16200/cs`.

After you log in, the configuration page opens. The Oracle UCM configuration files must be on a shared disk. This example assumes that the shared disk is at `/orashare/orcl/ucm`.

Set the following values on the configuration page to the values below:

- **Content Server Instance Folder:** /u01/app/oracle/admin/*domainName*/ucm\_cluster/cs
- **Native File Repository Location:** /u01/app/oracle/admin/*domainName*/ucm\_cluster/cs/vault
- **WebLayout Folder:** /u01/app/oracle/admin/*domainName*/ucm\_cluster/cs/weblayout

Verify that the **Is new Content Server Instance?** check box is not selected.

After making these updates, click **Submit** and then restart the WLS\_UCM2 managed server.

### 11.3.6.7 Configure the WLS\_URM2 Managed Server

Access the WLS\_URM2 configuration page at <http://ECMHOST2:16250/urm>.

After you log in, the configuration page opens. The Oracle URM configuration files must be on a shared disk. This example assumes that the shared disk is at /u01/app/oracle/admin/*domainName*/ucm\_cluster.

Set the following values on the configuration page to the values below:

- **Content Server Instance Folder:** /u01/app/oracle/admin/*domainName*/ucm\_cluster/urm
- **Native File Repository Location:** /u01/app/oracle/admin/*domainName*/ucm\_cluster/urm/vault
- **Web Layout Folder:** /u01/app/oracle/admin/*domainName*/ucm\_cluster/urm/weblayout

Verify that the **Is new Content Server Instance?** check box is not selected.

After making these updates, click **Submit** and then restart the WLS\_URM2 managed server.

## 11.3.7 Installing and Configuring WEBHOST2

This section describes the installation and configuration steps to perform on WEBHOST2.

### 11.3.7.1 Installing Oracle HTTP Server on WEBHOST2

This section describes how to install Oracle HTTP Server on WEBHOST2.

1. Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
2. Ensure that port 7777 is not in use by any service on WEBHOST2 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

On UNIX:

```
netstat -an | grep "7777"
```

On Windows:

```
netstat -an | findstr :7777
```

3. If the port is in use (if the command returns output identifying the port), you must free the port.

On UNIX:

Remove the entries for port 7777 in the `/etc/services` file and restart the services, or restart the computer.

On Windows:

Stop the component that is using the port.

4. Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
5. Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port number for Oracle HTTP Server):

```
The port for Oracle HTTP server
Oracle HTTP Server port = 7777
```

6. Start the Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:

On UNIX, issue this command: **runInstaller**.

On Windows, double-click **setup.exe**.

The `runInstaller` and `setup.exe` files are in the `../install/platform` directory, where *platform* is a platform such as Linux or Win32.

This displays the Specify Oracle Inventory screen.

7. On the Specify Inventory Directory screen, enter values for the Oracle Inventory Directory and the Operating System Group Name. For example:

**Specify the Inventory Directory:** `/u01/app/oraInventory`

**Operating System Group Name:** `oinstall`

A dialog box appears with the following message:

"Certain actions need to be performed with root privileges before the install can continue. Please execute the script  
`/u01/app/oraInventory/createCentralInventory.sh` now from another window and then press "Ok" to continue the install. If you do not have the root privileges and wish to continue the install select the "Continue installation with local inventory" option"

Log in as root and run the `"/u01/app/oraInventory/createCentralInventory.sh"`

This sets the required permissions for the Oracle Inventory Directory and then brings up the Welcome screen.

---

**Note:** The Oracle Inventory screen is not shown if an Oracle product was previously installed on the host. If the Oracle Inventory screen is not displayed for this installation, make sure to check and see:

1. If the `/etc/oraInst.loc` file exists
  2. If the file exists, the Inventory directory listed is valid
  3. The user performing the installation has write permissions for the Inventory directory
-

8. On the Welcome screen, click **Next**.
9. On the Select Installation Type screen, select **Install and Configure**, and click **Next**.
10. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.  
Click **Next**.
11. On the Specify Installation Location screen:
  - On WEBHOST2, set the **Location** to:  
`/u01/app/oracle/admin`Click **Next**.
12. On the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Select **Associate Selected Components with Weblogic Domain**.Click **Next**.
13. On the Specify WebLogic Domain screen:

Enter the location where you installed Oracle WebLogic Server. Note that the Administration Server must be running.

  - **Domain Host Name:** ECMHOST1
  - **Domain Port No:** 7001
  - **User Name:** weblogic
  - **Password:** \*\*\*\*\*Click **Next**.
14. On the Specify Component Details screen:
  - Enter the following values for WEBHOST2:
    - **Instance Home Location:** /u01/app/oracle/admin/ohs\_inst2
    - **Instance Name:** ohs\_inst2
    - **OHS Component Name:** ohs2Click **Next**.
15. On the Specify Webtier Port Details screen:
  - Select **Specify Custom Ports**. If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
  - Enter the Oracle HTTP Server port, for example, 7777.Click **Next**.
16. On the Oracle Configuration Manager screen, enter the following:
  - **Email Address:** Provide the email address for your My Oracle Support account
  - **Oracle Support Password:** Provide the password for your My Oracle Support account.

- **I wish to receive security updates via My Oracle Support:** Click this checkbox.
- 17. On the Installation Summary screen, ensure that the selections are correct. If they are not, click **Back** and make the necessary fixes. After ensuring that the selections are correct, click **Next**.
- 18. On the Installation Progress screen on UNIX systems, a dialog appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.  
Click **Next**.
- 19. On the Configuration Progress screen, multiple configuration assistants are launched in succession; this process can be lengthy.
- 20. On the Configuration Completed screen, click **Finish** to exit.

### 11.3.7.2 Configuring Oracle HTTP Server on WEBHOST2

To configure Oracle HTTP Server on WEBHOST2, follow the instructions in [Section 11.3.4.2, "Configuring Oracle HTTP Server on WEBHOST1"](#) on WEBHOST2.

## 11.3.8 Configuring the I/PM Managed Servers

This section describes how to configure a JMS persistence store for Oracle I/PM JMS, how to configure a default persistence store for Transaction Recovery, and how to configure the Oracle I/PM Managed Servers with the Oracle UCM Managed Servers. Finally, the steps for configuring Server Migration of the Oracle I/PM Managed Servers are described here, also.

### 11.3.8.1 Configuring a JMS Persistence Store for I/PM JMS

Configure the location for the JMS persistence stores as a directory that is visible from both nodes. By default, the JMS servers used by I/PM are configured with no persistent store and use WebLogic Server's store (`ORACLE_BASE/admin/domain_name/msserver/domain_name/servers/server_name/data/store/default`). You must change I/PM's JMS server persistent store to use a shared base directory as follows:

---

**Note:** See [Considerations for Using File Stores on NFS](#) for important information about JMS messages and transaction logs, as well as releasing locks on file stores.

---

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Services** node and then click the **Persistence Stores** node.
3. In the Change Center, click **Lock & Edit**.
4. Click **New**, and then **Create File Store**.
5. Enter a **name** (for example 'IPMJMSserver1Store', which allows you identify the service it is created for) and target WLS\_IPM1. Enter a **directory** that is located in shared storage so that it is accessible from both ECMHOST1 and ECMHOST2 (`ORACLE_BASE/admin/domain_name/ipm_cluster/jms`).
6. Click **OK** and activate the changes.

7. In the Domain Structure window, expand the **Services** node and then click the **Messaging >JMS Servers** node.
8. Click on the IpmJmsServer1 JMS Server (represented as a hyperlink) from the **Name** column of the table.
9. In the Change Center, click **Lock & Edit**.
10. In the Persistent Store drop-down list, select **IPMJMSStore1Store**.
11. Click **Save** and **Activate Changes**.
12. Repeat the steps and create **IPMJMSStore2Store** for IPMJMSStore2.

#### 11.3.8.2 Configuring a Default Persistence Store for Transaction Recovery

Each server has a transaction log which stores information about committed transactions that are coordinated by the server that may not have been completed. The WebLogic Server uses this transaction log for recovery from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to the server.

---

**Note:** Preferably, this location should be a dual-ported SCSI disk or on a Storage Area Network (SAN).

---

To set the location for the default persistence store for WLS\_IPM1:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node and then click the **Servers** node.
3. Click WLS\_IPM1 (represented as a hyperlink) in the Name column of the table.
4. Click the **Services** tab.
5. In the Change Center, click **Lock & Edit**.
6. In the Default Store section of the page, enter the path to the folder where the default persistent stores will store its data files. The path directory structure is:  
*ORACLE\_BASE/admin/domain\_name/ipm\_cluster\_name/tlogs*
7. Click **Save** and activate the changes.
8. Repeat these steps for the WLS\_IPM2 server.

---

**Note:** To enable migration of the Transaction Recovery Service, specify a location on a persistent storage solution that is available to other servers in the cluster. ECMHOST1 and ECMHOST2 must be able to access this directory, which must exist before you restart the server.

---

#### 11.3.8.3 Configuring Oracle I/PM with Oracle UCM

The Oracle IP/M Cluster should be configured to access the Oracle UCM Cluster.

##### 11.3.8.3.1 Enabling I/PM in UCM

To enable Oracle I/PM in Oracle Content Server:

1. Log into Oracle Content Server at `http://ECMHOST1:16200/cs`.
2. Open the Administration tray or menu, and choose Admin Server.



3. Click **Component Manager**.
4. Enable the IpmRepository component.
5. Restart all the managed servers in the UCM cluster.

**11.3.8.3.2 Upgrading the Default File Store** Perform these steps to upgrade the default file store in Oracle Content Server:

1. Log into Oracle Content Server at `http://ECMHOST1:16200/cs`.
2. Open the Administration tray or menu and choose Providers.
3. Click on the **Info** link of **DefaultFileStore**.
4. Click **Edit**.
5. Click **Update**.
6. Restart all the content servers in the cluster.

**11.3.8.3.3 Adding the I/PM Server Listen Addresses to the List of Allowed Hosts in UCM** To add the Oracle I/PM Server listen addresses of the WLS\_IPM1 and WLS\_IPM2 managed servers (ECMHOST1VHN1 and ECMHOST2VHN1, respectively) to the list of allowed hosts in UCM:

1. Edit the file `/u01/app/oracle/admin/domainName/ucm_cluster/config/config.cfg` and remove or comment out the following line:  
`SocketHostAddressSecurityFilter=127.0.0.1|0:0:0:0:0:0:0:1`
2. Add the following two lines to include the WLS\_IPM1 and WLS\_IPM2 listen address to the list of addresses that are allowed to connect to UCM:  
`SocketHostNameSecurityFilter=localhost|localhost.mydomain.com|ecmhost1vhn1|ecmhost2vhn1`  
`AlwaysReverseLookupForHost=Yes`

Restart the UCM servers for the change to take effect.

**11.3.8.3.4 Creating a Connection to the UCM System** Perform these steps to create a connection to the UCM system:

1. Log into the WLS\_IPM1 imaging console at `http://ECMHOST1:16200/imaging`.
2. In the left-hand pane, click **Manage Connections**, and then **Create Content Server Connection**.
3. Enter a name and description for the new connection, and then click **Next**.
4. In the Connection Settings screen, do the following:
  - Ensure that **Use Local Content Server** is checked.
  - Set the Content Server port to 4444.
  - Add two servers to the Content Server Pool:
    - ECMHOST1:4444
    - ECMHOST2:4444

Click **Next**.

5. In the Connection Security screen, leave the default selections for the WebLogic user and then click **Next**.

6. Review the connection details and click **Submit**.

#### 11.3.8.4 Configuring BPEL CSF Credentials

When connecting to a BPEL system from Oracle I/PM, you must configure the required credential to communicate with the SOA system. To add these credentials:

1. Change directory to the `common/bin` location under the ECM Oracle Home in ECMHOST1 (where your administration servers resides):

```
ECMHOST1>cd ORACLE_HOME/common/bin
```

(ORACLE\_HOME is the ECM home under MW\_HOME/ecm.)

2. Run the Oracle WebLogic Scripting Tool (WLST):

```
ECMHOST1>./wlst.sh
```

3. Run `connect()` and supply the username, password, and administration server URL (t3://ECMHOST1:7001).

```
wls:/offline> connect()
```

4. Create a CSF (Credential Store Framework) credential. This credential is the credential that I/PM will use to connect to the BPEL system. It should be a BPEL admin user. CSF credentials are username/password pairs that are keyed by an "alias" and stored inside a named "map" in the CSF. Because of its integration with OWSM web services, I/PM is currently leveraging the standard OWSM CSF map named "oracle.wsm.security". To create a credential, use the `createCred` WLST command:

```
wls:/ecm_domain/serverConfig> createCred(map="oracle.wsm.security",
key="basic.credential", user="weblogic", password="password_for_credential")
```

The "key" in the command is the "alias," which is used for the 'Credential Alias' property of the BPEL connection definition in the I/PM administration user interface (also the `Connection.CONNECTION_BPEL_CSFKEY_KEY` property in the API). The alias "basic.credential" is used in the example because it is a standard default name used by OWSM and BPEL. However, the alias can be anything as long as it is unique within the map.

5. Run the list credentials command to verify that the credential was created:

```
wls:/ecm_domain/serverConfig> listCred(map="oracle.wsm.security",
key="basic.credential")
{map=oracle.wsm.security, key=basic.credential}
Already in Domain Runtime Tree
```

```
[Name : weblogic, Description : null, expiry Date : null]
PASSWORD: password_for_credential
```

#### 11.3.8.5 Configuring the BPEL Connection

To configure the BPEL connection:

1. Log into the WLS\_IPM1 imaging console at `http://ECMHOST1VHN1:16000/imaging`.
2. In the Left pane, click **Manage Connections**, then **Create Workflow Connection**.
3. Click **Create BPEL Connection**.
4. Enter a BPEL machine name and, optionally, a description.

5. Click **Next**.
6. In the Connection Settings screen, do the following:
  - a. In the provider name field, enter your two SOA server listen addresses separated by a comma: SOAHOST1VHN2, SOAHOST2VHN1
  - b. Provide the **BPEL port**: 8001.
  - c. Select the **SSL** option if the target BPEL system requires it.
  - d. Provide the credential alias created earlier in the "Configuring BPEL CSF Credentials" section.
  - e. Click the **Test Connection** button to confirm the connection parameters and see what composites exist on that BPEL machine.
  - f. Click **Next**.
7. Modify the security grants if desired.
8. Click **Next** then **Submit**.

### 11.3.8.6 Setting the Front End HTTP Host and Port

You must set the front end HTTP host and port for the Oracle WebLogic Server I/PM cluster:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Clusters**.
5. Select **IPM\_Cluster**.
6. Click the **HTTP** tab.
7. Set the following values:
  - **Frontend Host**: ecm.mycompany.com
  - **Frontend HTTPS Port**: 443
  - **Frontend HTTP Port**: 80 (if not using SSL)
8. Click **Save**.
9. Click **Activate Changes** in the Change Center section of the Administration Console.
10. Restart the servers to make the front end host directive in the cluster take effect.

### 11.3.8.7 Configuring Server Migration for Oracle I/PM Instances

This section describes how to configure server migration for Oracle I/PM instances.

**11.3.8.7.1 About Configuring Server Migration** The following steps enable server migration for the Oracle I/PM managed servers. This allows an Oracle I/PM managed server to failover to another node in the case of server or process failure.

To configure server migration for the WebLogic managed servers:

- Step 1: [Setting Up a User and Tablespace for the Server Migration Leasing Table](#)
- Step 2: [Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console](#)

- Step 3: [Editing Node Manager's Properties File](#)
- Step 4: [Setting Environment and Superuser Privileges for the wlsifconfig.sh Script](#)
- Step 5: [Configuring Server Migration Targets](#)
- Step 6: [Testing the Server Migration](#)

**11.3.8.7.2 Creating a GridLink or Multi Data Source Using the Oracle WebLogic Administration Console** You create a data source to each of the Oracle RAC database instances during the process of setting up the multi data source, both for these data sources and the global leasing multi data source. When you create a data source:

- Ensure that this is a non-XA data source.
- Use Oracle's Driver (Thin) Version 9.0.1, 9.2.0, 10, 11.
- Data sources do not require support for global transactions. Therefore, do *not* use any type of distributed transaction emulation/participation algorithm for the data source (do not choose the **Supports Global Transactions** option, or the **Logging Last Resource, Emulate Two-Phase Commit**, or **One-Phase Commit** options of the **Supports Global Transactions** option), and specify a service name for your database.
- Target these data sources to the cluster(s).
- Make sure the data source's connection pool initial capacity is set to 0 (zero). To do this, select **Services, JDBC** and then **Datasources**. In the Datasources screen, click the **Datasource Name**, click the **Connection Pool** tab, and enter 0 (zero) in the **Initial Capacity** field.
- Ensure that an ONS daemon is running on your database servers at all times. You can start the ONS daemon on a database server by running the `onsctl` command:  
`start`

### Creating a GridLink Data Source

To create a GridLink Data Source:

1. Log in to the Oracle WebLogic Server Administration Console.
2. If you have not already done so, in the Change Center of the Administration Console, click **Lock & Edit**.
3. In the **Domain Structure** tree, expand **Services**, then select **Data Sources**.
4. On the **Summary of Data Sources** page, click **New** and select **GridLink Data Source**.
5. Enter the following information:
  - a. Enter a logical name for the datasource in the Name field. For example, **gridlink**
  - b. Enter a name for JNDI. For example, **jdbc/gridlink**.
  - c. Click **Next**.
6. In the **Transaction Options** page, de-select **Supports Global Transactions**, and click **Next**.
7. Select **Enter individual listener information** and click **Next**.
8. Enter the following connection properties:

- **Service Name:** Enter the service name of the database. **For a GridLink data source**, you must enter the RAC service name in lower-case letters followed by the domain name `us.oracle.com`. For example, `<mydbservice>.us.oracle.com`

---

**Note:** The Oracle RAC Service name is defined on the database; it is not a fixed name. Oracle recommends that you register/add the RAC service name with the database domain name, for example, `us.oracle.com`.

---

- **Host Name** - The DNS name or IP address of the server that hosts the database. For an Oracle GridLink service-instance connection, this must be the same for each data source in a given multi data source.
- **Port** - The port on which the database server listens for connections requests.
- **Database User Name:** The database user name. For example, `myDataBase`.
- **Password:** The password. For example, `myPassword1`.
- Confirm Password and click **Next**.

**Tip:** For more information, see the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

The console automatically generates the complete JDBC URL. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=left)(PORT=1234))(ADDRES
S=(PROTOCOL=TCP)(HOST=right)(PORT=1234))(ADDRESS=(PROTOCOL
=TCP)(HOST=center)(PORT=1234)))(CONNECT_DATA=(SERVICE_
NAME=myService)))
```

9. On the **Test GridLink Database Connection** page, review the connection parameters and click **Test All Listeners**.

Oracle WebLogic attempts to create a connection from the Administration Server to the database. Results from the connection test appear at the top of the page. If the test is unsuccessful, correct any configuration errors and retry the test.

Click **Next**.

10. In the **ONS Client Configuration** page, do the following:

- Select **Fan Enabled** to subscribe to and process Oracle FAN events.
- In ONS host and port, enter a comma-separate list of ONS daemon listen addresses and ports for receiving ONS-based FAN events. You can use Single Client Access Name (SCAN) addresses to access FAN notifications.
- Click **Next**.

11. On the **Test ONS client configuration** page, review the connection parameters and click **Test All ONS Nodes**.

Click **Next**.

12. In the **Select Targets** page, select **Dept1\_Cluster1** as the target and **All Servers in the cluster**.

13. Click **Finish**.

**14. Click *Activate Changes*.****Creating a Multi Data Source**

To create a multi data source:

1. Log into the Oracle WebLogic Server Administration Console using the Admin credentials.
2. In the Domain Structure window, expand the **Services** node then expand the **DataSource** node.
3. Click **Lock and Edit** in the Change Center.
4. Click **New** then click **Multi Data Sources**.
5. Enter `leasing` as the name.
6. Enter `jdbc/leasing` as the JNDI name.
7. Select **Failover as algorithm** (default).
8. Click **Next**.
9. Select the target cluster(s).
10. Click **Next**.
11. Select **non-XA driver** (the default).
12. Click **Next**.
13. Click **Create New Data Source**.
14. Enter `leasing-rac0` as the name. Enter `jdbc/leasing-rac0` as the JNDI name. Enter `oracle` as the database type. For the driver type, select Oracle Driver (Thin) for RAC server-Instance connection Version 10,11.

---

---

**Note:** When you create multi data sources for the leasing table, enter names in the format `<MultiDS>-rac0`, `<MultiDS>-rac1`, and so on.

---

---

15. Click **Next**.
16. Deselect **Supports Global Transactions**.
17. Click **Next**.
18. Enter the service name, database name (that is, the RAC Node instance name such as `racdb1`, `racdb2`), host port, and password for your leasing schema.
19. Click **Next**.
20. Click **Test Configuration** and verify that the connection works.
21. Click **Next**.
22. Target the data source to the cluster(s).
23. Select the data source and add it to the right screen.
24. Click **Create a New Data Source** for the second instance of your Oracle RAC database, target it to the cluster(s), repeating the steps for the second instance of your Oracle RAC database.
25. Add the second data source to your multi data source.
26. Save and click **Activate Changes**.

**11.3.8.7.3 Testing the Server Migration** The final step is to test the server migration. Perform these steps to verify that server migration is working properly:

**From ECMHOST1:**

1. Stop the WLS\_IPM1 managed server. To do this, run this command:

```
ECMHOST1> kill -9 pid
```

where *pid* specifies the process ID of the managed server. You can identify the pid in the node by running this command:

```
ECMHOST1> ps -ef | grep WLS_IPM1
```

---

**Note:** For Windows, the Managed Server can be terminated by using the `taskkill` command. For example:

```
taskkill /f /pid <pid>
```

Where *<pid>* is the process Id of the Managed Server.

To determine the process Id of the Managed Server run the following command and identify the pid of the WLS\_IPMn Managed Server.

```
MW_HOME\jrockit_160_20_D1.0.1-2124\bin\jps -l -v
```

---

2. Watch the Node Manager console. You should see a message indicating that WLS\_IPM1's floating IP has been disabled.
3. Wait for Node Manager to try a second restart of WLS\_IPM1. It waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Node Manager should now log a message indicating that the server will not be restarted again locally.

**From ECMHOST2:**


1. Watch the local Node Manager console. After 30 seconds since the last try to restart WLS\_IPM1 on ECMHOST1, Node Manager on ECMHOST2 should prompt that the floating IP for WLS\_IPM1 is being brought up and that the server is being restarted in this node.
2. Access the soa-infra console in the same IP.

**Verification From the Administration Console**

To verify migration in the Administration Console:

1. Log in to the Administration Console.
2. Click **Domain** on the left console.
3. Click the **Monitoring** tab and then the **Migration** subtype.

The Migration Status table provides information on the migration status (Figure 11–8).

**Figure 11–8 Migration Status Screen in the Administration Console**

**Note:** After a server migrates, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager starts the managed server on the machine it was originally assigned to.

### 11.3.9 Configuration of Inbound Refinery Instances

This section provides configuration information for Inbound Refinery instances.

#### 11.3.9.1 Inbound Refinery and Inbound Refinery Cluster Concepts

Inbound Refinery instances are not clusterable; they operate completely independently. Several Inbound Refinery instances can be added to the same WebLogic domain, but the following restrictions should be observed:

- No more than one Inbound Refinery per domain per machine can be installed.
- Inbound Refinery instances that are on separate machines must ensure that their configuration is all local and not on shared disk.

The local content requirement is important if you install Inbound Refinery instances on machines that have existing Content Server installations. While configuration in a Content Server cluster *must* be shared, configuration information of Inbound Refinery instances must *not* be shared with other Inbound Refinery instances.

#### 11.3.9.2 Content Server and Inbound Refinery Configuration

You can configure a Content Server with one or more Inbound Refinery instances and the same Inbound Refinery can act as a provider to one or more Content Servers. To configure Inbound Refinery instances, see *Oracle Fusion Middleware Administrator's Guide for Conversion*.



It is recommended that all of the Inbound Refineries which were created in the preceding Configuration be added to the UCM or URM Cluster. This will create a many-to-many relationship in which all members of the UCM/URM cluster can request conversions from any member of the IBR cluster.

Inbound Refinery instances can be installed in their own domain or as an extension of an existing domain.

#### **11.3.9.3 Inbound Refinery Instances and Oracle HTTP Server**

An Inbound Refinery only needs to be accessed once through HTTP to initialize its configuration. This can be done directly, at the Managed Server's listen address. Do not place an Inbound Refinery behind an HTTP Server.

All subsequent access to an Inbound Refinery is through the socket listener.



---

## Active-Passive Topologies for Oracle Fusion Middleware High Availability

This chapter describes how to configure and manage active-passive topologies. It contains the following sections:

- [Section 12.1, "Oracle Fusion Middleware Cold Failover Cluster Topology Concepts"](#)
- [Section 12.2, "Configuring Oracle Fusion Middleware for Active-Passive Deployments"](#)
- [Section 12.3, "Oracle Fusion Middleware Cold Failover Cluster Example Topologies"](#)
- [Section 12.4, "Transforming the Administration Server in an Existing Domain for Cold Failover Cluster"](#)

### 12.1 Oracle Fusion Middleware Cold Failover Cluster Topology Concepts

Oracle Fusion Middleware provides an active-passive model for all its components using Oracle Fusion Middleware Cold Failover Cluster (Cold Failover Cluster). In a Cold Failover Cluster configuration, two or more application server instances are configured to serve the same application workload but only one is active at any particular time.

You can use a two-node Cold Failover Cluster to achieve active-passive availability for Oracle Application Server middle-tier components. In a Cold Failover Cluster, one node is active while the other is passive, on standby. If the active node fails, the standby node activates and the middle-tier components continue servicing clients from that node. All middle-tier components are failed over to the new active node. No middle-tier components run on the failed node after the failover.

The most common properties of a Cold Failover cluster configuration include:

- **Shared Storage:** Shared storage is a key property of a Cold Failover Cluster. The passive Oracle Fusion Middleware instance in an active-passive configuration has access to the same Oracle binaries, configuration files, domain directory, and data as the active instance. You configure this access by placing these artifacts in storage that all participating nodes in the Cold Failover Cluster configuration can access. Typically the active node has shared storage mounted, while the passive node's is unmounted but accessible if the node becomes active. The shared storage can be a dual-ported disk device accessible to both the nodes or a device-based storage such as a NAS or a SAN. You can install shared storage on a regular file system. With Cold Failover Cluster, you mount the volume on one node at a time.

- **Virtual hostname:** In the Cold Failover Cluster solution, two nodes share a virtual hostname and a virtual IP. (The virtual hostname maps to the virtual IP and is used interchangeably in this guide.) However, only the active node can use this virtual IP at any one time. When the active node fails and the standby node is made active, the virtual IP moves to the new active node. The new active node now services all requests through the virtual IP. The Virtual Hostname provides a single system view of the deployment. A Cold Failover Cluster deployment is configured to listen on this virtual IP. For example, if the two physical hostnames of the hardware cluster are `node1.mycompany.com` and `node2.mycompany.com`, the name `cfcvip.mycompany.com` can provide the single view of this cluster. In the DNS, `cfcvip.mycompany.com` maps to the virtual IP, which floats between `node1` and `node2`. When a hardware cluster is used, it manages the failover of the virtual IP without the middle tier clients detecting which physical node is active and actually servicing requests.
- **Hardware Cluster:** Typically a Cold Failover Cluster deployment uses a hardware cluster. The hardware cluster addresses the management of shared storage and virtual IP in its architecture. It plays a role in reliable failover of these shared resources, providing a robust active-passive solution. Most Cold Failover Cluster deployment to hardware clusters that include the following:
  - Two nodes that are in the same subunit,
  - A high-speed, private interconnect between the two nodes,
  - Public network interfaces, on which the client requests are served and the virtual IP is enabled.
  - Shared storage accessible by the two nodes. This includes shared storage that acts as a quorum device and shared storage for Oracle Fusion Middleware and database installs.
  - Clusterware running to manage node and component failures
- **Planned Switchover and Unplanned Failover:** The typical Cold Failover Cluster deployment is a two-node hardware cluster. To maximize utilization, both of these nodes typically have some elements of the deployment running, with the other node acting as a backup node for the appropriate element if needed. For example, a deployment may have the application tier (WebLogic container) running on one node and the Web tier (Oracle HTTP Server) running on the other node. If either node is brought down for maintenance or if either node crashes, the surviving node hosts the services of the down node while continuing to host its current services.

The high-level steps for switch-over to the standby node are as follows:

1. Stop the middle-tier service on the primary node if the node is still available.
2. Fail over the virtual IP from the current active node to the passive node. Bring it down on the current node then enable it and bring it up on the passive node.
3. Fail over the shared disk from the current active node to the passive node. This involves unmounting the shared disk from the current node and mounting it on the passive node.
4. Start the middle-tier service on the passive node, which becomes active.

For failover management, there are two possible approaches:

- Automated failover using a cluster manager facility

The cluster manager offers services, which allows development of packages to monitor the state of a service. If the service or the node is found to be down

(either due to planned operation or unplanned operation), it automatically fails over the service from one node to the other node. The package can be developed to attempt to restart the service on a given node before failing over. Similarly, when the active node itself is down, the clusterware can detect the down state of the node and attempt to bring it up on the designated passive node. These can be automated using any clusterware to provide failover in a completely automated fashion.

Oracle Fusion Middleware provides this capability using Oracle Clusterware. Please refer to [Chapter 13, "Using Oracle Cluster Ready Services"](#) for details of managing a Cold Failover Cluster environment with Oracle Cluster-Ready Services. However, you can use any available clusterware to manage cluster failover.

– Manual failover

For this approach, the planned switchover steps are executed manually. Both the detection of the failure and the failover itself is manual, Therefore this method may result in a longer period of service unavailability.

In active-passive deployments, services are typically down for a short period of time. This is the time taken to either restart the instance on the same node, or to failover the instance to the passive node.

### Active-Passive Topologies: Advantages

- Increased availability

If the active instance fails for any reason or must be taken offline, an identically configured passive instance is prepared to take over at any time. This provides an increased level of availability than a normal single instance deployment. Active-passive deployments also are used to provide availability and protection against planned and unplanned maintenance operation of the hardware used. When a node needs to be brought down for planned maintenance such as a patch apply or OS upgrade or any other reason, the middle ware services can be brought up on the passive node. Switchback to the old node can be done at appropriate times.

- Reduced operating costs

In an active-passive configuration only one set of processes is up and serving requests. Management of the active instance is generally less costly than managing an array of active instances.

- A Cold Failover Cluster topology is less difficult to implement because it does not require a load balancer, which is required in active-active topologies

- A Cold Failover Cluster topology is less difficult to implement than active-active topologies because you are not required to configure options such as load balancing algorithms, clustering, and replication.

- Active-passive topologies better simulates a one-instance topology than active-active topologies.

- Application independence

Some applications may not be suited to an active-active configuration. This may include applications which rely heavily on application state or on information stored locally. Singleton application by very nature are more suitable for active-passive deployments. An active-passive configuration has only one instance serving requests at any particular time.

**Active-Passive Topologies: Disadvantages**

- Active-passive topologies do not scale as well as active-active topologies. You cannot add nodes to the topology to increase capacity.
- State information from HTTP session state and EJB stateful session beans is not replicated, and therefore gets lost when a node terminates unexpectedly. Such state can be persisted to the database or to the file system residing on a shared storage, however, this requires additional overhead that may impact performance of the single node Cold Failover Cluster deployment.
- Active-passive deployments have a shorter downtime than a single node deployment. However, downtime is much shorter in an active-active deployment.

## 12.2 Configuring Oracle Fusion Middleware for Active-Passive Deployments

Oracle Fusion Middleware components come in a variety of Java EE container-deployed components and non-Java EE components. Oracle Internet Directory, Oracle Virtual Directory, and Oracle Reports are system components. Oracle SOA Suite and Oracle WebCenter Portal are Java EE components that are deployed to Oracle WebLogic Server.

WebLogic Server Administration Console and Oracle Enterprise Manager Fusion Middleware Control are also deployed to the WebLogic container. Both Java EE and system components can be deployed to Cold Failover Cluster environments. They can co-exist on the same system or on different systems. When on the same system, you can configure them to failover as a unit, sharing the same virtual IP, or failover independently using separate virtual IPs. In most Oracle Fusion Middleware deployments, a database is used either for the component metadata created using Repository Creation Utility (RCU), or for application data. In many cases, a Cold Failover Cluster middle tier deployment uses a Cold Failover Cluster database, both deployed to the same cluster. The typical deployment has the two components configured as separate failover units using different VIPs and different shared disks on the same hardware cluster.

For Oracle Fusion Middleware, the recommended procedure to create an active-passive topology is:

- Install the component as a single instance configuration. If you planned to transform this instance to a Cold Failover Cluster deployment, install it using a shared disk. This means that the Middleware home, the Instance home, in the case of system components, and the domain directory, in case of a WebLogic deployment are on a shared disk. Everything that fails over as a unit should be on a shared disk.
- After the installation, you transform the deployment into a Cold Failover Cluster deployment, and configure it to listen on a Virtual IP. The Virtual IP is configured on the current active node. It fails over, along with the Oracle Fusion Middleware deployment to the passive node when failure occurs.

This general procedure applies to the Cold Failover Cluster Oracle database. For example, the Oracle database instance is installed as a single instance deployment and subsequently transformed for Cold Failover Cluster. A Cold Failover Cluster Oracle Fusion Middleware deployment can also use an Oracle Real Application Clusters (Oracle RAC) database.

The following sections describe the procedures for post-installation configuration to transform a single instance deployment to a Cold Failover Cluster deployment.

The rest of this chapter describes how to transform Cold Failover Cluster for each of the individual components in the Oracle Fusion Middleware suite. The first section details the procedure for the basic infrastructure components, and the subsequent section does so for the individual Oracle Fusion Middleware component. Any given deployment, for example an Oracle instance or domain, has more than one of these present in a given machine. To transform the entire instance or domain:

- Decide which components form a unit of failover.
- Deploy them on the same shared disk.

---

**Note:** For details about installing and deploying Oracle Fusion Middleware components, see the installation guide for the specific Fusion Middleware component.

---

- Determine a virtual IP to use for this unit of failover. Typically, a single virtual IP is used for all the components, but separate IPs can be used as long as all of them fail over together.
- Apply the transformation procedure to each of the individual components to transform the deployment as a whole. Since more than one of these sections will apply for Cold Failover Cluster transformation of an installation, the order of transformation should always be as follows:
  - Transform the Administration Server or Enterprise Manager instance (if applicable).
  - Transform all managed servers in the deployment.
  - Transform the Oracle instances (non-Java EE deployments).

## 12.2.1 General Requirements for Cold Failover Cluster

As described earlier, a Cold Failover Cluster deployment has at least two nodes in the deployment. The installation is done on one of these nodes. The other node is the passive node. The requirements on these two nodes are as follows:

- The nodes should be similar in all respects at the operating system level. For example, they should be the same operating system, the same version, and the same patch level.
- The nodes should be similar in terms of the hardware characteristics. This ensures predictable performance during normal operations and on failover. Oracle suggests designing each node for capacity to handle both its normal role, as well as the additional load required to handle a failover scenario. However, if the SLA agreement indicates that during outage scenarios, reduced performance is acceptable, this is not required.
- The nodes should have the same mount point free so that mounting of shared storage can occur to the same node during normal operations and failover conditions.
- The user ID and group ID on the two nodes are similar, and the user ID and group ID of the user owning the instance is the same on both nodes.
- The `oraInventory` location is the same on both nodes and has similar accessibility of the instance or domain owner. The location of the `oraInst.loc` file, as well as the `beahomelist` file should be the same.

- Since a given instance uses the same listen ports irrespective of the machine on which it is currently active, it is required to ensure that the ports used by the Cold Failover Cluster instance are free on both the nodes.

---

**Note:** Before beginning the transformation, back up of your entire domain. Oracle also recommends creating a local backup file before editing it. For details on backing up your domain, see the *Oracle Fusion Middleware Administrator's Guide*. Oracle recommends backing up:

- All domain directories
  - All Instance homes
  - Optionally, the database repository and the Middleware homes
  - For all the sections below, before changing any file, ensure that a local backup copy is made before editing the file.
- 

### 12.2.1.1 Terminology for Directories and Directory Environment Variables

The following list describes the directories and variables used in this chapter:

- **ORACLE\_BASE:** This environment variable and related directory path refers to the base directory under which Oracle products are installed.
- **MW\_HOME:** This environment variable and related directory path refers to the location where Fusion Middleware (FMW) resides.
- **WL\_HOME:** This environment variable and related directory path contains installed files necessary to host a WebLogic Server.
- **ORACLE\_HOME:** This environment variable and related directory path refers to the location where a specific Oracle FMW Suite (for example, SOA, WebCenter, or Identity Management) is installed.
- **ORACLE\_COMMON\_HOME:** The Oracle home that contains the binary and library files that are common to all the Oracle Fusion Middleware software suites. In particular, the Oracle Common home includes the files required for Oracle Enterprise Manager Fusion Middleware Control (which is used to manage the Fusion Middleware) and the Oracle Java Required Files (JRF).
- **DOMAIN Directory:** This directory path refers to the location where the Oracle WebLogic Domain information (configuration artifacts) is stored.
- **ORACLE\_INSTANCE:** An Oracle instance contains one or more system components, such as Oracle Web Cache, Oracle HTTP Server, or Oracle Internet Directory. An Oracle instance directory contains updatable files, such as configuration files, log files, and temporary files.
- **/localdisk:** /localdisk stands for the root of the directory tree when the FMW install (either MW\_HOME or DOMAIN HOME) is on a local disk. It is used to represent the MW\_HOME on local disk.
- **/sharedisk:** /sharedisk stands for root of the directory tree when the FMW install (either MW\_HOME or DOMAIN HOME) is on a shared storage system that is mounted by any one node of a CFC configuration. It is used to represent the MW\_HOME on shared disk.

The example values used and recommended for consistency for these directories are:

- **ORACLE\_BASE:** /u01/app/oracle



- MW\_HOME (Apptier): ORACLE\_BASE/product/fmw
- ORACLE\_COMMON\_HOME: MW\_HOME/oracle\_common
- WL\_HOME: MW\_HOME/wlserver\_10.3

The following table includes examples of Oracle home, domain home, and domain directory values used for some of the Oracle Fusion Middleware components:

| Component           | ORACLE_HOME     | DOMAIN_HOME    | Domain Directory                            |
|---------------------|-----------------|----------------|---------------------------------------------|
| Identity Management | MW_HOME/idm     | IDMDomain      | MW_HOME/user_projects/domains/IDMDomain     |
| Oracle SOA          | MW_HOME/soa     | SOADomain      | MW_HOME/user_projects/domains/SOADomain     |
| WebCenter           | MW_HOME/wc      | WCDomain       | MW_HOME/user_projects/domains/WCDomain      |
| Oracle Portal       | MW_HOME/portal  | PortalDomain   | MW_HOME/user_projects/domains/PortalDomain  |
| Oracle Forms        | MW_HOME/forms   | FormsDomain    | MW_HOME/user_projects/domains/FormsDomain   |
| Oracle Reports      | MW_HOME/reports | ReportsDomain  | MW_HOME/user_projects/domains/ReportsDomain |
| Oracle Discoverer   | MW_HOME/disco   | DiscoDomain    | MW_HOME/user_projects/domains/DiscoDomain   |
| Web Tier            | MW_HOME/web     | Not applicable | Not applicable                              |
| Directory Tier      | MW_HOME/idm     | Not applicable | Not applicable                              |

Location for Applications Directory: ORACLE\_BASE/admin/domain\_name/apps (for example)

Location for Oracle Instance: ORACLE\_BASE/admin/instance\_name (for example)

All entities on disk that failover as a unit should preferably be on the same mount point. They can, however, be on separate shared storage, on separate mount points as well. Oracle recommends that the mount point for the shared storage is ORACLE\_BASE. In most cases, this ensures that all the persistent bits of a failover unit are on the same shared storage. When more than one Cold Failover Cluster exists on a node, and each fails over independent of the other, different mount points will exist for each such failover unit.


## 12.2.2 Transforming Oracle Fusion Middleware Infrastructure Components

An Oracle Fusion Middleware deployment is made up of basic infrastructure components that are common across all the product sets. This section describes Cold Failover Cluster transformation steps for these components.

There are two Administration Server topologies supported for Cold Failover Cluster configuration. The following sections describe these two topologies and provide installation and configuration steps to prepare the Administration Server for Cold Failover Cluster transformation.

### 12.2.2.1 Administration Server Topology 1

Figure 12–1 illustrates the first supported topology for Oracle Cold Failover Cluster.

**Figure 12–1 Administration Server Cold Failover Cluster Topology 1**

In Figure 12–1, the Administration Server runs on a two-node hardware cluster: Node 1 and Node 2. The Administration Server is listening on the Virtual IP or hostname. The Middleware Home and the domain directory is on a shared disk that is mounted on Node 1 or Node 2 at any given point. Both the Middleware home and the domain directory should be on the same shared disk or shared disks that can fail over together. If an enterprise has multiple Fusion Middleware domains for multiple applications or environments, this topology is well suited for Administration Server high availability. A single hardware cluster can be deployed to host these multiple Administration Servers. Each Administration Server can use its own virtual IP and set of shared disks to provide high availability of domain services.

**12.2.2.1.1 Topology 1 Installation Procedure** To install and configure Cold Failover Cluster for the application server in this topology:

#### Install the Middleware Home

This installation includes the Oracle home, WebLogic home, and the Domain home on a shared disk. This disk should be mountable by all the nodes that act as the failover destination for the Administration Server. Depending on the storage sub-system used, the shared disk may be mountable only on one node at a time. This is the preferred configuration, even when the storage sub system allows simultaneous mounts on more than one node. This is done as a regular single-instance installation. Please refer to the component chapters for details on installing the Administration Server (and Enterprise Manager) alone. The overall procedure for each suite is as follows:

For Oracle SOA or Oracle WebCenter Portal:

1. Install the WebLogic Server software.  
See the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.
2. Install the Oracle Home for Oracle SOA or WebCenter.  
See the *Oracle Fusion Middleware Installation Guide for Oracle SOA Suite* or the *Oracle Fusion Middleware Installation Guide for Oracle WebCenter Portal*.
3. Invoke the Configuration Wizard and create a domain with just the Administration Server.

In the Select Domain Source screen, select the following:

- **Generate a domain configured automatically to support the following products**
- **Select Enterprise Manager and Oracle JRF.**

For Oracle Identity Management:

1. Install the WebLogic Server software.

See the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

2. Using Oracle Identity Management 11g Installer, install and configure the IDM Domain using the create domain option. In the Configure Components Screen, de-select everything except **Enterprise Manager** (selected by default)

See the *Oracle Fusion Middleware Installation Guide for Oracle Identity Management*.

For Oracle Portal, Forms, Reports and Discoverer:

1. Install the WebLogic Server software.

See the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

2. Using Oracle Fusion Middleware 11g Portal, Forms, Reports, and Discoverer Installer, install and configure the Classic Domain using the create domain option. In the Configure Components Screen, ensure that **Enterprise Manager** is also selected.

---

**Note:** In this case, at least one more Managed Server for the product components is also installed in this process (the Administration Server by itself cannot be installed). This Managed Server must also be transformed to CFC using the specific procedure for the component. It is part of the same failover unit as the Administration Server.

---

For Oracle Business Intelligence:

1. Install the WebLogic Server software.

See the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

2. Using Oracle Fusion Middleware 11g Business Intelligence Installer, install and configure the Business Intelligence domain using the create new BI system option.

---

**Note:** In this case, at least one more Managed Server for the product components is also installed in this process (the Administration Server by itself cannot be installed). This Managed Server must also be transformed to CFC using the specific procedure for the component. It is part of the same failover unit as the Administration Server.

---

### Configuring the Administration Server for Cold Failover Cluster

To configure the Administration Server for Cold Failover Cluster:

1. Provision the Virtual IP using the following commands as root user:

For Linux

```
/sbin/ifconfig interface:index IP_Address netmask netmask
/sbin/arping -q -U -c 3 -I interface IP_Address
```

Where *IP Address* is the virtual IP address and the *netmask* is the associated netmask. In the following example, the IP address is being enabled on the interface `eth0`.

```
ifconfig eth0:1 130.35.46.17 netmask 255.255.224.0
/sbin/arping -q -U -c 3 -I eth0 130.35.46.17
```

For Windows:

```
netsh interface ip add address interface IP Address netmask
```

Where *IP Address* is the virtual IP address and the *netmask* is the associated netmask. In the example below, the IP address is being enabled on the interface 'Local Area Connection'.

```
netsh interface ip add address "Local Area connection" 130.35.46.17
255.255.224.0
```

2. Transform the Administration Server instance to Cold Failover Cluster following the procedure in section [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)

3. Validate the Administration Server transformation by accessing the consoles on the virtual IP.

`http://cfcvip.mycompany.com:7001/console`

`http://cfcvip.mycompany.com:7001/em`

4. Failover the Administration Server manually to the second node using the following procedure:
  - a. Stop the Administration Server process (and any other process running out of a given Middleware Home)
  - b. Unmount the shared storage from Node1 where the Middleware Home and domain directory exists.
  - c. Mount the shared storage on Node2, follow storage specific commands.
  - d. Disable the Virtual IP on Node1 using the following command as root user:

For Linux:

```
/sbin/ifconfig interface:index down
```

Where *IP Address* is the virtual IP. In the example below, the IP address is being disabled on the interface `eth0`.

```
/sbin/ifconfig eth0:1 down
```

On Windows:

```
netsh interface ip delete address interface IP_Address
```

Where *IP Address* is the virtual IP address and the *netmask* is the associated netmask. In the example below, the IP address is being enabled on the interface Local Area Connection.

```
netsh interface ip delete address "Local Area connection" 130.35.46.17
```

- e. Enable the virtual IP on Node2 using similar commands as in Step 1.
- f. Start the Administration Server process.

```
DOMAIN_HOME/bin/startWebLogic.sh
```


Where *DOMAIN\_HOME* is the location of your domain directory.

- g. Validate access to both the Administration Server and Enterprise Manager console.

### 12.2.2.2 Administration Server Topology 2

Figure 12–2 illustrates the second supported Administration Server topology for Oracle Cold Failover Cluster.

**Figure 12–2 Administration Server Cold Failover Cluster Topology 2**



In Figure 12–2, the Administration Server runs on a two-node hardware cluster: Node 1 and Node 2. The Administration Server is listening on the Virtual IP or hostname. The domain directory used by the Administration Server is on a shared disk. This is mandatory. This shared disk is mounted on Node 1 or Node 2 at any given point. The Middleware Homes, which contain the software, (WebLogic Home and the Oracle Home) are not necessarily on a shared disk. They can be on the local disk as well. The Administration Server uses the Middleware Home on Node1 for the software when it is running on Node1 and it uses the Middleware Home on Node2 when it is running on Node2. These two Middleware Home must be maintained to be the same in terms of deployed products, Oracle Homes, and patches. In both cases, it uses the configuration available in the shared Domain Directory/Domain Home. Since this is shared, it ensures that the same configuration is used before and after failover.

This shared domain directory may also have other Managed Servers running. It may also be used exclusively for the Administration Server. If the domain directory is shared with other managed servers, appropriate consideration must be made for their failover when the Administration Server fails over. Some of these considerations are:

1. If the shared storage can be mounted as read/write on multiple nodes simultaneously, the Administration Server domain directory can be shared with other managed servers. In addition, it can be failed over independently of the Managed Server. The Administration Server can failover and Managed Servers can continue to run independently on their designated nodes. This is possible because

the Administration Server in this case requires only failover of the VIP, and does not require failover of the shared disk. The domain directory/domain home continues to remain available by the Managed Servers. Example of such storage include a NAS or a SAN/Direct attached storage with Cluster file system.

2. If only one node can mount the shared storage a time, sharing the Administration Server domain directory with a Managed Server implies that when the Administration Server fails over, the Managed Server that runs off the same domain directory must be shut down.

A hardware cluster may be used in this topology. The cluster helps automate failover (when used with properly configured clusterware). However, it is not required. A single hardware cluster can be deployed to host these multiple Administration Servers. Each Administration Server can use its own virtual IP and set of shared disks to provide domain services high availability.

This topology is supported for Oracle SOA Suite and Oracle WebCenter Portal Suite only.

---

**Note:** For the Oracle Identity Management, an alternate topology is also supported for Cold Failover Cluster. See the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management* for more details.

---

**12.2.2.2.1 Topology 2 Installation Procedure** To install and configure Cold Failover Cluster for the Administration Server in this topology:

#### **Install the Middleware Home**

Install the Middleware Home including the Oracle Home and WebLogic Home separately on the two nodes of the domain. The Administration Server domain directory is created on a shared disk. This disk should be mountable by all the nodes that act as the failover destination for the Administration Server. Depending on the storage sub-system used, the shared disk may be mountable only on one node at a time. This is a regular single-instance installation. Refer to the product suite for details on installing the Administration Server, and Enterprise Manager alone. To install the Middleware Home:

For Oracle SOA Suite or Oracle WebCenter Portal:

1. Install the Oracle WebLogic Server software on Node 1.
2. Install the Oracle Home for SOA or WebCenter on Node 1.
3. Repeat steps 1 and 2 on Node 2.
4. Start the Configuration Wizard on Node 1 and create a domain with just the Administration Server.

In the Select Domain Source screen, select the following:

- **Generate a domain configured automatically to support the following products.**
  - Select **Enterprise Manager** and **Oracle JRF**.
5. In the Specify Domain Name and Location screen, enter the domain name, and be sure the domain directory matches the directory and shared storage mount point.

## Configuring the Middleware Home for Cold Failover Cluster

To configure the Middleware Home for Cold Failover Cluster:

### 1. Provision the Virtual IP. For example:

For Linux:

```
ifconfig eth0:1 IP_Address netmask netmask
/sbin/arping -q -U -c 3 -I eth0 IP_Address
```

Where *IP\_Address* is the virtual IP address and the *netmask* is the associated netmask. In the following example, the IP address is being enabled on the interface *eth0*.

```
ifconfig eth0:1 130.35.46.17 netmask 255.255.224.0
/sbin/arping -q -U -c 3 -I eth0 130.35.46.17
```

For Windows:

```
netsh interface ip add address interface IP_Address netmask
```

Where *IP\_Address* is the virtual IP address and the *netmask* is the associated netmask. In the example below, the IP address is being enabled on the interface "Local Area Connection".

```
netsh interface ip add address "Local Area connection" 130.35.46.17
255.255.224.0
```

2. Transform the Administration Server instance to Cold Failover Cluster using the procedure in [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster."](#)
3. Validate the Administration Server by accessing the consoles on the virtual IP.  
<http://cfcvip.mycompany.com:7001/console>  
<http://cfcvip.mycompany.com:7001/em>
4. Failover the Administration Server manually to the second node:
  - a. Stop the Administration Server process (and any other process running out of a given Middleware Home).
  - b. Unmount the shared storage from Node1 where the Middleware Home or domain directory exists.
  - c. Mount the shared storage on Node2, following storage specific commands.
  - d. Disable the virtual IP on Node1:

For Linux:

```
ifconfig interface:index down
```

In the following example, the IP address is being disabled on the interface *eth0*.

```
ifconfig eth0:1 down
```

For Windows:

```
netsh interface ip delete address interface addr=IP_Address
```

Where *IP\_Address* is the virtual IP address. In the following example, the IP address is enabled on the interface 'Local Area Connection'.

```
netsh interface ip delete address 'Local Area connection' addr=130.35.46.17
DOMAIN_HOME/bin/startWebLogic.sh
```

Where *DOMAIN\_HOME* is the location of your domain directory.

- e. Enable the virtual IP on Node 2.
- f. Start the Administration Server process using the following command:

```
DOMAIN_HOME/bin/startWebLogic.sh
```

Where *DOMAIN\_HOME* is the location of your domain directory.

- g. Validate access to both Administration Server and Enterprise Manager console.

### 12.2.2.3 Transforming the Administration Server for Cold Failover Cluster

To transform the Administration Server installed on a shared disk from Node 1, follow the steps in this section. These steps transform the container, therefore, both the WebLogic Server Administration Console and Oracle Enterprise Manager Fusion Middleware Control, are transformed for Cold Failover Cluster. This results in other components such as OWSM-PM, deployed to this container, to become Cold Failover Cluster ready as well. The address for all of these services transforms `cfcvip.mycompany.com`. After installation, to transform a non Cold Failover Cluster instance, to a Cold Failover Cluster:

1. Log into the WebLogic Server Administration Console.
2. Create a Machine for the Virtual Host
  - a. Select **Environment**, and then **Machines**.
  - b. In the Change Center, click **Lock & Edit**.
  - c. Click **New**.
  - d. In the Name field, enter **cfcvip.mycompany.com**
  - e. Select the appropriate operating system.
  - f. Click **OK**.
  - g. On the Summary of Machines tab, click the machine name just created.
  - h. Click the **Servers** tab.
  - i. Click **Add**.
  - j. Select an existing server, and associate it with this machine.

In the select server drop down list ensure AdminServer is selected.
  - k. Click **Activate Changes**.
3. Configure the Administration Server to Listen on `cfcvip.mycompany.com`.
  - a. Select **Environment**, and then **Servers** from the Domain Structure menu.
  - b. In the Change Center, click **Lock & Edit**.
  - c. Click on the Administration Server (**AdminServer**)
  - d. Change the Listen Address to **cfcvip.mycompany.com**
  - e. Click **Save**.
  - f. Click **Activate Changes**.



- g. Restart the Administration Server.

---

**Note:** Since it is typically expected that transformation to Cold Failover Cluster for the Administration Server is done at domain creation time No other changes to other parts of the domain are expected. If this change happens post domain creation, and other components are installed in the domain, follow the steps in the following Administration Server transformation section.

---

### Changing Client Side Configuration for Administration Server

Any existing entities in the domain must communicate with the Administration Server using the new address. For example, when starting the Managed Servers manually, the Administration Server address should be specified as `cfcvip.mycompany.com`.

In the `instance.properties` file, located in the `INSTANCE_HOME/OPMN/opmn` directory, make the following change:

```
adminHost=cfcvip.mycompany.com
```

If the Oracle Instance is to be registered or re-registered with a Cold Failover Cluster Administration Server using the OPMN registration commands, the `AdminHost` location in the `opmnctl` command should reference the new location of the Administration Server (`cfcvip.mycompany.com`).

### Changing Client Side Configuration for Oracle Enterprise Manager

Since the Enterprise Manager is part of the same container where the Administration Server runs, transforming the Administration Server to Cold Failover Cluster also transforms the Enterprise Manager. If there are existing Enterprise Manager Agents configured to be part of the domain, these agent configurations must use the new location for the Enterprise Manager. To configure the new location for Enterprise Manager, use the following steps for each agent:

1. Set the directory to `ORACLE_INSTANCE/EMAGENT/emagent_asinst_1/sysman/config`.
2. In the `emd.properties` file, change `node1.mycompany.com` to `cfcvip.mycompany.com` in the following attributes:

- `REPOSITORY_URL`
- `EmdWalletSrcUrl`

3. Stop and restart the agent using the following commands:

```
cd INSTANCE_HOME/EMAGENT/emagent_dir/bin
./emctl stop agent
./emctl start agent
./emctl status agent
```

This shows the Repository URL and it should now point to the new host.

#### 12.2.2.4 Transforming Oracle WebLogic Managed Servers

All Oracle Fusion Middleware components are deployed to a Managed Server. An important step to convert an application or component that is deployed to Oracle WebLogic Server to Cold Failover Cluster is to change its listen address to the virtual IP being used. This change is done for the specific Managed Server to which the

component has been deployed. You can make this change using the WebLogic Server Administration Console or using WLST commands.

The following example describes the generic steps for Cold Failover Cluster transformation of a Managed Server named WLS\_EXMPL. These steps apply to any Managed Server in the Fusion Middleware components.

**12.2.2.4.1 Transforming an Oracle WebLogic Managed Server using the Fusion Middleware Administration Console** For this procedure, the WebLogic Server Administration Console must be running. In the following example, cfcvip.mycompany.com is the virtual IP used for the Cold Failover Cluster, and WLS\_EXMPL is the managed server to be transformed.

1. Log into the WebLogic Server Administration Console.
2. Create a machine for the virtual host:
  - a. Select **Environment > Machines**.
  - b. In the Change Center, click **Lock & Edit**.
  - c. Click **New**.
  - d. For the **Name** field, enter **cfcvip.mycompany.com**
  - e. For the **Machine OS** field, select the appropriate operating system.
  - f. Click **OK**.
  - g. Click the newly created Machine.
  - h. Click **Node Manager** tab.
  - i. Update Listen Address: **cfcvip.mycompany.com**.
  - j. Click **Save**.
  - k. Click **Activate Changes**.
3. Stop the WLS\_EXMPL Managed server:
  - a. Choose **Environment > Servers**.
  - b. Click **Control**.
  - c. Select **WLS\_EXMPL**.
  - d. Select **Force Shutdown Now** in the **Shutdown** drop-down menu.
4. Associate the WLS\_EXMPL Managed Server with the VirtualHost Machine:
  - a. Choose **Environment > Servers**.
  - b. In the Change Center, click **Lock & Edit**.
  - c. Click **Configuration**.
  - d. Select **WLS\_EXMPL**.
  - e. For Machine, assign the newly created Machine by assigning it from the pull down menu.
  - f. For **Listen Address**, enter **cfcvip.mycompany.com**.
  - g. Click **Save**.
  - h. Click **Activate Changes**.
5. Start the WLS\_EXMPL Managed Server:

- a. Choose **Environment > Servers**.
- b. Click **Control**.
- c. Select **WLS\_EXMPL**.
- d. Click **Start**.

#### 12.2.2.4.2 Transforming an Oracle WebLogic Managed Server using the WLST Command Line

You can transform an Oracle WebLogic managed server using WLST commands as well.

Oracle recommends shutting down the managed server you are transforming before performing these steps.

To transform a Managed Server using the WLST command line in online mode (with the WebLogic Server Administration Server up):

1. In the command line, enter:

```
WL_HOME/server/bin/setWLSEnv.sh
WL_HOME/common/bin/wlst.sh
```

2. In WLST, enter the following commands:

```
wls:/offline>connect(<username>,<password>,<AdminServer location>)
```

For example:

```
wls:/offline>connect('WebLogic', 'welcome1', 't3://admin.mycompany.com:7001')

wls:/DomainName/serverConfig> edit()
wls:/DomainName/edit> startEdit()
wls:/DomainName/edit !> create('cfcvip.mycompany.com','Machine')
wls:/DomainName/edit !>
cd('Machines/cfcvip.mycompany.com/NodeManager/cfcvip.mycompany.com')
wls:/DomainName/edit !> set('ListenAddress', 'cfcvip.mycompany.com')
wls:/DomainName/edit !>cd ('Servers')
wls:/DomainName/edit/Servers !>cd ('WLS_EXMPL')
wls:/DomainName/edit/Servers/WLS_EXMPL !>set('Machine',' cfcvip.mycompany.com
')
wls:/DomainName/edit/Servers/WLS_EXMPL !>set('ListenAddress','
cfcvip.mycompany.com ')
wls:/DomainName/edit/Servers/WLS_EXMPL !> save()
wls:/DomainName/edit/Servers/WLS_EXMPL !> activate()
wls:/DomainName/edit/Servers/WLS_EXMPL> exit()
```

Stop (if not already down) and start the Managed server.

Once the Managed server transformation is completed, all references to it should use the new Listen Address - cfcvip.mycompany.com. If Oracle HTTP Server serves as a front end to this Managed server, then any mod\_wls\_ohs configuration with mount points referring to applications in this Managed server should be changed to route to the new listening end point.

#### 12.2.2.5 Transforming Node Manager

Node Manager can be used in a Cold Failover Cluster environment. The possible configurations are:

- Using a Node Manager that listens on the virtual IP as well and fails over with the rest of the Cold Failover Cluster stack. With ASCRS based deployments, the Node Manager must be part of the same Middleware Home as where the Cold Failover

Cluster Fusion Middleware instance is running. This Node Manager is assumed to be dedicated for this Fusion Middleware instance. In this case, Oracle recommends having additional Network Channels listening on the localhost configured. Other Node Managers may co-exist on the box listening on other ports. For more details, see [Chapter 13, "Using Oracle Cluster Ready Services."](#)

- A Node Manager that does not failover with the rest of the Cold Failover Cluster stack. In this case, Node Manager is not configured for Cold Failover Cluster and listens on all IPs on the machine, and not specifically on the virtual IP for Cold Failover Cluster. The failover nodes also have a similarly configured Node Manager already available and configured. The Machine associated with the WebLogic instance communicates with the Node Manager on the localhost. For more details, see the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.

For Cold Failover Cluster in general, port usage should be planned so that there are no port conflicts when failover occurs.

To convert the Node Manager to Cold Failover Cluster:

1. If Node Manager is running, stop it.

The `nodemanager.properties` file is created only after the first start of Node Manager.

Restart the Node Manager if necessary.

2. In the `nodemanager.properties` file located in the `WL_HOME/common/nodemanager/` directory, set the `ListenAddress` to the virtual IP.

For example:

```
ListenAddress=cfcvip.mycompany.com
```

3. Restart the Node Manager using the `startNodeManager.sh` file, located in the `WL_HOME/server/bin` directory. For ASCRS based deployment, the Node Manager is started using `WL_HOME/server/bin/cfcStartNodemanager.sh`.

See [Chapter 13, "Using Oracle Cluster Ready Services."](#) for more details.

---

---

**Note:** If needed, start the Node Manager only using the `cfcStartNodemanager.sh` script instead of the `startNodeManager.sh` script.

---

---

---

---

**Note:** For WebLogic Managed Servers and Administration Servers, hostname verification may be enabled or disabled in a given installation. For CFC installation where hostname verification is enabled and Node Manager is managing these instances, the hostname verification step should use certificates for the virtual IP `cfcvip.mycompany.com` as part of these steps.

---

---

### 12.2.2.6 Transforming Oracle Process Management and Notification Server

Oracle Process Management and Notification Server (OPMN) is used for Process Management of system components and is part of the application server instance.

Oracle recommends keeping the default OPMN configuration in a Cold Failover Cluster environment. No further steps are necessary for Cold Failover Cluster transformation of the OPMN process itself.

If you are transforming an Oracle Instance for Cold Failover Cluster and it has already been registered with an Administration Server, make the following changes in these files:

1. In the `topology.xml` file in the `DOMAIN_HOME/opmn` directory of the Administration Server domain, change host name entries for this specific Oracle instance (being transformed to Cold Failover Cluster) to `cfcvip.mycompany.com`.

For example, for an Oracle HTTP Server instance transformed to Cold Failover Cluster, set the following in the `topology.xml` file

```
<property name="HTTPMachine" value="cfcvip.mycompany.com"/>
```

For the instance itself:

```
<ias-instance id="asinst " instance-home="/11gr1as3/MW/asinst"
host="cfcvip.mycompany.com" port="6701">
```

2. In the `opmn.xml` file in the `INSTANCE_HOME/config/OPMN/opmn` directory, replace the "`<physicalhostname>`" with `cfcvip.mycompany.com`:

```
</ias-component><ias-component id="ReportsServer_<physicalhostname>_asinst_2">
<process-set id="ReportsServer_<physicalhostname>_asinst_2"
restart-on-death="true" numprocs="1">
```

3. In the `instance.properties` file in the `INSTANCE_HOME/config/OPMN/opmn` directory, change "`adminHost=<physical hostname>`" to "`adminHost=<cfcvip.mycompany.com>`".

### 12.2.2.7 Transforming Oracle Enterprise Manager for an Oracle Instance

When an Oracle instance such as Oracle Internet Directory, Oracle Virtual Directory, Web Tier components, and Oracle Portal, Forms, Reports, and Discoverer components is transformed to Cold Failover Cluster, the Enterprise Manager agent that is part of this Oracle instance must be transformed to Cold Failover Cluster as well.

To transform the Enterprise Manager agent:

1. Stop the Enterprise Manager agent using the following command:

```
cd INSTANCE_HOME/EMAGENT/emagent_dir/bin
./emctl stop agent
```

2. Set the directory to `ORACLE_INSTANCE/EMAGENT/emagent_instance name/sysman/config`.
3. In the `emd.properties` file, change `node1.mycompany.com` to `cfcvip.mycompany.com` for the `emd_url` attribute.
4. Change the `targets.xml` file on the agent side:

```
cd INSTANCE_HOME/EMAGENT/emagent_dir/cp targets.xml targets.xml.org
```

Modify `targets.xml` so that it has only targets related to the host and `oracle_emd`. Remove all other entries. For example:

```
<Targets AGENT_TOKEN="ad4e5899e7341bfe8c36ac4459a4d569ddb03bc">
 <Target TYPE="oracle_emd" NAME=cfcvip.mycompany.com:port"/>
 <Target TYPE="host" NAME=cfcvip.mycompany.com DISPLAY_
```

```
NAME=cfcvip.mycompany.com/>
</Targets>
```

### 5. Stop and restart the agent

```
cd $INSTANCE_HOME/EMAGENT/emagent_dir/bin
./emctl start agent
```

Make the following changes for the Enterprise Manager server in the Administration Server domain directory:

1. Set your directory to `MW_HOME/user_projects/domains/domain_name/sysman/state`.
2. In the `targets.xml` file, located in `MW_HOME/user_projects/domains/domain_name/sysman/state` directory, modify the hostname from `node1.mycompany.com` to `cfcvip.mycompany.com`

## 12.2.2.8 Transforming Web Tier Components and Clients

The Web tier is made up of two primary components, Oracle HTTP Server and Oracle Web Cache. The next two sections describe how to transform Oracle HTTP Server and Oracle Web Cache for Cold Failover Cluster.

**12.2.2.8.1 Transforming Oracle HTTP Server** To transform Oracle HTTP Server for Cold Failover Cluster:

In `$INSTANCE_HOME/config/OHS/component_name/httpd.conf`, change the following attributes

```
Listen cfcvip.mycompany.com:port #OHS_LISTEN_PORT
Listen cfcvip.mycompany.com:port #OHS_PROXY_PORT
ServerName cfcvip.mycompany.com
```

Also, perform a single sign-on reregistration, as described in [Section 12.2.3.14, "Single Sign-On Re-registration \(If required\)."](#)

### Clients of Oracle HTTP Server

If an Oracle Web Cache instance is routing to Oracle HTTP Server that has been transformed to Cold Failover Cluster, in `$INSTANCE_HOME/config/WebCache/component_name/webcache.xml`, change the following attributes:

Change `node1.mycompany.com` to `cfcvip.mycompany.com`, where `node1.mycompany.com` is the previous address of the Oracle HTTP server before transformation.

```
HOST ID="h1" NAME="cfcvip.mycompany.com" PORT="8888" LOADLIMIT="100"
 OSSTATE="ON" />
<HOST ID="h2" NAME="cfcvip.mycompany.com" PORT="8890" LOADLIMIT="100" OSSTATE="ON"
 SSLENABLED="SSL" />
```

**12.2.2.8.2 Transforming Oracle Web Cache** To transform an Oracle Web Cache for Cold Failover Cluster:

1. Set up an alias to the physical hostname on both nodes of the cluster in `/etc/hosts`.

This is an alias to the IP address of the node. Set this in `/etc/hosts` for Linux and Windows location for Windows. The alias name is `wcprfx.mycompany.com`. For example, On node Node1, the `/etc/hosts` file (on UNIX), the entry would be **n.n.n.n node1 node1.mycompany.com wcprfx wcprfx.mycompany.com**

On the failover node Node2, the `/etc/hosts` file (on UNIX), the entry would be **n.n.n.m node2 node2.mycompany.com wcprfx wcprfx.mycompany.com**.

On Windows, this change should be done on all nodes in the file located at `C:\SystemRoot\system32\drivers\etc\hosts`. SystemRoot is either Winnt or Windows.

2. In `INSTANCE_HOME/config/WebCache/wc1/webcache.xml`:

- Change node `.mycompany.com` to `cfcvip.mycompany.com` where Oracle Web Cache was installed, and the host address it is listening on before transformation.

```
SITE NAME="cfcvip.mycompany.com"
```

- Change the Virtual Host Name entries to be `cfcvip.mycompany.com` for the SSL and non-SSL ports. For example:

```
<HOST SSLENABLED="NONE" ISPROXY="NO" OSSTATE="ON" NUMRETRY="5"
PINGINTERVAL="10" PINGURL="/" LOADLIMIT="100" PORT="8888"
NAME="cfcvip.mycompany.com" ID="h0"/>
 <HOST SSLENABLED="SSL" ISPROXY="NO" OSSTATE="ON" NUMRETRY="5"
PINGINTERVAL="10" PINGURL="/" LOADLIMIT="100" PORT="8890"
NAME="cfcvip.mycompany.com" ID="h3"/>
 <VIRTUALHOSTMAP PORT="8094" NAME="cfcvip.mycompany.com">
 <HOSTREF HOSTID="h3"/>
 </VIRTUALHOSTMAP>
 <VIRTUALHOSTMAP PORT="8090" NAME="cfcvip.mycompany.com">
 <HOSTREF HOSTID="h0"/>
 </VIRTUALHOSTMAP>
```

- Change cache name entries to be based of `wcprfx.mycompany.com` where `wcprfx.mycompany.com` is an alias created in `/etc/hosts` on all nodes of the cluster. For example:

```
<CACHE WCDEBUGON="NO" CAPACITY="30" VOTES="1" INSTANCENAME="asinst_1"
COMPONENTNAME="wc1" ORACLEINSTANCE="ORACLE_INSTANCE_HOME"
HOSTNAME="wcprfx.mycompany.com" ORACLEHOME="ORACLE_HOME"
NAME="wcprfx.mycompany.com-WebCache">
```

- In the `MULTIPORT` section, change `IPADDR` from **ANY** to **cfcvip.mycompany.com** for the following:

```
PORTTYPE="NORM"
SSLENABLED="SSL" PORTTYPE="NORM"
PORTTYPE="ADMINISTRATION"
PORTTYPE="INVALIDATION"
PORTTYPE="STATISTICS"
```

For example:

```
<MULTIPORT>
 <LISTEN PORTTYPE="NORM" PORT="8090"
IPADDR="stbdd01-vip.us.oracle.com"/>
 <LISTEN SSLENABLED="SSL" PORTTYPE="NORM" PORT="8094"
IPADDR="cfcvip.mycompany.com">
 <WALLET>ORACLE_INSTANCE/config/WebCache/wc1/keystores/
default</WALLET>
 </LISTEN>
 <LISTEN PORTTYPE="ADMINISTRATION" PORT="8091"
IPADDR="cfcvip.mycompany.com"/>
 <LISTEN PORTTYPE="INVALIDATION" PORT="8093" IPADDR="
```

```
cfcvip.mycompany.com"/>
 <LISTEN PORTTYPE="STATISTICS" PORT="8092"
IPADDR="cfcvip.mycompany.com"/>
</MULTIPOINT>
```

### 12.2.2.9 Instance-specific considerations

In a Cold Failover Cluster environment, a failover node (node2.mycompany.com) must be equivalent to the install machine (node1.mycompany.com) in all respects. To make the failover node equivalent to the installation node, perform the following procedure on the failover instance:

#### 12.2.2.9.1 UNIX Platforms

For UNIX platforms follow these steps:

1. Failover the Middleware Home from Node 1 (the installation node) to the failover node (Node 2), following the mount/unmount procedure described previously.
2. As root, do the following:
  - Create an `oraInst.loc` file located in the `/etc` directory identical to the file on Node1.
  - Run the `root.sh` file located in the `ORACLE_HOME` directory on node2, if it is required, and is available for the product suite.
3. Create the `oraInventory` on the second node, using the `attachHome` command located in the `ORACLE_HOME/oui/bin/attachHome.sh` directory.

#### 12.2.2.9.2 Windows Platform

For Windows, follow these steps:

#### For Application Server Instances (Web Tier, OID/OVD for IDM Installs, Oracle Portal, Form, Reports, and Discoverer):

1. To create the OPMN service on Machine 2, run the following commands:

```
sc create OracleProcessManager_instance_name binPath= "ORACLE_
HOME\opmn\bin\opmn.exe -S -I Instance_Home"
```

For example:

```
sc create OracleProcessManager_asinst binPath= "X:\Middleware\im_
oh\opmn\bin\opmn.exe -S -I X:\Middleware\asinst"
```

2. On both Machine 1 and 2, set the service `OracleProcessManager_instance_name` to be started manually.

```
sc config OracleProcessManager_instance_name start= demand
```

#### For Oracle Business Intelligence Installations:

When Oracle Business Intelligence is used in a Cold Failover Cluster, certain Windows registry entries used by Oracle Business Intelligence must be exactly the same on the active node and the passive node. To ensure that these registry entries are the same on the active and passive nodes, follow these steps for exporting registry entries on Node 1 and importing those registry entries on Node 2 (where Node 1 is the active node and Node 2 is the passive node).

#### Export the Oracle Business Intelligence Registry Entries on Node 1

Follow these steps to export registry entries for Oracle Business Intelligence on Node 1:



1. On Node 1, start the Windows Registry Editor. To do this, type the following from the Windows command line:

```
C:\>regedit
```

2. In the Registry Editor, navigate to and select the node HKEY\_LOCAL\_MACHINE\SOFTWARE\Oracle. Right-click the node and choose **Export**. Save these entries as a uniquely named .reg file.
3. Navigate to and select the node HKEY\_LOCAL\_MACHINE\SOFTWARE\ODBC. Right-click the node and choose **Export**. Save these entries as a uniquely named .reg file.
4. Navigate to and select the node HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet001\Services. Export each node within this node that begins with "Oracle" to a uniquely named .reg file.
5. Navigate to and select the node HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet002\Services. Export each node within this node that begins with "Oracle" to a uniquely named .reg file.
6. Navigate to and select the node HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services. Export each node within this node that begins with "Oracle" to a uniquely named .reg file.
7. Copy all the .reg files you created to Node 2.
8. Exit the Registry Editor.

### Import the Oracle Business Intelligence Registry Entries on Node 2

Follow these steps to import registry entries for Oracle Business Intelligence on Node 1 to Node 2:

1. On Node 2, double-click Oracle\_BI1\bifoundation\install\vc80\vcredist\_x86.exe.
2. Double-click on each .reg file you copied from Node 1. In each case, click **Yes** to import the contents into the registry.

For all installations:

If the system environment variable 'Path' on Machine 1 contains any instance specific information, set this information in the same variable on Machine 2. Also export the registry HKEY\_LOCAL\_MACHINE\SOFTWARE\Oracle from Machine 1 and import it to Machine 2.

To copy the Start Menu from Machine 1 to Machine 2, copy the files under C:\Document and Settings\All Users\Start Menu\Programs\<Menu> from Machine 1 to Machine 2. You can do this with any Windows Backup tool, or Zip utility. For example, using the Windows Backup Utility:

1. Invoke the Windows Backup Utility by selecting **Accessories, System Tools**, and then **Backup**.
2. Select C:\Document and Settings\All Users\Start Menu\Programs\<Menu>
3. Copy the backup file to the second node.
4. Restore the backup to the same location: C:\Document and Settings\All Users\Start Menu\Programs\

For Node Manager, use the standard WebLogic Server procedure to configure Node Manager as a service.

Node Manager is configured as a service using the `WL_HOME\server\bin\installNodeMgrSvc.cmd` command. To configure Node Manager:

- It must be set on both Machines.
- Node Manager should be set to start manually on both nodes.

To configure Node Manager on a Machine,

1. Ensure that the Middleware Home is available on the machine.
2. Install the service running the above command.
3. Set it to be started manually using the following command:

```
sc config nodemanager_service_name start= demand
```

## 12.2.3 Transforming Oracle Fusion Middleware Components

This section describes the considerations and the transformation steps for each of the Oracle product suites. For detailed explanation on the product components, see the appropriate component chapter in this guide.

### 12.2.3.1 Transforming Oracle Internet Directory and Its Clients

This section describes how to transform Oracle Internet Directory and its clients.

**12.2.3.1.1 Transforming Oracle Internet Directory** To transform an Oracle Internet Directory server from PS1:

1. In a text editor, create a file named `oidcfc.ldif` that sets up the virtual IP `cfcvip.mycompany.com` for the Oracle Internet Directory server:

```
dn: cn=oid1,cn=osldapd,cn=subconfigsubentry
changetype: modify
replace: orclhostname
orclhostname: cfcvip.mycompany.com
```

2. Run the following command:

```
ORACLE_HOME/bin/ldapmodify -p <oidPort> -h <oidHost> -D cn=orcladmin -w
<adminPasswd> -f oidcfc.ldif
```

OID must be up at the time of running this command. *oidHost* is the physical hostname listening end point installed with OID.

3. Stop and restart the Oracle Internet Directory server using `opmnctl`:

For example:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=oid1
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=oid1
```

4. Re-Register OID with the Administration Server

```
ORACLE_INSTANCE/bin/opmnctl updatecomponentregistration -adminHost
myAdminHost -adminPort 7001 -adminUsername weblogic -componentType OID
-componentName oid1 -Host cfcvip.mycompany.com
```

To upgrade Oracle Internet Directory from PS1, run the following:

```
ORACLE_HOME/opmn/bin/upgradenonj2eeapp.sh -oracleInstance
ORACLE_INSTANCE
```

### 12.2.3.1.2 Transforming Oracle Internet Directory Clients

To transform an Oracle Internet Directory client:

1. All clients of the Oracle Internet Directory server should use the virtual IP `cfcvip.mycompany.com` to access that Oracle Internet Directory server.
2. Oracle Directory Integration Platform installation that uses the Oracle Internet Directory server must access the Oracle Internet Directory server using the virtual IP. To do this:
  - a. In a text editor, open the `dip-config.xml` file located in the `DOMAIN_HOME/config/fmwconfig/servers/wls_ods1/applications/DIP_<version_number>/configuration` directory.

For example:

```
MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/
wls_ods1/applications/DIP_11.1.1.2.0/configuration
```

- b. Enter the following value to set the LDAP address to the virtual IP:
 

```
OID_NODE_HOST>cfcvip.mycompany.com</OID_NODE_HOST>
```
3. An Oracle Directory Services Manager instance managing the Oracle Internet Directory server in [Section 12.2.2.8.1, "Transforming Oracle HTTP Server"](#) must use the virtual IP to connect to the Oracle Internet Directory server.

For example, from the Oracle Directory Services Manager screen, verify that you can connect to Oracle Internet Directory using the virtual server by following these steps:

- a. Select the **Connect to a directory --> Create A New Connection** link in the upper right hand corner.
- b. In the New Connection screen, fill in the connection information below and click **Connect**:
  - \* **Directory Type:** OID
  - \* **Name:** OIDHA
  - \* **Server:** `cfcvip.mycompany.com`
  - \* **Port:** 389
  - \* **SSL Enabled:** Leave blank
  - \* **User Name:** `cn=orcladmin`
  - \* **Password:** `*****`
  - \* **Start Page:** Home (default)

### 12.2.3.2 Transforming Oracle Virtual Directory and Its Clients

This section describes how to transform Oracle Virtual Directory and its clients.

#### 12.2.3.2.1 Transforming Oracle Virtual Directory

Follow these steps to transform an Oracle Virtual Directory server:

1. In a text editor, open the `listeners.os_xml` file in the `ORACLE_INSTANCE/config/OVD/componentname` directory.
2. Enter the following value to set the LDAP address to the virtual IP:

```
<host>cfcvip.mycompany.com</host>
```

3. Restart the Oracle Virtual Directory server using `opmnctl`.

For example:

```
ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=ovd1
ORACLE_INSTANCE/bin/opmnctl startproc ias-component=ovd1
```

**12.2.3.2.2 Transforming Oracle Virtual Directory Clients** All clients of Oracle Virtual Directory must use the virtual IP `cfcvip.mycompany.com` to access Oracle Virtual Directory. For example, when using Oracle Directory Services Manager to administer a Cold Failover Cluster Oracle Virtual Directory instance, create a connection using `cfcvip.mycompany.com` as the location of the Oracle Virtual Directory instance.

### 12.2.3.3 Transforming Oracle Directory Integration Platform and Oracle Directory Services Manager and Their Clients

This section describes how to transform Oracle Directory Integration Platform, Oracle Directory Services Manager, and their clients.

**12.2.3.3.1 Transforming Oracle Directory Integration Platform and Oracle Directory Services Manager** Oracle Directory Integration Platform and Oracle Directory Services Manager are deployed to a Managed Server. The procedure for CFC transformation is to configure the Managed Server to which they are deployed to listen on the `cfcvip.mycompany.com` virtual IP. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the WLS\_ODS managed server to listen on the `cfcvip.mycompany.com` virtual IP.

**12.2.3.3.2 Transforming Oracle Directory Integration Platform and Oracle Directory Services Manager Clients** Follow these steps to transform Oracle Directory Integration Platform and Oracle Directory Services Manager clients:

1. Clients of Oracle Directory Integration Platform and Oracle Directory Services Manager must use the virtual IP `cfcvip.mycompany.com` to access these applications.
2. When Oracle HTTP Server is the front end for Oracle Directory Services Manager, the WebLogic configuration for Oracle Directory Services Manager must specify the virtual IP `cfcvip.mycompany.com` as the address for the WLS\_ODS Managed Server. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used by Oracle HTTP Server and Oracle Directory Services Manager. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file:

```
#Oracle Directory Services Manager
<Location /odsm>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>
```

### 12.2.3.4 Transforming Oracle Identity Federation and Its Client

This section describes how to transform Oracle Identity Federation and its clients.

**12.2.3.4.1 Transforming Oracle Identity Federation** Oracle Identity Federation is a component that is deployed to a Managed Server. The procedure for Cold Failover

Cluster transformation is to configure the Managed Server to which it is deployed to listen on the `cfcvip.mycompany.com` virtual IP. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the WLS\_OIF Managed Server to listen on the `cfcvip.mycompany.com` virtual IP. Since Oracle Identity Federation Cold Failover Cluster deployments are likely to be split into Service Provider and Identity Provider, more than one instance of WLS\_OIF is likely to exist in a given deployment. Use the same Cold Failover Cluster procedure for both WLS\_OIF instances.

After configuring the Managed Server to listen on the `cfcvip.mycompany.com` virtual IP, log into the Oracle Enterprise Manager Fusion Middleware Control and perform these steps:

1. Navigate to **Farm > Identity and Access > OIF**.
2. In the right frame, navigate to **Oracle Identity Federation > Administration** and then make these changes:
  - a. **Server Properties**: change the host to `cfcvip.mycompany.com`
  - b. **Identity Provider > Common**: change the **providerId** to `cfcvip.mycompany.com`
  - c. **Service Provider > Common**: change the **providerId** to `cfcvip.mycompany.com`
  - d. **Data Stores**: If LDAP is the data store, then replace the value of **Connection URL for User Data Store and Federation Data Store** with `cfcvip.mycompany.com`
  - e. **Authentication Engine > LDAP Directory**: Set the **ConnectionURL** to `cfcvip.mycompany.com`

The metadata needs to be generated after the above changes are done.

**12.2.3.4.2 Transforming Oracle Identity Federation Clients** Follow these steps to transform Oracle Identity Federation clients:

1. Clients of Oracle Identity Federation must use the virtual IP `cfcvip.mycompany.com` to access these applications.
2. When Oracle HTTP Server is the front end for Oracle Identity Federation, the WebLogic configuration for Oracle Directory Services Manager must specify the virtual IP `cfcvip.mycompany.com` as the address for the WLS\_OIF Managed Server. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used by Oracle HTTP Server and Oracle Directory Services Manager. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file:

```
#Oracle Identity Federation
<Location /oif>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>
```

### 12.2.3.5 Transforming an Oracle SOA Suite

Oracle SOA Suite is made up of Java EE components deployed to a managed server. The typical configuration has deployments of OWSM-PM applications and SOA applications. The SOA applications are always deployed to a separate managed server.

In many Cold Failover Cluster deployments OWSM-PM is deployed to the Administration Server, however, they can be deployed to a managed server by itself, for example, WLS\_OWSM, or to the SOA managed server. The WLS\_SOA container may have applications such as BPEL, UMS, and B2B deployed or may additionally have the BPM components deployed as well. In all cases, since these are Java EE components, the Cold Failover Cluster transformation procedure involves configuring the managed server to which they are deployed to listen on the virtual IP. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for information on transforming managed servers WLS\_OWSM and WLS\_SOA.

In addition, follow these steps to transform a SOA component managed server:

1. Set the front-end host for WLS\_SOA managed server to cfcvip.mycompany.com.
  - a. Log into Oracle WebLogic Server Administration Console.
  - b. In the **Environment** section, select **Servers**.
  - c. Select the name of the managed server.
  - d. Select **Protocols**, then select **HTTP**.
  - e. In the **Frontend Host** field, enter the host name as **cfcvip.mycompany.com**.
  - f. Set the Frontend Port to the HTTP port
  - g. Click **Save**.
  - h. Activate the changes.
  - i. Restart the Managed Server.
2. When using Oracle HTTP Server as the front end, the mod WebLogic configuration for the applications deployed to WLS\_OWSM and WLS\_SOA should provide the VIP cfcvip.mycompany.com as the address of these managed server. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used by SOA components. For example, use a text editor to make the following edits in the mod\_wl\_ohs.conf file:

```
#SOA soa-infra app
<Location /soa-infra>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

Worklist
<Location /integration/>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

B2B
<Location /b2b>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

UMS
<Location /sdpMessaging/ >
 SetHandler weblogic-handler
 WebLogicHost cfc.mycompany.com:port
```

```

</Location>

UMS WS
<Location /ucs/messaging/webservice >
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

WSM
<Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

workflow
<Location /workflow>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

#Required if attachments are added for workflow tasks
<Location /ADFAttachmentHelper>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

Required if SOA is being used for Oracle Image and Process Management
SOA inspection.wsil
<Location /inspection.wsil>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

```

### 12.2.3.6 Transforming Oracle Access Manager and Its Clients

This section describes how to transform Oracle Access Manager to work in a Cold Failover Cluster environment.

As with other managed servers created using Configuration Wizard, a Cold Failover Cluster set up that requires the managed server to listen on a virtual IP can be done during the initial creation as well by specifying the listen address of the managed server as the virtual hostname (cfcvip.mycompany.com). In this case, the explicit transformation step is not needed.

**12.2.3.6.1 Transforming Oracle Access Manager** Oracle Access Manager is deployed to a managed server (for example, WLS\_OAM1) and the procedure for CFC transformation is to configure this managed server to listen on the virtual IP cfcvip.mycompany.com. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the WLS\_OAM1 managed server to listen on the cfcvip.mycompany.com virtual IP. All other requirements related to the placement of the Middleware Home and other related domain artifacts on a shared storage that can be failed over apply, as described in [Section 12.2.1, "General Requirements for Cold Failover Cluster."](#)

The Oracle Access Manager console is deployed as part of the Administration Server for the domain. For Cold Failover Cluster configuration of this console, the whole Administration Server needs to be configured in Active Passive. The Administration Server could share the same virtual IP `cfcvip.mycompany.com`, or it could be configured to fail over independently using a separate virtual IP and shared disk. If they are using the same virtual IP, the Administration Server and the Oracle Access Manager managed server will fail over as a unit.

**12.2.3.6.2 Transforming Oracle Access Manager Clients** Follow these steps to transform Oracle Access Manager clients:

1. Clients of Oracle Access Manager must use the virtual IP `cfcvip.mycompany.com` to access these applications. Any wiring done with other components such as Oracle Identity Manager and Oracle Adaptive Access Manager should also use the virtual IP `cfcvip.mycompany.com` to access the applications.
2. When Oracle HTTP Server is the front end for Oracle Access Manager, the mod webLogic configuration for Oracle Access Manager must specify the virtual IP `cfcvip.mycompany.com` as the address for the Oracle Access Manager managed server. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file:

```
#Oracle Access Manager
<Location /oam>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:port
</Location>
```

3. When the Oracle Access Manager console as part of the Oracle Administration Server is also configured to be Active Passive, and the Administration Server is configured to be front-ended by Oracle HTTP Server, you must change the WebLogic host configuration in the webserver proxy plugin configuration files. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file:

```
#Oracle Access Manager Admin Console deployed to the Admin Server
<Location /oamconsole>
 SetHandler weblogic-handler
 WebLogicHost ADMINVHN
 WebLogicPort 7001
</Location>
```

### 12.2.3.7 Transforming Oracle Adaptive Access Manager and Its Clients

This section describes how to transform Oracle Adaptive Access Manager to work in a Cold Failover Cluster environment.

As with other managed servers created using Configuration Wizard, a Cold Failover Cluster set up that requires the managed server to listen on a virtual IP can be done during the initial creation as well by specifying the listen address of the managed server as the virtual hostname (`cfcvip.mycompany.com`). In this case, the explicit transformation step is not needed.

**12.2.3.7.1 Transforming Oracle Adaptive Access Manager** Oracle Adaptive Access Manager is deployed to managed servers, both deployed to fail over as a single unit, and therefore sharing the same virtual IP and the same shared storage. The procedure to convert the OAAM Admin managed server and the OAAM Server managed server for



CFC transformation is to configure these managed servers to listen on the virtual IP cfcvip.mycompany.com. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure these managed servers to listen on the cfcvip.mycompany.com virtual IP. All other requirements related to the placement of the Middleware Home and other related domain artifacts on a shared storage that can be failed over apply, as described in [Section 12.2.1, "General Requirements for Cold Failover Cluster."](#)

**12.2.3.7.2 Transforming Oracle Adaptive Access Manager Clients** Follow these steps to transform Oracle Adaptive Access Manager clients:

1. Clients of Oracle Adaptive Access Manager must use the virtual IP cfcvip.mycompany.com to access these applications. Any wiring done with other components such as Oracle Access Manager and Oracle Identity Manager should also use the virtual IP cfcvip.mycompany.com to access the applications.
2. When Oracle HTTP Server is the front end for Oracle Adaptive Access Manager, the mod webLogic configuration for Oracle Adaptive Access Manager must specify the virtual IP cfcvip.mycompany.com as the address for the Oracle Adaptive Access Manager managed servers. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used. For example, use a text editor to make the following edits in the mod\_wl\_ohs.conf file:

```
#Oracle Adaptive Access Manager
<Location /oaam_server>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:port
</Location>

#Oracle Adaptive Access Manager Admin Console
<Location /oaam_admin>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:port
</Location>
```

### 12.2.3.8 Transforming Oracle Identity Manager and Its Clients

This section describes how to transform Oracle Identity Manager to work in a Cold Failover Cluster environment.

As with other managed servers created using Configuration Wizard, a Cold Failover Cluster set up that requires the managed server to listen on a virtual IP can be done during the initial creation as well by specifying the listen address of the managed server as the virtual hostname (cfcvip.mycompany.com). In this case, the explicit transformation step is not needed.

**12.2.3.8.1 Transforming Oracle Identity Manager** Oracle Identity Manager is deployed to managed servers. The typical CFC deployment will have the Oracle Identity Manager managed server and the another managed server with Oracle SOA and Oracle Web Services Manager deployed to fail over as a single unit and therefore sharing the same virtual IP and the same shared storage. The procedure to convert both these managed server for CFC transformation is to configure this managed servers to listen on the virtual IP cfcvip.mycompany.com. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure these managed servers to listen on the cfcvip.mycompany.com virtual IP. All other requirements related to the placement of the Middleware Home and other related domain artifacts on a shared storage that

can be failed over apply, as described in [Section 12.2.1, "General Requirements for Cold Failover Cluster."](#)

**12.2.3.8.2 Transforming Oracle Identity Manager Clients** Follow these steps to transform Oracle Identity Manager clients:

1. Clients of Oracle Identity Manager must use the virtual IP `cfcvip.mycompany.com` to access these applications. Any wiring done with other components such as Oracle Access Manager and Oracle Adaptive Access Manager should also use the virtual IP `cfcvip.mycompany.com` to access the applications. Since SOA is also configured to fail over with Oracle Identity Manager, the wiring from Oracle Identity Manager to SOA should also use the common virtual IP.
2. When Oracle HTTP Server is the front end for Oracle Identity Manager, the `mod_webLogic` configuration for Oracle Identity Manager must specify the virtual IP `cfcvip.mycompany.com` as the address for the managed servers. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file or `oim.conf` file:

```
#Oracle Identity Manager
<Location /oim>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:port
</Location>
```

Refer to *Oracle Fusion Middleware System Administrator's Guide for Oracle Identity Manager* for more information about making configuration changes in Oracle Identity Manager.

### 12.2.3.9 Transforming an Oracle WebCenter Portal Suite

The WebCenter Suite is made up of Java EE components. The typical WebCenter Suite deployment has four managed servers, though more are possible if you develop WebCenter Portal applications. Typical managed server deployments include:

- WC\_SPACES
- WC\_PORTLETS
- WC\_COLLABORATION
- WC\_UTILITIES

These are Java EE components, therefore, the procedure for Cold Failover Cluster transformation is to configure the Managed Server they are deployed to so that it listens on the Virtual IP. In many Cold Failover Cluster deployments of Oracle WebCenter Portal, OWSM-PM may be deployed to WC\_PORTLETS and WC\_SPACES container. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for information on transforming Oracle WebCenter Portal Suite managed servers.

When you use Oracle HTTP Server as the front end, the `mod_weblogic` configuration for WebCenter Suite applications should provide the VIP `cfcvip.mycompany.com` as the managed servers address instead. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points that WebCenter components use.

To use a virtual host as the managed servers address:

1. Verify that the Virtual Host CFCVIP is enabled on the machine that the managed server runs on.

2. Verify that the managed server Listen Address is set to CFCVIP.
3. Use the CFCVIP address in the `mod_wl_ohs.conf` file.

### 12.2.3.10 Transforming Oracle Portal, Forms, Reports, and Discoverer

Cold Failover Cluster configuration varies for each Oracle Portal, Forms, Reports, and Discoverer component. This section documents the steps for Cold Failover Cluster transformation of Oracle Portal, Forms, Reports, and Discoverer.

**12.2.3.10.1 Transforming Oracle Forms for Cold Failover Cluster** Oracle Forms is made up of both Java EE and system components. To configure Oracle Forms for Cold Failover Cluster:

---

**Note:** The transformation process requires that you shutdown the domain before editing these files.

---

1. Transform the WLS\_FORMS managed server. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for this procedure. Start the Domain Administration Server to do the Managed server transformation, and shut it down once the Managed Server transformation is finished.
2. Transform the Oracle Forms OPMN. See [Section 12.2.2.6, "Transforming Oracle Process Management and Notification Server"](#) for this procedure.
3. Transform Enterprise Manager. See [Section 12.2.2.7, "Transforming Oracle Enterprise Manager for an Oracle Instance"](#) for this procedure.
4. Transform Oracle Forms Web tier components. See [Section 12.2.2.8, "Transforming Web Tier Components and Clients"](#) for this procedure.
5. Reregister Single Sign-On using the steps described in the [Section 14.6.4.9.4, "Enable Single Sign On."](#)
6. In the `forms.conf` file located in the `INSTANCE_HOME/config/OHS/ohs1/moduleconf` directory, change the `mod_weblogic` configuration.

Start the Administration Server and Managed Servers in the WebLogic Domain as well as the Oracle instances to validate the Cold Failover Cluster installation:

```
<Location /Forms>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 DynamicServerList OFF
</Location>
```

**12.2.3.10.2 Transforming Oracle Reports for Cold Failover Cluster** Oracle Reports is made up of both Java EE and system components. To configure Oracle Reports for Cold Failover Cluster:

1. Transform the WLS\_REPORTS managed server. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for this procedure. Start the Domain Administration Server to do the Managed Server transformation, and shut it down once the Managed Server transformation is finished.
2. Transform the Oracle Reports OPMN. See [Section 12.2.2.6, "Transforming Oracle Process Management and Notification Server"](#) for this procedure

In addition, do the following:

- a. In the `opmn.xml` file, located in the `INSTANCE_HOME/config/OPMN/opmn` directory, change the Reports server name in `ias-component` to reflect its new name.

For example:

```
<ias-component id="ReportsServer_virtual_hostname_instance_name">
 <process-set id="ReportsServer_virtual_hostname_instance_name"
 restart-on-death="true" numprocs="1">
 </process-set>
</ias-component>
```

For example:

```
<ias-component id="ReportsServer_cfcvip_asinst1">
 <process-set id="ReportsServer_cfcvip_asinst1" restart-on-death="true"
 numprocs="1">
 </process-set>
</ias-component>
```

- b. In `DOMAIN_HOME/opmn/topology.xml` of the Administration Server domain home, change all occurrences of the Report server name from `ReportsServer_physical_hostname_instance_name` (example: `ReportServer_node1_asinst`) to `ReportsServer_virtual_hostname_instance_name` (example: `ReportServer_cfcvip_asinst`).
3. Transform Oracle Enterprise Manager. See [Section 12.2.2.7, "Transforming Oracle Enterprise Manager for an Oracle Instance"](#) for this procedure.
4. Transform Oracle Reports Web tier components. See [Section 12.2.2.8, "Transforming Web Tier Components and Clients"](#) for this procedure.
5. Reregister Single Sign-On. To do this perform the steps described in the [Section 14.6.4.9.4, "Enable Single Sign On."](#)
6. In the `reports_ohs.conf` file located in the `INSTANCE_HOME/config/OHS/ohs1/moduleconf` directory, change the following:

```
<Location /reports>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>
```

7. In the directory `INSTANCE_HOME/config/ReportsServerComponent`, rename the following directory:

`ReportsServer_physical_hostname_instance_name` (for example: `ReportServer_node1_asinst`)

To the following:

`ReportsServer_VIP_instance_name` (for example: `ReportServer_cfcvip_asinst`)

---

---

**Note:** All new logs for the report server go to `ReportsServer_Virtual_Hostname_Instance_Name` after restart of the instance.

---

---

8. In the renamed directory, replace physical hostname with a virtual hostname in the following files:
  - `component-logs.xml`

- logging.xml
- 9. In the `targets.xml` file, located in the `INSTANCE_HOME/EMAGENT/emagent_dir/sysman/emd` directory, change the following:
  - a. Change all hostnames from `node1.mycompany.com` to `cfcvip.mycompany.com`
  - b. Change the report server name from `ReportsServer_physical_hostname_instance_name` (for example: `ReportServer_node1_asinst`) to `ReportsServer_VIP_instance_name` (for example: `ReportServer_cfcvip_asinst`) for the following two elements:
 

```
Target TYPE="oracle_repserv"
Target TYPE="oracle_repapp"
```
- 10. In the `INSTANCE_HOME/reports` directory, replace the physical hostname with a virtual hostname in `reports_install.properties`.
- 11. In the `INSTANCE_HOME/reports/server` directory, rename the following files:
  - Rename `reportsserver_physical_hostname_instance_name.dat` to `reportsserver_virtual_host_name_instance_name.dat`
  - Rename `rep_wls_reports_physical_hostname_instance_name.dat` to `rep_wls_reports_virtual_hostname_instance_name.dat`
- 12. In the `DOMAIN_HOME/servers/WLS_REPORTS/tmp/_WL_user/reports_<version_number>/1ww3jm/configuration` directory (for example, the `DOMAIN_HOME/servers/WLS_REPORTS/tmp/_WL_user/reports_11.1.1.3.0/1ww3jm/configuration` directory), replace physical hostname with virtual hostname in the `rwervlet.properties` file (including changes to any occurrence of the reports server name).
- 13. In the `DOMAIN_HOME/servers/WLS_REPORTS/tmp/_WL_user/reports_<version_number>/1ww3jm/META-INF` directory (for example, the `DOMAIN_HOME/servers/WLS_REPORTS/tmp/_WL_user/reports_11.1.1.3.0/1ww3jm/META-INF` directory), replace physical hostname with virtual hostname in the `mbeans.xml` file (including changes to any occurrence of the reports server name).
- 14. Start the Administration Server and Managed Servers in the WebLogic Domain, as well as the Oracle instances to validate the Cold Failover Cluster installation.

**12.2.3.10.3 Transforming Oracle Discoverer for Cold Failover Cluster** Oracle Discoverer is made up of both Java EE and system components. To configure Oracle Discoverer for Cold Failover Cluster:

1. Transform the WLS\_DISCO managed server. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for this procedure. Start the domain Administration Server to do the Managed Server transformation, and shut it down once the Managed Server transformation is finished.
2. Transform the Oracle Discoverer OPMN. See [Section 12.2.2.6, "Transforming Oracle Process Management and Notification Server"](#) for this procedure.
3. Transform Oracle Enterprise Manager. See [Section 12.2.2.7, "Transforming Oracle Enterprise Manager for an Oracle Instance"](#) for this procedure.
4. Transform Oracle Discoverer Web tier components. See [Section 12.2.2.8, "Transforming Web Tier Components and Clients"](#) for this procedure.

5. Reregister Single Sign-On. To do this perform the steps described in the [Section 14.6.4.9.4, "Enable Single Sign On."](#)
6. In the `reports_ohs.conf` file located in the `INSTANCE_HOME/config/OHS/ohs1/moduleconf` directory, change the following:

```
<IfModule mod_weblogic.c>
<Location /discoverer>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 DynamicServerList ON
</Location>
</IfModule>
```

**12.2.3.10.4 Transforming Oracle Portal for Cold Failover Cluster** Oracle Portal is made up of both Java EE and system components. To configure Oracle Portal for Cold Failover Cluster:

1. Transform the WLS\_PORTAL managed server. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for this procedure. Start the domain Administration Server to do the Managed Server transformation, and shut it down once the Managed Server transformation is finished.
2. Transform the Oracle Portal OPMN. See [Section 12.2.2.6, "Transforming Oracle Process Management and Notification Server"](#) for this procedure.
3. Transform Oracle Enterprise Manager. See [Section 12.2.2.7, "Transforming Oracle Enterprise Manager for an Oracle Instance"](#) for this procedure.
4. Transform Oracle Portal Web tier components. See [Section 12.2.2.8, "Transforming Web Tier Components and Clients"](#) for this procedure.
5. Reregister Single Sign-On. To do this, perform the steps described in the [Section 14.6.4.9.4, "Enable Single Sign On."](#)
6. In the `portal.conf` file located in the `INSTANCE_HOME/config/config/OHS/ohs1/moduleconf`, change the mod weblogic configuration, change the mod weblogic configuration:

```
WLS routing configuration
<Location /portal>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

<Location /portalTools>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

<Location /wsrp-tools>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

<Location /richtextportlet>
 SetHandler weblogic-handler
```

```

 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 </Location>

 <Location /jpd<
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 </Location>

 <Location /portalHelp>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 </Location>

 <Location /portalHelp2>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 </Location>

```

## 7. Rewire the Portal Repository:

- a. Log into the domain WebLogic Server Enterprise Manager using the following URL:

`http://cfcvip.mycompany.com:7001/em`

- b. Change InValidation and Administrator password.

In the Navigator Window Expand the **Web Tier** tree.

Click the component **wc1**.

From the drop-down list at the top of the page, select **Administration - Passwords**.

Enter a new invalidation password, confirm it, and click **Apply**.

Enter a new administrator password, confirm it and click **Apply**.

- c. Expand the **Fusion Middleware** menu in the left pane.

- d. Select **Classic**.

- e. Click **Portal**

The **Portal Domain** information page appears.

- f. Right-click on **Portal** and select **Settings**, and then **Wire Configuration**.

- g. Enter the following information for **Portal Midtier**

**Host:** Enter the Cold Failover Cluster Virtual IP name of the Web Cache host cfcvip.mycompany.com.

**Port:** Enter the Web Cache port being used (HTTP or non HTTP)

**SSL Protocol:** Enter this value if appropriate.

- h. Enter the following information for **Web Cache**:

**Host:** Enter the Cold Failover Cluster Virtual IP name of the Web Cache host cfcvip.mycompany.com mysite.mycompany.com

**Invalidation port:** Enter the Invalidation port, for example, **9401**.



**Invalidation User Name:** Enter the user name for Portal invalidations.

**Invalidation Password:** Enter the password for this account.

- i. Click **Apply** to start the rewire.
- j. When the rewire is complete, click the **Portal** menu option again, and ensure that the Portal URL is now the following:

`https://cfcvip.mycompany.com:WCHTTPPort/portal/pls/portal`

**8. Change Host Assertion in Oracle WebLogic Server.**

Because the Oracle HTTP Server acts as a proxy for WebLogic, by default certain CGI environment variables, including the host and port, are not passed through to WebLogic. TO ensure that Web Logic is ware that it is using a virtual site name and port so that it can generate internal URLs appropriately:

- a. Log in to the WebLogic Server Administration Console using the following URL:

`http://cfcvip.mycompany.com:7001/console`

- b. Select **WLS\_PORTAL** from the home page or select **Environment**, and then **Clusters** from the **Domain Structure** menu.

- c. In the Change Center, click **Lock & Edit**.

- d. Click **Protocol**, and select **HTTP**.

- e. Enter the following values:

Frontend Host: **Cfcvip.mycompany.com**

Frontend HTTP Port: **WCHTTPPort (8090)**

Frontend HTTPS Port: **WCHTTPSPort (8094)**

This ensures that any HTTPS URLs created from within WebLogic are directed to port 443 on the load balancer.

- f. Click **Activate Changes** in the **Change Center** window to enable editing.

- g. Restart the WLS\_PORTAL managed server

**12.2.3.10.5 Transforming Oracle Business Activity Management (BAM)** Oracle BAM is made up of Java EE components deployed to a managed server. Since these are Java EE components, the Cold Failover Cluster transformation procedure involves configuring the managed server to which they are deployed to listen on the virtual IP. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for information on transforming managed servers WLS\_BAM.

When using Oracle HTTP Server as the front end, the mod WebLogic configuration for the applications deployed to WLS\_BAM should provide the VIP `cfcvip.mycompany.com` as the address of these managed server. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used by SOA components. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file:

```
BAM Web Application
<Location /OracleBAM >
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>
```



**12.2.3.10.6 Transforming a Custom ADF Deployment** For a deployment that uses a custom ADF application, Cold Failover Cluster can be used in the same way as any of the Fusions Middleware deployment. The domain is created in this case using the installation from Oracle Application Developer DVD. Since this is primarily a Java EE components. the Cold Failover Cluster transformation procedure involves configuring the managed server to which they are deployed to listen on the virtual IP. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for information on transforming managed servers.

When using Oracle HTTP Server as the front end, the mod WebLogic configuration for the applications deployed to WLS\_ADF (the name of the managed server for the customer app) should provide the VIP cfcvip.mycompany.com as the address of these managed server. To do this, change the WebLogic host configuration in the webserver proxy plugin configuration files for the mount points used by SOA components. For example, use a text editor to make the following edits in the mod\_wl\_ohs.conf file:

```
BAM Web Application
<Location /ADFApplicationMountPoint >
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>
```

### 12.2.3.11 Transforming an Oracle WebCenter Content

The Oracle WebCenter Content has products such as Oracle Imaging and Process Management (Oracle I/PM), Oracle Universal Content Management (Oracle UCM), Oracle Universal Records Management (URM), and Oracle Inbound Refineries (Oracle IBR).

These managed servers are typically included when you install Oracle WebCenter Content:

- Oracle I/PM (WLS\_IPM)
- Oracle UCM (WLS\_UCM)
- Oracle URM (WLS\_URM)
- Oracle IBR (WLS\_IBR)

For Cold Failover deployments, the recommendation is to set up the Application tier and the Web tier as recommended for CFC deployments in previous sections of this chapter, and to follow the procedure for Cold Failover Cluster transformation for these managed servers. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) for information on transforming Oracle I/PM, UCM (or URM), and IBR managed servers.

Additional things to consider in Cold Failover Cluster deployments of Oracle WebCenter Content include:

- In many Cold Failover Cluster deployments of Oracle WebCenter Content, the likely deployment topologies include:
  - Both I/PM and UCM deployed together on the same node of a hardware cluster.
  - A deployment with I/PM on one node of a hardware cluster and UCM on the other node of a hardware cluster and configured for mutual failover.
- All I/PM related files that need to be available when the I/PM server fails over should also be on the shared disk. This includes files such as the input files and the

images. These should always be on the shared disk so they are accessible from all nodes of a CFC configuration. It is recommended that there is a separate volume for the input files and the image files.

- All persistence stores such as one used by JMS and TLogs should also be on the shared disk.
- All restrictions related to Inbound Refinery continue to apply in a CFC configuration. For information on configuring Inbound Refinery instances, refer to *Oracle Fusion Middleware Administrator's Guide for Conversion*.
- All content related folders for UCM should be on a shared disk as well, which will fail over as a unit along with the UCM server configuration. This includes folders such as:
  - Content server instance folder
  - Native File repository location
  - Web Layout folder
- When configuring UCM or URM using `http://cfcvip.mycompany.com:port/cs`, the **Webserver HTTP Address** is the hostname and port of the location of the HTTP server front-ending the ECM managed servers. Depending on the high availability configuration of the HTTP server this can be a load balancer address, physical host or a virtual host. In the case where Oracle HTTP Server is also in a CFC deployment, the above value should be set to the virtual IP of the Oracle HTTP Server. When Oracle HTTP Server is collocated on the same hardware cluster as the ECM servers, this is likely to be `cfcvip.mycompany.com`.
- In cases where UCM is configured to be in a Cold Failover Cluster (including the example above), when I/PM is wired to UCM, the recommendation is to use the virtual IP address of UCM. When doing this configuration using `http://cfcvip.mycompany.com:port/imaging`:
  1. In the left-hand pane, click **Manage Connections**, and then **Create Content Server Connection**.
  2. Enter a name and description for the new connection, and then click **Next**.
  3. In the Connection Settings screen, enter the following:
    - Content Server Version: CS 11g
    - Primary: `cfcvip.mycompany.com:4444`Click **Next**.
  4. In the connection security screen, leave the default selections for the WebLogic user, and then click **Next**.
  5. Review the connection details and click **Submit**.
- In addition, when Oracle HTTP Server is the front end for Oracle WebCenter Content applications, the `mod_weblogic` configuration should use the VIP `cfcvip.mycompany.com` as the address of the managed servers. To do this, change the WebLogic host configuration in the `webserver proxy` plugin configuration files for the mount points used by ECM components. For example, use a text editor to make the following edits in the `mod_wl_ohs.conf` file:

```
Content Server
<Location /cs>
 SetHandler weblogic-handler
```

```

 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 WLCookieName IDCCS_SESSIONID
 </Location>

 # URM
 <Location /urm>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 WLCookieName IDCURM_SESSIONID
 </Location>

 # IBR
 <Location /ibr>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 WLCookieName IDCIBR_SESSIONID
 </Location>

 # IPM
 <Location /imaging>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
 </Location>

```

### 12.2.3.12 Transforming Oracle Business Intelligence

Oracle Business Intelligence is an integrated business intelligence (BI) solution that provides the business user with a complete picture across the entire organization. Oracle Business Intelligence is designed to quickly and easily integrate diverse data sources, find information from the database, share the database information, and exploit the data to learn more about the business and its customers.

This section describes active-passive high availability for Oracle Business Intelligence Enterprise Edition, Oracle Business Intelligence Publisher and Oracle Real-Time Decisions.

**12.2.3.12.1 Transforming Oracle Business Intelligence Enterprise Edition and its Clients** This section describes how to transform Oracle Business Intelligence Enterprise Edition (Oracle BI EE) to work in a Cold Failover Cluster environment.

#### Transforming Oracle Business Intelligence Enterprise Edition

The following section describes how to transform Oracle BI EE to work in a Cold Failover Cluster environment.

#### Transform Managed Server

Oracle BI Enterprise Edition is deployed to a managed server (for example, BI\_SERVER1) and the procedure for Cold Failover Cluster transformation is to configure this managed server after install to listen on the virtual IP `cfcvip.mycompany.com`. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the BI\_SERVER1 managed server to listen on the `cfcvip.mycompany.com` virtual IP. After these steps, restart the managed server using the WebLogic Server Administration Console or the WLST command line.

All requirements for placing the Middleware Home and other related domain artifacts on shared storage that can be failed over apply as [Section 12.2.1, "General Requirements for Cold Failover Cluster"](#) describes. The specific requirement for Cold Failover Cluster deployments of the Oracle BI Enterprise Edition is to ensure that the following artifacts are also on the shared disk and available on the same mount points on both nodes of the failover cluster. When configured in the default install, this happens implicitly:

- BIEE Presentation Catalog
- BIEE Repository Publishing Directory (RPD)

To transform an Oracle BI EE component managed server:

1. Edit the `biee-domain.xml` file to replace the hostname with `cfcvip.mycompany.com`. The `biee-domain.xml` file is located in this directory:  
`MW_HOME/user_projects/domains/bifoundation_domain/config/fmwconfig`
2. Set the front-end host for the BI\_SERVER managed server to `cfcvip.mycompany.com`.
  - a. Log into Oracle WebLogic Server Administration Console.
  - b. In the **Environment** section, select **Servers**.
  - c. Select the name of the managed server.
  - d. Select **Protocols**, then select **HTTP**.
  - e. In the **Frontend Host** field, enter the host name as `cfcvip.mycompany.com`.
  - f. Set the Frontend Port to the HTTP port
  - g. Click **Save**.
  - h. Activate the changes.
  - i. Restart the Managed Server.

### Transforming System Components

To transform system components:

1. Log into Oracle Fusion Middleware Console for the domain.
2. Navigate to **Business Intelligence > Core Application > Capacity Management > Scalability**.
3. For the Oracle Instance (instance1), change the ListenAddress to `cfcvip.mycompany.com`.
4. Click **Lock and Edit Configuration** and change the Listen Address to `cfcvip.mycompany.com`.
5. Click **Apply**.
6. Click **Activate Changes**.
7. Restart both the Managed servers and the System components (using `opmnctl`).

### Transforming Oracle BI Enterprise Edition Clients

---

**Note:** Change data sources that refer to the BI Enterprise Edition instance before you transform the managed server to CFC. See step 4.

---

To transform Oracle BI Enterprise Edition clients:

1. Oracle BI Enterprise Edition clients must use the virtual IP `cfcvip.mycompany.com` to access these applications.
2. Since WSM-PM is deployed to the same managed server, ensure that the client side configuration changes required for WSM-PM are also complete.
3. Configure other components in the BI Enterprise Edition product suite, such as Oracle BI Publisher, that use Oracle BI Enterprise Edition services to take advantage of these services using the virtual hostname.
4. In BI Publisher, data sources that refer to this BI Enterprise Edition instance should change or be created (if new using the new virtual host). To change:
  - a. Navigate to the BI Publisher Administration Console.
  - b. Click **JDBC Connection** under Data Sources.
  - c. Edit any Data source for BI Enterprise Edition for this instance to connect to `cfcvip.mycompany.com`.
5. Where BI Publisher is integrated with BIEE deployment, configure BI Publisher to discover the Presentation Services at the new virtual address. Check details for these in the BI publisher section of this chapter.
6. When Oracle HTTP Server is the front end for Oracle BI Enterprise Edition, use a text editor to update the `mod_wl_ohs` file to specify the virtual IP `cfcvip.mycompany.com` as the address for the Oracle BI Server managed server. See the following example:

```
BIEE Analytics
<Location /analytics>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>

<Location /analytics-ws>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>

<Location /bimiddleware>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>
```

### Special Considerations for Windows

To achieve Cold Failover Cluster Deployment on a Windows platform:

1. Make the shared disk available on drive (for example, E:\).
2. Install BI EE suite (software-only install) on node1 of the Windows Cluster on E:\.
3. Delete the software-only install from E:\ by removing the Middleware Home.
4. Failover the same shared disk to node2 and make it available on node2 as the same drive (E:\).
5. Install BI EE suite (software-only install) on node2 of the Windows Cluster on E:\.

6. Invoke the OUI Configuration Wizard and create the WLS domain and the Oracle Instance on E:\.
  7. On Node 1:
    - a. Run regedit.
    - b. Locate the node HKEY\_LOCAL\_MACHINE\SOFTWARE\ORACLE.
    - c. Right click on this node, click **Export**, and enter a filename.
    - d. Locate the node HKEY\_LOCAL\_MACHINE\SOFTWARE\ODBC
    - e. Right click on this node and click **Export**. Enter a file name.
    - f. Locate the node HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet001\Services. Export each node within this node that starts with Oracle to a .reg file.
    - g. Locate the node HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet002\Services. Export each node within this node that starts with Oracle to a .reg file.
    - h. Locate the node HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services. Export each node within this node that begins "Oracle" to a .reg file
    - i. Copy all the reg files that you exported to node 2.
- On Node 2:
- a. Double click Oracle\_BI1\bifoundation\install\vc80\vcredist\_x86.exe
  - b. Double click on each reg file you copied from node 1. For each file, click **Yes** to import the contents into the registry.
8. Enable the Virtual IP on the current node (node2) and perform the transformation steps for this BI suite install in [Section 12.2.3.12.1, "Transforming Oracle Business Intelligence Enterprise Edition and its Clients."](#)

**12.2.3.12.2 Transforming Oracle Business Intelligence Publisher and its Clients** This section describes how to transform Oracle Business Intelligence Publisher (Oracle BI Publisher) to work in a Cold Failover Cluster environment.

### Transforming Oracle Business Intelligence Publisher

---

**Note:** You must hand edit the BIP connection to BI Presentation server before transforming the managed server to CFC. Change the hostname in user\_projects/domains/bifoundation\_domain/config/bipublisher/repository/Admin/Configuration/xmlp-server-config.xml, property SAW\_SERVER.

---

To transform Oracle BI Publisher to work in a Cold Failover Cluster environment:

1. Open the BI Publisher application at `http://hostname:port/xmlpserver` and login.
2. Change the BI Scheduler 's JMS configuration:
  - a. Click **Administration**.
  - b. Click **Scheduler Configuration** (under System Maintenance)

- c. Change Weblogic JNDI URL to `t3://cfc.mycompany.com:9704`
- d. Click **Apply**.

### Transform Managed Server

Oracle BI Publisher is deployed to a managed server, for example, `BI_SERVER1`, and the procedure for Cold Failover Cluster transformation is to configure this managed server after install to listen on the virtual IP `cfcvip.mycompany.com`. See [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the `BI_SERVER1` managed server to listen on the `cfcvip.mycompany.com` virtual IP.

Then restart the managed server using the WebLogic Server Administration Console or the WLST command line.

All requirements related to placing the Middleware Home and other related domain artifacts on shared storage that can be failed over apply as described in [Section 12.2.1, "General Requirements for Cold Failover Cluster."](#) The specific requirement for Cold Failover Cluster deployments of the Oracle BI Publisher is to ensure that the following artifacts are also on the shared disk and available on the same mount points on both failover cluster nodes. When configured in the default install, this happens implicitly:

- Transaction logs and JMS persistence store
- BI Publisher configuration folder
- BI Publisher Catalog repository
- BI Publisher Scheduler temp directory

### BI EE Integration

When you deploy BI Publisher and BI EE together, BI Publisher is integrated to use BI EE Presentation Services. In this setup, the BI EE deployment is typically also a Cold Failover Cluster deployment listening on the same virtual hostname, `cfc.mycompany.com`. In this case, you must change the BI Publisher configuration related to BI EE integration to ensure that BI Publisher discovers the BI EE Presentation Services on this new address. To do this:

1. Invoke the BI Publisher application URL (`http://hostname:port/xmlpserver`) and log in.
2. Change the BI Scheduler JMS configuration:
  - a. Click **Administration**.
  - b. Click **Scheduler Configuration** (under System Maintenance).
  - c. Change Weblogic JNDI URL to `t3://cfc.mycompany.com:9704`
  - d. Click **Apply**.
3. Restart the Managed server using Admin server console or the command line.

All requirements for placing the Middleware Home and other related domain artifacts on a shared storage that can be failed over applies as [Section 12.2.1, "General Requirements for Cold Failover Cluster"](#) describes. The requirement for CFC deployments of Oracle Business Intelligence Publisher ensures that the following artifacts are also on the shared disk and available on the same mount points on both nodes of the failover cluster:

- Transaction Logs and JMS persistence Store
- BI publisher Configuration Folder
- BI Publisher Catalog repository

- BI Publisher Scheduler Temp Directory

### Transforming Oracle BI Publisher Clients

To transform Oracle BI Publisher clients:

1. Oracle BI Publisher clients must use the virtual IP `cfcvip.mycompany.com` to access these applications.
2. Since WSM-PM is deployed to the same managed server, ensure that the client side configuration changes required for WSM-PM are done as well.
3. When Oracle HTTP Server is the front end for Oracle BI Publisher, use a text editor to update the `mod_wl_ohs` file to specify the virtual IP `cfcvip.mycompany.com` as the address for the Oracle BI Publisher managed server, as shown in the following example:

```
BI Publisher
<Location /xmlpservlet>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>

WSM-PM
<Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1:9704, APPHOST2:9704
</Location>
```

**12.2.3.12.3 Transforming Oracle Real-Time Decisions and its Clients** This section describes how to transform Oracle Real-Time Decisions (Oracle RTD) to work in a Cold Failover Cluster environment.

### Transforming Oracle RTD

To transform Oracle RTD to a Cold Failover Cluster deployment, follow the instructions in this section.

Oracle RTD is deployed to a managed server (for example, `BI_SERVER1`) and the procedure for Cold Failover Cluster transformation is to configure this managed server after install to listen on the virtual IP `cfcvip.mycompany.com`. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the `BI_SERVER1` managed server to listen on the `cfcvip.mycompany.com` virtual IP.

All requirements related to the placement of the Middleware Home and other related domain artifacts on shared storage that can be failed over apply, as described in [Section 12.2.1, "General Requirements for Cold Failover Cluster."](#)

### Transforming Oracle RTD Clients

To transform Oracle RTD clients:

1. Oracle RTD clients must use the virtual IP `cfcvip.mycompany.com` to access these applications.
2. Since WSM-PM is deployed to the same managed server, ensure that the client side configuration changes required for WSM-PM are done as well.
3. When Oracle HTTP Server is the front end for Oracle RTD, use a text editor to update the `mod_wl_ohs` file to specify the virtual IP



cfcvip.mycompany.com as the address for the Oracle RTD managed server, as shown in the following example:

```
<Location /rtis>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>

<Location /schemas>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>

<Location /ws>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com:<port>
</Location>

WSM-PM
<Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST:9704, APPHOST2:9704
</Location>

RTD
<Location /ui>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1:9704, APPHOST2:9704
</Location>
```

### 12.2.3.13 Transforming Oracle BI for Microsoft Office and its Clients

This section provides instructions to transform Oracle BI for Microsoft Office to work in a cold Failover Cluster environment.

Oracle BI for Microsoft Office is deployed to a managed server, for example, BI\_Server1, and the procedure for Cold Failover Cluster transformation is to configure this server to listen on the virtual IP cfcvip.mycompany.com. Follow the steps in [Section 12.2.2.4, "Transforming Oracle WebLogic Managed Servers"](#) to configure the BI\_SERVER1 managed server to listen on the cfcvip.mycompany.com virtual IP.

All requirements related to Middleware Home placement and other domain artifacts on shared storage that can be failed over apply as [Section 12.2.1, "General Requirements for Cold Failover Cluster"](#) describes.

---

**Note:** For BI Office to work correctly in a CFC configuration, you must change the local host value in the biooffice.xml file to the new VIP value. See `/bidomain/servers/bi_server1/tmp/_WL_user/biooffice_11.1.1/cvsibb/war/WEB-INF` for the biooffice.xml file.

---

### Transforming Oracle BI for Microsoft Office Clients

All client invocations of BI Office must use cfcvip.mycompany.com. For OHS frontending the BIOFFICE, use a text editor to update the mod\_wl\_ohs file to specify the virtual IP cfcvip.mycompany.com as the address for the Oracle BI for Microsoft Office managed server. See the following example:

```
BI Office
```

```
<Location /biooffice>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>

<Location /bioofficeclient>
 SetHandler weblogic-handler
 WebLogicHost cfcvip.mycompany.com
 WebLogic port
</Location>
```

#### 12.2.3.14 Single Sign-On Re-registration (If required)

Single Sign-On (SSO) re-registration typically applies only to Oracle Portal, Forms, Reports, and Discoverer. Once the front end listening endpoint on Oracle HTTP Server for this tier has been changed to the Virtual IP, it becomes necessary to do SSO re-registration so that the URL to be protected is configured with the virtual IP. To re-register SSO, perform these steps on the 10.1.x installation of Identity Management where the SSO server resides:

1. Set the `ORACLE_HOME` variable to the SSO `ORACLE_HOME` location.
2. Execute `ORACLE_HOME/sso/bin/ssoreg.sh` (`ssoreg.bat` for Windows) with the following parameters:

```
-site_name cfcvip.mycompany.com:port
-mod_osso_url http://cfcvip.mycompany.com
-config_mod_osso TRUE
-oracle_home_path ORACLE_HOME
-config_file /tmp/osso.conf
-admin_info cn=orcladmin
-virtualhost
-remote_midtier
```

3. Set the `ORACLE_HOME` variable to the SSO `ORACLE_HOME` location.
4. Copy `/tmp/osso.conf` file to the mid-tier home location:  
`ORACLE_INSTANCE/config/OHS/ohs1`
5. Restart Oracle HTTP Server by issuing the following command from the `ORACLE_HOME/opmn/bin/directory`:

```
opmnctl restartproc process-type=OHS
```

6. Log into the SSO server through the following URL:

```
http://login.mycompany.com/pls/orasso
```

7. In the **Administration** page and then **Administer Partner** applications, delete the entry for **node1.mycompany.com**.

## 12.2.4 Transforming an Oracle Database

In a typical Cold Failover Cluster deployment of Oracle Fusion Middleware, the database is also deployed as a cold failover cluster. This section describes how to transform a single instance Oracle database to a Cold Failover Cluster database. Perform this transformation before seeding the database using RCU and subsequent Fusion Middleware installations that use this seeded database.

To enable the database for Cold Failover Cluster:

1. Change the listener configuration that the database instance uses in the `listener.ora` file.

Ensure that the `HOST` name in the listener configuration has the value of the virtual hostname. In addition, ensure that no other process (Oracle or third party) uses the listener port.

```
<listener_name> =
 (DESCRIPTION_LIST =
 (DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP) (HOST = <virtual hostname>) (PORT = port))
)
)
```

For example:

```
LISTENER_CFCDB =
 (DESCRIPTION_LIST =
 (DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP) (HOST = cfcdhdbhost.mycompany.com) (PORT =
1521))
)
)
```

2. Change the `tnsnames.ora` file.

Change an existing TNS service alias entry or create a new one:

```
<tns alias name> =
 (DESCRIPTION =
 (ADDRESS_LIST =
 (ADDRESS = (PROTOCOL = TCP) (HOST = <virtual hostname>) (PORT = port))
)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = <db service name>)
 (INSTANCE_NAME = <db instance name>)
)
)
```

For example:

```
CFCDB =
 (DESCRIPTION =
 (ADDRESS_LIST =
 (ADDRESS = (PROTOCOL = TCP) (HOST = cfcdhdbhost.mycompany.com) (PORT = 1521))
)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = cfcdhdb)
 (INSTANCE_NAME = cfcdhdb)
)
)
```

3. Change the local `sp` file to update the `local_listener` parameter of the instance.

Log in as `sysdba` using `SQL*Plus`:

```
SQL> alter system set local_listener='<tns alias name>' scope=both;
)
```

For example:

```
SQL> alter system set local_listener='CFCDDB' scope=both;
```

4. Shutdown and restart the listener.
5. Shutdown and restart the database instance.
6. Create the database service for the application server.

Oracle recommends a dedicated service separate from the default database service used with the Oracle Application server. To create this service, execute the following SQL\*Plus command:

```
SQL> execute DBMS_SERVICE.CREATE_SERVICE(
 '<cfc db service name>' , '<cfc db network name>') \
```

For example:

```
SQL> execute DBMS_SERVICE.CREATE_SERVICE(
 'cfcd_basservice' , 'cfcd_basservice')
```

```
SQL> execute DBMS_SERVICE.START_SERVICE('cfcd_basservice')
```

You can set additional parameters for this service depending on the needs of the installation. See the *Oracle Database PL/SQL Packages and Types Reference* for details about the DBMS\_SERVICE command.

#### 12.2.4.1 Database Instance Platform-Specific Considerations

Consider the following procedures for Unix and Windows platform database instances:

##### For Unix:

1. Manually Failover the Database Oracle Home from Node 1 (the installation node) to the failover node (Node 2) following the mount/unmount procedure described earlier.
2. As root, do the following:
  - Create an oraInst.loc file located in the /etc directory identical to the file on Node1.
  - Create an oratab file located in the /etc directory, identical to the file on Node1.
  - Run the oracleRoot.sh file located in the ORACLE\_HOME directory on node2, if it is required, and is available for the product suite.
3. Create the oraInventory file on the second node, using the attachHome command located in ORACLE\_HOME/oui/bin/attachHome.sh directory.

##### For Windows:

If the system environment variable 'Path' on Node 1 contains any information pertaining to this database, ensure this information is set in the same variable on Node 2. Export the registry HKEY\_LOCAL\_MACHINE\SOFTWARE\Oracle from Node 1 and import it to Node 2. Also, ensure the user 'system' is also in the ORA\_DBA group on Node 2.

Using the SC tool from the Windows Resource Kit, create services for the following:

##### ■ db service

```
sc create OracleService<oracle_sid> start= demand binPath= "ORACLE_
```

```
HOME\bin\ORACLE.EXE <oracle_sid>"
```

For example:

```
sc create OracleServiceORCL start= auto
binPath= "C:\Oracle\db\product\11.1.0\bin\ORACLE.EXE <oracle_sid>"
```

---

---

**Note:** oracle\_sid should be in upper case

---

---

- **listener**

```
sc create Oracle<home name>TNSListener start= demand
binPath= "ORACLE_HOME\bin\TNSLSNR"
```

For example:

```
sc create OracleINFRAOracleTNSListener start= demand
binPath= "C:\Oracle\db\product\11.1.0\bin\TNSLSNR"
```

- **cluster services**

```
sc create OracleCSService start= auto
binPath= "ORACLE_HOME\bin\ocssd.exe service"
```

For example:

```
sc create OracleCSService start= auto
binPath= "INFRAHOME\bin\ocssd.exe service"
```

- **database console**

```
sc create OracleDBConsole<oracle_sid> start= auto
binPath= "ORACLE_HOME\bin\nmesrv service"
```

For example:

```
sc create OracleDBConsoleorcl start= auto
binPath= " C:\Oracle\db\product\11.1.0\bin\nmesrv service"
```

- **job scheduler**

```
sc create OracleJobScheduler<oracle_sid> start= demand
binPath= "ORACLE_HOME\bin\extjob.exe <oracle_sid>"
```

For example:

```
sc create OracleJobSchedulerORCL start= auto
binPath= "C:\Oracle\db\product\11.1.0\INFRAHOME\bin\extjob.exe <oracle_sid>"
```

---

---

**Note:** oracle\_sid should be in upper case.

---

---

- **Vss service**

```
sc create OracleVssWriter<oracle_sid> start= demand
binPath= "ORACLE_HOME\bin\OraVSSW.exe <oracle_sid>"
```

For example:

```
sc create OracleJobSchedulerORCL start= auto
binPath= "C:\Oracle\db\product\11.1.0\INFRAHOME\bin\ OraVSSW.exe <oracle_sid>"
```

---

**Note:** `oracle_sid` should be in upper case.


---

## 12.3 Oracle Fusion Middleware Cold Failover Cluster Example Topologies

This section illustrates sample Cold Failover Cluster topologies. Since there are many possible combinations of topologies, these topologies are illustrative only. To achieve these topologies, more than one of the transformation steps apply. Refer to the steps mentioned earlier to configure the transformation.


### 12.3.1 Example Topology 1

[Figure 12-3](#) shows an Oracle WebCenter Portal Cold Failover Cluster deployment. Both the Administration Server and the WebCenter Managed Servers are in the domain, and failover as unit. Therefore, they share the same virtual IP and are installed together on the same shared disk. There may be an Oracle HTTP Server front ending this topology. It is on a separate node in the example topology. It can also be on the same node, and can be part of the Cold Failover Cluster deployment. In this example, the database is also on a separate node. However, it is equally likely that the database is on the same cluster and is also Cold Failover Cluster-based (using its own virtual IP and shared disk).

**Figure 12-3 Cold Failover Cluster Example Topology 1**

### 12.3.2 Example Topology 2

Figure 12-4 shows an example SOA Cold Failover Cluster deployment. In this example, only the SOA instance is deployed as Cold Failover Cluster, and the Administration Server is on a separate node. The database is also on a separate node in this example topology. Oracle HTTP Server in this case is part of the Cold Failover Cluster deployment, and part of the same failover unit as the SOA Managed Servers. Important variants of this topology include a Cold Failover Cluster Administration Server on the same hardware cluster. It may share the same virtual IP and shared disk as the SOA Managed Servers (SOA and Administration Server are part of the same failover unit) or use a separate virtual IP, and shared disk (Administration Server fails over independently). Similarly, depending on the machine capacity, the database instance can also reside on the same hardware cluster.


**Figure 12–4 Cold Failover Cluster Example Topology 2**

### 12.3.3 Example Topology 3

Figure 12–5 shows an Oracle Identity Management deployment. In this example topology, all components are on a two-node hardware cluster. Identity Management fails over as a unit, and both the Java EE (Administration Server and WLS\_ods Managed Server) and system components are part of the same failover unit. They share the same virtual IP and shared disk (`cfcvip1.mycompany.com`). The database is also on the same hardware cluster. It uses a different virtual IP (`cfcvip2.mycompany.com`), and a different set of shared disks. During normal operations, the database runs on Node2 and the IDM stack runs on Node1. The other node acts as a back up for each.

This topology is the recommended topology for most Cold Failover Cluster deployments. The example is for Identity Management, but this is true for the Oracle SOA, Oracle WebCenter Portal, and Oracle Portal, Forms, Reports, and Discoverer suites as well. In this recommended architecture, Oracle Fusion Middleware runs as one node of the hardware cluster. The Oracle database runs on the other node. Each node is a backup for the other. The Oracle Fusion Middleware instance and the database instance failover independently of each other, using different shared disks and different VIPs. This architecture also ensures that the hardware cluster resources are optimally used.



**Figure 12-5 Cold Failover Cluster Example Topology 3**

## 12.4 Transforming the Administration Server in an Existing Domain for Cold Failover Cluster

This section describes the steps for transforming an Administration Server in an existing domain to Cold Failover Cluster.

---

**Note:** After Administration Server transformation, client side changes for Administration Server and Enterprise Manager, as mentioned in [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster,"](#) may be required.

---

### Assumptions

The procedures in this section assume that:

- All the Fusion Middleware components are in a nostage deployment.
- The starting topology is an active-active cluster of the product suite (Node1 and Node2).
- The administration server is on Node1 to start.
- It shares the same domain home as the Managed Server on Node1.
- The `MW_HOME` path is the same on both Node1 and Node2.

### Start Topologies

Refer to the following sections to achieve the start topologies:


- Oracle SOA Suite: [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines"](#)
- Oracle ADF: [Section 6.2, "Configuring an Oracle ADF High Availability Deployment"](#)
- Oracle WebCenter Portal: [Section 6.4, "Configuring High Availability for WebCenter Portal"](#)

- Oracle Identity Manager: [Chapter 8, "Configuring High Availability for Identity Management Components"](#)

These procedures also apply to customer applications deployed to Oracle WebLogic cluster installations, provided the previously stated assumptions are met.

[Figure 12–6](#) shows an example start topology before transforming the Administration Server.


**Figure 12–6 Cold Failover Cluster Example Start Topology**



### 12.4.1 Destination Topologies

[Figure 12–7](#) shows the possible destination topology after transforming the Administration Server for Cold Failover Cluster with the following characteristics:

- The Administration Server Domain Home is moved out onto a shared disk that is mountable by both Node1 and Node2, but is mounted by either one of the two at any given point in time.
- It continues to use the original Middleware Home available on Node1 and Node2.
- The Listen Address of the Administration Server is moved to a virtual IP.

**Figure 12-7 Possible Destination Topologies**

## 12.4.2 Cold Failover Cluster Transformation Procedure

To transform the Administration Server in an existing domain:

1. Shutdown the cluster of the component Managed Servers and Administration Server.
2. Shutdown the Node Manager process on each of the nodes (if it is running).
3. Back up of the entire domain.
4. Provision the Virtual IP on Node1. For example:

For Linux:

```
ifconfig eth0:1 IP_Address netmask netmask
/sbin/arping -q -U -c 3 -I eth0 IP_Address
```

Where *IP\_Address* is the virtual IP address and the *netmask* is the associated netmask.

In the example below, the IP address is enabled on the interface "Local Area Connection."

```
ifconfig eth0:1 130.35.46.17 netmask 255.255.224.0
/sbin/arping -q -U -c 3 -I eth0 130.35.46.17
```

For Windows:

```
netsh interface ip add address interface IP_Address netmask
```

Where *IP\_Address* is the virtual IP address and the *netmask* is the associated netmask.

In the example below, the IP address is enabled on the interface "Local Area Connection."

```
netsh interface ip add address "Local Area connection" 130.35.46.17
255.255.224.0
```

5. Start the Administration Server from the local Domain Home.

```
cd DOMAIN_HOME/bin
./startWeblogic.sh
```

6. Transform the Administration Server instance to Cold Failover Cluster:

Log into the Oracle WebLogic Server Administration Console.

Create a Machine for the Virtual Host.

- a. Select **Environment**, and then **Machines**.
- b. In the Change Center, click **Lock & Edit**.
- c. Click **New**.
- d. In the **Name** field, enter **cfcvip.mycompany.com**.

---

**Note:** Keep the Listen Address field set to localhost; the CFC solution relies on this setting. Do not change it to the VIP address or any other value.

---

- e. Select the appropriate operating system and click **OK**.
- f. Select the machine you just created.
- g. Click the **Servers** tab then click **Add**.
- h. Select an existing server, and associate it with this machine.
- i. In the **Select Server** drop-down list, ensure **AdminServer** is selected.
- j. Click **Finish**.
- k. Click **Activate Changes**.

Configure the administration server to listen on cfcvip.mycompany.com.

- a. Select **Environment**, and then **Servers** from the Domain Structure menu.
- b. In the Change Center, click **Lock & Edit**.
- c. Click on the Administration Server (**AdminServer**).
- d. Change the **Listen Address** to **cfcvip.mycompany.com**
- e. Click **Save**.

Stop the Administration Server from the Administration Console.

- a. Select **Environment**, and then **Servers** from the Domain Structure menu.
- b. Click **Control**.
- c. Select **Adminserver** by clicking on the checkbox next to it.
- d. Shut down Adminserver by selecting **Force Shutdown Now** under **Shutdown** pull-down menu.
- e. Click **Yes**.

Ensure that the VIP is enabled on the system and start the Administration Server from the command line:

```
cd DOMAIN_HOME/bin
./startWeblogic.sh
```

7. Validate the Administration Server by accessing the consoles on the virtual IP:
  - `http://cfcvip.mycompany.com:7001/console`
  - `http://cfcvip.mycompany.com:7001/em`
8. Shut down the Administration Server on Node1 using the Administration Console.
9. Ensure that the shared disk is provisioned and mounted on Node1.
10. Pack the entire domain on Node1:

```
cd ORACLE_COMMON_HOME/common/bin
./pack.sh -managed=false -domain=/localdisk/user_projects/domains/domain_name
-template=cfcdomaintemplate_all.jar -template_name=cfc_domain_template_all
```

11. Unpack it on the shared disk:

```
./unpack.sh -domain=/sharedisk/user_projects/domains/domain_name
-template=cfcdomaintemplate_all.jar
-app_dir=/sharedisk/user_projects/apps -server_start_mode=prod
```

12. Product suite-specific changes:

In the Oracle Identity Management Product Suite:

- Back up the `config.xml` file, located in the `/sharedisk/user_projects/domains/domain_name/config/` directory in this example.
- Change ECM Socket Host Filter to VIP in the following files:
  - `{domain}/ucm/ibr/config/config.cfg`
  - `{domain}/ucm/urm/config/config.cfg`
  - `{domain}/ucm/cs/config/config.cfg`
- Edit the `config.xml` file, located in the `/sharedisk/user_projects/domains/domain_name/config` directory in this example, and make the following changes to source-path:

For `dipapp`, change source path to `ORACLE_HOME/ldap/odi/dipapp/dipapps.ear`.

For example:

```
<app-deployment>
 <name>DIP#11.1.1.2.0</name>
 <target>cluster_ods</target>
 <module-type>ear</module-type>
 <source-path>ORACLE_HOME/ldap/odi/dipapp/dipapps.ear</source-path>
 <security-dd-model>DDOnly</security-dd-model>
 <staging-mode>nostage</staging-mode>
</app-deployment>
```

For the `ODSM` app, change the source path to `ORACLE_HOME/ldap/odsm/odsm.ear`

For example:

```
<app-deployment>
 <name>odsm#11.1.1.2.0</name>
 <target>cluster_ods</target>
 <module-type>ear</module-type>
 <source-path>ORACLE_HOME/ldap/odsm/odsm.ear</source-path>
 <security-dd-model>DDOnly</security-dd-model>
```

```
<staging-mode>nostage</staging-mode>
</app-deployment>
```

- For the OIF:

Change the source path of OIF-APP to *ORACLE\_HOME*/fed/install/oif.ear:

```
<app-deployment>
 <name>OIF#11.1.1.2.0</name>
 <target>cluster_oif</target>
 <module-type>ear</module-type>
 <source-path>ORACLE_HOME/fed/install/oif.ear</source-path>
 <security-dd-model>Advanced</security-dd-model>
 <staging-mode>nostage</staging-mode>
</app-deployment>
```

Change the source path of oif-libs to *ORACLE\_HOME*/lib/java/shared/oracle.idm.oif/11.1.1.0.0/oif-libs.ear:

```
<library>
 <name>oif-libs#11.1.1.2.0@11.1.1.2.0</name>
 <target>cluster_oif</target>
 <module-type>ear</module-type>
 <source-path>ORACLE_HOME/lib/java/shared/oracle.idm.oif/11.1.1.0.0/oif-libs.ear</source-path>
 <security-dd-model>DDOnly</security-dd-model>
</library>
```

13. Start the Administration Server:

```
cd /shareddisk/user_projects/domains/domain_name/bin
./startWeblogic.sh
```

14. Validate the Administration Server by accessing the consoles on the virtual IP.

- <http://cfcvip.mycompany.com:7001/console>
- <http://cfcvip.mycompany.com:7001/em>

15. Shut down the Administration Server.

16. Perform the following on Node1:

```
mv /localdisk/user_projects/domains/domain_name
 /localdisk/user_projects/domains/domain_name_old
mv /localdisk/user_projects/applications/domain_name
 /localdisk/user_projects/applications/domain_name_old
cd ORACLE_COMMON_HOME/common/bin
./pack.sh -managed=true -domain=/shareddisk/user_projects/domains/domain_name
-template=cfcdomaintemplate_mngd.jar -template_name=cfc_domain_template_mngd
./unpack.sh -domain=/localdisk/user_projects/domains/domain_name
-template=cfcdomaintemplate_mngd.jar
```

---

---

**Note:** These commands assume an applications directory exists under *user\_projects*.

---

---

17. Copy the template to Node2:

```
scp ORACLE_COMMON_HOME/common/bin/cfcdomaintemplate_mngd.jar
```

```
Node2:ORACLE_COMMON_HOME/common/bin
```

## 18. Log in to Node2, and unpack on Node2

```
mv /localdisk/user_projects/domains/domain_name
/localdisk/user_projects/domains/domain_name_old
mv /localdisk/user_projects/applications/domain_name
/localdisk/user_projects/applications/domain_name_old
cd ORACLE_COMMON_HOME/common/bin
./unpack.sh -domain=/localdisk/user_projects/domains/domain_name
-template=cfcdomaintemplate_mngd.jar
```

---

**Note:** These commands assume an applications directory exists under `user_projects`.

---

## 19. For Oracle Identity Manager Suite, the following addition steps are required:

For DIP:

- a. Locate the applications directory in the Oracle WebLogic Server domain directory on Node1:

```
MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers/wls_ods1/applications
```

- b. Copy the applications directory and its contents on Node1 to the same location in the domain directory on Node2.

```
scp -rp MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig/servers
/wls_ods1/applications
user@IDMHOST2:MW_HOME/user_projects/domains/IDMDomain/config/fmwconfig
/servers/wls_ods2/applications
```

For OIF:

- a. Locate the applications directory in the Oracle WebLogic Server domain directory on Node1:

```
MW_HOME/user_projects/domains/OIFDomain/config/fmwconfig/servers/wls_oif1/applications
```

- b. Copy the applications directory and its contents on Node1 to the same location in the domain directory on Node2.

```
scp -rp MW_HOME/user_projects/domains/OIFDomain/config/fmwconfig/servers
/wls_oif1/applications
user@IDMHOST2:MW_HOME/user_projects/domains/OIFDomain/config/fmwconfig
/servers/wls_oif2/applications
```

## 20. Start the Administration Server:

```
cd /sharedisk/user_projects/domains/domain_name/bin
./startWeblogic.sh
```

## 21. Start the node manager (if used) on Node1 and Node2:

```
cd WL_HOME/server/bin
./startNodeManager.sh
```

22. Start the component Managed Servers on Node1 and Node2 from the Administration Server Console (if Node Manager used) or from the command line.

23. Validate the deployment using component-specific functional tests.

24. Test Administration Server failover.

Failover the Administration Server manually to the second node:

- a. Stop the Administration Server process (and any other process running out of a given Middleware Home).
- b. Unmount the shared storage from Node1 where the Middleware Home or domain directory exists.
- c. Mount the shared storage on Node2, following storage specific commands.
- d. Disable the virtual IP on Node1:

For Linux:

```
ifconfig interface:index down
```

In the following example, the IP address is disabled on the interface eth0:

```
ifconfig eth0:1 down
```

For Windows:

```
netsh interface ip delete address interface addr=IP Address
```

Where IP Address is the virtual IP address.

In the following example, the *IP address* is enabled on the interface Local Area Connection.

```
netsh interface ip delete address 'Local Area connection' addr=130.35.46.17
```

- e. Enable the virtual IP on Node2.
- f. Start the Administration Server process using the following command:

```
DOMAIN_HOME/bin/startWebLogic.sh
```

Where *DOMAIN\_HOME* is the location of your domain directory.

Validate access to both the Administration Server and Oracle Enterprise Manager Administration Console.

Validate with component-specific tests.

25. After validation, fail back the Administration Server to the node where it will normally run (this could be Node1 or Node2) and run normal operations.



---

## Using Oracle Cluster Ready Services

This chapter describes conceptual information as well as configuration procedures for Oracle Cluster Ready Services. It contains the following sections:

- [Section 13.1, "Introduction to Oracle Clusterware"](#)
- [Section 13.2, "Cluster Ready Services and Oracle Fusion Middleware"](#)
- [Section 13.3, "Installing and Configuring Oracle Clusterware with CRS"](#)
- [Section 13.4, "Using ASCRS to Manage Resources"](#)
- [Section 13.5, "Example Topologies"](#)
- [Section 13.6, "Troubleshooting Oracle CRS"](#)

### 13.1 Introduction to Oracle Clusterware

Oracle Clusterware manages the availability of user applications and Oracle databases in a clustered environment. In an Oracle Real Application Clusters (Oracle RAC) environment, Oracle Clusterware manages all of the Oracle database processes automatically. Anything managed by Oracle Clusterware is known as a cluster resource, which could be a database instance, a listener, a virtual IP (VIP) address, or an application process.

Oracle Clusterware was initially created to support Oracle RAC. Its benefits, however, are not limited to Oracle RAC; it can also be used to manage other applications through add-on modules (or action scripts). It is this flexibility and extensibility in Oracle Clusterware that forms the basis of a high availability solution for Oracle Fusion Middleware.

For more information about Oracle Clusterware, see the *Oracle Clusterware Administration and Deployment Guide*.

### 13.2 Cluster Ready Services and Oracle Fusion Middleware

Oracle Clusterware includes a high availability framework that offers protection to any resource with the help of resource-specific add-on modules. In Oracle Clusterware terminology, a resource refers to an object that Oracle Clusterware creates to identify the entity to be managed, such as an application, a virtual IP or a shared disk. Oracle Clusterware monitors resource availability by frequently checking its state and trying to restart it if it is down. If restarting fails, the resource starts on a new node, a process called failover. Resource switchover, an intentional switch of the operating environment of resources, is also allowed through the proper user interface.

With this high availability framework, Oracle Clusterware manages resources through user-provided add-on modules. For example, to create a resource for an application running in a single process, a module has to be supplied by the user to start and stop the process, and check its state. If this application fails, Oracle Clusterware attempts to restart it using this module. If the node on which this application is currently running fails, Oracle Clusterware attempts to restart it on another node if the application and resource are configured properly. You can configure the monitoring frequency for a resource and define its relationships to other resources.

Application Server Cluster Ready Services (ASCRS) relieves you from writing your own add-on modules for your critical and complex Oracle Fusion Middleware application environment, thereby giving you easy access to Oracle Clusterware's high availability features.

---

**Note:** ASCRS is a solution for managing an Oracle Fusion Middleware environment that is already CFC enabled. Refer to [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability"](#) for information on enabling CFC for the various Fusion Middleware components.

---

ASCRS consists of a frontend and a backend. The frontend is a command line interface, *ascrsctl*, with which you can perform administrative tasks such as resource creation, deletion, update, start, stop or switchover between cluster nodes. The backend is logic for the life cycle management of the various Fusion Middleware resources. The frontend and backend have their own separate log files.

On Unix platforms and Windows Server 2008, ASCRS supports virtual IPs, shared disks, database listeners, database instances, OPMN managed instances, and WebLogic servers. You can create Oracle Clusterware resources from these middleware components, allowing Oracle Clusterware and ASCRS to maintain their high availability within the cluster.

Oracle Clusterware and ASCRS provide a means to improve the survivability of the various resources when their hosting environment is corrupted or lost. However, they do not prevent disk corruption or application malfunctioning caused by disk corruption.

ASCRS supports Oracle Clusterware version 10.2.0.4, 11.1.0.7 and higher, or 11.2.0.2 and higher. Open online help using the following command:

```
ascrsctl help -c command -t resource_type
```

For example, the following command shows the help for creating a virtual IP resource:

```
ascrsctl help -c create -t vip
```

To view the full contents of ascrsctl online help, see [Appendix E, "ascrsctl Online Help."](#)

## 13.3 Installing and Configuring Oracle Clusterware with CRS

As an extension of CRS, ASCRS must be installed within each CRS home on each node of the cluster and configured separately before it can be used.

With the ASCRS command line tool *ascrsctl*, you can give ASCRS control of various middleware components. Once a component is controlled by CRS, its runtime state is closely monitored and CRS takes the proper actions if the component fails. With

ascsctl, you can create a CRS resource, and once a resource is created, you can perform start, stop, update, switch, status, and delete operations on it.

### 13.3.1 Upgrading Older Versions of ASCRS to the Current ASCRS Version

If an older version of ASCRS is already in use, following these steps to upgrade to the latest ASCRS version:

1. Use the ASCRS `stop` command to take all ASCRS resources offline.
2. Take note of all Fusion Middleware resource settings by viewing their status with ASCRS `status` command.
3. Delete these resources.
4. Remove the `CRS_HOME/ascrs` directory.
5. Install the new version of ASCRS as described in [Section 13.3.2, "Installing ASCRS"](#) and finish with the `setup` command.
6. Using the information from step 2 above, use the ASCRS `create` command to recreate the Fusion Middleware resources.

### 13.3.2 Installing ASCRS

Install ASCRS on each node of the cluster. To successfully install ASCRS on a particular node, check the following:

- Operating system: ASCRS is supported on Unix platforms, and Windows Server 2008. The system version and patch level should be compatible to the CRS version supported on that platform.
- CRS and version: CRS is installed on this system, started, and functioning correctly. The CRS version must be 10.2.0.4 or higher. For information about installing Oracle Clusterware and CRS, see the *Oracle Clusterware Installation Guide for Linux*.
- User account: The ASCRS installation user should be the same as the owner of the CRS home. It should also be the same as the owner of the application resource being managed. On Windows, this user must have administrator privilege.

---

**Note:** To install ASCRS for CRS 10.2.0.4, Sun JDK (or JRE) 1.5 or higher must be installed on the local system. It is needed by ascsctl, the command line tool.

---

To install ASCRS:

1. Log in as the CRS owner.
2. Insert the Oracle Fusion Middleware Companion CD and run the following commands to unzip and install the ascrs.zip file:

```
cd CRS_HOME
unzip Disk1/ascrs/ascrs.zip
cd ascrs/bin
setup
```

If the CRS version is 10.2.0.4, after JDK(JRE) 1.5+ or higher is installed, run the following command:

```
setup -j JDK/JRE_HOME
```

When installation is complete, the following ASCRS directory structure appears:

```
CRS_HOME/ascrs/bin
CRS_HOME/ascrs/config
CRS_HOME/ascrs/lib
CRS_HOME/ascrs/log
CRS_HOME/ascrs/public
CRS_HOME/ascrs/perl
CRS_HOME/ascrs/sql
CRS_HOME/ascrs/wlst
```

### 13.3.3 Configuring ASCRS with Oracle Fusion Middleware

After you install ASCRS, it is ready for use with the default configuration. To customize logging locations, logging levels, or the default CRS properties, edit the `config.xml` file located in the `CRS_HOME/ascrs/config` directory.

---

---

**Note:** ASCRS requires an available listen port in the `config.xml` file. Depending on the Weblogic server release, after converting the AdminServer to Cold Failover Cluster configuration, the `config.xml` file does not include `listen-port` directive explicitly when the AdminServer listen port is the default (7001). Verify that the file `DOMAIN_HOME/config/config.xml` explicitly specifies the `listen-port` directive before you configure ASCRS. For example:

```
<server>
 <name>AdminServer</name>

 <listen-port>7001</listen-port>
 <listen-port-enabled>true</listen-port-enabled>
 <listen-address>cfc.mycompany.com</listen-address>

</server>
```

---

---

The `config.xml` file contains the configuration for `ascrsctl` ASCRS agent logging. To change either of their locations, specify an existing path name or a path name within this CRS home using the `ORACLE_HOME` prefix. The available logging levels in the decreasing order of verbosity are ALL, FINEST, FINER, FINE, INFO, WARNING and SEVERE. Each resource has its own agent log file that rolls over after its size exceeds `rollover_size` bytes.

CRS properties are configured for policies. A policy name describes the characteristics of the CRS property values under that policy. A policy can be normal or fast. Policy 'fast' means more frequent resource health checking, and less delay in failover.

The following represents the default `config.xml` file that ships with ASCRS:

```
<?xml version="1.0" ?>
<config>
 <ascrsctl>
 <display level="normal"/>
 <log path="{ORACLE_HOME}/ascrs/log/ascrsctl.log" level="FINER"/>
 <resource-params target="vip" policy="normal">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="7"/>
 <param name="FAILOVER_DELAY" value="5"/>
 <param name="FAILURE_INTERVAL" value="50"/>
 </resource-params>
 </ascrsctl>
</config>
```

```

 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="2"/>
 <param name="SCRIPT_TIMEOUT" value="120"/>
 <param name="START_TIMEOUT" value="120"/>
 <param name="STOP_TIMEOUT" value="120"/>
</resource-params>

<resource-params target="vip" policy="fast">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="5"/>
 <param name="FAILOVER_DELAY" value="4"/>
 <param name="FAILURE_INTERVAL" value="30"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="2"/>
 <param name="SCRIPT_TIMEOUT" value="120"/>
 <param name="START_TIMEOUT" value="120"/>
 <param name="STOP_TIMEOUT" value="120"/>
</resource-params>

<resource-params target="disk" policy="normal">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="7"/>
 <param name="FAILOVER_DELAY" value="5"/>
 <param name="FAILURE_INTERVAL" value="50"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="2"/>
 <param name="SCRIPT_TIMEOUT" value="120"/>
 <param name="START_TIMEOUT" value="120"/>
 <param name="STOP_TIMEOUT" value="120"/>
</resource-params>

<resource-params target="disk" policy="fast">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="5"/>
 <param name="FAILOVER_DELAY" value="4"/>
 <param name="FAILURE_INTERVAL" value="30"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="2"/>
 <param name="SCRIPT_TIMEOUT" value="120"/>
 <param name="START_TIMEOUT" value="120"/>
 <param name="STOP_TIMEOUT" value="120"/>
</resource-params>

<resource-params target="dblsnr" policy="normal">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="50"/>
 <param name="FAILOVER_DELAY" value="20"/>
 <param name="FAILURE_INTERVAL" value="300"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="4"/>
 <param name="SCRIPT_TIMEOUT" value="120"/>
 <param name="START_TIMEOUT" value="120"/>
 <param name="STOP_TIMEOUT" value="120"/>
</resource-params>

<resource-params target="dblsnr" policy="fast">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="40"/>
 <param name="FAILOVER_DELAY" value="20"/>
 <param name="FAILURE_INTERVAL" value="250"/>

```

```

 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="4"/>
 <param name="SCRIPT_TIMEOUT" value="120"/>
 <param name="START_TIMEOUT" value="120"/>
 <param name="STOP_TIMEOUT" value="120"/>
 </resource-params>

 <resource-params target="db" policy="normal">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="120"/>
 <param name="FAILOVER_DELAY" value="30"/>
 <param name="FAILURE_INTERVAL" value="700"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="4"/>
 <param name="SCRIPT_TIMEOUT" value="300"/>
 <param name="START_TIMEOUT" value="300"/>
 <param name="STOP_TIMEOUT" value="300"/>
 </resource-params>

 <resource-params target="db" policy="fast">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="60"/>
 <param name="FAILOVER_DELAY" value="20"/>
 <param name="FAILURE_INTERVAL" value="400"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="4"/>
 <param name="SCRIPT_TIMEOUT" value="300"/>
 <param name="START_TIMEOUT" value="300"/>
 <param name="STOP_TIMEOUT" value="300"/>
 </resource-params>

 <resource-params target="as" policy="normal">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="50"/>
 <param name="FAILOVER_DELAY" value="20"/>
 <param name="FAILURE_INTERVAL" value="350"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="4"/>
 <param name="SCRIPT_TIMEOUT" value="600"/>
 <param name="START_TIMEOUT" value="600"/>
 <param name="STOP_TIMEOUT" value="600"/>
 </resource-params>

 <resource-params target="as" policy="fast">
 <param name="AUTO_START" value="1"/>
 <param name="CHECK_INTERVAL" value="40"/>
 <param name="FAILOVER_DELAY" value="10"/>
 <param name="FAILURE_INTERVAL" value="300"/>
 <param name="FAILURE_THRESHOLD" value="5"/>
 <param name="RESTART_ATTEMPTS" value="4"/>
 <param name="SCRIPT_TIMEOUT" value="600"/>
 <param name="START_TIMEOUT" value="600"/>
 <param name="STOP_TIMEOUT" value="600"/>
 </resource-params>
</ascrsctl>
<agent>
 <log path="{ORACLE_HOME}/ascrs/log" level="FINER" rollover_size="5242880"/>
</agent>
</config>

```

---

**Note:** The parameter values for a resource should fall within proper ranges when creating it. If a parameter is not configured for a policy, an internal default value is assumed.

The internal default parameter values are listed in [Table E-1](#).

---

Consult Oracle Clusterware documentation for the definitions of these parameters before editing their values.

Since computing environments vary in speed, Oracle recommends measuring the application's start and stop latency before setting the script, start, and stop timeout values. These values may be twice as much as the observed latencies.

## 13.4 Using ASCRS to Manage Resources

With the `ascrsctl` command line you manage CRS resources created for Fusion Middleware components. With this tool you can create, update, start, stop, switch and delete resources.

As mentioned in a previous section, a resource refers to an object that is created by CRS to identify the entity to be managed, such as an application, a virtual IP, or a shared disk. If the auto start for a resource is set to 1, CRS ensures this resource starts when CRS starts. Since Fusion Middleware resources depend on each other, start or stop of one resource may affect other resources. Resource dependency is enforced in the resource creation through `ascrsctl` syntax. At runtime, CRS uses this dependency knowledge for start/stop.

### 13.4.1 Creating CRS Managed Resources

On Unix and Windows 2008, ASCRS supports WebLogic Server, OPMN managed instance, Oracle database, Oracle database listener, virtual IP and shared disk. After a CRS managed resource is created for one of these components, it can be managed by CRS.

CRS resources created with the `ascrsctl` command line follow a naming convention. Follow this naming convention to ensure that the resources function correctly. To avoid unexpected errors, Oracle recommends using the CRS installation exclusively for Oracle Fusion Middleware, so that all the CRS managed resources are created with `ascrsctl`.

Under this naming convention, the canonical name for a resource has the following format:

`ora.name.cfctype`

Where *name* refers to the short name of the resource, for example, **sharedisk**, or **myvip**, and *type* refers to one of the resource types, such as **vip**, **disk**, **db**, **dblsnr** or **as**.

For example, on Linux, the following command creates a virtual IP resource named `ora.myvip.cfcvip` from the IP address 192.168.1.10 on network interface `eth0` with netmask 255.255.255.0:

```
ascrsctl create -name myvip -type vip -ipAddr 192.168.1.10 -netmask 255.255.255.0
-interface eth0
```

On Windows:

```
ascrsctl create -name myvip -type vip -ipAddr 192.168.1.10 -netmask 255.255.255.0
```

```
-interface "Public network"
```

---

**Note:** If a resource has a dependent resource, set the check interval for the dependent resource higher or at least equal to that of the resource on which it depends.

---

#### 13.4.1.1 Creating a Virtual IP Resource

The following information is required for creating a virtual IP resource:

- A valid virtual host name or IP address and it is not used by any host in the network
- A valid netmask number for this host name or IP
- One or more valid network interface names and they are present on all the cluster nodes where this IP is used. On Windows, the interface name is the same as the network connection name.
- On Windows, the interface name refers to the network connection name, such as *Public network*.

Since a virtual IP resource is a system resource, on Unix, the create or update command generates a script that must be executed by root user to complete the operation.

The following command creates a virtual IP resource named ora.myvip.cfcvip on Linux:

```
ascrsctl create -name myvip -type vip -ipAddr 192.168.1.10 -netmask 255.255.255.0
-interface eth0
```

#### 13.4.1.2 Creating a Shared Disk Resource

In a Fusion Middleware environment, shared disks are those disk storages that are used to hold Oracle database software, the database data files, WebLogic servers, OPMN managed components, and their Oracle homes. Shared disks allow the use of the same data when application resources are switched among the nodes within a cluster.

When creating a shared disk resource, carefully consider the following:

On Unix:

- Before creating a shared disk resource, create an empty signature file named `.ascrssf` on the root of the shared disk. The owner of the CRS home should own this file. This file is used by CRS after the resource is created.
- You can specify `nop` for either the mount or unmount command. You can use it for the mount command if the shared disk is never offline. If the disk does go off line for some reason, CRS will detect it and mark it as down. The `nop` command can be used for the unmount command if the disk does not need to be unmounted by CRS. In such a case, be absolutely sure that the disk does not need to be unmounted. There are potential disk corruption issues if the shared disk is mounted on two nodes without protection. Again, the signature file is always needed on the shared disk.
- The unmount command, may fail if there are active processes using the shared disk. To prevent this command failure, avoid accessing this disk from other applications while this disk resource is in online state.



- For complex mount and unmount commands, encapsulate the logic in executable scripts and specify the full path of these scripts as the mount and unmount commands. A proper unmount script is capable of killing other processes that are using this disk to ensure a successful and clean disk unmount. If the unmount command is in a script, do some basic file system checking, such as running a `fsck` command. Such a script should return 0 for success and 1 for failure.
- A shared disk resource is a system resource. Create, update, or delete commands generate scripts that must be executed as root to complete the create operation. Follow the instructions from the screen output.
- If the signature file is at the mount point of the shared disk, the start/stop operation may fail. Having the signature file on the mount point signals ASCRS that the disk is mounted, even if its not.
- Validate the mount/unmount command before using it in the `mc` or `umc` parameters or in the script file. There is no validation from ASCRS for the commands.
- If the shared disk is not protected by a cluster file system, it could be corrupted if it is mounted from multiple nodes. To avoid this, before creating the ASCRS resource, mount the disk only on the node where you create the resource.

On Windows Server 2008:

- Open Microsoft Disk Management and take note of the shared disk number. A disk number is a non-negative integer, such as 0, 2, or 5.
- Create an empty mount directory on the system drive on each cluster node, such as `c:\oracle\asdisk`.
- Ensure this disk is no longer used by any application on any node.
- From one node, in **Disk Management**, right click the drive and, online it, remove all partitions on it, create a single partition on this hard drive, and format it with NTFS. Remove any drive letter that may be assigned to it, and mount it to the directory you just created. Right click the drive again and offline it.
- On each other node, open Microsoft Disk Management, online this drive, remove the drive letter, if any, and mount it to the directory you just created. Right click the drive and offline it.
- Go to the node where you will create the disk resource, online the disk.
- This disk root should be accessible from the mount directory.
- Create an empty signature file named `.ascrssf` on the root of the shared disk. The CRS home owner should own this file. This file is used by CRS after the resource is created.
- The mount command is "diskmgr online disknumber" and unmount command is "diskmgr offline disknumber", where diskmgr is an ASCRS built-in command. There is no need to install any additional software to use the diskmgr command.

To create a shared disk resource, on Unix, run the following `ascrsctl` command that includes a valid mount point, a mount command, and an unmount command:

```
ascrsctl create -n sharedisk -type disk -path /asdisk -mc "/bin/mount
/dev/sda /asdisk" -umc "/bin/umount /asdisk"
```

To create a shared disk resource on Windows Server 2008, run an `ascrsctl` command similar to the following:

```
ascrsctl create -n sharedisk -type disk -path c:\oracle\asdisk -mc "diskmgr online
```

```
2" -umc
"diskmgr offline 2"
```

After a resource is created, start it explicitly on the node where the create command is executed to make sure the already mounted disk is under ASCRS control:

```
ascrsctl start -n ora.sharedisk.cfcdisk -node <disk resource creation node>
```

---

---

**Note:** On Windows Server 2008, mapped drives cannot be used as shared disk resources, and no general cluster file systems have been certified for this purpose.

---

---

### 13.4.1.3 Creating an Oracle Database Listener Resource

The following information is required for creating an Oracle listener resource:

- A valid Oracle database home
- The listener name
- The virtual IP resource name for the listen address
- The disk resource name for the Oracle home

Before creating the database listener resource, carefully check the following:

- The database listener home is installed on a shared disk. A CRS resource has been created for the shared disk with an ascrsctl command, and the resource is started.
- A CRS resource has been created for the virtual IP with an ascrsctl command, and the resource is started.
- The listener Cold Failover Cluster (CFC) enabled. See [Section 12.2.4, "Transforming an Oracle Database"](#) for details.
- Ensure that the listener name and Oracle home are valid database sid and database home. ASCRS does not do exhaustive validation on this type of information for this release.
- On Windows, ensure that the start method of the listener Windows service is manual.

Here is a syntax example for creating the resource:

```
ascrsctl create -n mydblsnr -type dblsnr -loh /cfcdb -ln LISTENER -disk ohdisk
-vip myvip
```

For online help information for creating an Oracle database listener resource, use the following command:

```
ascrsctl help -c create -t dblsnr
```

### 13.4.1.4 Creating an Oracle Database Resource

A database resource is a resource of any one of the following:

- Oracle database instance
- Oracle Database Console (dbconsole) process
- Oracle job scheduler process for the Windows platform
- Oracle Volume Shadow Copy Service for the Windows platform

Creating an Oracle database instance resource requires the following:

- A valid Oracle database home
- The database sid name
- Disk resource name for the Oracle home
- Disk resource name(s) if data files reside on different shared disk(s)
- The listener resource name

Before creating the Oracle database instance resource, carefully check the following:

- On Windows, ensure the built-in user system is in DBA\_GROUP and the start method of the corresponding Windows service is 'manual.'
- The database home is installed on a shared disk. The data files of this database are on the same or different shared disk(s). CRS resources have been created for all these shared disks with ascrsctl and started.
- A CRS resource has been created for the database listener with an ascrsctl command, and the resource is started.
- The database is CFC enabled. See [Section 12.2.4, "Transforming an Oracle Database"](#) for details.
- Ensure the database sid and Oracle home are valid. ASCRS does not do extensive validation of this information for this release.

The following is a syntax example for creating the database instance resource:

```
ascrsctl create -n mydb -type db -oh /cfcdb -lsnr mydblsnr -oraSID sid -disk
ohdisk datadisk
```

Creating all the other database resources requires the following:

- A valid Oracle database home.
- The database SID name.
- The disk resource name for the Oracle home.
- A valid virtual IP resource name. This is required for the Oracle Database Console (dbconsole) database resource only.
- The database is CFC enabled.

For online help information for creating an Oracle database resource, use the following command:

```
ascrsctl help -c create -t db
```

### 13.4.1.5 Creating a Middleware Resource

OPMN instances and WebLogic servers are collectively called Application Server (AS) components and are managed by separate resources. Specifically, all OPMN managed components have to be managed by one resource and all servers under a WebLogic domain have to be managed by a different resource.

**13.4.1.5.1 Creating a Resource for OPMN Managed Components** The following information is needed for creating a resource for an OPMN managed instance:

- A valid instance home for the OPMN managed components.
- A disk resource name for the instance home

- A disk resource name for the instance's Oracle home if it is on a different shared disk
- The names of the OPMN managed applications for inclusion in the resource. If you plan to include only a subset of all the components, the other remaining components won't be managed by CRS and they shouldn't be started outside CRS. By default, all the components are included.

Before creating the OPMN resource, carefully check the following:

- The Oracle home is installed on a shared disk. The OPMN instance is on the same or different shared disk. CRS resources have been created for all these shared disks with `ascrsctl` and started. Shutdown all OPMN managed applications.
- The OPMN server and its managed Oracle Fusion Middleware components have been CFC enabled. Refer to [Section 12.2.2.6, "Transforming Oracle Process Management and Notification Server"](#) through [Section 12.2.2.9, "Instance-specific considerations"](#) and to [Section 12.2.3, "Transforming Oracle Fusion Middleware Components"](#) for the steps to enable CFC for Oracle Fusion Middleware components.
- On Windows, ensure the start method of the OPMN Windows service pertaining to this instance is 'Disabled.'

The following is a syntax example for creating the resource (Oracle home and instance home are on the same disk. All components are included.):

```
ascrsctl create -n myopmn -type as -ch /cfcas -disk ohdisk
```

For online help information for creating an OPMN instance resource, use the following command:

```
ascrsctl help -c create -t as
```

**13.4.1.5.2 Creating a Resource for WebLogic Servers** Creating a CRS resource for a WebLogic domain requires more preparation than other resource types. Due to its complexity, the procedure is divided into the following sections:

- Basic Setup
- Node Manager Setup
- WebLogic Server Setup
- Creating the Resource

### Basic Setup

Before starting the basic setup, be sure that WebLogic is installed on shared disk(s). WebLogic Server software and the domain instance can be installed on either the same or separate shared disk.

In addition, before proceeding to the following Node Manager and Server Setup, ensure that the WebLogic Server environment is CFC enabled. See [Section 12.2.2.3, "Transforming the Administration Server for Cold Failover Cluster"](#) through [Section 12.2.2.5, "Transforming Node Manager"](#) for details on enabling CFC for the Oracle WebLogic Server environment. Once CFC is enabled, you can manually start and stop the server, the original node, and the failover node(s) without noticeable difference.

To create the dependency resources:

1. Create a CRS resource for each shared disk and start it on the node on which it was created.

2. Create a CRS resource for the virtual IP with the `ascrsctl` command and start it on the same cluster node.

### Node Manager Setup

To set up the Node Manager

1. For Windows Server 2008, on each node, create Node Manager Windows service if it does not already exist, by executing the following command from the `WL_HOME/server/bin` directory:

```
installNodeMgrSvc.cmd
```

From Windows Service Manager, make sure this service is in manual start mode.

2. If you have not yet done so, change Node Manager's username and password. The initial password is randomly generated. To change the Node Manager password, in the WebLogic Server Administration Console, select **Domain, Security, General**, and then **Advanced**. Enter the new password and click **Save**.
3. If you have changed anything in steps 1 or 2, restart the Node Manager.

On Unix, using the following command from the `WL_HOME/server/bin` directory:

```
startNodemanager.sh
```

On Windows, start Node Manager from the service manager.:

4. Start the WebLogic scripting tool in the `WL_HOME/common/bin` directory. To persist Node Manager's user login information in the `ascrscf.dat` and `ascrskf.dat` files, use the following commands:

```
nmConnect('nmUser','nmPasswd','hostname','nmPort','domainName','domainDir')
storeUserConfig('WL_HOME/common/nodemanager/ascrscf.dat',
 'WL_HOME/common/nodemanager/ascrskf.dat','true')
nmDisconnect()
exit()
```

5. For Unix platforms, copy `CRS_HOME/ascs/public/cfcStartNodeManager.sh` to the `WL_HOME/server/bin` directory, and make the script executable.

---

**Note:** To keep the setup consistently in sync, step 4 must be performed whenever the Node Manager passwords or usernames are changed.

---

After you have started Node Manager for the first time, you can edit the `nodemanager.properties` file to set the `StartScriptEnabled` property. The `nodemanager.properties` file does not exist until Node Manager is started for the first time.

In the `WL_HOME/common/nodemanager` directory, set the `StartScriptEnabled` property in the `nodemanager.properties` file to `true`.

```
StartScriptEnabled=true
```

Check the `nodemanager.properties` file to ensure no value is assigned to `ListenAddress`, and that a valid port number is assigned to `ListenPort`.

When this property is set in the `nodemanager.properties` file, you no longer need to define it in the `JAVA_OPTIONS` environment variable.

## WebLogic Server Setup

---

**Note:** The following steps apply to all WebLogic server setups that Oracle clusterware manages.

---

1. All WebLogic servers listen on the virtual IP. To ensure this is configured correctly, log in to the WebLogic Server Administration Console and navigate to **Environment > Servers > server\_name > Configuration > General** page and verify that the virtual IP and the port number are both set correctly and click **Save**.
2. Each server must also listen on the localhost. To ensure this is configured correctly, log into the WebLogic Server Administration Console and do the following:
  - a. In the Domain tree, select **Environment, Servers, server\_name, Protocols**, and then **Channels**.
  - b. In the Change Center, click **Lock & Edit**.
  - c. Click **New**, select a channel name, and protocol **t3**, and continue to the next screen.
  - d. Enter **localhost** for both the **Listen Address** and **External ListenAddress**.
  - e. Enter the port number to the **Listen Port** and **External ListenPort**. This port number must be exactly the same as the virtual IP port number.
  - f. Continue to the next screen and verify that **Enabled** is selected.
  - g. Click **Finish**.
  - h. Click **Activate Changes**.
3. If this is the Administration Server, ensure the `DOMAIN_HOME/servers/serverName/security` directory exists. This directory should contain the `boot.properties` file. If this file does not exist, create it and include the following properties:

```
username=<admin server user name>
password=<admin server user password>
```
4. If this domain does not have an Administration Server, ensure the `DOMAIN_HOME/servers/serverName/security` directory exists. This directory should contain the `boot.properties` file. If this file does not exist, create it and include the following properties:

```
username=<admin server user name>
password=<admin server user password>
```
5. If `DOMAIN_HOME/servers/serverName/data/nodemanager/startup.properties` exists, ensure the property `AutoRestart` defined in this file is set to `false`.
6. Repeat the previous steps for all servers, and then restart all servers in the Administration Server console.
7. Shutdown all WebLogic processes and stop the Node Manager.

---

**Note:** To keep the setup consistently in sync, Step 3 must be performed whenever the Administration Server password is changed.

---

### Creating the Resource

After Basic Setup, Node Manager Setup, and Server Setup, create the CRS resource using the following command:

```
ascrsctl create -n mywls -type as -ch /cfcas -disk sharedisk -vip myvip
```

---

---

**Note:** The WebLogic component home argument (-ch) must be valid. ASCRS does not perform extensive validation for this information.

---

---



---

---

**Note:** By default, the ASCRS agent checks WebLogic Server health by periodically initiating a TCP connection to the server. To alter this behavior, refer to [Section 13.4.8, "Configuring and Using Health Monitors."](#)

---

---

## 13.4.2 Updating Resources

You can update resources created with ascrsctl using the update command. Depending on the resource type, you can update the resource profile by specifying the appropriate parameter through the update command line. You can perform updates only when the resource is in the offline state.

For example, to update the virtual IP resource created in the last section with a new IP address and a different interface, use the following command:

```
ascrsctl update -n myvip -type vip -ip 192.168.1.20 -if eth1
```

---

---

**Note:** If you want to change the set of nodes hosting a particular resource, you must stop all dependent resources and then update the cluster nodes for each resource with the same node set and ordering. To find out related resources, run ascrsctl status command for this resource.

---

---

## 13.4.3 Starting Up Resources

When a resource is started, it is put under the control of CRS and its runtime status is monitored continuously by CRS. If the resource depends on other resources, starting this resource automatically starts the dependent resources. Refer to Oracle Clusterware documentation for information about the role of resource placement policy during resource start up. The ascrsctl start command maps to the CRS command.

For example, to start the virtual IP resource, use the following command:

```
ascrsctl start -n ora.myvip.cfcvip
```

---

---

**Note:** If a resource depends on more than one resource, while starting that resource, be sure that the resources, if online, are targeted on the same node.

---

---

### 13.4.4 Shutting Down Resources

When a resource is stopped, it is brought down and put in offline state and CRS stops monitoring its runtime status. If the resource depends on other resources that are in online status, the dependent resources are not stopped unless you confirm the prompt or (-np) option is specified. Refer to Oracle Clusterware documentation for more information about the implications of resource dependency during resource stop. The `ascrsctl stop` command maps to CRS command `crs_stop`.

For example, to stop the virtual IP resource, run the following command:

```
ascrsctl stop -n ora.myvip.cfcvip
```

### 13.4.5 Resource Switchover

Resource switchover is a process of shutting down the resource on the node on which it is running and restarting it on another node. The new node, if not specified, is determined by CRS, based on the placement policy. If the resource to be switched depends on other resources, or there are resources that are online and depend on it, this resource must be switched with -np flag.

To switch over a resource to another available node in the cluster, run the following command:

```
ascrsctl switch -n ora.myvip.cfcvip
```

### 13.4.6 Deleting Resources

You can delete a resource from CRS control. After a resource is deleted, the corresponding application or component's functionality is not affected, but CRS no longer monitors that resource. If a resource has dependent resources, it can not be removed.

To delete a resource from CRS control, run the following command:

```
ascrsctl delete -n ora.myvip.cfcvip
```

---

---

**Note:** If you delete a resource from CRS, the log directory and the log files for the deleted resource are NOT automatically removed. If you don't plan to reuse them in the future, you should delete them manually. The log files are located in the `ORA_CRS_HOME/ascrs/log` directory by default.

---

---

### 13.4.7 Checking Resource Status

You can check resource status with the `ascrsctl status` command. With this command, you can view the states of all resources and their dependents. If a particular resource is specified, the status command show its CRS profile, its direct and indirect dependency relationships, and its current state information.

For example, to check the status of a resource, run the following command:

```
ascrsctl status -n ora.myvip.cfcvip
```

Assuming the virtual IP resource is used by a database listener resource and the listener resource is in turn required by a database instance resource, all the



dependency information is shown in a tree structure, along with other status information in the following status output:

```

Basic information
-----+-----
Name | ora.myvip.cfcvip
Type | Virtual IP
Target state | ONLINE
Resource state | ONLINE on stjz11
Restart count | 0
Failure count | 0
Hosting members | stjz11, stjz12
-----+-----

Common CRS parameters
-----+-----
Auto start | Yes
Check interval | 7 sec
Failover delay | 5 sec
Failure interval | 50 sec
Failure threshold | 5
Restart attempts | 2
Script timeout | 30 sec
Start timeout | 30 sec
Stop timeout | 30 sec
-----+-----

Resource specific parameters
-----+-----
Interface(s) | eth2
Netmask | 255.255.252.0
Virtual IP address | 140.87.27.48
-----+-----

Resource dependency tree(s)
-----+-----
ora.mydb.cfcd
|
+-->ora.mydisk.cfcdisk
|
+-->ora.mylsnr.cfcdbsnr
|
| +-->ora.mydisk.cfcdisk
| |
| +-->ora.myvip.cfcvip
|
+-->ora.myvip.cfcvip

ora.myopmn.cfcd
|
+-->ora.mydisk.cfcdisk
|
+-->ora.myvip.cfcvip

ora.dbc.cfcd
|
+-->ora.mydisk.cfcdisk
|
+-->ora.myvip.cfcvip

ora.mywls.cfcd
|
+-->ora.mydisk.cfcdisk

```

```
|
+-->ora.myvip.cfcvip
```

### 13.4.8 Configuring and Using Health Monitors

Resource health state is the most important information in a CRS restart or failover decision. However, precisely determining the true state of a server is not an easy task. One situation is that the server process is still running, but its service is actually down. Thus, a simple TCP ping of a WebLogic server does not necessarily tell its true functional state. As an alternative, ASCRS provides some limited support for checking WebLogic Server functional state through user-defined monitors.

All monitors are defined in `ORA_CRS_HOME/ascrs/config/mconfig.xml`. Here is an example of this file:

```
<monitors>

 <!-- HTTP code monitor -->
 <monitor name="http_cm">
 <method type="ping" protocol="http" url="/index.html"/>
 <result type="code" code="200"/>
 </monitor>

 <!-- HTTP response content monitor that checks the response of a particular
URL -->
 <monitor name="http_rcm">
 <method type="ping" protocol="http" url="/index.html"/>
 <result type="exact_content" file="/crs/ascrs/public/index_content.txt"/>
 </monitor>

 <!-- HTTP response content monitor that checks the desired pattern in any line
of the returned content -->
 <monitor name="http_regex">
 <method type="ping" protocol="http" url="/cqi"/>
 <result type="regex_content"><![CDATA[.+Welcome.+]]></result>
 </monitor>

 <!-- WebLogic callout script monitor -->
 <monitor name="wls_sm">
 <method type="script" command="/var/scripts/wlsrv3_checker.sh"/>
 <result type="code" code="0"/>
 </monitor>
</monitors>
```

For each monitor, both "method" and "result" need to be defined. Method type can be either "ping" or "script." For a "ping" method, only the "http" protocol is supported and a valid URL has to be present. The result type is one of "code," "exact\_content," and "regex\_content." For a ping result with the "http" protocol, the code refers to the expected HTTP return code, while "exact\_content" and "regex\_content" specifies the exact return content and expected pattern in a line.

To use a monitor for a WebLogic Server, you only need to assign the monitor name to the server name with the `-m` option during its creation or update.

For example, assuming the WebLogic resource has two servers, AdminServer and wlsapp, you can customize their monitors with the command:

```
ascrsctl create -n mywls -type as -ch /cfcas -disk sharedisk -vip myvip -m
AdminServer=http_cm wlsapp=http_rcm
```

---

**Note:** When using the `exact_content` method, use the `CRS_HOME/ascrs/bin/mu` utility to generate the file for the expected response content, since the content saved from the browser is not usually exactly the same text as sent from the server.

---

## 13.5 Example Topologies

The following two examples illustrate how to use ASCRS to manage Fusion Middleware resources.

**Figure 13–1 CRS Topology Example 1**




Figure 13–1 illustrates the CRS Example 1 topology. In this example, Oracle HTTP Server and SOA are installed in a two-node cluster. Oracle HTTP Server is managed by OPMN. The SOA installation has a WebLogic server running that it hosts four Java EE applications.

Assumptions:

- **Operating Environment:** This is a Linux, two-node cluster with `node1.company.com` and `node2.company.com` as its members. Node 1 is designated as the primary node and node2 is the failover node. CRS has been installed on both nodes in the `/crshome` directory and is started. ASCRS has been installed on both nodes and has been configured.
- **One shared disk** has been allocated for both Oracle HTTP Server (Oracle home and instance home), WebLogic installation (server software and the domain home of the SOA server). It is a SCSI drive identified with `/dev/sda1` and has ext2 file system on it. It is mounted on `/sharedisk1`. Assume this shared disk is not used for other purposes.

- Both Oracle HTTP Server and WebLogic use virtual IP 192.168.1.10 for its public listen address. One each node, two network interface controllers, eth0 and eth1 are available for binding the virtual IP. The netmask is 255.255.255.0.
- Oracle HTTP Server Oracle home is /sharedisk1/ohsoh. The instance home is /sharedisk1/ohinst.
- WebLogic Server software is installed in the /sharedisk1/fmw directory, and the domain directory is /sharedisk1/fmw/user\_projects/domains/asdomain.

Under these assumptions, the following procedure describes the Cold Failover Clusters automation setup:

1. Installing WebLogic software and enabling Cold Failover Clusters:
  - a. If the shared disk /dev/sda1 is mounted on Node 2, unmount it. Mount the shared disk on /sharedisk1 on Node 1 if not yet mounted.
  - b. Create an empty signature file .ascrssf in /sharedisk1. Create this file only after this shared disk is mounted.
  - c. Bind the virtual IP to eth0 using /sbin/ifconfig.
  - d. Install SOA and OHS on the shared disk. Perform the CFC enabling procedure for both SOA and OHS using the virtual IP. After CFC enabling, shutdown all processes belonging to SOA server and OHS instance. To do a basic checking of the CFC enabling, unmount the shared disk on Node 1 and mount on Node 2 and try to start SOA and OHS on Node 2. If start fails, fix it before you proceed.
  - e. After step 3 is done, shutdown all OHS and SOA processes and unmount the disk on Node 2 and mount it on Node 1.
  - f. Follow the procedure in [Section 13.4.1.5.2, "Creating a Resource for WebLogic Servers,"](#) except for the "Creating the Resource" step, to configure the WebLogic Server in SOA install for ASCRS.
  - g. Shutdown any OHS and SOA processes.
  - h. Unbind the virtual IP.
2. Create a CRS resource.
  - a. On Node 1, cd to the /CRS\_HOME/ascrs/bin directory.
  - b. Create the virtual IP resource:

```
ascrsctl create -n asvip -t vip -if "eth0|eth1" -ip 192.168.1.10 -nm 255.255.255.0
```
  - c. Create the shared disk resource:

```
ascrsctl create -n asdisk -t disk -path /sharedisk1
-mc "/bin/mount /dev/sda1/sharedisk1"
-umc "/bin/umount /sharedisk1"
```
  - d. Create the SOA WebLogic server resource:

```
ascrsctl create -n soa -t as -vip asvip -disk asdisk
-ch /sharedisk1/fmw/user_projects/domains/asdomain
```
  - e. Create the Oracle HTTP Server resource with SOA as its dependency:

```
ascrsctl create -n ohs -t as -vip asvip -disk asdisk -as soa -ch /sharedisk1/ohsinst
```

- f. Start all resources on Node 1. Since the Oracle HTP Server resource depends on all other resources, starting this resource automatically starts other resources as well.

```
ascrsctl start -n ora.ohs.cfcas -node node1
```

**Figure 13–2 CRS Topology Example 2**




Figure 13–2 illustrates the CRS Example 2 topology. In this example topology, WebLogic Server and the Oracle database are installed on a two-node cluster with the following characteristics:

- The WebLogic Administration Server is the only WebLogic server running on this cluster. The WebLogic software and the domain home reside on the same shared disk.
- The Administration Server, along with Oracle Enterprise Manager, run in a WebLogic Java EE container with the first node as its primary node.
- The database software and its data files reside on two other shared disks with the second node as its primary node.

The goal of this topology is to provide a failover solution for both the WebLogic Administration Server and the database instance.

Assumptions:

- Operating Environment: This is a Linux two-node cluster with node1.company.com and node2.company.com as its members. Node 1 is designated as the primary node for WebLogic Server and Node 2 as its failover node. Node 2 is designated as the primary node for the Oracle database, and node 1 as its failover node. CRS is installed on both nodes in the /crshome directory and is started. ASCRS is installed on both nodes and is configured.
- Three shared disks are allocated. They are all SCSI drives identified with /dev/sda1, /dev/sda2, /dev/sda3 and have ext2 file system on them. /dev/sda1 is used for WebLogic software.
- Domain home: /dev/sda2 is used for Oracle database software. /dev/sda3 is used for the database data files. They are mounted on /sharedisk1, /sharedisk2, and /sharedisk3, respectively.

- WebLogic Server uses virtual IP 192.168.1.10 for its public listen address, and 7001 for its listen port. On each node, two network interface controllers, eth0 and eth1 are available for binding this virtual IP.

The database listener uses virtual IP 192.168.1.20 for its listen address, and 1521 for its listen port. On each node, network interface controller eth2 are used for binding this virtual IP.

The netmask is 255.255.255.0 for both virtual IPs.

- WebLogic Server is installed on the shared disk in the /sharedisk1/fmw directory, and the domain directory is /sharedisk1/fmw/user\_projects/domains/asdomain.
- The database is installed on the shared disk in the /sharedisk2/dbhome directory, and the data files are created in the /sharedisk3/dbdata directory. Assume orcl is the Oracle SID name and LISTENER is the listener name.

Under these assumptions, the following describes the procedure for automating Cold Failover Clusters:

1. Install WebLogic software and enable Cold Failover Clusters:
  - a. If the shared disk /dev/sda1 is mounted on Node 2, unmount it. Mount the shared disk on /sharedisk1 on Node 1 if it is not yet mounted.
  - b. Create an empty signature file .ascrssf in /sharedisk1. Create this file only after this shared disk is mounted.
  - c. Bind the virtual IP 192.168.1.10 to eth0 using /sbin/ifconfig.
  - d. Install WebLogic on the shared disk. Perform the Cold Failover Clusters enabling procedure for this installation using this virtual IP.
  - e. Follow the procedure in [Section 13.4.1.5.2, "Creating a Resource for WebLogic Servers"](#) to configure the WebLogic server for ASCRS.
  - f. Shutdown the WebLogic server.
  - g. Unbind the virtual IP.
2. Create the WebLogic-related resource.
  - a. Go to Node 1, cd to the /CRS\_HOME/ascs/bin directory.
  - b. Create the virtual IP resource:

```
ascrsctl create -n asvip -t vip -if "eth0|eth1" -ip 192.168.1.10 -nm 255.255.255.0
```
  - c. Create the shared disk resource:

```
ascrsctl create -n asdisk -t disk -path /sharedisk1
-mc "/bin/mount /dev/sda1 /sharedisk1" -umc "/bin/umount /sharedisk1"
```
  - d. Create the WebLogic server resource:

```
ascrsctl create -n adminserver -t as -vip asvip -disk asdisk
-ch /sharedisk1/fmw/user_projects/domains/asdomain
```
  - e. Start all WebLogic related resources on Node 1. Since this WebLogic resource depends on the disk and virtual IP resource, starting WebLogic resource automatically starts the other two as well.

```
ascrsctl start -n ora.adminserver.cfcas -node node1
```
3. Install the Oracle database software and enable Cold Failover Clusters.

- a. If the shared disk `/dev/sda2` is mounted on Node 1, unmount it. Mount the shared disk on `/sharedisk2` on Node 2 if it is not yet mounted.
  - b. Create an empty signature file `.ascrssf` in `/sharedisk2`. Create this file only after this shared disk is mounted.
  - c. If the shared disk `/dev/sda3` is mounted on Node 1, unmount it. Mount the shared disk on `/sharedisk3` on Node 2 if it is not yet mounted.
  - d. Create an empty signature file `.ascrssf` in `/sharedisk3`. Create this file only after this shared disk is mounted.
  - e. Bind the virtual IP 192.168.1.20 to `eth2` using `/sbin/ifconfig`.
  - f. Install Oracle database in directory `/sharedisk2/dbhome` and put the data files in `/sharedisk3/dbdata`.
  - g. Perform the CFC enabling procedure for this install for this virtual IP.
  - h. Shutdown the database.
  - i. Unbind the virtual IP.
4. Create the Oracle database-related resource.
    - a. In Node 2, cd to the `/CRS_HOME/ascrs/bin` directory.
    - b. Create the virtual IP resource:
 

```
ascrsctl create -n dbdisk -t disk -path /sharedisk2
 -mc "/bin/mount /dev/sda2
 /sharedisk2" -umc "/bin/umount /sharedisk2"
```
    - c. Create the shared disk resource for the database software:
 

```
ascrsctl create -n dbdisk -t disk -path /sharedisk2
 -mc "/bin/mount /dev/sda2
 /sharedisk2" -umc "/bin/umount /sharedisk2"
```
    - d. Create the shared disk resource for the database data files:
 

```
ascrsctl create -n dfdisk -t disk -path /sharedisk3
 -mc "/bin/mount /dev/sda3
 /sharedisk3" -umc "/bin/umount /sharedisk3"
```
    - e. Create the database instance resource:
 

```
ascrsctl create -n asdblsnr -t lsnr -vip dbvip
 -disk dbdisk -loh /sharedisk2/dbhome -ln LISTENER
```
    - f. Create the database listener resource:
 

```
ascrsctl create -n asdb -t db -lsnr asdblsnr -sid orcl
 -oh /sharedisk2/dbhome -disk dbdisk dfdisk
```
    - g. Start all the database related resources on Node 2. Since the database resource depends on all the other resources directly or indirectly, starting the database resource automatically starts the others as well.
 

```
ascrsctl start -n ora.asdb.cfcd -node node2
```

## 13.6 Troubleshooting Oracle CRS

This section includes troubleshooting information for Oracle CRS.

### 13.6.1 OPMN Resource Depends on a Virtual IP Resource

The OPMN resource depends on a virtual IP resource, but during create or update operations, ASCRS does not validate whether the virtual IP resource that the OPMN instance depends on actually exists. If an incorrect virtual IP resource is provided, the OPMN resource startup fails, and the OPMN instance may be left in an inconsistent state. To fix this problem, follow these steps:

1. Stop the OPMN resource using the `-f` option.
2. Ensure that all relevant processes are stopped. If not, stop all remaining stray processes manually.
3. Update the resource with the correct virtual IP resource.
4. Start the resource.

### 13.6.2 ASCRS Logging

ASCRS relies on logging for diagnosing unexpected issues. To get more diagnostic information, you can increase the log level verbosity by changing the ASCRS configuration file `config.xml`. See [Section 13.3.3, "Configuring ASCRS with Oracle Fusion Middleware"](#) for information on how to change the ASCRS logging level.

In addition, you can also check CRS daemon logs for basic CRS issues.



---

## Configuring High Availability for Oracle Portal, Forms, Reports, and Discoverer

This chapter describes high availability concepts and configuration procedures for Oracle Portals, Forms, Reports, and Discoverer. This chapter includes the following topics:

- [Section 14.1, "Overview of Oracle Portal, Forms, Reports, and Discoverer"](#)
- [Section 14.2, "Oracle Portal and High Availability Concepts"](#)
- [Section 14.3, "Oracle Reports and High Availability Concepts"](#)
- [Section 14.4, "Oracle Forms and High Availability Concepts"](#)
- [Section 14.5, "Oracle Discoverer and High Availability Concepts"](#)
- [Section 14.6, "Configuring Oracle Portal, Forms, Reports, and Discoverer for High Availability"](#)

### 14.1 Overview of Oracle Portal, Forms, Reports, and Discoverer

Oracle Portal offers a complete portal framework for building, deploying, and managing portals that are tightly integrated with Oracle Fusion Middleware. Oracle Portal provides a rich, declarative environment for creating a portal Web interface and accessing dynamic data with an extensible framework for Java EE-based enterprise application access.

With Oracle Portal, you can enhance your deployments with enterprise content management capabilities through Oracle Universal Content Management. In addition, you can add Enterprise 2.0 capabilities to existing portals using Oracle WebCenter Portal Services. Oracle Portal has certified interoperability with both of these complementary solutions, as well as with Oracle enterprise applications and other Oracle Fusion Middleware components.

When Oracle Portal is implemented, it often becomes the entry point into an organization internet or intranet, and as such, it is crucial that it is as highly available as possible.

Oracle Forms is Oracle's long established technology to design and build enterprise applications quickly and efficiently.

Oracle Reports Services is the reports publishing component of Oracle Fusion Middleware. It is an enterprise reporting service for producing high quality production reports that dynamically retrieve, format, and distribute any data, in any format, anywhere. You can use Oracle Reports Services to publish in both Web-based and non-Web-based environments.


Oracle Discoverer is a business intelligence tool for analyzing data. It is a key component of Oracle Fusion Middleware. Discoverer provides an integrated business intelligence solution that comprises intuitive ad-hoc query, reporting, analysis, and Web publishing functionality. These tools enable non-technical users to gain immediate access to information from data marts, data warehouses, multidimensional (OLAP) data sources, and online transaction processing systems. Discoverer integrates seamlessly with Oracle Portal and Oracle WebCenter Portal, enabling rapid deployment of Discoverer workbooks and worksheets to Web portals.

Oracle Portal, Forms, Reports and Discoverer can be installed individually or collectively.

### 14.1.1 Oracle Portal, Forms, Reports, and Discoverer Architecture

Figure 14-1 Illustrates a single-instance Oracle Portal, Forms, Reports, and Discoverer deployment.

**Figure 14-1 Oracle Portal, Forms, Reports, and Discoverer Architecture**



#### Oracle Portal, Forms, Reports, and Discoverer Common Components

Figure 14-1 includes the following common components:

- **Oracle Web Cache** performs two functions. Its primary function is to serve static Web content from its cache, much faster than could be achieved by the Oracle HTTP Server alone. If Oracle Web Cache does not have a cacheable page in its cache, or that page is not current, it requests the page from the attached Oracle HTTP Servers.

The second function of Oracle Web Cache is used in high availability environments. It receives a request, and if it cannot service that request from its own cache, it can load balance it between several Oracle HTTP Servers.

Oracle Web Cache is optional, but can be used in conjunction with Oracle Forms, Reports and Discoverer.

- **Oracle HTTP Server** is responsible for assembling requested pages. Page assembly is not always straightforward. Depending on how the page is made up, the Oracle HTTP Server performs one of the following:
  - If the page is a simple HTML document, then the Web tier finds and returns the document.
  - If the Web page needs to be assembled by executing a Java EE application, the Oracle Web Tier routes the request to Oracle WebLogic Server, which, after processing the request, sends the result back to the user through the Oracle Web Tier.
  - If the Web page needs to be assembled by executing some other application such as PLSQL or CGI, the Oracle Web Tier routes the request to the appropriate application, and once that application has processed the request, it sends the result back to the user through the Oracle Web Tier.
  - If the requested page is security controlled, the Oracle Web Server invokes Oracle Identity Management to ensure that the user is authorized to view the page.

The Oracle HTTP server can be used as stand-alone or in conjunction with Oracle Web Cache.

In Oracle Portal, Forms, Reports and Discoverer deployments, the Oracle HTTP Server uses an Apache module called `mod_wl_ohs` to route requests to the Oracle WebLogic Managed Servers. When WebLogic Managed Servers are clustered together, `mod_wl_ohs` load balances requests among all Managed Servers in a given cluster.

- **Oracle WebLogic Managed Servers** are the Java EE runtime containers for Oracle Portal, Forms, Reports and Discoverer Applications.
- **Oracle Process Manager and Notification Server (OPMN)** is used to start, stop, and monitor Oracle Web Cache and Oracle HTTP Server. It periodically polls Oracle Web Cache and the Oracle HTTP Server to ensure that they are functioning. If they are not functioning, OPMN takes appropriate action to restart the failed component to ensure that service is maintained.
- **Oracle Node Manager** is used to start, stop, and monitor Oracle WebLogic Managed Servers including the Administration Server. It periodically polls the WebLogic Managed Servers to ensure that they are functioning. If they are not functioning, the Oracle Node Manager takes appropriate action to restart the failed component to ensure that service is maintained.
- The **Oracle WebLogic Administration Server** contains both the WebLogic Administration Console and the Oracle Fusion Middleware Enterprise Manager. It is active only once within a WebLogic domain. If the server hosting the Administration Server fails, it must be manually restarted elsewhere.

- **Oracle Single Sign-On** is Oracle's Enterprise authentication mechanism. Oracle Single Sign-On is integrated into Oracle HTTP Server for each of the product components. When a request which requires authentication comes in to the Oracle HTTP Server, it determines whether the user has been authenticated through the Oracle Single Sign-On server. If the user has been authenticated, the request is processed. If however, the user has not been authenticated, the Oracle Single Sign-On server is contacted to gain authorization.

## 14.1.2 Common Log Files

The following table lists, describes, and provides the location for generic log files used by Oracle Portal, Forms, Reports, and Discoverer:

| File       | Location                                                              | Description                                                          |
|------------|-----------------------------------------------------------------------|----------------------------------------------------------------------|
| access_log | <i>ORACLE_</i><br><i>INSTANCE</i> /diagnostics/logs/O<br>HS/ohs1      | Lists each access to the Oracle HTTP Server and the HTTP Return code |
| ohs1.log   | <i>ORACLE_</i><br><i>INSTANCE</i> /diagnostics/logs/O<br>HS/ohs1      | HTTP Server error log                                                |
| access_log | <i>ORACLE_</i><br><i>INSTANCE</i> /diagnostics/logs/We<br>bCache/web1 | Lists each access to Oracle WebCache and the HTTP Return code        |
| access_log | <i>DOMAIN_HOME</i> /servers/WLS_<br>PORTAL/logs                       | Lists access requests being received by the WebLogic Managed Server  |
| opmn.log   | <i>ORACLE_</i><br><i>INSTANCE</i> /diagnostics/logs/OP<br>MN          | OPMN log files                                                       |

## 14.1.3 Common Component Failures and Expected Behaviors

This section describes high availability concepts that apply to Oracle Portal, Forms, Reports, and Discoverer.

### 14.1.3.1 Oracle Web Cache and Oracle HTTP Server Process Failures

Oracle HTTP Server, Oracle Web Cache, and Oracle Discoverer preference server processes are protected by the Oracle Process Manager and Notification system (OPMN). If one of these processes fails, OPMN automatically restarts the process.

Oracle WebLogic Managed Servers are started and monitored by Oracle Node Manager. If an Oracle WebLogic Managed Server fails, Oracle Node Manager restarts the process.

### 14.1.3.2 Common Component Node Failures

If a node that is not fronted by Oracle Web Cache fails, the Load balancer automatically reroutes requests to a surviving node.

If Oracle Web Cache fails, the Load Balancer automatically reroutes requests to a surviving web cache node.

If a node with Oracle HTTP Server fronted by Oracle Web Cache fails, Oracle Web Cache reroutes the request to a surviving Oracle HTTP Server.

If Oracle WebLogic Server fails, Oracle HTTP Server reroutes requests to another WebLogic cluster member.

Oracle Portal, Forms, Reports and Discoverer requests being processed by the failed node must be restarted.

### 14.1.3.3 Common Component WebLogic Managed Server Failures

In a high availability configuration, Oracle WebLogic Managed Servers are clustered together. If one of the managed servers fails, `mod_wl_ohs` automatically redirects requests to one of the surviving cluster members. If the application stores state, state replication is enabled within the cluster, which allows redirected requests access to the same state information.

### 14.1.3.4 Common Component Database Failures

Databases are recommended to be implemented using high availability technologies such as Oracle Real Application Clusters. If one of the database nodes fails, the database as a whole remains available. In some cases you may have to resubmit the request.

In a multi database node environment, if a user session is connected to the database node that fails, the following occurs:

- Oracle Portal: The user is required to resubmit the request.
- Oracle Forms: The user is required to resubmit the request.
- Oracle Reports: If the database connect string is configured using Oracle Transparent Application Failover, no action is required (unless the report writes to a database during its execution).

If the database containing the reports queue loses a node, then the user is required to resubmit the report request.

- Oracle Discoverer: The user is required to resubmit the request.

## 14.1.4 Oracle Portal, Forms, Reports, and Discoverer Cluster-Wide Configuration Changes

Oracle Web Cache instances are clustered. Once a configuration change is made through the Oracle Fusion Middleware Console or through the Oracle Web Cache Administration utility, these changes are propagated to other cluster members. Propagation is done manually using these tools.

Oracle HTTP Servers are not clustered. The Oracle HTTP server configuration is file-based. As a result, changes made to one Oracle HTTP Server must be manually copied to other Oracle HTTP Servers in the configuration. This also applies to static HTML files stored in the `htdocs` directory.

Configure Oracle Portal, Forms, Reports and Discoverer using a series of configuration files. Any changes to these files must be manually applied to all members in the architecture.

WebLogic Managed Servers are clustered and share resources at the cluster level. Changes to these resources can be made once without the need for propagation. These resources include:

- Data sources
- Application redeployments

- State replication

### 14.1.5 Common Component Log File Information

Cluster wide log consolidation is not offered for Oracle Web Cache, Oracle HTTP Server, OPMN, and WebLogic Managed Servers. For information about the status of an Oracle HTTP Server application, refer to the log files on each Oracle HTTP Server node. To For information about the status of an failed application, refer to the log files on each Server node.

## 14.2 Oracle Portal and High Availability Concepts

This section describes single-instance information, as well as high availability concepts specific to Oracle Portal. This section guides you through the concepts and considerations necessary for creating a successful high availability Oracle Portal deployment.

### 14.2.1 Oracle Portal Single-Instance Characteristics

For information about the single-instance architecture of Oracle Portal, see the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Portal*.

- Understanding the Oracle Portal Components - this section introduces the components of the Oracle Fusion Middleware. It describes how these components work with Oracle Portal.
- Understanding the Oracle Portal Architecture - this section describes the Portal architecture. Read the following topics in this section:
  - How Does Oracle Portal Integrate with Other Components?
  - How Are Pages Assembled in Oracle Portal?
  - How Does Communication Flow in Oracle Portal?
- Understanding Caching in Oracle Portal - this section describes the caching configurations you can implement to increase the availability and scalability of medium to large deployments.
- Understanding WSRP and JPS - this section provides an introduction to the Web Services for Remote Portlets (WSRP) specifications and Java Portlet Specification (JPS). These two standards enable the development of portlets that interoperate with different Portal products, thereby increasing the availability of portlets within an organization.

#### 14.2.1.1 Oracle Portal Request Flow

For information about request flow in Oracle Portal, see the following sections in the *Oracle Fusion Middleware Administrator's Guide for Portal Guide*:

- How Are Pages Assembled in Oracle Portal?
- How Does Communication Flow in Oracle Portal?

#### 14.2.1.2 Oracle Portal Component Characteristics

The following lists the characteristics of Oracle Portal components:

- Oracle Portal application runs inside a WebLogic container (WLS\_PORTAL).

- Oracle Portal repository stores metadata, documents, customizations, and personalizations.
- Oracle Portal has a strong dependency on Oracle Web Cache to process and cache content.
- Oracle Portal depends on Oracle HTTP server to route traffic to WebLogic Server, service product images, and rewrite URLs.

#### 14.2.1.3 Oracle Portal Startup and Shutdown of Processes and Lifecycle

Oracle Portal does not have any process of its own. It relies on standard Oracle tools and utilities to manage Oracle Web Cache, Oracle HTTP Server, WebLogic Managed Server WLS\_PORTAL, and the database.

The following lists the configuration and monitoring interfaces for Oracle Portal components:

- Enterprise Manager for managing Oracle Web Cache, Oracle HTTP Server, Oracle Portal, and WLS\_PORTAL
- Oracle WebLogic Administration Console for managing WLS\_PORTAL
- Oracle WebLogic Scripting Tool (WLST) for adding or editing Portal-specific configuration parameters
- The `opmnctl` command line interface for managing Oracle Web Cache and Oracle HTTP Server

#### 14.2.1.4 Oracle Portal Deployment Artifacts

The Portal application is deployed as a staged application (`portal.ear`) to WebLogic Server (WLS\_PORTAL). The configuration files are available in `DOMAIN_HOME/servers/WLS_PORTAL/stage/portal/portal/configuration`

#### 14.2.1.5 Oracle Portal Configuration Information

Some of the Portal configuration is stored in the Portal repository. This information includes site details (site host and port), Web Cache details (webcache host, invalidation port, and invalidation password), and Oracle Internet Directory details (Oracle Internet Directory host and port). Such information can be changed by accessing Enterprise Manager or by using WLST commands for any of the containers.

The remaining Portal configuration is available in configuration files that are located inside the staged location of the Portal application. Such configuration files are located in `DOMAIN_HOME/config/fmwconfig/servers/WLS_PORTAL/applications/portal/configuration/`. In the HA environment, if a configuration file is modified, the same change must be repeated on each WLS\_PORTAL instance either by repeating the action on each node or by copying over the configuration files from one node to another.

**Table 14–1 Portal Configuration Files**

| Configuration File         | Description                      |
|----------------------------|----------------------------------|
| <code>appConfig.xml</code> | Portal Page Engine configuration |

**Table 14–1 (Cont.) Portal Configuration Files**

| Configuration File             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>portal_dads.conf</code>  | Portal Database Access Descriptor (DAD) configuration<br><br>This file contains properties related to High Availability (HA). It contains connect string information to the database. If Oracle RAC nodes are added or removed, this file must be updated to reflect the final set of nodes. The file contains a configuration property, which can be enabled to test a pooled database connection before using it. At a minimal cost, this property can be enabled to ensure better results when a database node goes down. |
| <code>portal_cache.conf</code> | Portal Cache configuration                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <code>portal_plsql.conf</code> | Portal global settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

#### 14.2.1.6 Oracle Portal Logging and Log Configuration

Portal log files are generated in the `DOMAIN_HOME/servers/WLS_PORTAL/logs` directory. The log file is named `WLS_PORTAL-diagnostic.log`. Log rotation ensures that older logs are archived with a similar name. The configuration of log files in Portal is controlled by the `logging.xml` file, which is located in `DOMAIN_HOME/config/fmwconfig/servers/WLS_PORTAL`.

**14.2.1.6.1 Oracle Portal Log Files** The following table lists and describes log files used by Oracle Portal:

| File                        | Location                                         | Description                                                                                                           |
|-----------------------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <code>WLS_PORTAL.log</code> | <code>DOMAIN_HOME/servers/WLS_PORTAL/logs</code> | Log file for the WebLogic Managed Server                                                                              |
| <code>WLS_PORTAL.out</code> | <code>DOMAIN_HOME/servers/WLS_PORTAL/logs</code> | Supplemental log file for the WebLogic Managed Server. This is generally the first place to look for WebLogic issues. |

#### 14.2.1.7 Oracle Portal External Dependencies


Oracle Portal requires an Oracle HTTP Server to service requests. Optionally these requests may be cached by Oracle Web Cache, in which case it also acts in the capacity of a load balancer.

Oracle Portal requires a database to store information. This database is pre-seeded with schemas using the Oracle Repository Creation Assistant.

### 14.2.2 Oracle Portal Protection from Failures and Expected Behavior

Figure 14–2 illustrates a high availability configuration for Oracle Portal.



**Figure 14–2 Oracle Portal High Availability Deployment**

The Oracle Portal high availability setup includes two or more Oracle Web Cache instances. These Oracle Web Cache instances are clustered together to provide cache consistency and failover. Each Oracle Web Cache is configured to route traffic to two or more Oracle HTTP servers, and these servers load balance incoming requests.

Each Oracle HTTP Server is configured to route Oracle Portal requests to two or more Oracle WebLogic managed servers, which host the Oracle Portal application. Oracle HTTP Server is a web server provided by Oracle Fusion Middleware. It incorporates an OpenSSL module to support Secure Sockets Layer (SSL) and HTTP Secure Sockets Layer (HTTPS). Oracle HTTP Server also provides a mod connector to route traffic to Oracle WebLogic Server, which runs Java servlets and J2EE applications. Oracle Portal relies on Oracle HTTP Server to service product images and rewrite URLs rewriting, in some scenarios.

WLS\_PORTAL is the WebLogic Server that runs the Oracle Portal application. WebLogic Server provides a complete Java EE environment that includes a JSP translator, a JSP servlet engine (OJSP), and an Enterprise JavaBeans (EJB) container. It provides a fast, lightweight, highly scalable, easy-to-use, complete Java EE environment. It is written entirely in Java and executes on the standard Java Development Kit (JDK) Virtual Machine (JVM).

The Portal application running inside WLS\_PORTAL works closely with Oracle Web Cache, using its Edge Side Includes (ESI) processing, to service dynamic Portal pages in a secure fashion. Oracle Portal uses ESI to securely cache various pieces of metadata and content in Oracle Web Cache. As content changes, Oracle Portal invalidates any cached copies in Oracle Web Cache. Such content is regenerated in a subsequent request. ESI allows content to be cached at a more granular level, thereby increasing the cache hit ratio and enabling a more granular invalidation of Portal content. For most content/metadata, Oracle Portal also backs up the in-memory content in Web Cache into the Portal Cache, which is based on file system.

Oracle Portal stores its content and metadata information in a database. To ensure that the database is not a single point of failure, the database is built using Oracle Real Application Clusters.

Finally, a load balancer front ends the Oracle Web Cache instances to provide a single virtual access point to Oracle Portal. Besides load balancing external traffic coming from the browser, the load balancer also provides load balancing for internal traffic generated by Oracle Portal (referred to as Portal Loopback Requests). The load balancer also provides load balancing of invalidation messages generated from the Portal metadata repository. Such invalidation requests are sent to Oracle Web Cache to invalidate any stale content cached in Oracle Web Cache. The clustering feature in Oracle Web Cache ensures cache consistency across the Oracle Web Cache instances.

---

**Note:** If you are deploying your own producers, ensure that there is redundancy for the Producers and the producers are front-ended by a load balancer. This step ensures that there is no single-point of failure in the system.

---

#### 14.2.2.1 Oracle Portal Process Failures

Due to redundancy in each component, if a particular process goes down, Oracle Portal continues to work. Note that Oracle Portal does not attempt to retry a request. Therefore, if a failure occurs while processing a particular request, that request fails.

#### 14.2.2.2 Oracle Portal Node Failures

For information about Oracle Portal Node failure, see [Section 14.1.3.2, "Common Component Node Failures."](#)

#### 14.2.2.3 Oracle Portal WebLogic Managed Server Failures

For information about Oracle Portal WebLogic Managed Server failure, see [Section 14.1.3.3, "Common Component WebLogic Managed Server Failures."](#)

#### 14.2.2.4 Oracle Portal Protection from Database Failures

For information about Oracle Portal database failure, see [Section 14.1.3.4, "Common Component Database Failures."](#)

## 14.3 Oracle Reports and High Availability Concepts

This section describes single-instance information, as well as high availability concepts specific to Oracle Reports. This section guides you through the concepts and considerations necessary for creating a successful high availability Oracle Reports deployment.

---

**Note:** Active-active configurations support Web clients only. The command-line client, `rwclient`, is not supported. `rwclient` connects to a specific Reports Server by name and there can be only one named instance of a specific Reports Server.

---

### 14.3.1 Oracle Reports Single-Instance Characteristics

Oracle Reports Services is the reports publishing component for Oracle Fusion Middleware. It is an enterprise reporting service for producing high quality production reports that dynamically retrieve, format, and distribute any data, in any format, anywhere. Oracle Reports Services can be used to publish in both Web-based and non-Web-based environments.

The following components are specific to Oracle Reports:

**Oracle Reports Server**

The Reports Server (rwserver) processes client requests, including ushering them through its various services, such as authentication and authorization checking, scheduling, caching, and distribution (including distribution to custom-or pluggable-output destinations). The Reports Server also spawns runtime engines for generating requested reports, fetches completed reports from the Reports Server cache, and notifies the client that the job is ready.

The reports server can be run stand-alone or in-process.

**Oracle Reports Engine**

The Reports Engine includes components for running SQL-based and pluggable data source-based reports, fetches requested data from the data source, formats the report, sends the output to cache, and notifies the Reports Server that the job is complete.

**Oracle Reports Cache**

The Reports Cache is used to store the output of reports, which have been successfully run. By using the reports cache it is not necessary to generate the same report repeatedly.

**Oracle Reports Queue**

The Reports Queue contains a list of the report requests. When processing capacity becomes available, the next report in the queue is executed.

**Oracle Reports Customer Databases**

Typically, reports are compiled using information stored in a variety of databases. These are referred to as customer databases.

**Oracle Single Sign On**

Although advisable it is not a mandatory requirement to restrict access to Oracle Reports using Oracle Single Sign-On (SSO).

**14.3.1.1 Oracle Reports State Information**

Oracle Reports only state information is job metadata. For example, which reports are to be run, and the parameters with which they are run. This is often referred to as the reports queue. This information is stored in operating system files `servername.dat`, or alternatively in a database.

**14.3.1.2 Oracle Reports External Dependencies**

Oracle Reports requires Oracle HTTP Server to service requests. Optionally these requests may be cached by Oracle Web Cache, in which case it also acts in the capacity of a load balancer.

**14.3.1.3 Oracle Reports Specific Configuration Files**

The following table lists and locates configuration files used by Oracle Reports:

| File             | APPHOST1 Location                                                                                                     | APPHOST2 Location                                                                                                       |
|------------------|-----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| rwserver.conf    | <i>DOMAIN_<br/>HOME/config/fmwconfig/servers/WLS_<br/>REPORTS/applications/reports_<br/>11.1.1.2.0/configuration/</i> | <i>DOMAIN_<br/>HOME/config/fmwconfig/servers/WLS_<br/>REPORTS1/applications/reports_<br/>_11.1.1.2.0/configuration/</i> |
| reports_ohs.conf | <i>ORACLE_<br/>INSTANCE/config/OHS/ohs1/<br/>moduleconf</i>                                                           | <i>ORACLE_<br/>INSTANCE/config/OHS/ohs1/<br/>moduleconf</i>                                                             |

#### 14.3.1.4 Oracle Reports Connection Retry

This section describes the following connection retries:

- Oracle Portal Database Connection Retry:  
If the connection from the Reports Server to the Oracle Portal database schema is dropped, the Reports Server tries to reestablish the connection before generating an error. If reconnection is successful, there is no need to restart the Reports Server.
- Oracle Internet Directory Connection Retry:  
If the Oracle Internet Directory connection becomes stale, the Reports servlet and the Reports Server try to reestablish the connection before generating errors. If reconnection is successful, there is no need to restart the Reports Server.
- Oracle Metadata Repository and Oracle Identity Management Outage:  
The outage of the Oracle Metadata Repository (which stores security metadata) does not bring down the Reports Server. If the Oracle Metadata Repository is unavailable, the Reports Server rejects new requests as a result of the component being unavailable. When the Oracle Metadata Repository is brought back on-line, the Reports Server recovers itself and begins to receive and process new requests.  
  
If Oracle Identity Management components become unavailable, the Reports Server also reject new requests, much like the outage of the Oracle Metadata Repository.
- Reports Server Timeout:  
The Reports Server has a configurable timeout for waiting for requests to be returned from the database. This timeout must be set to a high enough value to allow valid reports to run but not so high as to cause excessively long waits.

#### 14.3.1.5 Oracle Reports Process Flow

The various components of Oracle Reports Services contribute to the process of running a report as follows:

1. The client requests a report by contacting a server through a URL (Web)
2. The Reports Server processes the request as follows:  
  
If the request includes a TOLERANCE option, the Reports Server checks its cache to determine whether it already has output that satisfies the request. If it finds acceptable output in its cache, then it immediately returns that output rather than rerunning the report.  
  
If the request is the same as a currently running job, it reuses the output from the current job rather than rerunning the report.  
  
If neither of these conditions is met:

- If Oracle Reports Servlet (rwservlet) is SSO-enabled, it checks for authentication. A secure Reports Server then authorizes the user using Oracle Internet Directory. If Oracle Reports Servlet (rwservlet) is not SSO-enabled, a secure Reports Server authorizes and authenticates the user.
  - If the report is scheduled, the Reports Server stores the request in the scheduled job queue, and the report is run according to schedule. If the report is not scheduled, it is queued in the current job queue for execution when a Reports Engine becomes available.
3. At runtime, the Reports Server spawns a Reports Engine and sends the request to that engine to be run.
  4. The Reports Engine retrieves and formats the data.
  5. The Reports Engine populates the Reports Server cache.
  6. The Reports Engine notifies the Reports Server that the report is ready.
  7. The Reports Server accesses the cache and sends the report to output according to the runtime parameters specified in either the URL, the command line, or the keyword section in the `cgicmd.dat` file (URL requests only).

If a report server dies while running a request the request must be resubmitted. Once resubmitted, one of the surviving reports servers processes the request.


#### 14.3.1.6 Oracle Reports Log Files

The following table shows log files used by Oracle Reports:

| File                     | Location                                     | Description                                                                                                           |
|--------------------------|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| WLS_REPORTS.log          | <i>DOMAIN_HOME</i> /servers/WLS_REPORTS/logs | Log file for the WebLogic Managed Server                                                                              |
| WLS_REPORTS.out          | <i>DOMAIN_HOME</i> /servers/WLS_REPORTS/logs | Supplemental log file for the WebLogic Managed Server. This is generally the first place to look for WebLogic issues. |
| rwservlet_diagnostic.log | <i>DOMAIN_HOME</i> /servers/WLS_REPORTS/logs | Log information relating to the Reports servlet.                                                                      |
| rwserver_diagnostic.log  | <i>DOMAIN_HOME</i> /servers/WLS_REPORTS/logs | Log information relating to the Reports server                                                                        |

### 14.3.2 Oracle Reports Protection from Failure and Expected Behavior

Figure 14–3 illustrates an example Oracle Reports high availability deployment.

**Figure 14–3 Oracle Reports High Availability Deployment**

As shown in [Figure 14–3](#), the Oracle Real Application Clusters database provides a high availability repository for the reports queue. In high availability configurations, each Reports server has access to the same reports queue. A shared reports queue ensures that any request is only processed once, but it can be processed by any Reports server in the deployment.

Typically, reports are compiled using information stored in a variety of databases, the diagram above refers to these databases as customer databases.

In order to make full use of the Reports cache, the cache should be available to any of the Reports servers in the implementation. This way reports generated by one Reports server can be reused or served by any Reports server. In [Figure 14–3](#), the Reports cache is made available using a shared directory.

Single Sign-On (SSO) is not required to restrict access to Oracle Reports, however, many organizations do restrict access to Oracle Reports using Oracle Single Sign-on. This chapter includes the steps required to protect the deployment using Oracle Single Sign On.

The in-process Reports server is recommended for high availability reports deployments. The in-process reports server runs inside Oracle WebLogic Server and the Reports servers themselves can be clustered to ensure that high availability is maintained.

When using Web Cache in a highly available distributed configuration, it is important that contents of all of the distributed caches are consistent. For example, the same cached copy of a given report is available from any Oracle Web Cache instance. When cached content becomes invalid, it is essential that it becomes invalid in all of the web caches. To achieve this goal web cache clustering is used.

The Diagram above includes the Web Logic Administration Server and shows how it can be positioned on APPHOST2. To determine how to do this refer to the appropriate chapter elsewhere in this guide.

#### **14.3.2.1 Oracle Reports Process Failures**

If the Reports Server hangs or fails to respond, but the WebLogic Managed Server does not, the WebLogic Managed Server must be restarted manually.

#### **14.3.2.2 Oracle Reports Node Failures**

For information about Oracle Reports Node failure, see [Section 14.1.3.2, "Common Component Node Failures."](#)

#### **14.3.2.3 Oracle Reports WebLogic Managed Server Failures**

For information about Oracle Reports WebLogic Managed Server failure, see [Section 14.1.3.3, "Common Component WebLogic Managed Server Failures."](#)

#### **14.3.2.4 Oracle Reports Database Failures**

For information about Oracle Reports database failure, see [Section 14.1.3.4, "Common Component Database Failures."](#)

## **14.4 Oracle Forms and High Availability Concepts**

This section describes single-instance information, as well as high availability concepts specific to Oracle Forms. This section guides you through the concepts and considerations necessary for creating a successful high availability Oracle Forms deployment.

### **14.4.1 Oracle Forms Single-Instance Component Characteristics**

In the Oracle Fusion Middleware Forms Services architecture there is only one connection between the client and the HTTP Listener, much like any Web-based application. The HTTP Listener routes the request to the Forms Listener Servlet, which controls routing the requests from the Forms client to the Forms runtime.

The communication between the Forms client and the Forms runtime always goes through the HTTP Listener, leaving the application with only one port open to the network.

The client sends HTTP requests and receives HTTP responses from the HTTP Listener process. With the HTTP Listener acting as the network endpoint for the client, the other server machines and ports are not exposed at the firewall.

#### **14.4.1.1 Oracle Forms State Information**

Oracle Forms employs a stateful architecture. Each Runtime process keeps the state for the client it serves in working memory. No state is ever, or can ever be, serialized or shared between runtime processes

#### **14.4.1.2 Oracle Forms Database Requirements**

Oracle Forms only requires access to the databases with which it will interact. There are no special requirements for Oracle Forms itself.

### 14.4.1.3 Oracle Forms Request Flow

The Oracle Forms request flow is as follows:

1. The user chooses a Link from a Web page, or types a URL directly in the Browser.
2. The HTTP Listener interprets the URL that is passed and displays an HTML page containing an `<EMBED>` or `<OBJECT>` tag (depending on which browser is used) that describes the Forms Java Client to the Browser. The URL that is passed calls the Forms Servlet to create an HTML page dynamically based on the `base.html` files located on the web server.
3. The Client receives the HTML file served by the HTTP Listener. The tag in the HTML file supplies the information required to locate the Java Class files that make up the Forms Java Client. Within the tag in the HTML file you supply information about the Form that should run, and any other parameters that you want to pass to your Forms session, such as the Login information. The tag definition also contains instructions on what Forms Services to run and many parameters which help to customize aspects of the Java Client, including the look-and-feel, and color schemes.

The HTML file might also contain other HTML attributes such as those to tell the browser to run this particular applet using a particular version of the JRE on the client.

4. The Browser then asks the HTTP Listener for the Java Class files from the location specified in the HTML file. The `CODEBASE` parameter in the HTML file is used to define this. The files may be downloaded individually or as an "Archive". This archive has an extension of `.JAR` and can be best thought of as a `.ZIP` file containing all of individual `CLASS` files required by the applet. The use of a `JAR` file speeds up the download of the Java Client and enables caching on the client for subsequent calls. The `ARCHIVE` parameter defines which (if any) `.JAR` file should be used. The JRE plug-in carries out the additional step of checking the version of the Forms Client Java code available on the HTTP Listener and only downloads it if it turns out to be newer than any version that the plug-in currently has cached.
5. The `CLASS` or `JAR` files are downloaded (if not already present) to the Browser and the Java applet starts.
6. The Java Client applet sends a request to start a Forms session through the HTTP Listener to the Forms Listener Servlet. The Forms Listener Servlet is defined by the `serverURL` parameter in the HTML file's tag.
7. After receiving the connection request from the Java Client, the Forms Listener Servlet starts a new Forms Runtime process for this client. You can define user specific environments for each runtime process by setting the Servlet initialization `envFile` parameter in the `formsweb.cfg` configuration file to a specific environment file.
8. The Forms Runtime process allocated to this client, loads the module specified in the HTML file and any libraries and menus that are required by that form. All communication between the Forms Client and the Forms Runtime process is passed through the Forms Listener Servlet.
9. The user is prompted for database login information, if this had not already been supplied, and the connection to the database server is established.
10. The user is now ready to work.



#### 14.4.1.4 Oracle Forms Configuration Persistence

Oracle Forms uses several files for startup configuration:

- `formsweb.cfg` holds the Forms specific startup parameters for the runtime process that used to be expressed on the command line. It consists of a default section that is used if no specific section is used and one or more user sections that holds parameters specific to that section. These sections correspond to what is referred to as an application, a set of forms that belong to the same logical unit.

`formsweb.cfg` is located in the following directory:

```
DOMAIN_HOME/config/fmwconfig/servers/<FORMS_MANAGED_SERVER>/applications/formsapp_11.1.1/config
```

- `default.env` holds startup parameters that used to be expressed in environment variables. On Windows these variables can also be put in the system registry.

This file is located in the same directory as `formsweb.cfg`:

```
DOMAIN_HOME/config/fmwconfig/servers/<FORMS_MANAGED_SERVER>/applications/formsapp_11.1.1/config
```

It is possible to create customized `.env` files with different names. If such files exist, they are specified in the `formsweb.cfg` and are thus discoverable.

- `base.html` and `basejpi.htm` or a customized version of either, holds the template structure for the HTML code that launches the Forms client applet.

These files are located in the following directory:

```
ORACLE_INSTANCE/config/FormsComponent/forms/server
```

- `Registry.dat` holds the default location and search paths for fonts, icons, and images. Find this file in:

```
DOMAIN_HOME/config/fmwconfig/servers/<FORMS_MANAGED_SERVER>/applications/formsapp_11.1.1/config/forms/registry/oracle/forms/registry
```

- `frmweb.res` holds key binding definitions for the Forms client. The bindings define the relationships between keyboard keys and the internal Forms functions they can trigger. For UNIX this bindings file is located at:

```
ORACLE_INSTANCE/config/FormsComponent/forms/admin/resource/<language>
```

Where `<language>` is a two character language denominator. For Windows this file is available at:

```
ORACLE_INSTANCE/config/FormsComponent/forms
```

On Windows, the language is defined by putting the language abbreviation at the end of the file name: `frmwebf.res` for French for example.

- `jvmcontroller.cfg` is the file for configuring JVM Controller(s). Any controllers configured for the system will have been defined in this file. This file is located in the following directory:

```
ORACLE_INSTANCE/config/FRComponent/frcommon/tools/jvm/
```

#### 14.4.1.5 Oracle Forms Runtime Considerations

The Oracle Fusion Middleware Forms Services is made up of three components: a Forms Client that is downloaded automatically to the end user's client machine and cached, the Forms Listener Servlet, and the Forms Runtime, on the middle tier.

##### Oracle Forms Client (Java Applet)

When a user runs a Forms session, the Forms Client, a thin 100 percent Java Applet, dynamically downloads from the Oracle Fusion Middleware Application Server. This generic Java Applet provides the classes for rendering the user interface for the associated Forms Runtime process on the middle tier, and handles user interaction and visual feedback, such as that generated by navigating between items or checking a checkbox. The same Java applet is used for all Forms application, therefore it is downloaded only once and cached on the client and so is available for subsequent Forms applications.

In order to run a Java applet in a browser, it is necessary to have a Java Runtime Engine (JRE) installed. The JRE is installed on the client and is platform dependent.

##### Oracle Forms Runtime Process

The Forms Runtime process is the process that maintains a connection to the database on behalf of the Forms Client. The process is created when a user accesses a page containing a Forms application. The process is automatically stopped when the user closes the Forms application or terminates the browser window.

##### Oracle Forms Listener Servlet

The Forms Listener Servlet manages:

- The creation of a Forms runtime process for each client when a user requests to run a Forms application.
- The Forms Listener Servlet is also in charge of stopping the runtime process as the user closes the Forms application or terminates the browser window.
- Network communications between the client and its associated Forms runtime process

#### 14.4.1.6 Oracle Forms Process Flow

The various components of Oracle Forms Services contribute to the process of running and Oracle Form as follows:

1. Client requests a form by contacting a server through a URL.
2. Oracle HTTP Server routes the request to Oracle WebLogic Server
3. Oracle WebLogic Server creates a forms Runtime process to run the form.

If a middle tier server crashes or a servlet session is interrupted, recover from either failure by restarting the application. A new Forms Runtime is created by doing so and any unsaved data can be reentered. The unsaved data does not cause database corruption since Forms uses atomic transactions which guarantees that database saves happen in an orderly and defined manner.

#### 14.4.1.7 Oracle Forms Configuration Files

The following table shows configuration files used by Oracle Forms:

| File              | APPHOST1 Location                                                                                                                                             | APPHOST2 Location                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| formsweb.cfg      | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS/applications/formsapp_<br>11.1.1/config                                                                | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS1/applications/formsapp_<br>_11.1.1/config                                                              |
| default.env       | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS/applications/formsapp_<br>11.1.1/config                                                                | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS1/applications/formsapp_<br>_11.1.1/config                                                              |
| base(jpi).htm     | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS/applications/formsapp_<br>11.1.1/config                                                                | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS1/applications/formsapp_<br>_11.1.1/config                                                              |
| Registry.dat      | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS/applications/formsapp_<br>11.1.1/config/forms/registry/oracle/forms/registry                           | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>FORMS1/applications/formsapp_<br>_11.1.1/config/forms/registry/oracle/forms/registry                         |
| frmweb.res        | UNIX:<br><br>ORACLE_<br>INSTANCE/config/FormsComponent/forms/admin/resource/<language><br><br>Windows:<br><br>ORACLE_<br>INSTANCE/config/FormsComponent/forms | UNIX:<br><br>ORACLE_<br>INSTANCE/config/FormsComponent/forms/admin/resource/<language><br><br>Windows:<br><br>ORACLE_<br>INSTANCE/config/FormsComponent/forms |
| jvmcontroller.cfg | ORACLE_<br>INSTANCE/config/FRComponent/frcommon/tools/jvm/                                                                                                    | ORACLE_<br>INSTANCE/config/FRComponent/frcommon/tools/jvm/                                                                                                    |
| forms.conf        | ORACLE_<br>INSTANCE/config/OHS/ohs1/moduleconf                                                                                                                | ORACLE_<br>INSTANCE/config/OHS/ohs1/moduleconf                                                                                                                |

#### 14.4.1.8 Oracle Forms External Dependencies

Oracle Forms requires an Oracle HTTP Server to service requests. Optionally these requests may be cached by Oracle Web Cache, in which case it also acts as a load balancer.

Oracle Forms uses Sun's Java plug-in (JRE) on the client to run the Forms Java Client.

#### 14.4.1.9 Oracle Forms Log Files


The following table shows log files used by Oracle Forms:

| File          | Location                           | Description                              |
|---------------|------------------------------------|------------------------------------------|
| WLS_FORMS.log | DOMAIN_HOME/servers/WLS_FORMS/logs | Log file for the WebLogic Managed Server |

| File          | Location                                           | Description                                                                                                                                                                                                                          |
|---------------|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WLS_FORMS.out | <i>DOMAIN_HOME</i> /servers/WLS_FORMS/logs         | Supplemental log file for the WebLogic Managed Server. This is generally the first place to look for WebLogic issues.                                                                                                                |
| Various       | <i>ORACLE_INSTANCE</i> /FormsComponent/forms/trace | Forms trace files, generated if a Forms runtime process fails.<br><br>The filename has the format:<br><br><forms_runtime_process>_dump_<process_id><br><br>Forms runtime process does not write to a log file under normal operation |
| Various       | <i>DOMAIN_HOME</i> /servers/WLS_FORMS/logs         | Listener Servlet logs                                                                                                                                                                                                                |

### 14.4.2 Oracle Forms Protection from Failover and Expected Behavior

Oracle Forms has no serializable state, which means that transparent failover is not possible. If a runtime process fails, the client process served by that server process also fails. Oracle Forms employs atomic database transactions. With atomic transactions, a set of data, a record, or set of records, is either fully saved or not saved at all. As a result, the failure causes data not yet saved to be lost in a defined and precise manner. The data must be re-entered by the user when the application is restarted. If the process failed as a result of the machine it ran on failing, a substantial part, or all of the existing capacity to service the user base may be unavailable. In this scenario the application as a whole may be unavailable. To ensure continuity one option is to have redundant capacity either in an active-active or active-passive configuration.

**Figure 14–4 Oracle Forms High Availability Deployment**

#### 14.4.2.1 Oracle Forms N+1 Redundancy

N+1 refers to an approach to redundant capacity that is based on the idea that hardware tends to break in units rather than in groups, meaning that a network of computers is more likely to lose one of its components rather than lose several at the same time. The principal of N+1 is to have as many machines as needed to service the entire user base at peak load plus one additional unit of equal capacity as the machine with the largest capacity in the set. This is often referred to as active-passive failover, since the failover happens on the machine level.

If you had six servers in your deployment that can handle the entire user base and are identically configured, they all have an HTTP Server and a Java Runtime and the Forms Runtime installed. There is a hardware (or software) load balancer to ensure that no single server is over utilized. The load balancer is aware of one machine labeled Standby, which has the same capacity as the machine with the largest capacity among the rest of the set, and is identically configured to all of them, but does not route to it. In an active-passive scenario, the standby machine doesn't have to be running. In an active-active scenario it could be an active part of the set, and provide spare capacity in an ongoing basis.

Now let's assume one of the machines fails:

If one of the machines fails, the standby machine is brought online and the load balancer has been reconfigured to route new requests to that machine as well, and to ignore the failed machine (some hardware load balancers do these two steps automatically). If the standby machine was already running and serving the user base in an active-active deployment, the downtime may be as small as the time it takes to restart the browser on the client. If the Standby machine was not running in an

active-passive deployment, and was required to be started, and the load balancer had to be reconfigured manually, the downtime would be longer.

#### **14.4.2.2 Oracle Forms N+M Redundancy**

For added failover capacity more than one standby machine can be used. That is usually referred to as N+M. The chance of more than one machine failing at the same time (or close to the same time) is significantly smaller, but if there is a requirement to be able to handle that situation, a N+M setup is possible and perhaps called for.

For redundancy in the case of a condition that affects all the machines in the deployment, perhaps a natural disaster that destroys all machines in one location, this deployment can be duplicated in two different geographical locations.

#### **14.4.2.3 Oracle Forms Virtual Machines**

In a datacenter that uses virtual machines, the standby machine can be brought online using any spare hardware and thus be implemented very cheaply.

#### **14.4.2.4 Oracle Forms Configuration Cloning**

The significance of cloning the environment becomes apparent when studying the scenarios described in the previous sections. All changes done to one machine's environment must also be present on all other machines, including the standby machine. This can present a practical problem, especially if the spare machine is virtual. The image that is the virtual machine must be kept in sync with the changes made to the other machines.

#### **14.4.2.5 Oracle Forms Process Failures**

For information about Oracle Forms process failure, see [Section 14.1.3.1, "Oracle Web Cache and Oracle HTTP Server Process Failures."](#)

#### **14.4.2.6 Oracle Forms Node Failures**

For information about Oracle Forms Node failure, see [Section 14.1.3.2, "Common Component Node Failures."](#)

#### **14.4.2.7 Oracle Forms WebLogic Managed Server Failures**

For information about Oracle Forms WebLogic Managed Server failure, see [Section 14.1.3.3, "Common Component WebLogic Managed Server Failures."](#)

#### **14.4.2.8 Oracle Forms Database Failures**

For information about Oracle Forms database failure, see [Section 14.1.3.4, "Common Component Database Failures."](#)


## **14.5 Oracle Discoverer and High Availability Concepts**

This section describes single-instance information, as well as high availability concepts specific to Oracle Discoverer. This section guides you through the concepts and considerations necessary for creating a successful high availability Oracle Discoverer deployment.

## 14.5.1 Oracle Discoverer Single-Instance Characteristics

When you install Discoverer, you create a Discoverer topology within a single instance, as shown in [Figure 14–5](#). After installation, you can configure other types of topologies for Discoverer.


**Figure 14–5 Discoverer Topology for a Single Instance**



### 14.5.1.1 Oracle Discoverer Runtime Considerations

[Figure 14–6](#) illustrates the runtime interaction between Discoverer components.

**Figure 14–6 Runtime Interaction Between Discoverer Components**



The **Discoverer Servlet** is a Java EE application that is deployed on a managed server, which can be started and shutdown through the Oracle WebLogic Server administration console and by using the Oracle Enterprise Manager Fusion Middleware Control. The Discoverer Servlet is a stateless process.

The **Discoverer Session Server** is an OPMN-managed CORBA component that performs Discoverer operations such as connecting to the database or opening a workbook. The Session Server provides the link between the Discoverer Servlet and the database. There is one Session Server component per active user login session. The Discoverer Session Server is a stateful process; the state is stored only in memory.

The **Discoverer Preference Server** is an OPMN-managed component that provides preference settings (both default and user-defined) for all Discoverer users. These preferences control the Discoverer behavior. For information about starting and stopping the Preference Server, see the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

The **Discoverer Services Status** is a dummy process managed by OPMN. This process must be running for the Session Server to be spawned.

---

**Note:** The Discoverer Servlet and Session Server must run on the same machine.

---

**Oracle Discoverer External Dependencies**

Discoverer requires an Oracle HTTP Server and Enterprise Manager Fusion Middleware Control.

The database schema for Discoverer and the portlet schema must be loaded (before installing Discoverer) by using the Repository Creation Utility (RCU).

Discoverer interacts with customer databases that contain Discoverer workbooks and the Discoverer End User Layer and other data sources. For more information, see the "About the Discoverer database tier" section in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

Oracle Web Cache can be used to configure load balancing for Discoverer.

The Portlet Provider component of Discoverer can provide portlets to Oracle Portal and Oracle Web Center.

The Web Services component of Discoverer can be used by external clients (such as Oracle BI Publisher and Oracle BI Enterprise Edition) to obtain Discoverer connections and workbooks.

It is recommended (but not mandatory) to restrict access to Oracle Discoverer by using Oracle Single Sign-On (SSO).

**Discoverer Process Flow**

See the "How does Discoverer work?" section in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

**Connection Protocols**

Client browsers send HTTP requests to the Discoverer Java EE application.

The Discoverer Java EE application accesses the database schema by using WebLogic Server data sources.

The Discoverer Session Server (non-Java EE component) uses the OCI layer to connect to data sources.

The Discoverer Java EE application and the Discoverer Session Server processes communicate by using CORBA.

**14.5.1.2 Oracle Discoverer Viewer and Web Cache**

You can improve Discoverer Viewer performance and availability by using Oracle Web Cache. For information about when and how to use Web Cache, see the "Using Discoverer Viewer with Oracle Web Cache" chapter in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

**14.5.1.3 Oracle Discoverer Configuration Considerations**

The environment and behavior of Oracle Discoverer are controlled by Discoverer preferences and configuration parameters.

- For information about the list of preferences and the procedure to change preference settings, see Managing Discoverer Preferences in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*
- Configuration parameters are stored in the `configuration.xml` file. For information about configuration parameters and how you can define settings for those parameters, see Managing and Configuring Discoverer in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.



For information about the location of the files in which preferences and configuration parameters are stored, see the "Discoverer Configuration Files" section in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

#### 14.5.1.4 Oracle Discoverer Deployment Considerations

The Discoverer Java EE application deployment and contains basic deployment descriptors.

- The `discoverer.ear` file consists of the following files: `application.xml`, `jazn-data.xml`, and `weblogic-application.xml`.

These files are located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_DISCO/applications/discoverer_version/configuration/` directory, where `DOMAIN_HOME` is the name of the domain directory.

- The `discoverer.war` file contains the following files: `custom-laf.xml`, `plus_versions.properties`, `uix-config.xml`, `portlet.xml`, `weblogic.xml`, `struts-config.xml`, and `web.xml`.

These files are located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_DISCO/applications/discoverer_version/configuration/` directory, where `DOMAIN_HOME` is the name of the domain directory.

#### 14.5.1.5 Oracle Discoverer Log File Locations

You can search for and view Discoverer log files in the Enterprise Manager Fusion Middleware Control.

For more information, see the Enterprise Manager Fusion Middleware Control online help.

#### 14.5.1.6 Discoverer Log Files


The following table shows log files used by Oracle Discoverer:

| File          | Location                                                                     | Description                                                                                                           |
|---------------|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| WLS_DISCO.log | <code>DOMAIN_HOME/servers/WLS_DISCO/logs</code>                              | Log file for the WebLogic Managed Server                                                                              |
| WLS_DISCO.out | <code>DOMAIN_HOME/servers/WLS_DISCO/logs</code>                              | Supplemental log file for the WebLogic Managed Server. This is generally the first place to look for WebLogic issues. |
| Various       | <code>ORACLE_INSTANCE/diagnostics/logs/Discoverer/Discoverer_Instance</code> | Supplemental Discoverer log files including Preference store.                                                         |

### 14.5.2 Oracle Discoverer Protection from Failures and Expected Behavior

Figure 14–7 shows a Discoverer topology that consists of two Discoverer instances.

The Discoverer Java EE application in each managed server communicates with a single Preference Server. Therefore, the same preferences are applied to both the Discoverer instances. The Discoverer instances can exist either on the same machine or on different machines.

**Figure 14–7 Discoverer Topology for Multiple Instances**

The following sections discuss specific high-availability considerations relevant for Discoverer. After making configuration changes to provide high availability, you must restart the Oracle HTTP Server and the managed servers in which the Discoverer Java EE applications are deployed.

#### 14.5.2.1 Preference Server Failover

Failure detection and restart of the Discoverer Java EE application in a cluster is managed by WebLogic Server.

The Preference Server is a singleton process and runs only one instance in the entire cluster.

If the Preference Server process goes down, OPMN brings it up again automatically. If OPMN too is down, the administrator must either start the Preference Server manually or move the files containing the preferences (`reg_key.dc` and `pref.txt`) to another machine where the Preference Server is configured and then start the Preference Server on that machine.

If the Preference Server is started in another node in the cluster, the Discoverer Java EE application must be configured to recognize the new Preference Server.

#### 14.5.2.2 Session State Replication and Failover

When a server or machine that is handling requests pertaining to a particular Discoverer HTTP session is down, subsequent request are diverted to other managed servers in the cluster. The session state can be replicated in the new server. The necessary information to achieve this replication is available in the HTTP request (as GET/POST).

---

**Note:** Discoverer requires a Session Server (C++ process) to be available to complete any request. Therefore, in a failover situation, the response time would be marginally higher because a C++ process must be spawned and brought to the required state.

---

### 14.5.2.3 Performance Recommendation

The Discoverer Java EE Servlet spawns a Session Server process for each HTTP request that it receives. If multiple Discoverer instances (managed server and the associated oracle instance) exist in a single machine, the load on the CPU might become very high because all Discoverer instances might spawn C++ processes simultaneously. Adding Discoverer instances on different machines yields better performance than adding them on the same machine.

### 14.5.2.4 Propagation of Configuration Changes Across the Cluster

Any configuration change must be implemented individually in each managed server in the cluster. You can achieve this by copying the `configuration.xml` file to all the managed servers.

For information about the location of the configuration files, see the "Discoverer Configuration Files" section in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

### 14.5.2.5 Cluster-Wide Application Deployment

Each Discoverer Java EE application instance (managed server) should be associated with an Oracle instance that contains Discoverer components. These components are created during the configuration phase of the installation process. So adding managed servers to a cluster and deploying the Discoverer Java EE application from the WebLogic Server administration console are not supported.

All the managed servers defined in the cluster are targets for deployment of the Discoverer Java EE application.

### 14.5.2.6 Online Application Deployment

Most of the Discoverer Java EE application-dependent jar files are available through shared libraries. Patching these shared libraries does not require the Discoverer Java EE application to be restarted.

Changes to the following jar files, which are available through the system classpath, require the Discoverer Java EE application to be restarted.

```
WL_HOME/server/lib/weblogic.jar
ORACLE_HOME/modules/oracle.jrf_11.1.1/jrf.jar
ORACLE_HOME/opmn/lib/nonj2eebeans.jar
ORACLE_HOME/server/lib/ojdbc6.jar
ORACLE_HOME/opmn/lib/optic.jar
ORACLE_HOME/opmn/lib/iasprovision.jar
ORACLE_HOME/opmn/lib/ons.jar
ORACLE_HOME/modules/oracle.adf.share_11.1.1/oracle-el.jar
ORACLE_HOME/jlib/share.jar
ORACLE_HOME/jlib/jewt4.jar
```

### 14.5.2.7 Oracle Discoverer Process Failures

For the Discoverer Java EE application to access databases, an entry for each database must be created in the `ORACLE_INSTANCE/config/tnsnames.ora` file. For more information, see the "About configuring the tnsnames.ora file" section in the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

#### 14.5.2.8 Oracle Discoverer Node Failures

For information about Oracle Discoverer Node failure, see [Section 14.1.3.2, "Common Component Node Failures."](#)

#### 14.5.2.9 Oracle Discoverer WebLogic Managed Server Failures

For information about Oracle Discoverer WebLogic Managed Server failure, see [Section 14.1.3.3, "Common Component WebLogic Managed Server Failures."](#)

#### 14.5.2.10 Oracle Discoverer Database Failures

For information about Oracle Discoverer database failure, see [Section 14.1.3.4, "Common Component Database Failures."](#)

## 14.6 Configuring Oracle Portal, Forms, Reports, and Discoverer for High Availability

This section describes procedures for configuring Oracle Portal, Forms, Reports, and Discoverer for a high availability deployment. This section contains information on the following topics

---

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

- [Section 14.6.1, "Prerequisites"](#)
- [Section 14.6.2, "Assumptions"](#)
- [Section 14.6.3, "Creating the Metadata Repository"](#)
- [Section 14.6.4, "Install and Configure Application Tier on APPHOST1"](#)
- [Section 14.6.5, "Install and Configure Application Tier on APPHOST2"](#)
- [Section 14.6.6, "Scaling Out the Deployment"](#)

---

---

**Note:** The hostnames and ports listed in these sections are, for illustrating this example deployment. When configuring your actual deployment, you can specify your own hostnames and ports.

---

---

### 14.6.1 Prerequisites

Before installing Oracle Portal, Forms, Reports and Discoverer, review the prerequisites described in this section.

#### 14.6.1.1 Dependencies

If you are using Oracle Single Sign-On, it is assumed that a high availability Identity Management Framework is available. Oracle Portal, Forms, Reports, and Discoverer uses Single Sign-On, which is available in Oracle Identity Management 10g.

#### 14.6.1.2 Network Requirements

This section describes network requirements.

**14.6.1.2.1 Load Balancer** In order to distribute requests across the Oracle Web servers, a load balancer is required. This external load balancer should have the following features:

- Virtual server name and port configuration
- Process failure detection
- Monitoring of ports (HTTP, HTTPS) for Oracle HTTP and HTTPS
- SSL Translation (if required)

**14.6.1.2.2 Load Balancer Configuration - Virtual Server Names and Ports** If you are using a load balancing router, it must be configured to enable the following:

**If using SSL traffic terminated at the load balancer**

A virtual IP address (VIP1) that listens for requests to `mysite.mycompany.com` on port 443 (an HTTPS listening port), and balances them to the application tier Oracle Web Cache, running on WEBHOST1 and WEBHOST2, port 7777 (an HTTP listening port). You must configure the Load Balancing Router to perform protocol conversion.

**If using site SSL**

A virtual IP address (VIP1) that listens for requests to `mysite.mycompany.com` on port 443 (an HTTPS listening port), and balances them to the application tier, Oracle Web Caches running on WEBHOST1 and WEBHOST2 port 443 (an HTTPS listening port).

**If using HTTP traffic terminated at the load balancer**

A virtual IP address (VIP1) that listens for requests to `mysite.mycompany.com` on port 80 (an HTTP listening port), and balances them to the application tier, Oracle Web Caches running on WEBHOST1 and WEBHOST2 port 7777 (an HTTP listening port).

**Web Cache**

- The virtual IP address VIP1 listens for requests to `mysite.mycompany.com` on port 9401 (an HTTP listening port), and balances them to the application tier Oracle Web Cache on WEBHOST1 and WEBHOST2 port 9401 (an HTTP listening port).

---

**Note:** For security reasons, ports 9401, 9402 and 9403 on the load balancing router should not be visible to external users.

---

- HTTP/HTTPS monitoring of Oracle Web Cache: The Load Balancing Router must be configured to detect an inoperative computer and stop routing requests to it until it is functioning again. Two Oracle Web Cache ports must be monitored: the HTTP request port and the invalidation port.

To monitor port 7777, use the following URL in the Load Balancing Router configuration:

```
hostname:port/_oracle_http_server_webcache_static_.html
```

For example:

```
http://webhost1.mycompany.com:7777/_oracle_http_server_webcache_static_.html
```

If the Load Balancing Router receives a response from this URL, then the Oracle Web Cache instance is running. If it does not receive a response, the process or the

server is down, and the load balancing router forwards all requests to the active computer.

To monitor port 9401, use the following URL in the Load Balancing Router configuration:

`http://hostname.domain.com:9401/x-oracle-cache-invalidate-ping`

For example:

`http://apphost1.mycompany.com:9401/x-oracle-cache-invalidate-ping`

The load balancing router sends an HTTP request to this URL; the response header resembles the following:

`HTTP/1.0`

The load balancing router must be configured to detect the string HTTP in the first line of the response header. Therefore, when the load balancing router detects HTTP in the first line of the response header, the invalidation port is available. If it does not, all invalidation requests are routed to the active computer.

---

**Note:** The `sqlnet.ora` file must be updated to prevent connection timeouts related to the load balancing router and firewall.

---

To summarize, the load balancer requires the following configuration:

| Virtual Host         | Virtual Port | Server Pool        | Server   | Port | Comments                                                             |
|----------------------|--------------|--------------------|----------|------|----------------------------------------------------------------------|
| mysite.mycompany.com | 443/80       | UserRequest        | WEBHOST1 | 7777 | Protocol conversion required if terminating SSL at the load balancer |
| mysite.mycompany.com | 443/80       | UserRequest        | WEBHOST2 | 7777 | Protocol conversion required if terminating SSL at the load balancer |
| mysite.mycompany.com |              | Cache Invalidation | WEBHOST1 | 9401 | Invisible to the external clients                                    |
| mysite.mycompany.com |              | Cache Invalidation | WEBHOST2 | 9401 | Invisible to the external clients                                    |

### 14.6.1.3 Databases

Different products use databases in different ways. Below is a summary of the database requirements for Oracle Portal, Forms, Reports and Discoverer:

**Oracle Portal**

Oracle Portal stores both its own metadata and user content in a database. The majority of the Portal application logic is also placed into this database in the form of PLSQL. Using a combination of the metadata, user content, and PLSQL, Portal generates Web pages.

Oracle Portal requires a high availability database, which has been pre-seeded with database objects required by the Portal application. For information about creating the Metadata Repository, see [Section 14.6.3, "Creating the Metadata Repository."](#)

**Oracle Forms**

Oracle Forms interacts with the database and therefore requires access to it. There are no special database requirements of Oracle Forms itself.

**Oracle Reports**

Oracle Reports requires a highly available database, which contains the reports queue. It also requires access to various customer databases, which it uses to compile report content.

**Oracle Discoverer**

Oracle Discoverer requires a database for portlets. This should be a high availability database, and should be seeded using the Oracle Repository Creation Utility (RCU).

In addition, Oracle Discoverer interacts with a number of customer databases. These databases are used to store Discoverer End User Layers, and work books as well as the data on which it reports.

**14.6.1.4 Shared Directories**

Shared directory requirements vary depending on the product. These requirements are described in the following table:

| Product           | Shared Directory Requirement                                                       |
|-------------------|------------------------------------------------------------------------------------|
| Oracle Portal     | None                                                                               |
| Oracle Forms      | Not mandatory, but useful to have a shared directory for Oracle Forms executables. |
| Oracle Reports    | Reports Cache                                                                      |
| Oracle Discoverer | Portlet Preference Store                                                           |

**14.6.1.5 Managed Port Numbers**

Many Oracle Fusion Middleware components and services use ports. As an administrator, it is important to know the port numbers used by these services, and to ensure that the same port number is not used by two services on the same host.

Port numbers can either be automatically or manually assigned at installation time.

**14.6.1.6 Site Names**

In order to configure a Web site a site name is required. This site name, for example `mysite.mycompany.com`, must be defined in DNS and be associated with the Virtual IP address assigned to the load balancer.

A site name is also required for the Single Sign-On server, which is set up as part of the high availability Single Sign-On installation.

## 14.6.2 Assumptions

For the example high availability configuration described in this chapter, consider the following assumptions:

### 14.6.2.1 Ports

The following table lists the typical ports required by an Oracle Portal, Forms, Reports, and Discoverer implementation.

| Purpose                   | Host(s)              | Port | Comment                             |
|---------------------------|----------------------|------|-------------------------------------|
| mysite.mycompany.com      | Load balancer        | 80   | HTTP port on the load balancer      |
| Web Cache HTTP            | APPHOST1<br>APPHOST2 | 7777 | Web Cache HTTP port                 |
| Web Cache HTTPS           | APPHOST1<br>APPHOST2 | 4443 | Web Cache HTTPS port                |
| Web Cache Invalidation    | APPHOST1<br>APPHOST2 | 9401 | Web Cache invalidation port         |
| Web Cache Admin           | APPHOST1<br>APPHOST2 | 9400 | Web Cache Administration port       |
| HTTP Server (OHS) - HTTP  | APPHOST1<br>APPHOST2 | 7778 | OHS HTTP listening port             |
| HTTP Server (OHS) - HTTPS | APPHOST1<br>APPHOST2 | 4444 | OHS HTTPS listening port            |
| WebLogic Admin port       | APPHOST1             | 7001 | WebLogic Administration Server port |
| WLS_PORTAL                | APPHOST1             | 7050 | WebLogic Managed Server port        |
| WLS_PORTAL1               | APPHOST2             | 7050 | WebLogic Managed Server port        |
| WLS_REPORTS               | APPHOST1             | 7051 | WebLogic Managed Server port        |
| WLS_REPORTS1              | APPHOST2             | 7051 | WebLogic Managed Server port        |
| WLS_DISCO                 | APPHOST1             | 7052 | WebLogic Managed Server port        |
| WLS_DISCO1                | APPHOST2             | 7052 | WebLogic Managed Server port        |
| WLS_FORMS                 | APPHOST1             | 7053 | WebLogic Managed Server port        |
| WLS_FORMS1                | APPHOST2             | 7053 | WebLogic Managed Server port        |
| Internet Directory        | SSOHOST              | 389  | Oracle Internet Directory HTTP port |
| Single Sign On            | SSOHOST              | 7777 | Single Sign On listening port       |



## 14.6.3 Creating the Metadata Repository

Before installing Oracle Portal and Oracle Discoverer, create and "seed" a repository with metadata that these applications require to function, by running the Repository Creation Utility (RCU).

---

**Note:** A metadata repository is not required for Oracle Forms and Reports.

---

### 14.6.3.1 Install the Repository Creation Utility (RCU)

Install the latest version of the Repository Creation Utility, as described in the *Oracle Fusion Middleware Repository Creation Utility User's Guide*.

### 14.6.3.2 Run Repository Creation Utility

After installing RCU, use it to populate the database. Start the Repository Creation Utility using the following commands:

For UNIX:

```
rcu
```

For Windows:

```
rcu.exe
```

These commands are located in the RCU ORACLE\_HOME/bin directory.

1. In the Create Repository screen, select **Create** and then click **Next**.
2. In the Database Connection Details screen, specify the following values:
  - **Database Type:** Oracle Database
  - **Host Name:** Enter *one* of the Oracle RAC nodes (use the VIP name), if using Oracle RAC.
  - **Port:** Enter the listener port.
  - **Service Name:** Enter the service name of the Oracle Real Application Clusters database.
  - **User Name:** Enter *sys*.
  - **Password:** Enter the *sys* user password.
  - **Role:** Select SYSDBA.
3. In the Check Prerequisites screen, click **OK** when the prerequisites have been validated.
4. In the Select Components screen, specify the following values:
  - **Create New Prefix:** Enter a prefix to be added to database schemas, for example: MYS
  - **Components:** Portal and BI > Portal and Discoverer
 Click **Next**.
5. In the Check Prerequisites screen, click **OK** when the prerequisites have been validated.

6. In the Schema Passwords screen, enter passwords for each of the portal schemas or use the same password for all schemas.  
Click **Next**.
7. In the Map Tablespaces screen, click **Next** to accept the defaults.
8. In the Create Tablespaces screen, click **Yes** to allow RCU to create any missing tablespaces.
9. In the Creating Tablespaces screen, click **OK** to acknowledge tablespace creation.
10. In the Summary screen, click **Create** to begin the creation process.

## 14.6.4 Install and Configure Application Tier on APPHOST1

The following steps are applicable for Oracle Portal, Forms, Reports and Discoverer installations.

- [Section 14.6.4.1, "Install Oracle WebLogic Server"](#)
- [Section 14.6.4.2, "Install Oracle Portal, Forms, Reports, and Discoverer"](#)
- [Section 14.6.4.3, "Configure Oracle Portal, Forms, Reports, and Discoverer Software"](#)
- [Section 14.6.4.4, "Validation"](#)
- [Section 14.6.4.5, "Generic Configuration"](#)
- [Section 14.6.4.6, "Configure Oracle Portal for High Availability"](#)
- [Section 14.6.4.7, "Configure Oracle Forms for High Availability"](#)
- [Section 14.6.4.8, "Configure Oracle Reports for High Availability"](#)
- [Section 14.6.4.9, "Configure Oracle Discoverer for High Availability"](#)

### 14.6.4.1 Install Oracle WebLogic Server

To install Oracle WebLogic Server binaries, see the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

### 14.6.4.2 Install Oracle Portal, Forms, Reports, and Discoverer

To install the Oracle Portal, Forms, Reports, and Discoverer binaries, see *Oracle Fusion Middleware Installation Guide for Oracle Portal, Forms, Reports and Discoverer*.

---

**Note:** Before you run the Configuration Wizard by following the instructions in [Section 14.6.4.3, "Configure Oracle Portal, Forms, Reports, and Discoverer Software,"](#) ensure that you applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home so that you have the latest version of Oracle Fusion Middleware.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

---

### 14.6.4.3 Configure Oracle Portal, Forms, Reports, and Discoverer Software

Start the Oracle Fusion Middleware Configuration Wizard by running this command:

```
ORACLE_HOME/bin/config.sh
```

---

---

**Note:** Before starting the configuration, ensure that the following environment variables (UNIX) are not set:

- LD\_ASSUME\_KERNEL
  - ORACLE\_BASE
  - LD\_LIBRARY\_PATH
- 
- 

Continue with the following steps:

1. In the Welcome screen, click **Next**.
2. In the Prerequisite Checks screen, once all the prerequisites have been validated, click **Next**.
3. In the Create Domain screen, select **Create New Domain** and enter these values:
  - **User Name:** Name of the user to log into the WebLogic domain.
  - **User Password:** Password for the domain.
  - **Confirm Password:** Type the same password again.
  - **Domain Name:** Name for the domain. For example: MyDomain
 Click **Next**.
4. In the Specify Security Updates screen, if required, enter a username and password to receive Oracle Security updates for Oracle Support.
5. In the Specify Installation Location screen, enter these values:
  - **WebLogic Server Location:** Enter the installation directory for Oracle WebLogic Server. This directory should be MW\_HOME/wlserver\_10.3, for example:  

```
/u01/app/oracle/product/FMW/wlserver_10.3
```
  - **Oracle Instance Location:** Enter the directory where the Oracle configuration files are to be placed. This should be outside of the Oracle home directory. This directory will be known as the ORACLE\_INSTANCE. For example:  

```
/u01/app/oracle/admin/fmw1
```
  - **Oracle Instance Name:** fmw1
 Click **Next**.
6. In the Configure Components screen, choose which products to install and configure.  
 Ensure that **Management Components - Enterprise Manager** is also selected.  
 Select **Clustered**.  
 Click **Next**.
7. In high availability implementations, it is not mandatory for all of the ports used by the various components to be are synchronized across hosts, but Oracle

strongly recommends it. Oracle Allows the bypassing of Automatic Port Configuration by specifying ports to be used in a file.

Create a file with the ports you wish to assign by Select a file name, and click **View/Edit**.

You can find a sample `staticports.ini` file on installation Disk1 in the `stage/Response` directory.

This file should have entries for the following:

```
[Domain Port No = 7001
Oracle HTTP Server Port No = 7778
WebCache Port No = 7777
WebCache Administration Port No = 9400
WebCache Statistics Port No = 9402
Web Cache Invalidation Port = 9401
Portal Managed Server Port = 7050
Reports Managed Server Port = 7051
Discoverer Managed Server Port = 7052
Forms Managed Server Port = 7053]
```

Save the file and click **Next**.

8. In the Specify Proxy Details screen (Reports only), if a proxy server is in use, enter the details in this screen.

Otherwise, click **Next**.

9. In the Specify Schema screen (Portal and Discoverer only), specify the following values:

- **Database Connect String** in the format:

For an Oracle RAC database:

```
racnode1-vip:ListenerPort:racnode2-vip:ListenerPort@mydb.mycompany.com
```

For other databases:

```
host:port:mydb.mycompany.com
```

- **Portal Schema Name:** `MYS_PORTAL`
- **Portal Schema Password:** Enter password entered in RCU.
- **Discoverer Schema Name:** `MYS_DISCOVERER`
- **Discoverer Schema Password:** Enter password entered in RCU.

Click **Next**.

10. In the Specify Portlet Schema screen (Portal and Discoverer only), specify the following values:

- **Portlet Schema Name:** `MYS_PORTLET`
- **Portlet Schema Password:** Enter password entered in RCU.

Click **Next**.

11. In the Specify Application Identity Store screen, specify the following values:

- **Hostname:** Name of the Oracle Internet Directory server, for example: `myoid.mycompany.com`
- **Port:** Oracle Internet Directory port, for example: 389

- **User Name:** cn=orcladmin
- **Password:** Password for Oracle Internet Directory's orcladmin account

12. In the Summary screen, click **Install** to begin the creation process.

#### 14.6.4.4 Validation

Validate the initial installation by performing the following tests:

| Test       | URL                                                           | Result                            |
|------------|---------------------------------------------------------------|-----------------------------------|
| Portal     | http://apphost1.mycompany.com:7777/portal/pls/portal          | Portal Home page opens            |
| Forms      | http://apphost1.mycompany.com:7777/forms/frmservlet           | Forms Servlet Home page opens     |
| Reports    | http://apphost1.mycompany.com:7777/reports/rwservlet/showjobs | Job Queue opens                   |
| Discoverer | http://apphost1.mycompany.com:7777/discoverer/viewer          | Discoverer Viewer Home page opens |

#### 14.6.4.5 Generic Configuration

The following steps are required regardless of the product being configured:

**14.6.4.5.1 Set Admin Server Listen Address** In servers where multiple network cards exist, it is important to bind the Administration Server to the network card that you wish to use.

To do this:

1. Log in to the WebLogic Console using the URL  
http://apphost1.mycompany.com:7001/console
2. In the **Domain Structure** menu, select **Environment**, and then **Servers**.
3. Click **AdminServer (admin)**
4. In the Change Center, click **Lock & Edit**.
5. Set the listen address to the DNS name referring to the network card you wish to use. This is generally the public server name.
6. Click **Save**.
7. Click **Activate Changes** from the change center.
8. Stop the Administration Server using the stopWebLogic.sh script, located in the DOMAIN\_HOME/bin directory. Start the Administration Server using the startWebLogic.sh script, also located in DOMAIN\_HOME/bin directory.

**14.6.4.5.2 Configure Virtual Hosts** For the site to function properly with the load balancer, you must add two virtual hosts. You can configure the virtual hosts using Oracle Enterprise Manager Fusion Middleware Control, or you can edit the file ORACLE\_INSTANCE/config/OHS/ohs1/httpd.conf or create a file called virtual\_hosts.conf in the directory ORACLE\_INSTANCE/config/OHS/ohs1/moduleconf.

To add two virtual hosts to the virtual\_hosts.conf file, in a text editor, add the following entries to the file:

```
NameVirtualHost *:7778
<VirtualHost *:7778>
 ServerName http://mysite.mycompany.com:80
```

```
RewriteEngine On
RewriteOptions inherit
UseCanonicalName On
</VirtualHost>

<VirtualHost *:7778>
 ServerName apphost1.mycompany.com:7777
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>
```

---

**Note:** Where:

7778 is the Oracle HTTP Server Listening Port

80 is the Load Balancer Listening Port

7777 is the Web Cache Listening Port.

Name of the server on which the Oracle HTTP server resides:  
apphost1

Use ServerName HTTP if your site works only in the http protocol.

---

Restart the Web Tier using these commands:

```
opmnctl stopall
opmnctl startall
```

**14.6.4.5.3 Create boot.properties File** Create a `boot.properties` file for the Administration Server on APPHOST1. The `boot.properties` file enables the Administration Server to start without prompting you for the administrator username and password.

In a text editor, create a file called `boot.properties` in the directory `DOMAIN_HOME/servers/AdminServer/security`, and enter the following lines in the file:

```
username=adminuser
password=password
```

Restarting the Administration Server encrypts the values in this file, for that reason it is recommended that the Administration Server is restarted on each node that can host it.

Stop the Administration Server using the script `stopWebLogic.sh`, which is located in `DOMAIN_HOME/bin`, or by using the script `startWebLogic`, also located in `DOMAIN_HOME/bin`.

**14.6.4.5.4 Configure sqlnet.ora** Create a file called `sqlnet.ora` in the directory `ORACLE_INSTANCE/config/` and add the following entry to the file:

```
TCP.CONNECT_TIMEOUT=10
```

This ensures that database connections time out after a reasonable time.

**14.6.4.5.5 Configure Web Cache** Web Cache is optional for Oracle Forms, Reports and Discoverer, but mandatory for Oracle Portal. If the environment being configured does not use Web Cache, ignore the following Web Cache configuration steps:

## Log Into Enterprise Manager Console

Log into the Oracle Fusion Middleware Enterprise Manager Console using the following URL:

`http://apphost1.mycompany.com:7001/em`

The default **User Name** and **Password** are the same as the WebLogic domain username and password entered during the installation.

## Create Site

1. In the Navigator window, expand the Web Tier tree.
2. Click on the component wc1.
3. From the list at the top of the page, select **Administration - Sites**.
4. Select **Create Site**.
5. Enter the following information to add a site.

| Field                  | Value                                                                                                                                                                                                                    |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Host Name              | mysite.mycompany.com                                                                                                                                                                                                     |
| Port                   | Specify the port number on the load balancer where requests are received, for example, 80 for HTTP requests and 443 for HTTPS request.<br><br>If using a load balancer, this is the listening port on the load balancer. |
| Default site           | Yes                                                                                                                                                                                                                      |
| Site Wide Compression  | Yes                                                                                                                                                                                                                      |
| Site Alias - Host Name | mysite.mycompany.com                                                                                                                                                                                                     |
| Site Alias - Port      | 7777                                                                                                                                                                                                                     |
| Site Alias - Host Name | mysite.mycompany.com                                                                                                                                                                                                     |
| Site Alias - Port      | 80 <sup>1</sup>                                                                                                                                                                                                          |

<sup>1</sup> Required when the load balancer can send requests to port 80 as well as port 443.

6. Do not change any other defaults.
7. Select **Submit**.
8. Select **OK** to save each entry.

## Create Site-to-Server Mapping

1. On the same page, select **Create** in the Site-to-Server Mapping section.
2. Enter the following information to add the site:

| Field                   | Value                       |
|-------------------------|-----------------------------|
| Host Pattern            | mysite.mycompany.com        |
| Port Pattern            | 80                          |
| Selected Origin Servers | apphost1.mycompany.com:7778 |

3. Click **OK** to store the site.
4. Ensure that the site `mysite.mycompany.com:443/80` appears first in the list of site-to-server mappings.
5. Click **Apply** to save the changes.

### Enable Session Binding

The session binding feature in Oracle Web Cache is used to bind user sessions to a given origin server to maintain state for a period of time. Although almost all components running in a default Oracle Fusion Middleware mid-tier are stateless, session binding is required for Oracle Portal.

Enabling session binding forces all the user requests to go to a specific Oracle Portal middle tier, resulting in a better cache hit ratio for the portal cache.

Follow these steps to enable session binding:

1. Select **Administration - Session Configuration** at the top of the page.
2. Select the site `mysite.mycompany.com:443/80` from the drop down list.
3. In the Session Binding section, select **Cookie based Session Binding with any Set Cookie**.
4. Click **Apply** to save the changes.

**14.6.4.5.6 Change the Web Cache Passwords** Web Cache passwords are randomly generated, however they are required in later stages. It is therefore recommended that the Web Cache passwords be changed from the default value to a new known value.

To change the default password, follow these steps:

1. In the Navigator window, expand the **Web Tier** tree.
2. Click on the component `wc1`.
3. From the drop down list at the top of the page, select **Administration - Passwords**.
4. Enter new invalidation and administration passwords, confirm them, and click **Apply**.

**14.6.4.5.7 Restart Web Tier (Oracle HTTP Server and Web Cache)** After making the previous changes, restart the Web Tier components using these commands:

```
opmnctl stopall
opmnctl startall
```

**14.6.4.5.8 Register with Single Sign-On Server** Perform these steps from the Single Sign-On server:

1. Set the `ORACLE_HOME` variable to the Single Sign-On Server `ORACLE_HOME` location.
2. Execute `ORACLE_HOME/sso/bin/ssoreg.sh` (`ssoreg.bat` on Windows) with the following parameters:

```
-site_name mysite.mycompany.com
-mod_osso_url http://mysite.mycompany.com
-config_mod_osso TRUE
-oracle_home_path ORACLE_HOME
-config_file /tmp/osso.conf
-admin_info cn=orcladmin
```



```
-virtualhost
-remote_midtier
```

3. Copy `/tmp/osso.conf` to the mid-tier home location, which is:

```
ORACLE_INSTANCE/config/OHS/ohs1
```

4. Restart Oracle HTTP server on APPHOST1 using the following command:

```
opmnctl restartproc process-type=OHS
```

5. Log into the Single Sign-On Server using the following URL:

```
http://login.mycompany.com/pls/orasso
```

6. Go to the Administration page and then Administer Partner applications. Delete the entry for `apphost1.mycompany.com`.

**14.6.4.5.9 Enabling the WebLogic Plug-in** For security purposes, turn on the WebLogic plug-in enabled flag for the domain. To do this, follow these steps:

The Oracle web tier sits in front of the Oracle WebLogic server, routing requests from the Oracle HTTP Servers to the Oracle WebLogic managed servers. In this scenario, you must tell WebLogic that the Oracle WebLogic plug-in is being used. To do this:

1. Log on to the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, click on the domain name.
3. Click on the **Web Applications** tab.
4. Click **Lock & Edit**.
5. Select the **WebLogic Plugin Enabled** box.
6. Save and activate the changes.
7. Restart the Administration Server.

**14.6.4.5.10 Change Host Assertion in WebLogic** Because the Oracle HTTP server acts as a proxy for WebLogic, by default certain CGI environment variables are not passed through to WebLogic. These include the host and port. WebLogic must be told that it is using a virtual site name and port so that it can generate internal URLs appropriately.

1. Log into the Oracle WebLogic Server Administration Console using the following URL:

```
http://apphost1.mycompany.com:7001/console
```

2. Select **Clusters** from the home page or choose **Environment -> Clusters** from the Domain structure menu.
3. In the Change Center, click **Lock & Edit**.
4. Click on the Cluster Name (`cluster_portal`, `cluster_forms`, `cluster_reports`, `cluster_disco`).
5. Select **HTTP** and enter the following values:

| Parameter          | Value                |
|--------------------|----------------------|
| Frontend Host      | mysite.mycompany.com |
| Frontend HTTP Port | 80                   |

| Parameter           | Value |
|---------------------|-------|
| Frontend HTTPS Port | 443   |

This ensures that any HTTPS URLs created from within WebLogic are directed to port 443 or 80 on the load balancer.

6. Click **Activate Changes** in the Change Center window to save the changes.
7. Restart the WLS\_PORTAL, WLS\_FORMS, WLS\_REPORTS, WLS\_DISCO Managed Servers using the following steps:
  - a. Select **Servers** from the Home page or **Environment > Servers** from the **Domain structure** menu.
  - b. Select the **Control** tab.
  - c. Select the box next to each Managed Server.
  - d. Choose **Shutdown > Force Shutdown Now**.
  - e. Click **Yes** to shut down the Managed Server.
  - f. Once the Managed Server is shut down, select the box next to the Managed Server.
  - g. Click **Start**.
  - h. Click **Yes** to start the Managed Server.

#### 14.6.4.6 Configure Oracle Portal for High Availability

The following steps are required to configure Oracle Portal only.

**14.6.4.6.1 Rewire Portal Repository** Follow these steps to rewire the Oracle Portal repository:

1. Log into the domain using Oracle Fusion Middleware Enterprise Manager using the following URL:  
<http://apphost1.us.oracle.com:7001/em>
2. Expand the **Fusion Middleware** menu on the left hand side.
3. Expand the **Portal** menu (under **Fusion Middleware** menu).
4. Click **Portal**.
5. Right-click **Portal** and select settings **Wire Configuration**.
6. Enter the following information for Portal midtier:

| Parameter    | Value                                                                                                                       |
|--------------|-----------------------------------------------------------------------------------------------------------------------------|
| Host         | Enter the DNS name of the load balancer. For example: mysite.mycompany.com                                                  |
| Port         | Enter the port that the load balancer is listening on. For example, 443 for SSL Termination/Site Enabled SLL or 80 for HTTP |
| SSL Protocol | Select this if SSL Termination/Site Enabled SSL is being used.                                                              |

7. Enter the following information for Web Cache:

| Parameter              | Value                                                                                                                                         |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Host                   | Enter the DNS name of the load balancer. For example: mysite.mycompany.com                                                                    |
| Invalidation Port      | Enter the Portal Invalidation port as configured at the load balancer, for example: 9401                                                      |
| Invalidation User Name | Enter the user name to be used for Portal invalidations, for example: invalidator                                                             |
| Invalidation Password  | Password for the above account.<br>If you do not know the value of this password, it can be changed in Enterprise Manager as described above. |

8. Click **Apply** to start the rewire.
9. After the rewire completes, click the **Portal** menu option again.
10. Ensure that the Portal URL now shows:

`http://mysite.mycompany.com:80/portal/pls/portal`

**14.6.4.6.2 Configure Parallel Page Engine Loop-Back with Load Balancer** The purpose of the Parallel Page Engine (PPE Servlet) is to construct pages that have been requested by users. It does this by receiving the page request from a user, making its own new requests to fetch all the pieces of the page in parallel, assembling these pieces into a single page file and then sending the page content back to the end user (or back to the client browser).

These internal requests should be kept inside of the organization, and be served using the HTTP protocol.

The following steps are required only if you use a load balancer.

1. Log into the Oracle Fusion Middleware Enterprise Manager using this URL:  
`http://apphost1.us.oracle.com:7001/em`
2. Select **Fusion Middleware > Classic -> Portal** from the object browser on the left.
3. Right-click **Portal**, and select **Settings > Page Engine**.
4. In the Advanced Properties section, add the following information:

| Parameter  | Value                                                                                                                                        |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| UsePort    | Select the internal loopback port number, for example: 7777. This port will have been configured at the load balancer for internal requests. |
| Use Scheme | Change to the value of HTTP                                                                                                                  |

5. Click **Apply** to save the settings.
6. Restart the WebLogic Managed Server from the Oracle WebLogic Administration Console as follows:
  - a. Select **Servers** from the Home page or **Environment > Servers** from the **Domain structure** menu.
  - b. Select the **Control** tab.

- c. Select the box next to the **WLS\_PORTAL** Managed Server.
- d. Choose **Shutdown > Force Shutdown Now**.
- e. Click **Yes** to shut down the Managed Server.
- f. Once the Managed Server is shut down, select the box next to the Managed Server.
- g. Click **Start**.
- h. Click **Yes** to start the Managed Server.

**14.6.4.6.3 Database Wallets and Portal** If you are using SSL, Oracle Portal requires that the database in which the portal schema resides has a wallet. In this wallet is stored the certificate of the load balancer or the site. Once the wallet has been created, Oracle Portal must be informed of its location.

---

**Note:** This step is required only if you are using SSL Termination or Site SSL.

---

#### Create a Database Wallet

Before starting this process, copy the certificate to the database servers.

Each browser does this in a slightly different way. For Firefox browsers:

1. Use a browser to access the URL `https://mysite.mycompany.com`.
2. Go to **Firefox > Preferences > Advanced > Encryption > view certificates**.
3. Highlight the certificate for `mysite.mycompany.com` select export and give the file a name.
4. Copy this file to the database server.
5. Save the certificate if requested.
6. Now that you have obtained a copy of the certificate, create a wallet on each of the database servers, and import this certificate using the Oracle Wallet Manager from the database server. This must be done for all of the Oracle RAC nodes by following these steps:
  - a. Type **owm** to invoke the Oracle Wallet Manager.
  - b. Select **Wallet > New**.
  - c. Select **No** so that you do not create the wallet in the default location.
  - d. Enter a password for the wallet (remember this password - it is needed later)
  - e. Set the Wallet Type to **Standard**.
  - f. Select **No** to the question Do you want to create a certificate at this time?.
  - g. In Oracle Wallet Manager, select **Operations > Import Trusted certificate**.
  - h. Select **Select a file that contains the certificate** and click **OK**.
  - i. Select the certificate file selected above and click **Import**.
  - j. Select **Wallet** and **Save As**.
  - k. Select a location for the wallet. For example:

`ORACLE_BASE/admin/DB_NAME/wallet`

- I. Repeat this procedure for other Oracle RAC database nodes.

You must use the same directory paths on all database nodes.

### Identify the Wallet to Portal

Now that the certificate is stored inside the database wallet, store the location of the wallet within the portal repository by following these steps:

1. Run the SQL\*Plus script `secwc.sql`, which is located in the following directory

`ORACLE_HOME/portal/admin/plsql/wwc`

It may be necessary to create a database entry in the file `tnsnames.ora` located in `ORACLE_HOME/network/admin`:

```
SQL> @secwc 'file:$ORACLE_BASE/admin/DB_NAME/wallet' 'walletpassword'
```

In the command above:

- Use the absolute path to the wallet. Do not use environment variables.
- `walletpassword` is the password for the wallet.
- Use the path to the wallet directory, not the wallet file itself. The keyword `file` is required.

**14.6.4.6.4 Restart All Components** After making the changes in previous sections, the Web tier components must be restarted. This section describes the procedure for restarting the components.

### Restart Web Components

Restart the Oracle HTTP Server using these commands:

```
opmnctl stopall
opmnctl startall
```

### Restart WLS\_Portal

Restart the Managed Server `WLS_PORTAL` by logging into the WebLogic Administration Console and following these steps:

1. Select **Servers** from the Home page or **Environment > Servers** from the **Domain structure** menu.
2. Select the **Control** tab.
3. Select the box next to the `WLS_PORTAL` Managed Server.
4. Choose **Shutdown > Force Shutdown Now**.
5. Click **Yes** to shut down the Managed Server.
6. Once the Managed Server is shut down, select the box next to the `WLS_PORTAL` Managed Server.
7. Click **Start**.
8. Click **Yes** to start the Managed Server.

**14.6.4.6.5 Post-installation Step for Portal Installation with Oracle RAC** Oracle recommends setting initial-capacity for individual data sources created for Portal and Portlet components to **0**. Please use Administration Console to set initial-capacity to 0 for these data sources.

**14.6.4.6.6 Validate Configuration** To validate the configuration, perform the following tests:

| Test                                | URL                                           | Result                  |
|-------------------------------------|-----------------------------------------------|-------------------------|
| Test Portal using the load balancer | http://mysite.mycompany.com/portal/pls/portal | Portal Home page opens. |

### Troubleshooting Single Sign-On Errors

If a Single Sign-On error message appears on the Portal front screen at this stage in the configuration, the database wallet may not have been created properly, or Oracle Portal may not have been notified of its correct location.

If this error appears, repeat the steps to create the database wallet and the subsequent identification to portal.

### 14.6.4.7 Configure Oracle Forms for High Availability

Use the following steps to configure Oracle Forms only.

**14.6.4.7.1 Create TNSNAMES Entries for Customer Databases** If the application will access one or more databases, an entry for each database being accessed must be placed into the file `tnsnames.ora`.

to place an entry for each database into the file, follow these steps:

1. Edit the file `ORACLE_INSTANCE/config/tnsnames.ora` and add an entry similar to this:

```
mydb.mycompany.com =
(DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode1-vip) (PORT = 1521))
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode2-vip) (PORT = 1521))
 (LOAD_BALANCE = yes)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = mydb.mycompany.com)
)
)
```

---

---

**Note:** This is an Oracle RAC database connect string.

---

---

2. Save the file and test that the database connection is configured correctly using the command:

```
tnsping mydb.mycompany.com
```

---

---

**Note:** Set the environment variable `TNS_ADMIN` before issuing the **tnsping** command above.

---

---

**14.6.4.7.2 Restart WLS\_FORMS** Restart the Managed Server `WLS_FORMS` by logging into the WebLogic Administration Console and following these steps:

1. Select **Servers** from the Home page or **Environment > Servers** from the **Domain structure** menu.

2. Select the **Control** tab.
3. Select the box next to the WLS\_FORMS Managed Server.
4. Choose **Shutdown > Force Shutdown Now**.
5. Click **Yes** to shut down the Managed Server.
6. Once the Managed Server is shut down, select the box next to the WLS\_FORMS Managed Server.
7. Click **Start**.
8. Click **Yes** to start the Managed Server.

**14.6.4.7.3 Validate Configuration** To validate the configuration, the following tests should be performed:

| Test               | URL                                          | Result                        |
|--------------------|----------------------------------------------|-------------------------------|
| Test load balancer | http://mysite.mycompany.com/                 | Home page opens               |
| Forms              | http://mysite.mycompany.com/forms/frmservlet | Forms Servlet Home page opens |

#### 14.6.4.8 Configure Oracle Reports for High Availability

The following steps are required to configure Oracle Reports only.

**14.6.4.8.1 Create Reports Queue in Database** To maintain a consistent reports queue across multiple Reports server instances and to be resilient to the failure of a reports server, create the reports queue in a high availability Real Application Clusters database.

To create the Reports queue, run the SQL\*Plus script `rw_server.sql` against the database.

The script is located in this directory:

`ORACLE_HOME/reports/admin/sql`

Follow these steps:

1. Create a user to hold the report queue in the database using the following commands:

```
sql> create user report_queue identified by MYSpasword;
sql> grant connect, resource,create view to report_queue;
```

2. Connect to the Reports user and execute the following script:

```
cd ORACLE_HOME/reports/admin/sql
sqlplus reports_queue/MYpasswd@rw_server.sql
```

**14.6.4.8.2 Create a TNSNAMES Entry for Reports Queue** Oracle Reports uses entries in the `tnsnames.ora` file to determine database connection information. Place an entry in this file for the reports queue database by following these steps:

1. Edit the file `ORACLE_INSTANCE/config/tnsnames.ora` file and add an entry similar to this:

```
myrepq.mycompany.com =
 (DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP)(HOST = mydbnode1-vip)(PORT = 1521))
```

```

 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode2-vip) (PORT = 1521))
 (LOAD_BALANCE = yes)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = myrepq.mycompany.com)
)
)
)

```

---

**Note:** This is an Oracle RAC database connect string.

---

2. Save the file and test that the database connection is configured correctly using the following command:

```
tnsping myrepq.mycompany.com
```

---

**Note:** Set the environment variable TNS\_ADMIN before using the **tnsping** command.

---

**14.6.4.8.3 Create a Security Key for the Reports Queue** Oracle Reports security is performed using an indirect model. Before the reports server can be configured to use the database reports queue, make an entry in WebLogic to hold the reports queue password by following these steps:

1. Log into Oracle Fusion Middleware Enterprise Manager using the following URL, entering the WebLogic administrator user and password when prompted:

```
http://apphost1.mycompany.com:7001/em
```

2. In the navigation tree on the left hand side, expand **WebLogicDomain**, and click on the name of the domain, for example: `ReportsDomain`.

The ReportsDomain overview page appears.

3. From the drop-down menu at the top of the page select **Security > Credentials**.
4. Click **Create Key**.
5. Provide the following information:

| Parameter   | Value                                     |
|-------------|-------------------------------------------|
| Select Map  | reports                                   |
| Key         | queuePassword                             |
| Type        | Password                                  |
| User Name   | report_queue                              |
| Password    | Password for the reports queue account    |
| Description | Description for the reports queue account |

6. Click **OK** when finished.

**14.6.4.8.4 Configure the Database Job Repository for In-Process Reports Servers** After the reports queue has been placed into the database, the reports server needs to be told how to access it.



To do this, edit the `rwserver.conf` file.

The file is located in the following directory:

```
DOMAIN_HOME/config/fmwconfig/servers/WLS_REPORTS/applications/reports_
11.1.1.2.0/configuration/
```

There is a fixed order for entries in the server configuration file. Add the following element immediately after the `<jobStatusRepository>` element in the `rwserver.conf` file:

```
<jobRepository>
 <property name="dbuser" value="dbuser"/>
 <property name="dbpassword" value="csf:reports:dbpasswdKey"/>
 <property name="dbconn" value="dbconn"/>
</jobRepository>
```

where:

- *dbuser* is the name of the schema where the reports queue resides
- *dbconn* references a database entry in the `tnsnames.ora` file in the `ORACLE_INSTANCE/config` directory

For example:

```
<jobRepository>
 <property name="dbuser" value="report_queue"/>
 <property name="dbpassword" value="csf:reports:queuePassword"/>
 <property name="dbconn" value="myrepq.mycompany.com"/>
</jobRepository>
```

**14.6.4.8.5 Configure the Reports Server to Access Shared Output Directory** Add the `CacheDir` or `JOCCacheDir` property to the `<cache>` element the file `rwserver.conf` file, located in the following directory:

```
DOMAIN_HOME/config/fmwconfig/servers/WLS_REPORTS/applications/reports_
11.1.1.2.0/configuration/
```

For example:

```
<property name="JOCCacheDir" value="folder_name"/>
<property name="CacheDir" value="folder_name"/>
```

On UNIX:

```
<cache class="oracle.reports.cache.RWCache">
 <property name="cacheSize" value="50"/>
 <!--property name="cacheDir" value="your cache directory"/-->
 <!--property name="maxCacheFileNumber" value="max number of cache files"/-->
 <property name="JOCCacheDir" value="/share/reports"/>
 <property name="CacheDir" value="/share/reports"/>
</cache>
```

**14.6.4.8.6 Restart WLS\_REPORTS** To restart the Managed Server `WLS_REPORTS`, log into the WebLogic Administration Console, and follow these steps:

1. Select **Servers** from the Home page or **Environment > Servers** from the **Domain structure** menu.
2. Select the **Control** tab.

3. Select the box next to the WLS\_REPORTS Managed Server.
4. Choose **Shutdown > Force Shutdown Now**.
5. Click **Yes** to shut down the Managed Server.
6. Once the Managed Server is shut down, select the box next to the WLS\_REPORTS Managed Server.
7. Click **Start**.
8. Click **Yes** to start the Managed Server.

**14.6.4.8.7 Validate Configuration** Validate the initial Reports configuration by performing these tests:

| Test          | URL                                                                                                                       | Result                                             |
|---------------|---------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Reports Queue | <a href="http://mysite.mycompany.com/reports/rwervlet/showjobs">http://mysite.mycompany.com/reports/rwervlet/showjobs</a> | Single Sign On screen and then Reports Queue opens |

For an alternate test, download a sample report from Oracle Technology Network, and run it. The URL for Oracle Technology Network is:

<http://www.oracle.com/technology/index.html>

#### 14.6.4.9 Configure Oracle Discoverer for High Availability

Follow these steps to configure high availability for Oracle Discoverer only.

**14.6.4.9.1 Create TNSNAMES Entries for Customer Databases** If the application accesses one or more databases, place an entry for each database being accessed into the `tnsnames.ora` file, by following these steps:

1. Edit the file `ORACLE_INSTANCE/config/tnsnames.ora`, and add an entry similar to this:

```
mydb.mycompany.com =
(DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode1-vip) (PORT = 1521))
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode2-vip) (PORT = 1521))
 (LOAD_BALANCE = yes)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = mydb.mycompany.com)
)
)
```

---

**Note:** This is an Oracle RAC database connect string.

---

2. Save the file and test that the database connection is configured correctly using the following command:

```
tnsping mydb.mycompany.com
```

---

**Note:** Set the environment variable `TNS_ADMIN` before issuing the `tnsping` command.

---

**14.6.4.9.2 Update configuration.xml** The configuration.xml stores information about the Discoverer configuration. The file is located in the following directory:

```
DOMAIN_HOME/config/fmwconfig/servers/WLS_DISCO/applications/discoverer_
version/configuration/
```

Update the line beginning with `applicationURL=` and change the URL to:

```
http://mysite.mycompany.com/discoverer
```

For example:

```
applicationURL="http://mysite.mycompany.com/discoverer">
```

**14.6.4.9.3 Discoverer Viewer and Web Cache** By default, Discoverer Viewer is not configured to make full use of Oracle Web Cache. When enabled, significant performance gains can be attained. However, it is not always appropriate to enable this functionality.

For details on when and how to enable Discoverer Viewer with Oracle Web Cache, see the "Using Discoverer Viewer with Oracle Web Cache" section of *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Discoverer*.

**14.6.4.9.4 Enable Single Sign On** By default, Discoverer is not protected by single sign-on. To secure Discoverer Plus and Viewer, use the following steps:

1. Edit the file `mod_osso.conf` file, located in the following directory:

```
ORACLE_INSTANCE/config/OHS/ohs1/moduleconf
```

2. Add the following lines to the file before the line that begins with `</IfModule>`:

```
<Location /discoverer/plus>
 require valid-user
 AuthType Osso
</Location>

<Location /discoverer/viewer>
 require valid-user
 AuthType Osso
</Location>

<Location /discoverer/app>
 require valid-user
 AuthType Osso
</Location>
```

**14.6.4.9.5 Restart All Components** After making the changes in previous sections, restart the Web tier components using the procedures described in this section

### Restart Web Components

Restart the Oracle HTTP Server using the following commands:

```
opmnctl stopall
opmnctl startall
```

**Restart WLS\_DISCO**

Restart the Managed Server WLS\_DISCO by logging into the WebLogic Administration Console and following these steps:

1. Select **Servers** from the Home page or **Environment > Servers** from the **Domain structure** menu.
2. Select the **Control** tab.
3. Select the box next to the WLS\_DISCO Managed Server.
4. Choose **Shutdown > Force Shutdown Now**.
5. Click **Yes** to shut down the Managed Server.
6. Once the Managed Server is shut down, select the box next to the WLS\_DISCO Managed Server.
7. Click **Start**.
8. Click **Yes** to start the Managed Server.

**14.6.4.9.6 Validate Configuration** Validate the initial Discoverer configuration by performing these tests:

| Test               | URL                                           | Result                                          |
|--------------------|-----------------------------------------------|-------------------------------------------------|
| Test load balancer | http://mysite.mycompany.com/                  | Home page opens                                 |
| Discoverer         | http://mysite.mycompany.com/discoverer/viewer | Single Sign-On and then Discoverer Viewer opens |

**14.6.5 Install and Configure Application Tier on APPHOST2**

The following sections describe how to install and configure the Application Tier on APPHOST2.

- [Section 14.6.5.1, "Install Oracle WebLogic Server"](#)
- [Section 14.6.5.2, "Install Oracle Portal, Forms, Reports, and Discoverer Software"](#)
- [Section 14.6.5.3, "Configure Oracle Portal, Forms, Reports, and Discoverer Software"](#)
- [Section 14.6.5.4, "Generic Configuration"](#)
- [Section 14.6.5.5, "Configure Oracle Portal for High Availability"](#)
- [Section 14.6.5.6, "Configure Oracle Forms for High Availability"](#)
- [Section 14.6.5.7, "Configure Oracle Reports for High Availability"](#)
- [Section 14.6.5.8, "Configure Oracle Discoverer for High Availability"](#)

**14.6.5.1 Install Oracle WebLogic Server**

To install Oracle WebLogic Server binaries, see the *Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server*.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the version of Oracle WebLogic Server to use with the latest version of Oracle Fusion Middleware.

### 14.6.5.2 Install Oracle Portal, Forms, Reports, and Discoverer Software

To install the Oracle Portal, Forms, Reports, and Discoverer binaries, see *Oracle Fusion Middleware Installation Guide for Oracle Portal, Forms, Reports and Discoverer*.

---

**Note:** Before you run the Configuration Wizard by following the instructions in [Section 14.6.5.3, "Configure Oracle Portal, Forms, Reports, and Discoverer Software,"](#) ensure that you applied the latest Oracle Fusion Middleware patch set and other known patches to your Middleware Home.

See "Understanding Your Installation Starting Point" in *Oracle Fusion Middleware Installation Planning Guide* for the steps you must perform to get the latest version of Oracle Fusion Middleware.

---

### 14.6.5.3 Configure Oracle Portal, Forms, Reports, and Discoverer Software

Start the Oracle Fusion Middleware Configuration Wizard by running this command:

```
ORACLE_HOME/bin/config.sh
```

---

**Note:** Before starting the configuration ensure that the following environment variables (UNIX) are not set:

- LD\_ASSUME\_KERNEL
  - ORACLE\_BASE
  - LD\_LIBRARY\_PATH
- 

Continue with the following steps:

1. In the Welcome screen, click **Next**.
2. In the Prerequisite Checks screen, once all the prerequisites have been validated, click **Next**.
3. In the Create Domain screen, select **Expand Cluster** and enter these values:
  - **Host Name:** Name of the host running the WebLogic Administration Server, for example: APPHOST1.mycompany.com
  - **Port:** The Administration Server port. For example: 7001
  - **User Name:** Administration Server administrator account name.
  - **User Password:** Password for the Administration Server administrator account.

Click **Next**.

4. In the Specify Security Updates screen, if required, enter a username and password to receive Oracle Security updates for Oracle Support.
5. In the Specify Installation Location screen, enter these values:
  - **WebLogic Server Location:** Enter the installation directory for Oracle WebLogic Server. This should be MW\_HOME/wlserver\_10.3, for example:  
/u01/app/oracle/product/FMW/wlserver\_10.3

- **Oracle Instance Location:** Enter the directory where the Oracle configuration files will be placed. This should be outside of the Oracle home directory. This directory will be known as the ORACLE\_INSTANCE. For example:

```
/u01/app/oracle/admin/fmw2
```

- **Oracle Instance Name:** `fmw2`

Click **Next**.

6. In the Configure Components screen, choose which products to install and configure. If placing Oracle Web Cache or the Oracle HTTP Server onto this node, ensure that they are selected.

---

**Note:** This should be the same list as selected on APPHOST1.

---

Click **Next**.

7. Select the same `staticports.ini` file used for APPHOST1.
8. In the Specify Application Identity Store screen, specify the following values:
  - **Hostname:** Name of the Oracle Internet Directory server, for example: `login.myoid.com`
  - **Port:** Oracle Internet Directory port, for example: `389`
  - **User Name:** `cn=orcladmin`
  - **Password:** Password for Oracle Internet Directory's `orcladmin` account
9. In the Summary screen, click **Install** to begin the creation process.

#### 14.6.5.4 Generic Configuration

The following steps are required regardless of the component you are configuring.

**14.6.5.4.1 Copy Configuration Information from APPHOST1** After the Expand Cluster operation has created a new WebLogic Managed Server and associated machine, it is necessary to copy some configuration information from APPHOST1 to APPHOST2.

Copy the following files located on APPHOST1 to the APPHOST2 locations shown in the following table:

| File                         | APPHOST1 Location                                       | APPHOST2 Location                                       |
|------------------------------|---------------------------------------------------------|---------------------------------------------------------|
| <code>mod_oradav.conf</code> | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> |
| <code>mod_osso.conf</code>   | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> |
| <code>plsql.conf</code>      | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> |
| <code>portal.conf</code>     | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> | <code>ORACLE_INSTANCE/config/OHS/ohs1/moduleconf</code> |

| File               | APPHOST1 Location                                             | APPHOST2 Location                                             |
|--------------------|---------------------------------------------------------------|---------------------------------------------------------------|
| forms.conf         | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf |
| reports_ohs.conf   | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf |
| module_disco.conf  | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf |
| virtual_hosts.conf | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/moduleconf |
| osso.conf          | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/           | <i>ORACLE_</i><br><i>INSTANCE</i> /config/OHS/ohs1/           |
| sqlnet.ora         | <i>ORACLE_INSTANCE</i> /config                                | <i>ORACLE_INSTANCE</i> /config                                |
| tnsnames.ora       | <i>ORACLE_INSTANCE</i> /config                                | <i>ORACLE_INSTANCE</i> /config                                |

**14.6.5.4.2 Configure Virtual Hosts** For the site to function properly with the load balancer, you must add two virtual hosts. You can configure the virtual hosts using Oracle Enterprise Manager Fusion Middleware Control, edit the file *ORACLE\_INSTANCE*/config/OHS/ohs1/httpd.conf, or edit the *virtual\_hosts.conf* file located in the *ORACLE\_INSTANCE*/config/OHS/ohs1/moduleconf directory (copied from APPHOST1).

In a text editor, update the file to change APPHOST1 to APPHOST2:

```
NameVirtualHost *:7778
<VirtualHost *:7778>
 ServerName http://mysite.mycompany.com:80
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>

<VirtualHost *:7778>
 ServerName apphost2.mycompany.com:7777
 RewriteEngine On
 RewriteOptions inherit
 UseCanonicalName On
</VirtualHost>
```

**14.6.5.4.3 Update Oracle HTTP Server Configuration to be Cluster Aware** When the installation was first created, it was configured so that all Web Logic requests were directed to the initially created Managed Servers residing on APPHOST1. Now that APPHOST2 exists, both the Oracle HTTP Servers on APPHOST1 and APPHOST2 must be made aware of all of the Oracle WebLogic Managed Servers.

To make the Oracle HTTP Server aware of APPHOST2, edit files located in the following directory:

*ORACLE\_INSTANCE*/config/OHS/ohs1/moduleconf

There are different files to edit for Portal, Forms, Reports, and Discoverer, as shown in the following table:

| Product    | File              |
|------------|-------------------|
| Portal     | portal.conf       |
| Forms      | forms.conf        |
| Reports    | reports_ohs.conf  |
| Discoverer | module_disco.conf |

Edit the files as follows:

1. Remove the lines that begin with `WebLogicHost` and `WebLogicPort` and add the following lines:

```
WebLogicCluster apphost1:9001,apphost2:9001
```

2. Change any lines that begin with `WebLogicCluster` to look similar to this:

```
WebLogicCluster apphost1:9001,apphost2:9001
```

For example, change the following:

```
<Location /portal>
 SetHandler weblogic-handler
 WebLogicHost apphost1.mycompany.com
 WebLogicPort PORT
</Location>
```

to this:

```
<Location /portal>
 SetHandler weblogic-handler
 WebLogicCluster apphost1.mycompany.com:PORT,apphost2.mycompany.com:PORT
</Location>
```

Change the following:

```
<Location /Forms>
 SetHandler weblogic-handler
 WebLogicCluster apphost1.mycompany.com:PORT
</Location>
```

to this:

```
<Location /Forms>
 SetHandler weblogic-handler
 WebLogicCluster apphost1.mycompany.com:PORT,apphost2.mycompany.com:PORT
</Location>
```

Repeat these steps on APPHOST2.

**14.6.5.4.4 Change the Web Cache Passwords** Web Cache passwords are randomly generated, however they are required in later stages. It is therefore recommended that the Web Cache passwords be changed from the default value to a new known value.

To change the default password, follow these steps:

1. In the Navigator window, expand the **Web Tier** tree.
2. Click on the component `wc1`.
3. From the drop down list at the top of the page, select **Administration - Passwords**.



4. Enter new invalidation and administration passwords, confirm them, and click **Apply**.

---

**Note:** The passwords used in this procedure must be the same as those used on APPHOST1. This is required for the web caches to be clustered together in later procedures.

---

**14.6.5.4.5 Configure Web Cache** Web Cache is optional for Oracle Forms, Reports and Discoverer, but mandatory for Oracle Portal. If Web Cache was not configured for APPHOST1, ignore the following Web Cache configuration steps:

#### Log into the Enterprise Manager Console

Log into the Enterprise Manager Console using the following URL:

`http://apphost1.mycompany.com:7001/em`

The default User Name and Password are the same as the WebLogic domain user name and password entered during the installation.

#### Create the Origin Server

To create the origin server:

1. In the Navigator Window, expand the Web Tier tree.
2. Click on the component `wc1` (make sure it is the one associated with APPHOST1).
3. From the drop down list at the top of the page select **Administration - Origin Servers**.
4. Click **Create**.
5. Enter the following information to add the origin server:

| Field              | Value                  |
|--------------------|------------------------|
| Host               | APPHOST2.mycompany.com |
| Port               | 7778                   |
| Capacity           | 100                    |
| Protocol           | HTTP                   |
| Failover Threshold | 5                      |
| Ping URL           | /                      |
| Ping Interval      | 10                     |

6. Click **OK** to save the changes.
7. Click **Apply** to save the changes.

#### Add Origin Server to Existing Site to Server Mapping

To add the origin server site to server mapping:

1. In the Navigator Window, expand the Web Tier tree.
2. Click on the component `wc1` (make sure it is the one associated with APPHOST1).
3. From the drop down list at the top of the page select **Administration - Sites**.

4. In the Site to Server Mapping section, click on the Host:Port, for example:  
`mysite.mycompany.com:80`
5. Click **Edit**.
6. Select the origin server `APPHOST2.mycompany.com:7778` and move it to the selected Origin servers list.
7. Click **OK** to save the changes.
8. Click **Apply** to save the changes.

### Cluster Web Cache on APPHOST1 and APPHOST2

Follow these steps to cluster Oracle Web Cache on APPHOST1 and APPHOST2:

---

**Note:** For Oracle Web Cache clustering to function properly, the administration password of each of the Web Caches clustered together must be the same.

---

1. In the Navigator Window, expand the Web Tier tree.
2. Click on the component `wc1` (make sure it is the one associated with APPHOST1).
3. From the drop down list at the top of the page select **Administration - Cluster**.
4. Click **Add**.  
The Web Cache from APPHOST2 is automatically added.
5. Click **Apply** to apply the changes.
6. Click on the newly created Web Cache entry (be sure not to click on the URL part of the entry).
7. Click **Synchronize** to copy the configuration to the Web Cache on APPHOST2.
8. Click **Yes** when prompted to confirm that you wish you perform the operation.
9. Click **Apply** to apply the new configuration.

**14.6.5.4.6 Restart Web Processes on APPHOST1 and APPHOST2** After making the previous changes, restart the Web Tier components on APPHOST1 and APPHOST2 using the following commands on each of the servers:

```
opmnctl stopall
opmnctl startall
```

### 14.6.5.5 Configure Oracle Portal for High Availability

This section describes the procedure for configuring only Oracle Portal for high availability.

**14.6.5.5.1 Copy Configuration Information from APPHOST1** Even though the Expand Cluster has created a new WebLogic Managed Server and associated machine, it is still necessary to copy some configuration information from APPHOST1 to APPHOST2.

Copy the following files located on APPHOST1 to the APPHOST2 locations shown in the following table:

| File or Directory  | APPHOST1 Location                                                                              | APPHOST2 Location                                                                                   |
|--------------------|------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| ORACLE_HOME/portal | ORACLE_HOME/portal                                                                             | ORACLE_HOME/portal                                                                                  |
| appConfig.xml      | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>PORTAL/applications/portal/configurati<br>on/ | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS<br>-<br>PORTAL1/applications/portal/configura<br>tion/ |
| portal_cache.conf  | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>PORTAL/applications/portal/configurati<br>on/ | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS<br>-<br>PORTAL1/applications/portal/configura<br>tion/ |
| portal_dads.conf   | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>PORTAL/applications/portal/configurati<br>on/ | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS<br>-<br>PORTAL1/applications/portal/configura<br>tion/ |
| portal_plsql.conf  | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS_<br>PORTAL/applications/portal/configurati<br>on/ | DOMAIN_<br>HOME/config/fmwconfig/servers/WLS<br>-<br>PORTAL1/applications/portal/configura<br>tion/ |

**14.6.5.5.2 Create Portal Directories** Create the following directories on APPHOST2 to allow the storage of the Oracle Portal Cache:

```
ORACLE_INSTANCE/portal/cache
ORACLE_INSTANCE/diagnostics/logs/portal
```

**14.6.5.5.3 Update Instance Paths** Edit the hard coded entries in the following two copied files to reflect the paths above.

- In the portal\_cache.conf file, change PlsqlCacheDirectory.
- In the portal\_plsql.conf file, change PlsqlLogDirectory.

These files are located in the DOMAIN\_HOME/config/fmwconfig/servers/WLS\_PORTAL1/applications/portal/configuration directory.

**14.6.5.5.4 Restart the Web Processes** Restart the Web Tier components on APPHOST2 using the following commands on each of the servers:

```
opmnctl stopall
opmnctl startall
```

**14.6.5.5.5 Start WLS\_PORTAL1** Now that the application files have been copied across. Start the managed server, WLS\_PORTAL1:

1. Log into the Administration Server on APPHOST1 using the URL:  
`http://APPHOST1.mycompany.com:7001/console`
2. Provide the WebLogic Administration Console login credentials.
3. Select **Environment -> Servers** from the Domain structure menu.
4. Select the **Control** tab.

5. Select the box next to the Managed Server WLS\_PORTAL1, and click **Shutdown - Force Shutdown Now**.
6. Click on **Yes** to confirm the operation.

This resets the server's status.

#### 14.6.5.5.6 Validate the Configuration To validate the configuration:

| Test                                | URL                                            | Result                                           |
|-------------------------------------|------------------------------------------------|--------------------------------------------------|
| Test Portal via Load Balancer       | http://mysite.mycompany.com/portals/pls/portal | Portal Home Page Displayed                       |
| Test Portal Login via Load Balancer | http://mysite.mycompany.com/portals/pls/portal | Should be able to log in using account orcladmin |

**14.6.5.5.7 Best Practices** By default configuration information is not automatically propagated to all server instances. If configuration files are changed on one portal server, propagate the configuration to all of the servers.

#### 14.6.5.6 Configure Oracle Forms for High Availability

The section describes the procedure for configuring only Oracle Forms for high availability.

**14.6.5.6.1 Create a TNSNAMES entries for Customer Databases** If the application is to access one or more databases, place an entry for each database being accessed into the file `tnsnames.ora`.

Edit the file `ORACLE_INSTANCE/config/tnsnames.ora`, and add an entry similar to the one below:

```
mydb.mycompany.com =
(DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP)(HOST = mydbnode1-vip)(PORT = 1521))
 (ADDRESS = (PROTOCOL = TCP)(HOST = mydbnode2-vip)(PORT = 1521))
 (LOAD_BALANCE = yes)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = mydb.mycompany.com)
)
)
```

This is an Oracle RAC database connect string.

**14.6.5.6.2 Copy Forms Configuration Files** Even though the expand cluster has created a new WebLogic Managed Server and associated machine. It is still necessary to copy some configuration information from APPHOST1 to APPHOST2.

Copy the following files located on APPHOST1:

| File | Location of APPHOST1                                                                            | Location of APPHOST2                                                                             |
|------|-------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| ALL  | <code>DOMAIN_HOME/config/fmwconfig/servers/WLS_FORMS/applications/formsapp_11.1.1/config</code> | <code>DOMAIN_HOME/config/fmwconfig/servers/WLS_FORMS1/applications/formsapp_11.1.1/config</code> |

| File | Location of APPHOST1                                | Location of APPHOST2                        |
|------|-----------------------------------------------------|---------------------------------------------|
| ALL  | ORACLE_<br>INSTANCE/config/FormsCom<br>ponent/forms | ORACLE_INSTANCE/config/FormsComponent/forms |
| ALL  | ORACLE_<br>INSTANCE/config/FRCompon<br>ent/frcommon | ORACLE_INSTANCE/config/FRComponent/frcommon |

**14.6.5.6.3 Update default.env** Having copied the above files, update the file `default.env`, located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_FORMS1/applications/formsapp_11.1.1/config` directory with the correct values for APPHOST2. In particular, the following entries must be changed:

- ORACLE\_INSTANCE
- TNS\_ADMIN
- FORMS\_PATH
- WEBUTIL\_CONFIG

Typically these entries have the following values:

```
ORACLE_INSTANCE=/u01/app/oracle/admin/Forms1
TNS_ADMIN=/u01/app/oracle/admin/Forms1/config
FORMS_
PATH=/u01/app/oracle/product/FMW/Forms/forms:/u01/app/oracle/admin/Forms1/FormsCom
ponent/forms
WEBUTIL_
CONFIG=/u01/app/oracle/admin/Forms1/config/FormsComponent/forms/server/webutil.cfg
```

And must be changed to:

```
ORACLE_INSTANCE=/u01/app/oracle/admin/Forms2
TNS_ADMIN=/u01/app/oracle/admin/Forms2/config
FORMS_
PATH=/u01/app/oracle/product/FMW/Forms/forms:/u01/app/oracle/admin/Forms2/FormsCom
ponent/forms
WEBUTIL_
CONFIG=/u01/app/oracle/admin/Forms2/config/FormsComponent/forms/server/webutil.cfg
```

**14.6.5.6.4 Restart WLS\_FORMS1** After making the changes in the previous section, restart the WebLogic managed servers WLS\_FORMS and WLS\_FORMS1 by logging into the Oracle WebLogic Administration Console using the following URL

`http://apphost1.mycompany.com:7001/console`

1. Select **Environment -> Servers** from the Domain Structure menu.
2. Select the **Control** tab.
3. Select the boxes next to the server WLS\_FORMS1.
4. Select **Shutdown -> Force shutdown now**.
5. Select **Yes** when the confirmation dialogue opens.
6. When the server is shutdown, restart it by selecting the boxes next to the server WLS\_FORMS1.
7. Click **Start**.

**14.6.5.6.5 Validate the Configuration** To validate the configuration, use the following tests:

| Test               | URL                                          | Result                             |
|--------------------|----------------------------------------------|------------------------------------|
| Test Load Balancer | http://mysite.mycompany.com/                 | AS Home Page Displayed             |
| Forms              | http://mysite.mycompany.com/forms/frmservlet | Forms Servlet Home Page displayed. |

**14.6.5.6.6 Best Practices** WebLogic Application Server does not replicate configuration information automatically between nodes. It is important to manually propagate any changes to any of the following to the other servers in the deployment:

- ORACLE\_INSTANCE
- DOMAIN\_HOME/config/fmwconfig/servers/WLS\_FORMS/applications/formsapp\_11.1.1/config

### 14.6.5.7 Configure Oracle Reports for High Availability

This section describes the procedures for configuring only Oracle Reports for high availability.

**14.6.5.7.1 Create TNSNAMES Entries for Customer Databases** Oracle Reports uses entries in the `tnsnames.ora` file to determine database connection information. Place an entry into this file for the Oracle Reports queue database.

Edit the file `ORACLE_INSTANCE/config/tnsnames.ora` and add an entry similar to the following:

```
mydb.mycompany.com =
(DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode1-vip) (PORT = 1521))
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode2-vip) (PORT = 1521))
 (LOAD_BALANCE = yes)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = mydb.mycompany.com)
)
)
```

Save the file and test that the database connection is configured correctly using the following command:

```
tnsping myreportq.mycompany.com
```

**14.6.5.7.2 Configure the Reports Server to Access Shared Output Directory** Add the `CacheDir` or `JOCCacheDir` property to the `<cache>` element of each of the `rwserver.conf` files, which are located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_REPORTS1/applications/reports_11.1.1.2.0/configuration/` directory.

For example:

```
<property name="JOCCacheDir" value="folder_name"/>
<property name="CacheDir" value="folder_name"/>
```

For UNIX

```
<cache class="oracle.reports.cache.RWCache">
 <property name="cacheSize" value="50"/>
```

```

<!--property name="cacheDir" value="your cache directory"/-->
<!--property name="maxCacheFileNumber" value="max number of cache files"/-->
<property name="JOCCacheDir" value="/share/reports"/>
<property name="CacheDir" value="/share/reports"/>
</cache>

```

**14.6.5.7.3 Configure the Database Job Repository for In-process Reports Servers** Now that the reports queue has been placed into the database, notify the Reports server how to access it using the by editing the file `rwserver.conf` file, located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_REPORTS1/applications/reports_11.1.1.2.0/configuration` directory.

When editing the file, the order in which different entries appear inside the server configuration file is fixed. As a result, the following element must be added immediately after the `<jobStatusRepository>` element in the server configuration file:

```

<jobRepository>
 <property name="dbuser" value="dbuser"/>
 <property name="dbpassword" value="csf:reports:dbpasswdKey"/>
 <property name="dbconn" value="dbconn"/>
</jobRepository>

```

Where `dbuser` is the name of the schema where the reports queue resides.

Also, `dbconn` references a database entry in the `tnsnames.ora` file located in `ORACLE_INSTANCE/config` directory. `dbpasswdKey` is the name of the entry created.

For example:

```

<jobRepository>
 <property name="dbuser" value="reports_queue"/>
 <property name="dbpassword" value="csf:reports:queuePassword"/>
 <property name="dbconn" value="myrepq.mycompany.com"/>
</jobRepository>

```

The value of `clusternodes` is the value that appears in the `<server>` tag in the file `rwervlet.properties` file on `APPHOST1`.

The `clusternodes` parameter should list all of the reports servers in the cluster (comma separated) **EXCEPT** the local report server.

**14.6.5.7.4 Creating an Oracle Reports Server Cluster** By creating a Reports cluster with a database reports queue it is possible to link all of the Reports servers to the same queue. The benefit of this procedure is that when a server has spare capacity, it can execute the next report in the queue, distributing the load. It also ensures that if a cluster member becomes unavailable, another Reports server can detect this and run any reports on which the failed server was working.

Create a Reports cluster by adding a cluster entry to the `rwervlet.properties` file on both `APPHOST1` and `APPHOST2`.

#### Cluster APPHOST1

Edit the `rwervlet.properties` file located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_REPORTS/applications/reports_11.1.1.2.0/configuration` directory.

Add the following line:

```

<cluster clustername="cluster_reports" clusternodes="rep_wls_reports1_APPHOST2_

```

```
reports2"/>
```

---

**Note:** The value of `clusternodes` is the value which appears in the `<server>` tag in the `rwsvrlet.properties` file located on `APPHOST2`.

---

---

**Note:** The `clusternodes` parameter should list all of the Reports servers in the cluster (comma separated) EXCEPT the local Reports server.

---

### Cluster APPHOST2

Edit the `rwsvrlet.properties` file located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_REPORTS1/applications/reports_11.1.1.2.0/configuration` directory.

Add the following line:

```
<cluster clustername="cluster_reports" clusternodes="rep_wls_reports_APPHOST1_reports1"/>
```

---

**Note:** The value of `clusternodes` is the value which appears in the `<server>` tag in the `rwsvrlet.properties` file located on `APPHOST1`.

---

---

**Note:** The `clusternodes` parameter should list all of the Reports servers in the cluster (comma separated) EXCEPT the local Reports server.

---

**14.6.5.7.5 Restart WLS\_REPORTS and WLS\_REPORTS1** Restart the Web Logic managed servers `WLS_REPORTS` and `WLS_REPORTS1` by logging into the WebLogic Administration Console using the following URL:

`http://apphost1.mycompany.com:7001/console`

1. Select **Environment -> Servers** from the Domain Structure menu.
2. Select the **Control** tab.
3. Select the boxes next to the server `WLS_REPORTS` and `WLS_REPORTS1`.
4. Select **Shutdown -> Force shutdown now**.
5. Select **Yes** when the confirmation dialogue opens.
6. When the server is shutdown, restart it by selecting the boxes next to the server `WLS_REPORTS` and `WLS_REPORTS1`.
7. Click **Start**.

**14.6.5.7.6 Validate the Configuration** To validate the configuration, run the following tests:



| Test               | URL                                                    | Result                                               |
|--------------------|--------------------------------------------------------|------------------------------------------------------|
| Test Load Balancer | http://mysite.mycompany.com/                           | AS Home Page Displayed                               |
| Test Reports Queue | http://mysite.mycompany.com/reports/rwservlet/showjobs | Single Sign On and Reports Queue pages are displayed |

**14.6.5.7.7 Managing Connection Availability for Oracle Reports Services** You must tune the time out settings to manage idle connections for the Load Balancing Router, firewall, Oracle HTTP Server, and WebLogic Server according to their Oracle Reports Services application requirements. For example, if the Load Balancing Router and firewall connection timeout is set to ten minutes, and the execution time for a report exceeds ten minutes, then the browser connection is timed out by the Load Balancing Router, and the Oracle Reports Services output does not reach the browser. The time out value for the connection between Oracle Reports Services clients and the Oracle Reports server is governed by the `idleTimeout` attribute of the connection element in the Oracle Reports Services Server configuration file `rwserver.conf`.

### 14.6.5.8 Configure Oracle Discoverer for High Availability

This section describes the procedures for configuring high availability for only Oracle Discoverer.

**14.6.5.8.1 Create TNSNAMES Entries for Customer Databases** If the application is to access one or more databases, place an entry for each database being accessed into the `tnsnames.ora` file.

Edit the file `ORACLE_INSTANCE/config/tnsnames.ora` and add an entry as follows:

```
mydb.mycompany.com =
 (DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode1-vip) (PORT = 1521))
 (ADDRESS = (PROTOCOL = TCP) (HOST = mydbnode2-vip) (PORT = 1521))
 (LOAD_BALANCE = yes)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = mydb.mycompany.com)
)
)
```

This is an Oracle RAC database connect string.

Save the file and test that the database connection is configured correctly using the following command:

```
tnsping myreportq.mycompany.com
```

**14.6.5.8.2 Copy Discoverer Configuration Files.** The expand cluster procedure has created a new WebLogic Managed Server and associated machine. However, you must still copy some configuration information from APPHOST1 to APPHOST2.

Prior to performing this step make a copy of the `configuration.xml` file located in the `DOMAIN_HOME/config/fmwconfig/servers/WLS_DISCO/applications/discoverer_version/configuration/` directory.

Copy this file from APPHOST1 to the same directory in APPHOST2.

**14.6.5.8.3 Update configuration.xml** Update the `configuration.xml` file you just copied, replacing the following values with the values in the `configuration.xml.orig` file.

```
<oracleInstance>
<discovererComponentName>
```

**14.6.5.8.4 Changing the Preference Store** In an enterprise deployment, there should only be one Discoverer preference store active at one time. To change the preference store, change Oracle Discoverer on APPHOST2 to point to the preference store on APPHOST1.

#### Identify Preference Store Host Name and Port

The nominated preference server in this example is configured to run on APPHOST1. When installed onto the server, the Preference server is assigned a port. Locate this value before proceeding to the following next steps:

1. On APPHOST1, open the `opmn.xml` file is located in the `ORACLE_INSTANCE/config/OPMN/opmn` directory.
2. Search the file for the entry `PREFERENCE_PORT` the value= shows the port assigned to the preference server.

#### Specifying a Discoverer Preferences Server on other Machines

Having identified the host name and port number of the machine that is going to run the Preferences component (that is, the Discoverer Preferences server machine), you must now ensure that other machines use the Preferences component on the Discoverer Preferences server machine.

To modify the `opmn.xml` file of other machines to use the Discoverer Preferences server machine, perform the following steps on every other machine in the installation.

1. On each machine except the Preferences server machine, open the `opmn.xml` file in a text editor (or XML editor).
2. Locate the `PREFERENCE_HOST` variable, and change its value to the host name of the Discoverer Preferences server machine, as follows:

```
<variable id="PREFERENCE_HOST" value="hostname">
```

3. Locate the `PREFERENCE_PORT` variable, and change its value to the port number of the Discoverer Preferences server machine, as follows:

```
<variable id="PREFERENCE_PORT" value="port">
```

4. Locate the `PreferenceServer` process type, and change its status to disabled, as follows:

```
<process-type id="PreferenceServer" module-id="Disco_PreferenceServer"
working-dir="$DC_LOG_DIR" status="disabled">
```

5. Save the `opmn.xml` file.

**14.6.5.8.5 Restart WLS\_DISCO and WLS\_DISCO1** Restart the Web Logic managed servers `WLS_DISCO` and `WLS_DISCO1` by logging into the WebLogic Administration Console using the following URL:

`http://apphost1.mycompany.com:7001/console`

1. Select **Environment -> Servers** from the Domain Structure menu.
2. Select the **Control** tab.
3. Select the boxes next to the server WLS\_DISCO and WLS\_DISCO1.
4. Select **Shutdown -> Force shutdown now**.
5. Select **Yes** when the confirmation dialogue opens.
6. When the server is shutdown, restart it by selecting the boxes next to the server WLS\_DISCO and WLS\_DISCO1.
7. Click **Start**.

**14.6.5.8.6 Validate the Configuration** Validate the configuration using the following tests:

| Test               | URL                                           | Result                                  |
|--------------------|-----------------------------------------------|-----------------------------------------|
| Test Load Balancer | http://mysite.mycompany.com/                  | AS Home Page Displayed                  |
| Discoverer         | http://mysite.mycompany.com/discoverer/viewer | Discoverer Servlet Home Page displayed. |

**14.6.5.8.7 Failover of the Preference Server** If the server hosting the preference server fails, the preference server must be started on one of the surviving nodes using the following procedure:

#### Identify Preference Store Host Name and Port

The nominated preference server in this example is configured to run on APPHOST1. When installed onto the server, the Preference server is assigned a port. Locate this value using the following procedure:

1. On APPHOST1, open the opmn.xml file is located in the *ORACLE\_INSTANCE/config/OPMN/opmn* directory.
2. Search the file for the entry PREFERENCE\_PORT the value= shows the port assigned to the preference server.

#### Enable the Preference Store on APPHOST2

Now that APPHOST2 is using the preference store on APPHOST1, disable the preference store on APPHOST2 by editing the opmn.xml file located in the *ORACLE\_INSTANCE/config/OPMN/opmn* directory

Locate the following line:

```
<process-type id="PreferenceServer" module-id="Disco-PreferenceServer"
working-dir="$DC_LOG_DIR" status="disabled">
```

Change status=disabled to **status=enabled**.

For example:

```
<process-type id="PreferenceServer" module-id="Disco-PreferenceServer"
working-dir="$DC_LOG_DIR" status="enabled">
```

#### Start the Preference Server on APPHOST2

Start the preference server on APPHOST2 using the following command:

```
opmnctl startproc process-type=PreferenceServer
```

**Change Surviving Servers to Use the Preference Store on APPHOST2**

Now that the preference server has been started on APPHOST2, amend all discoverer instances to use this preference server, including APPHOST2.

The preference store being used is determined at the startup of the WebLogic Managed Server WLS\_DISCO1. In order to get WLS\_DISCO1 to use a different preference store the startup parameters must be changed, by logging into the WebLogic Administration Console using the following URL:

`http://apphost1.mycompany.com:7001/console`

Continue with the following steps:

1. Select **Environment -> Servers** from the Domain Structure menu.
2. In the Change Center, click **Lock & Edit**.
3. Click on the server WLS\_DISCO1
4. Click on the **Configuration** Tab and the **Server Start** sub tab.
5. In the arguments field append the following

```
-Doracle.disco.activation.preferencePort=<portno>
-Doracle.disco.activation.preferenceHost=<hostname>
```

the port number is the value of PREFERENCE\_PORT defined previously.

The hostname is the name of the host on which the preference server is started, for example, APPHOST2.

**14.6.5.8.8 Setting up Discoverer WSRP Portlet Producer in a Clustered Environment** The portlet preference store is used for persisting consumer registration handles and portlet preference data.

Discoverer WSRP portlet producer uses a file-based preference store and the location of preference store is defined by the value of the `discoWsrpPrefStoreSharedPath` variable of the Discoverer deployment plan (an XML file). The default value of the `discoWsrpPrefStoreSharedPath` variable is `portletData`.

For more information about deployment plan and its contents, see the "Understanding Deployment Plan Contents" section in *Oracle Fusion Middleware Deploying Applications to Oracle WebLogic Server*.

In a clustered environment, for the file-based preference store, all Discoverer WSRP portlet producers running within the same Oracle WebLogic Server must use the same path for the `discoWsrpPrefStoreSharedPath` variable in the deployment plan. Therefore, the value of the `discoWsrpPrefStoreSharedPath` variable must be changed to a shared path in the deployment plan as described in the section ["Setting Up the Preference Store"](#).

When you change the `discoWsrpPrefStoreSharedPath` variable and if you want to migrate the preference store content from the existing preference store to a shared path, you must run the migration utility to transfer preference data from the source path to the destination path as described in the section ["Using the Migration Utility to Transfer Preference Store Content"](#).

**Setting Up the Preference Store**

To view and edit the Discoverer deployment plan by using WebLogic Server Administration Console:

1. Log in to Oracle WebLogic Server Administration Console.

2. In the left pane, click **Deployments**.
3. In the right pane, select the Discoverer application for which you want to update the deployment plan. The Discoverer: Overview page appears.

The Deployment Plan field on the Overview page displays the path of the Discoverer application deployment plan (an XML file). Modify the deployment plan as described the following procedure:

1. Open the deployment plan from the location (as displayed on the Discoverer: Overview page) in an XML editor.
2. Navigate to the **variable-definition** section.

```
...
 <variable-definition>
 <variable>
 <name>discoWsrpPrefStoreSharedPath</name>
 <value>portletData</value>
 </variable>
 </variable-definition>
 ...
```

---

**Note:** The discoWsrpPrefStoreSharedPath variable is defined in the **variable-definition** and **variable-assignment** sections of the deployment plan. You should modify only the discoWsrpPrefStoreSharedPath variable which is defined in the **variable-definition** section.

---

3. Change the **value** of the discoWsrpPrefStoreSharedPath variable to a *SharedPath* which all Managed Servers can access.

```
...
 <variable-definition>
 <variable>
 <name>discoWsrpPrefStoreSharedPath</name>
 <value>SharedPath</value>
 </variable>
 </variable-definition>
 ...
```

---

**Note:** You must specify a absolute path for the **value**. A relative path for the **value** is interpreted relative to the *ORACLE\_HOME/portal* directory.

---

4. Save your changes and close the XML file.

Once you update the deployment plan with the new value for the discoWsrpPrefStoreSharedPath variable, you must update the Discoverer application as described in the "Update (redeploy) an Enterprise application" section of *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

### Using the Migration Utility to Transfer Preference Store Content

The WSRP container preference store migration utility, PersistenceMigrationTool, enables you to migrate existing data between different preference stores.

The syntax of the PersistenceMigrationTool is:

```
java oracle.portlet.server.containerimpl.PersistenceMigrationTool
-sourceType file -destType file -sourcePath dir -destPath dir
```

`sourceType` indicates that the source store is in a file.

`destType` indicates that the destination store is in a file.

---

---

**Note:** Discoverer WSRP producer uses the file-based preference store. Therefore, you must specify `file` for the `sourceType` and `destType` arguments.

---

---

`sourcePath` and `destPath` are locations of file-based preference stores.

---

---

**Note:** You must set the classpath when running this tool to include `wsrp-container.jar`, `oracle.portlet-api.jar` and `ojdbc6.jar`. For example:

```
java -classpath $ORACLE_
HOME/webcenter/modules/oracle.portlet.server_
11.1.1/oracle-portlet-api.jar:$ORACLE_
HOME/webcenter/modules/oracle.portlet.server_
11.1.1/wsrp-container.jar:$WL_HOME/server/lib/ojdbc6.jar
```

---

---

**14.6.5.8.9 Best Practices** Oracle WebLogic Server does not replicate configuration information automatically between nodes. Therefore, it is important that any changes to any of the following be manually propagated to the other servers in the deployment:

- `ORACLE_INSTANCE`
- `DOMAIN_HOME/servers/WLS_DISCO1/applications/discoverer_`  
`version/configuration/configuration.xml`

#### Preference Server

The preference server contains details about user preferences, which can be updated on a frequent basis. If the Oracle Discoverer preference server fails, you must restart the preference server on another server. You would need to copy the preference information to these potential hosts on a regular basis.

The following file needs to be propagated to other servers, which could host the preference server on a regular basis:

`ORACLE_INSTANCE/config/PreferenceServer/Discoverer_Instance/.reg_key.dc`

## 14.6.6 Scaling Out the Deployment

To scale out the deployment to an additional node, follow the steps in [Section 14.6.5, "Install and Configure Application Tier on APPHOST2."](#) Be sure to use the same directory structure used for previous nodes.

---

## Configuring High Availability for Oracle Business Intelligence and EPM

Oracle Business Intelligence is an integrated business intelligence (BI) solution that provides a business user with a complete picture across the entire organization. Oracle Business Intelligence is designed to quickly and easily integrate diverse data sources, find information from the database, share the database information, and exploit the data to learn more about the business and its customers.

Oracle's Enterprise Performance Management (EPM) system integrates a modular suite of performance management applications with comprehensive business intelligence capabilities for reporting and analysis. It increases user productivity, lowers cost of ownership, and reduces business risk.

This chapter provides topics that describe high availability for the following:

- [Section 15.1, "High Availability for Oracle Business Intelligence Enterprise Edition and Enterprise Performance Management"](#)
- [Section 15.2, "High Availability for Oracle Business Intelligence Publisher"](#)
- [Section 15.3, "High Availability for Oracle Real-Time Decisions"](#)

This chapter assumes that you are familiar with Oracle Business Intelligence and EPM. If not, please refer to the Oracle Business Intelligence and EPM documentation.

### 15.1 High Availability for Oracle Business Intelligence Enterprise Edition and Enterprise Performance Management

This section provides an introduction to Oracle Business Intelligence Enterprise Edition (Oracle BI EE) and Enterprise Performance Management (EPM). This section also describes how to design and deploy a high availability environment for Oracle BI EE and EPM.

Oracle BI EE is part of Oracle Business Intelligence, a comprehensive collection of enterprise BI functionality that provides the full range of BI capabilities including interactive dashboards, full ad hoc, proactive intelligence and alerts, enterprise and financial reporting, real-time predictive intelligence, and more. The Oracle BI EE platform is based on a proven, modern web services-oriented architecture that delivers true next-generation BI capabilities. Oracle BI EE is a robust application that allows you to collect and organize disparate data and present the organized data to users who analyze, interpret, and act upon this content. For additional introductory information about Oracle BI EE, see *Oracle Fusion Middleware User's Guide for Oracle Business Intelligence Enterprise Edition*.

Enterprise Performance Management applications integrate strategy, planning, and execution into a seamless process. Oracle's EPM applications are a modular suite of integrated applications that integrate with both Oracle and non-Oracle transactional systems. You can deploy each application independently to deliver a high degree of value but they work better together, integrating strategic, financial, and operational management processes while delivering a low cost of deployment and ownership.

This section includes the following topics:


- [Section 15.1.1, "Oracle BI EE Component Architecture"](#)
- [Section 15.1.2, "Oracle BI EE High Availability Concepts"](#)
- [Section 15.1.3, "Oracle Essbase Component Architecture"](#)
- [Section 15.1.4, "Oracle Essbase High Availability Concepts"](#)
- [Section 15.1.5, "Oracle Hyperion Provider Services Component Architecture"](#)
- [Section 15.1.6, "Oracle Hyperion Provider Services High Availability Concepts"](#)
- [Section 15.1.7, "Oracle EPM Workspace Component Architecture"](#)
- [Section 15.1.8, "Oracle EPM Workspace High Availability Concepts"](#)
- [Section 15.1.9, "Oracle Hyperion Financial Reporting Component Architecture"](#)
- [Section 15.1.10, "Oracle Hyperion Financial Reporting High Availability Concepts"](#)
- [Section 15.1.11, "Allocation Manager Component Architecture and Characteristics"](#)
- [Section 15.1.12, "Oracle BI EE High Availability Configuration Steps"](#)

### 15.1.1 Oracle BI EE Component Architecture

[Figure 15–1](#) shows the Oracle BI EE components in a single instance architecture. All communication between Java and system components is via SOAP messages over HTTP.



**Figure 15–1 Oracle BI EE Single Instance Architecture**



Note that the port numbers shown in [Figure 15-1](#) are the default port numbers for a simple installation with one component of each component type on a single machine. When you use Oracle Enterprise Manager Fusion Middleware Control to scale out, you can choose a port number range. See [Using Fusion Middleware Control to Scale System Components](#).

#### 15.1.1.1 Oracle BI EE Component Characteristics

The following sections provide more information about the characteristics of the Oracle BI EE components. This information can help you better understand Oracle BI EE high availability requirements.

The following list summarizes the functionality of the Oracle BI EE system components.

- Oracle BI Presentation Services (Presentation Services): This is a C++ process that generates the user interface pages and renders result sets on behalf of the Oracle BI Scheduler. You can configure multiple Presentation Services processes, which share the load. No session replication takes place between the Presentation Services processes.

Presentation Services is almost stateless. The only significant state is the client authentication. If Oracle Business Intelligence is configured to use single sign on for authentication purposes, then users do not have to reauthenticate after a failover. For all other authentication schemes, when failover occurs, clients will have to reauthenticate. The client sees an interruption of service and is redirected to a login page.

- Oracle Business Intelligence Server (BI Server): This is a C++ process that does the data manipulation and aggregates data from data sources. You can configure multiple BI Server processes, which share the load. No session replication takes place between the BI Server processes.

The BI Server does not maintain user session state. For high availability deployments, query results are cached in the global cache.

- JavaHost: This is a Java process that includes resource-intensive graph and PDF rendering. You can configure multiple JavaHost processes, which share the load. No session replication takes place between the JavaHost processes.

The JavaHost is a stateless process.

- Cluster Controller: This is a C++ process which manages the population of BI Servers and Oracle BI Schedulers.
- Oracle Business Intelligence Scheduler (Oracle BI Scheduler): This is a C++ process that runs jobs according to a timed schedule. Jobs may be agents created in the Oracle BI Presentation Catalog, or jobs created by the job manager.

The following list summarizes the functionality of the Oracle BI EE Java components:

- Oracle BI Presentation Services plug-in servlet: Presentation Services runs as a process, not as a Web server, and does not communicate using any Web server plug-in API. The Oracle BI Presentation Services Plug-in forwards HTTP requests to Presentation Services. The HTTP requests are requests from the browser-based user interface, or SOAP requests.
- Web services:
  - Action execution service: Executes actions on behalf of Presentation Services and Oracle BI Scheduler. Actions may be invocations of third party web services, or invocations of user supplied Java code executed as EJBs.
  - Action registry web service: Provides a list of actions.
  - Action location service: Provides an indirection level so that actions can be developed in a customer test environment, and then deployed to a production environment without every action definition having to be changed to point to production sources.
  - Web services for SOA: Provides a web service interface to the contents of the Oracle BI Presentation Catalog. The tree of objects in the Oracle BI Presentation Catalog is exposed as a tree of web services, defined by a WSIL tree with WSDL leaves.
  - Security web service: Provides standards-based authentication and population services.

**15.1.1.1.1 Process Lifecycle** OPMN controls each Oracle BI EE system component. To start and stop, use the Oracle Enterprise Manager Fusion Middleware Control Oracle BI EE administration pages.

Each Oracle BI EE Java component is hosted by a managed Oracle WebLogic Server and controlled using the Oracle WebLogic Server Administration Console. Each of the managed Oracle WebLogic Servers is monitored and, if necessary, restarted by the local WebLogic Node Manager.

---

**Note:** See Ensuring That Shared Network Files Are Accessible in Windows Environments if you run OPMN processes using a named user.

---

**15.1.1.1.2 External Dependencies** Oracle BI EE requires the following:

- An SMTP mail server
- An LDAP provider, such as Oracle Internet Directory
- External web services to support actions (configured by the user)

The following clients depend on Oracle BI EE:

- ADF Desktop Integration: ADF uses both Presentation Services web services and direct invocation of BI Server via the ODBC/JDBC port. This is the only case where the Oracle BI Server is invoked directly, without going through Presentation Services.
- BI Publisher: Uses Presentation Services web services.
- BPEL: Uses web services for SOA web services and WSIL browsing servlet.

**15.1.1.1.3 Configuration Artifacts** Oracle BI EE provides three different configuration methods:

- Central configuration is provided by BI EE pages in Oracle Enterprise Manager Fusion Middleware Control. Using these pages, additional instances can be provisioned, providing high availability and scalability. Key configuration values can also be centrally administered, ensuring that all instances have consistent configuration values.
- Configuration of local system components. Local configuration is only required when making advanced configuration changes.
- Configuration of local Java components. As with configuration of system components, local configuration is only required for advanced options.

### Central Configuration

Oracle BI EE provides a central configuration mechanism, which is exposed in the Enterprise Manager Fusion Middleware Control.

Centrally managed configuration covers:

- Changes in the number of components.
- Listening port ranges. Ports used by each component are automatically assigned from this range.
- Basic SSL configuration: SSL can be enabled for BI EE components using the System MBean browser.
- Mail server configuration

- Principal tuning options
- Location of shared files.

Advanced configuration can be done locally by directly editing the local configuration files. When a central configuration is committed, only the configuration options that are set by central configuration are distributed to the local machines.

### Configuration of Local System Components

Each running system component sees a combination of shared and component-specific configuration files. The shared configuration files, such as `ClusterConfig.xml`, will be under this directory:

```
ORACLE_INSTANCE/config/OracleBIAApplication/coreapplication
```

The component-specific configuration files will be under:

```
ORACLE_INSTANCE/config/ComponentType/ComponentInstanceName
```

For example:

```
ORACLE_INSTANCE/config/OracleBIPresentationServicesComponent/
coreapplication_obips1
```

Since each component sees local files, they are not dependent on access to files held centrally on the Administration Server. The centrally managed configuration values are distributed to these local configuration files when central configuration is committed. These centrally managed configuration values are marked in the local files by comments warning the user not to edit them locally. The remaining values are available for local editing.

The `rpd` file contains references to data sources. Typically development environments use a different database for the fact tables than production environments. When the `rpd` file moves from a development to production environment, you must use the BI Administration Tool to edit the `rpd` file and change the data source references.

### Local Java Component Configuration

Java components are configured in the WebLogic directory tree at:

```
DOMAIN_HOME/config/fmwconfig/biinstances/coreapplication
```

**15.1.1.1.4 Deployment Artifacts** In general, Oracle BI EE deploys Java applications in nostage mode. WebLogic Server explodes the `.ear` file, putting it into a private temp directory. An exception to this is MapViewer, which has configuration requirements which can only be performed in a pre-exploded deployment directory.

All the Managed Server applications (excluding administration applications such as the Enterprise Manager Fusion Middleware Control and central configuration application) deploy to the whole Oracle BI EE cluster, not to individually named systems.

For information on the location of the log files when the Oracle BI EE system components are deployed, see [Section 15.1.1.1.5, "Log File Locations."](#)

**15.1.1.1.5 Log File Locations** [Table 15–1](#) shows the log files for the Oracle BI EE system components. You can access all of these log files from the Enterprise Manager Fusion Middleware Control.

**Table 15–1 Log Files for Oracle BI EE System Components**

| BI Component          | Log File                               | Log File Directory                                                                            |
|-----------------------|----------------------------------------|-----------------------------------------------------------------------------------------------|
| OPMN                  | debug.log                              | ORACLE_INSTANCE/diagnostics/logs/OPMN/opmn                                                    |
| OPMN                  | opmn.log                               | ORACLE_INSTANCE/diagnostics/logs/OPMN/opmn                                                    |
| BI Server             | nqserver.log                           | ORACLE_INSTANCE/diagnostics/logs/OracleBIServerComponent/coreapplication_obis1                |
| BI Cluster Controller | nqcluster.log                          | ORACLE_INSTANCE/diagnostics/logs/OracleBIClusterControllerComponent/coreapplication_obiccs1   |
| Oracle BI Scheduler   | nqscheduler.log                        | ORACLE_INSTANCE/diagnostics/logs/OracleBISchedulerComponent/coreapplication_obisch1           |
| Presentation Services | sawlog*.log (for example, sawlog0.log) | ORACLE_INSTANCE/diagnostics/logs/OracleBIPresentationServicesComponent/coreapplication_obips1 |
| BI JavaHost           | jh.log                                 | ORACLE_INSTANCE/diagnostics/logs/OracleBIJavaHostComponent/coreapplication_objh1              |

### 15.1.2 Oracle BI EE High Availability Concepts

This section provides conceptual information about using Oracle BI EE in a high availability deployment.

#### Requirements for Making Oracle BI EE Highly Available

Oracle BI EE achieves high availability through a combination of process replication and highly available storage (database and shared file system).

To provide a highly available system, Oracle BI EE requires the following external services:

- A fault tolerant HTTP load balancer
- A highly available shared file system
- A highly available database for Oracle BI Scheduler and fact tables

The following system components must be replicated (have at least two instances):

- Presentation Services
- Cluster Controller
- Oracle BI Scheduler
- BI Server
- JavaHost

The Enterprise Manager Fusion Middleware Control configuration UI will warn the administrator if the system components do not have at least two instances to meet the high availability requirement.

The following persistent data sources must be placed on the highly available shared file system:

- RPD publishing directory: The server metadata is contained in the repository file (.rpd file) that is local to each BI Server. One BI Server is designated as a Master. Online changes to the rpd file are made on the Master BI Server, and these changes are replicated to other members of the cluster.
- Oracle BI Presentation Catalog


- Global cache (optional, but advisable for higher performance). The global cache capability offers support for a common query cache that stores cache seeding and purging events. The global cache is visible to all BI Servers in the cluster.
- Scheduler scripts

Also, a WebLogic cluster with at least two members is required to host the Oracle BI EE Java components such as the Oracle BI Presentation Services Plug-in and the authentication service.

#### **15.1.2.1 Oracle BI EE and EPM High Availability Architecture**

[Figure 15–2](#) shows a highly available Oracle BI EE and EPM deployment.

**Figure 15–2 A Highly Available Oracle BI EE and EPM Deployment**





In [Figure 15-2](#), incoming requests are received by the hardware load balancer, which routes them to WEBHOST1 or WEBHOST2 in the web tier. These hosts have Oracle HTTP Server installed.

Oracle WebLogic Server is installed on APPHOST1 and APPHOST2 in the application tier. The Administration Server can be made active on APPHOST2 if APPHOST1 becomes unavailable, as described in [Chapter 3, "High Availability for WebLogic Server."](#) The Oracle BI EE, Oracle EPM, and Oracle BI Publisher Java components are deployed on the BI\_SERVER1 and BI\_SERVER2 managed servers on these hosts (note that Financial Reporting, Provider Services, and Workspace are Oracle EPM components). These managed servers are configured in an Oracle WebLogic cluster that allows them to work in an active-active manner. Whole server migration is configured for the managed servers.

The BI domain uses virtual hostnames as the listen addresses for the BI managed servers. You must enable VIP mapping for each of these hostnames on the two BI Machines, VIP1 on APPHOST1 and VIP2 on APPHOST2, and must correctly resolve the virtual hostnames in the network system that the topology uses (by DNS Server or by hosts resolution). [Section 15.1.12.22, "Configuring Server Migration for the BI\\_SERVERn Servers"](#) provides more details on configuring server migration for the BI servers.

Oracle BI Presentation Services, JavaHost, Oracle BI Cluster Controller, Oracle BI Scheduler, and Oracle BI Server are system components installed on APPHOST1 and APPHOST2 and configured as a cluster. The Cluster Controller and Scheduler on APPHOST2 are passive (they are started but do not service requests) and are made active only if APPHOST1 components fail.

In the data tier, shared external storage is set up to store Catalog, query cache, RPD, and Scheduler script data. An Oracle RAC database is used to store the BI EE schema metadata, BI Publisher repository data, and is also recommended for the enterprise database.

The following subsections provide high availability information for several of the Oracle BI EE components.

**15.1.2.1.1 Web Server High Availability Considerations** The Oracle BI Presentation Services Plug-in is the initial point of contact for any HTTP client, such as the browser user interface. This may in turn be fronted by an Oracle HTTP Server or other supported Web server. You must replicate this HTTP stack to avoid a single point of failure. To avoid presenting multiple HTTP URLs to users, set up an external load balancer, which routes user requests to one of the replicated Oracle BI EE URLs.

If internal HTTP requests route directly to the local WebLogic Managed Server, there is a reduced degree of high availability. If the WebLogic instance fails, Node Manager will restart it. WebLogic restart may take several minutes. During this period, the local security service is unavailable, preventing login. Contrast this with the recommended configuration, where the internal HTTP requests are routed through the load balancer. The alternative security service is quickly available.

**15.1.2.1.2 Oracle BI Presentation Services Plug-in High Availability Considerations** The Oracle BI Presentation Services Plug-in is deployed as the analytics .ear application. It routes session requests to Presentation Services instances using a native RPC protocol over TCP. It also performs load balancing between the multiple Presentation Services instances, with session affinity maintained in general through a session cookie or URL argument. Presentation Services and the Plug-in run as separate processes.

Each Oracle BI Presentation Services Plug-in instance is entirely independent - there is no communication between Oracle BI Presentation Services Plug-in instances. Each

makes its own decision on load balancing based on its experience of response times and number of outstanding requests.

It is recommended to configure session affinity on the load balancer.

The Oracle BI Presentation Services Plug-in is one of the Oracle Business Intelligence Java components and is scaled for high availability when you add additional Managed Servers to the deployment.

**15.1.2.1.3 Presentation Services High Availability Considerations** Presentation Services receives requests from the Oracle BI Presentation Services Plug-in on a port assigned from the port range defined in the Scalability tab of the Capacity Management page in Fusion Middleware Control. Although an initial user session request can go to any Presentation Services instance in the cluster, each user is then bound to a specific Presentation Services instance.

To process user requests, Presentation Services must communicate with the BI Servers. In a clustered environment, the first point of contact to the BI Servers is through the Cluster Controller. Presentation Services uses the Oracle BI ODBC data source that is configured for the clustered environment to identify the primary and secondary Cluster Controllers and the ports on which they listen. The centralized BI domain configuration file configures automatically connection parameters in the ODBC Data Source Name (DSN) when centralized configuration is changed. Do not modify the Oracle BI ODBC DSN connection parameters locally.

Cluster Controller assigns the Oracle BI Server instance that Presentation Services connects to. The connection to the Oracle BI Server establishes over ODBC, and subsequent requests in the same session go directly from Presentation Services to the assigned Oracle BI Server. The ODBC session between Presentation Services and the Oracle BI Server is stateful and must maintain affinity for the session lifetime.

To communicate with Oracle BI Scheduler, Presentation Services first contacts the Cluster Controller, which relays the active Oracle BI Scheduler instance. Presentation Services then establishes a session with the appropriate Oracle BI Scheduler instance. the Availability tab of the Capacity Management page in Fusion Middleware Control specifies the host and Oracle instance names for the Cluster Controllers and Oracle BI Scheduler.

**15.1.2.1.4 BI Cluster Controller High Availability Considerations** The Cluster Controller is the first point of contact for a new BI Server or Oracle BI Scheduler session from Presentation Services and other clients. The Cluster Controller deploys in active-passive configuration:

- Primary Cluster Controller: The active cluster controller.
- Secondary Cluster Controller: Assumes the role of active Cluster Controller if the Primary Cluster Controller is unavailable.

By default, the first Cluster Controller that you configure in your Oracle Business Intelligence installation is the primary Cluster Controller.

The Cluster Controller determines which Oracle BI Server in the cluster should receive incoming requests based on Oracle BI Server availability and load. It also monitors the server operation in the cluster, including the Oracle BI Scheduler instances.

The Cluster Controller supports:

- Detection of server and Oracle BI Scheduler failures.
- Failover for ODBC clients if their servers fail.

The Cluster Controller also determines the active Oracle BI Scheduler instance at run time. Note that Cluster Controller does not control Presentation Services instances.

It is important to avoid a split brain configuration, where both Cluster Controllers consider themselves to be active. To avoid a split brain configuration, do not configure Cluster Controllers on separate LANs connected across a potentially unreliable WAN. Fault tolerance against site failures should use the disaster recovery mechanisms, not the clustering mechanisms.

**15.1.2.1.5 BI Server High Availability Considerations** You can install and configure multiple BI Servers to create a BI Server cluster. The Cluster Controller dispatches requests from clients such as Presentation Services to an active member of this cluster. BI Servers listen for client requests on a port assigned from the port range that is defined in the Scalability tab of the Capacity Management page in Fusion Middleware Control.

---

**Note:** You must synchronize the clock on each server participating in a cluster. Unsynchronized clocks can skew reporting.

---

The Master BI Server is the server that the Oracle BI Administration Tool connects to in order to perform online metadata changes in the RPD file. The metadata changes then propagate out to the other servers when they restart. You can view which Oracle BI Server is the master using the Cluster Manager in the Administration Tool.

### **Communication from BI Server to Databases**

BI Server supports communication with many database types, including Oracle databases. The BI Server data repositories hold the raw fact tables. Oracle BI EE does not use the BI Server data repositories to store metadata; the RPD file or the Oracle BI Presentation Catalog (for Presentation Services) stores all metadata.

The Oracle BI Scheduler stores data in the Oracle BI Scheduler schema. Since the Oracle BI Scheduler database connection configuration uses the same libraries as BI Server, the fault tolerance characteristics are the same. The only difference is that the Oracle BI Scheduler interaction is much more write intensive. BI Server data sources are primarily read-only.

The main points of BI Server interaction with the database are:

- Source database: Source databases are used for Oracle BI EE reporting. They are primarily read-only. There are many supported options for source databases.
- Event polling: Keeps track of cache removal or creation from the file system.
- Usage tracking
- Write back facility: A user-driven capability.

The implementation provides the following characteristics:

- All database read/write operations are performed by C++ code.
- For Oracle databases, OCI is used. For other databases or data sources, native database capabilities are used.
- Oracle and non-Oracle database support is possible for not just the Source database, but also for Oracle BI Scheduler.

**15.1.2.1.6 Administration Tool High Availability Considerations** The Administration Tool uses the same Oracle BI ODBC DSN that Presentation Services uses. It is directed to

the Master BI Server by the Cluster Controller. If the Master BI Server is unavailable, the Administration Tool cannot perform RPD file writes.

**15.1.2.1.7 Oracle BI Scheduler High Availability Considerations** The Oracle BI Scheduler instances operate on an active-passive model. Only one Oracle BI Scheduler is active and processing requests at any one time to ensure that it allocates unique job ids to new jobs. The inactive Oracle BI Scheduler instance remains idle and does not process jobs until called on to take over if an active Oracle BI Scheduler fails.

Do not configure multiple Oracle BI installations to use the same scheduler tables. Doing so creates multiple active Oracle BI Schedulers, breaking the unique job ID constraint.

The Oracle BI Scheduler uses the same database technology as the BI Server.

If Oracle BI Scheduler is up and the database is down, Oracle BI Scheduler queues up failed transactions in its memory and retries them until they are successful. Therefore, if you shut down a back-end database source for maintenance you do not need to restart Oracle BI EE processes when you restart the database. The processes are retried on the same count as the periodic purge in the Oracle BI Scheduler, and therefore use the same configuration values. Oracle BI EE only retries statements that are required to be in Oracle BI EE metadata. Oracle BI EE does not retry queries, only inserts and deletes (except for deletes that would be reexecuted). Configuration values are in the `instanceconfig.xml` file for Oracle BI Scheduler (`PurgeIntervalMinutes`).

The Oracle BI Scheduler listens for Cluster Controller communication and for client requests on ports assigned from the port range defined in the Scalability tab of the Capacity Management page in Fusion Middleware Control.

**15.1.2.1.8 BI JavaHost High Availability Considerations** The Oracle BI JavaHost component provides services that enable Oracle BI Presentation Services to support various components such as Java tasks for Oracle BI Scheduler, Oracle BI Publisher, and graph generation. As with all other components, the administrator can use Fusion Middleware Control to control how many instances are deployed and where. The default configuration of JavaHost clients is to use all available JavaHosts; this differs from the previous 10g default of only using the local JavaHost.

The JavaHost service receives requests from Presentation Services and Oracle BI Scheduler on a port that is assigned from the port range that is defined in the Scalability tab of the Capacity Management page in Fusion Middleware Control.

### 15.1.2.2 Shared Files and Directories

In Oracle BI EE, process state is shared through shared file systems and databases rather than through distributed communication protocols. In particular, Oracle BI EE does not store state internal to processes or Java components. Therefore, Oracle BI EE does not need to serialize component state so it can be distributed within a Java cluster.

The fact tables and Oracle BI Scheduler tables are shared between BI Server and Oracle BI Scheduler instances, respectively. File system sharing is less standard, as the end of this section describes. You can use a shared storage device such as NAS or SAN.

**15.1.2.2.1 Oracle BI Presentation Catalog Shared Files** The Presentation Services instances in a cluster share a common Oracle BI Presentation Catalog. The Oracle BI Presentation Catalog should be placed on a shared NAS or SAN device. All instances of Presentation Services must have read and write access to the shared device. Concurrent access is managed by the file system locking semantics.

Because the Oracle BI Presentation Catalog consists of a large number of heavily accessed small files, be aware that in some cases the number of files may exceed the file limits for shared file systems. Refer to the documentation for your shared file system for information on setting the file limits for your system.

**15.1.2.2.2 Repository Publishing Directory Shared Files** All BI Servers participating in a cluster share the repository publishing directory, which holds the master copies of repositories edited in online mode. The Master BI Server must have read and write access to this directory; all other BI Servers must have read access.

The clustered BI Servers examine this directory on startup for any repository changes. On startup, a BI Server instance copies the repository in the publishing directory to its local repository directory, for example:

```
ORACLE_INSTANCE/bifoundation/OracleBIServerComponent/coreapplication_
obis1/repository
```

Because this copy occurs only at restart, a cluster with a writable rpd file may show inconsistent behavior (for example, metadata does not match requests from analyses) until all the BI Server instances restart. Metadata changes are normally made offline in a development environment with a read only rpd file in the production environment.

In Fusion Middleware Control, the **Shared Location** parameter on the Repository tab of the Deployment page specifies the repository publishing directory location.

**15.1.2.2.3 Global Cache Shared Files** The global cache is shared by all BI Servers participating in a cluster. It resides on a shared file system storage device and stores purging events, seeding events (often generated by agents), and result sets associated with seeding events. Each Oracle BI Server maintains its own local query cache for regular queries.

Shared files requirements for the global cache are:

- All BI Servers must have read and write access to the global cache directory.
- The global cache parameters are configured in the `NQSConfig.INI` file for each BI Server node participating in the cluster. The global cache is controlled centrally via Fusion Middleware Control.
- The BI Server maintains a query cache, a local, disk-based cache of query result sets. The query cache allows a BI Server to potentially meet many query requests without accessing back-end databases, dramatically decreasing query response time. Query cache entries become obsolete as updates occur on the back-end databases and must purge periodically. For more information on the query cache, refer to the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

**15.1.2.2.4 Oracle BI Scheduler Scripts Shared Files** You must create a network share for the Oracle BI Scheduler scripts. To create a network share, update the `SchedulerScriptPath` and `DefaultScriptPath` elements of the Oracle BI Scheduler `instanceconfig.xml` file (located in `ORACLE_INSTANCE/config/OracleBISchedulerComponent/coreapplication_obischn`). The Oracle BI Scheduler servers must have read and write access to this share.

### 15.1.2.3 Starting and Stopping the Cluster

Because there are dependencies between the processes, you must start a cluster in a particular order. The startup sequence is:

1. Start all of the WebLogic Managed Servers. The system components look to their local Managed Server for the authentication service.
2. Start other processes using the **Start All** button in Fusion Middleware Control. Using the **Start All** button ensures that the Java applications and the system components start in the correct startup order.

#### 15.1.2.4 Cluster-Wide Configuration Changes

Configuration changes, especially changes to the population of the cluster, are critical for high availability. Without the ability to add new members to a cluster, a failure cannot be repaired, leaving the cluster vulnerable to a subsequent failure (this is why mean time to repair is just as important in a high availability system as mean time between individual failures).

The two technology stacks, Java and system, are managed separately.

WebLogic configuration changes are made in the WebLogic Administration Console. Standard WebLogic mechanisms are used to lock configuration while it is being changed, and to propagate the changes to all Managed Servers.

System component configuration changes are made centrally using Fusion Middleware Control. After you enter configuration changes and click **Save**, the configuration changes propagate to all local configuration files. This propagation is important; if the configuration is corrupt or propagation fails, the next propagation corrects the local files. An Oracle JMX extension performs file distribution similar to the WebLogic config file distribution. Delivery at the remote site is announced by a JMX MBean call. Since MBeans must be hosted in an MBean server, this means that every machine which hosts Oracle BI EE system components must also have a WebLogic Managed Server even if that machine does not host any Java Oracle BI EE components.

All components access configuration changes at their next restart only; there is no online configuration reload mechanism.

Configuration changes are triggered by:

- An administrator committing changes made in the Fusion Middleware Control.
- An administrator committing changes made by the MBeans API (which directly mirrors the Fusion Middleware Control user interface).

An enterprise installation that registers a new installation does not trigger a configuration push. This changes the population of machines which the administrator may pick from in the Fusion Middleware Control - new component instances are not automatically added to the configuration.

#### 15.1.2.5 Protection From Failures and Expected Behaviors

This section describes different types of failures that can occur and the expected behavior when each type of failure occurs.

**15.1.2.5.1 Machine Failure** If a machine suffers a transient failure, the cluster detects the restarted components and starts using them again. If a machine fails permanently, you must replace it. Install Oracle BI EE on the new machine using an enterprise install to extend the existing cluster.

For information about recovering from Oracle BI EE machine failures, refer to the section on recovering Oracle BI EE to a different host in *Oracle Fusion Middleware Administrator's Guide*.

**15.1.2.5.2 WebLogic Administration Server Failure** The WebLogic Administration Server is a singleton and only one instance can be up and running at any given time. For information on protecting the WebLogic Administration Server, see [Chapter 12, "Active-Passive Topologies for Oracle Fusion Middleware High Availability."](#) In Oracle BI EE, the WebLogic Administration Server is used only for administrative tasks.

While the WebLogic Administration Server is down, you cannot reconfigure Oracle BI EE or use Fusion Middleware Control to look at log files.

**15.1.2.5.3 WebLogic Managed Server Failure** The WebLogic Node Manager detects failure of Managed Servers. Because all Java components are stateless, Oracle BI EE does not take Java cluster state maintenance into consideration.

All communications from system components to Managed Servers, for example, from Presentation Services to action services, use HTTP to the local Managed Server. If the local Managed Server fails, the Node Manager should restart it.

The following lists the effect on system components running on the same machine if the Managed Server cannot restart but the machine does not fail:

- BI Server: Taken out of service until the local Managed Server works to ensure that login (done through the BI Server) is not affected. When the BI Server detects that the local Managed Server is not running, it directs the Cluster Controller to send login requests to other BI Servers.
- Presentation Services: Unable to execute actions.
- Oracle BI Scheduler: Scheduled action executions fail if the active scheduler runs on this machine.

New Scheduler connections (for example, from Job Manager) will fail. The Oracle BI Scheduler on a host where the local Managed Server is not running is unable to authenticate new users.

If you must take a Managed Server offline, first stop all system components running on the machine. If the Managed Server on a particular machine becomes unavailable, stop all system components running on that machine.

**15.1.2.5.4 Oracle BI Scheduler Failure** The Cluster Controller monitors and manages the Oracle BI Scheduler. If the Oracle BI Scheduler is unavailable, the Cluster Controller determines the next Oracle BI Scheduler instance to activate. If the previous primary Oracle BI Scheduler becomes available again, the primary role doesn't revert to it.

When the active Oracle BI Scheduler fails, any open client connections do not receive an error because the Oracle BI Scheduler protocol is stateless and seamlessly fails over.

- Agents: Agent executions maintain state in the Oracle BI Scheduler tables. When the next instance of Oracle BI Scheduler becomes active, it reads the state of all job instances that were in progress, and executes them. A scheduled agent only delivers to those recipients that it did not deliver to before the failure of the active instance.

Successful delivery is recorded in the Oracle BI Scheduler tables after completion of the delivery. If Oracle BI Scheduler fails over after delivery is complete, but before recording that fact, then the delivery is repeated. Because of this, a small number of repeat deliveries are possible.

Note that Agent failover is only supported for agents that are scheduled. For example, if you click **Run Agent Now** on the toolbar in Answers, and then the primary Oracle BI Scheduler fails, that Agent will not continue to run on the

Secondary Scheduler and an error message will be displayed in the Run Now results dialog box.

- **Java, Command Line, or Script Job:** The jobs are re-executed from the beginning with a new job instance.

---

**Note:** Any job instance can be manually rerun from the Job Manager. For an agent, this only delivers to those users that did not have successful deliveries. For example, if the mail server goes down halfway through an agent's execution, the rerun of the instance delivers only to those recipients who did not receive e-mail due to the mail server crash.

---

**15.1.2.5.5 Cluster Controller Failure** The Cluster Controller supports detection of BI Server or Oracle BI Scheduler failures and failover for clients of failed servers.

The Cluster Controller works on an active-passive model. All clients first attempt to connect to the primary Cluster Controller. If the primary Cluster Controller is unavailable, clients then connect to the secondary Cluster Controller. The secondary Cluster Controller then directs requests to BI Servers based on load and availability, and to the active Oracle BI Scheduler instance. If the primary later becomes available, then all requests go to the primary again.

The secondary Cluster Controller monitors the session count on each BI Server (just like the primary), but the secondary Cluster Controller does not determine which Oracle BI Scheduler is the active Scheduler unless the primary Cluster Controller is down.

The primary and secondary Cluster Controllers monitor each other's life cycle. A Split-Brain failure can occur if communication is down between the Cluster Controller instances but each is up and can communicate with other clients. In these cases, BI Servers are not affected but the Oracle BI Scheduler may have two active instances at once. In rare cases, this situation can lead to double execution of jobs. When the line of communication comes back up, the primary Cluster Controller dictates to the cluster that only one Oracle BI Scheduler should be active. The Cluster component must reside on the same LAN, minimizing the possibility of a Split-Brain failure.

If both Cluster Controllers are unavailable, Presentation Services returns an error to any new user attempting to login. Existing sessions are not affected.

Cluster Controllers track the number of active sessions on each BI Server instance. When new ODBC connections are made, the Cluster Controller allocates the new session to keep the sessions for each BI Server in balance. If a BI Server instance goes down temporarily and comes back up, then all new ODBC sessions route to this node until its session count matches that of the others in the cluster.

When a Cluster Controller fails, the first person who tries to create a new session sees a delay of about six seconds as the secondary Cluster Controller picks up sessions. All other connections are then routed to the secondary Cluster Controller.

**15.1.2.5.6 Presentation Services Failure** A Presentation Services failure affects:

- **Web clients:** Although an initial user session request can go to any Presentation Services instance, each user is then bound to a specific Presentation Services instance. Loss of that Presentation Services instance disconnects the session, and an error is relayed back to the browser. Any work in progress during the loss of the server that was not saved to disk is lost. The user must log in again to establish a new connection to an available Presentation Services instance. If user login is



taking place using a Single Sign-On system such as Oracle Single Sign-On, then this relogin occurs automatically. The new Presentation Services session will create a new BI Server session.

---

**Note:** When a Presentation Services instance fails, there is a small interval of time before the system recognizes that the instance has failed and before users are migrated to a new Presentation Services instance. If a single sign-on solution is being used, then there is no need for users to reauthenticate. If single sign-on is not being used, then users must reauthenticate.

---

- Agents: An error is relayed to the Oracle BI Scheduler, which logs the failure and then retries the job. The retry establishes a new connection to an available Presentation Services.

**15.1.2.5.7 BI Server Failure** When a BI Server failure occurs, an ODBC error goes to the appropriate client:

- Presentation Services: Each Web user has requests that one BI Server serves. If the BI Server is unavailable, the user might see an error. However, a browser refresh makes a new session establish with an available BI Server. Note that the Presentation Services component performs this arbitration on behalf of its users.
- Administration Tool: The Administration Tool relays the ODBC error when the BI Server that it connects to becomes unavailable, and then closes the connection. The administrator must use the Administration Tool to reconnect.
- Agents: When BI Server failure occurs, the error relays to the Oracle BI Scheduler, which logs the failure and retries the job. A connection then establishes with an available BI Server.
- Third-party clients: Third-party clients use ODBC to connect the BI Server. When BI Server failure occurs, ODBC relays the failure and the session closes and reopens according to the ODBC standard.
- Master BI Server failure: If the Master BI Server is unavailable, online metadata changes cannot be performed. This is an administration operation; it does not affect run-time availability.

If the Master BI Server is permanently unavailable, you must appoint one of the other servers as the new master:

- Scale in the Master BI Server using the Scalability tab of the Capacity Management page in Fusion Middleware Control.
- Restart the Oracle BI EE system components.

The system assigns a new Master BI Server.

**15.1.2.5.8 Troubleshooting Oracle BI EE** Recovery issues typically involve failure to start a process. The suggested troubleshooting sequence is:

1. Check Fusion Middleware Control to see which processes are running. This reflects the state reported by OPMN.
2. Check the OPMN log files. They can be accessed from Fusion Middleware Control. They can also be accessed in the locations specified in [Table 15-1](#).
3. When a failed process is identified, check its log file, which can be accessed from Fusion Middleware Control or in the locations specified in [Table 15-1](#).

Certain processes shut down if they cannot communicate with their peers. To diagnose this problem:

1. Identify the failed process. Find the low-level socket error in the log file, which identifies the port and server address. Check in the local configuration files to see which process is configured to listen on that port. Simple deployments use the default port allocations, as the component descriptions in the subsections of [Section 15.1.2.1, "Oracle BI EE and EPM High Availability Architecture"](#) describe. You can see the ports that are currently configured in the Availability tab of the Capacity Management page in Fusion Middleware Control.
2. Check to see if the peer is running. If the peer is running, check its log file for an entry rejecting the communication. This can happen if SSL configuration is inconsistent or if there is a protocol failure.


Cluster Controller managed processes will fail if their configuration fails to match the Cluster Controller. While some failover events are not logged, the Cluster Controller log file records crashes of any Oracle BI Scheduler or BI Server instances.

Review the log files after initial start up. If a BI Server or Oracle BI Scheduler instance has not been configured correctly and as expected by the Cluster Controller, then the instance, though it may not shut down, will not be added to the cluster. The log files will record such failures.

### 15.1.3 Oracle Essbase Component Architecture

[Figure 15–3](#) shows the Oracle Essbase single instance architecture.

**Figure 15–3 Oracle Essbase Single Instance Architecture**



#### 15.1.3.1 Oracle Essbase Component Characteristics

Essbase component architecture comprises four tiers:

- **Web:** Oracle HTTP Server
- **Application:** Oracle Hyperion Provider Services, a mid-tier enterprise application deployed on a WebLogic application server

Provider Services communicates with Essbase Agent and Essbase Servers using TCP/IP.

- **Essbase Service:** Essbase Analytical Services, comprising OPMN, Essbase Agent, and Essbase application servers

OPMN communicates with Essbase Agent using Oracle Net Services. Essbase Agent acts as a coordinator between clients and Essbase application servers. Essbase Agent has a pipe communication channel with the servers. It also communicates using TCP/IP.

- **Storage:** The Hyperion Registry is a repository that stores configuration information in a relational database. Essbase Agent and Essbase application servers communicate with the back-end database using JDBC or ODBC.

**15.1.3.1.1 State Information** The Oracle Essbase session state is in memory, and there is no session failover. If a client session is lost, the client must log on again.

**15.1.3.1.2 Runtime Processes** Essbase Analytical Services comprises these processes at runtime:

- **Oracle Process Manager and Notification Server (OPMN):** A daemon that monitors Oracle Fusion Middleware components, including Essbase

OPMN-managed components are specified in `opmn.xml`.

You can use a command-line interface to OPMN to start, stop, restart, and monitor Essbase Agent; see [Section 15.1.1.1.1, "Process Lifecycle."](#) It does not start or stop the Essbase Application Server process directly. OPMN communicates with Essbase Agent using Oracle Net Services.

- **Essbase Agent (ESSBASE):** The process responsible for coordinating user requests for applications.

Essbase Agent is directly managed by OPMN. Essbase Agent manages application servers.

- **Essbase Application Server (ESSSVR):** The Essbase Application Server process handles all the data storage, caching, calculations, and data security.

**15.1.3.1.3 Process Lifecycle** Use the OPMN command `opmnctl` to start, stop, and monitor Essbase Agent.

### Start

To start OPMN and all its managed components, use `$ORACLE_INSTANCE/bin/opmnctl startall`.

To start Essbase Agent only, with OPMN already running, use `$ORACLE_INSTANCE/bin/opmnctl startproc ias-component=essbaseserver1`, where `essbaseserver1` is the name of the managed Essbase instance specified in `opmn.xml`.

### Stop

To stop OPMN and all its managed components, use `$ORACLE_INSTANCE/bin/opmnctl stopall`.

To stop Essbase Agent only if Essbase is configured in failover cluster mode and OPMN is already running, use `$ORACLE_INSTANCE/bin/opmnctl stopproc ias-component=essbaseserver1 process-type=Essbase`, where

essbaseserver1 is the name of the managed Oracle Essbase instance specified in `opmn.xml`.

### Monitor

To get status information about the managed process, use `$ORACLE_INSTANCE/bin/opmnctl status`.

**15.1.3.1.4 Request Flow** Communication with Essbase can be through a cluster name (resolved by Provider Services), a URL (resolved by the application), or a physical host and port.

**15.1.3.1.5 External Dependencies** Essbase uses Hyperion Registry, a database repository, to store and retrieve configuration information.

You can configure Essbase to work with an external directory service such as the Oracle Internet Directory (OID). Essbase communicates with OID using LDAP protocol.

**15.1.3.1.6 Configuration Artifacts** The following table describes Essbase configuration artifacts.

**Table 15–2 Essbase Configuration Artifacts**

| File                                                 | Location                                        | Description                                                                                            |
|------------------------------------------------------|-------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Essbase configuration file, <code>essbase.cfg</code> | <code>\$ARBORPATH/bin</code>                    | Contains configurable parameter values for Essbase.                                                    |
| The Essbase security file, <code>essbase.sec</code>  | <code>\$ARBORPATH/bin</code>                    | Contains information on security and applications                                                      |
| OPMN configuration file, <code>opmn.xml</code>       | <code>\$ORACLE_INSTANCE/config/OPMN/opmn</code> | Lists all the managed components and the environment variables to be passed to each managed component. |

**15.1.3.1.7 Deployment Artifacts** Deployment artifacts required for correct setup of Essbase include JDBC and ODBC, `reg.properties`, security artifacts, and so forth. Deployment artifacts must be passed to OPMN, which can then provision Essbase.

**15.1.3.1.8 Log Files** Diagnostic log files reside, by default, under `$ORACLE_INSTANCE/diagnostics`.

Essbase Agent writes messages to `Essbase.log`:

- Default location: `$ORACLE_INSTANCE/diagnostics/logs/Essbase/<name of the Essbase instance>/essbase`
- Configured location: `$HYPERION_LOGHOME/essbase`

Essbase application server writes messages to application server log files:

- Default location: `$ORACLE_INSTANCE/diagnostics/logs/Essbase/<name of the Essbase instance>`
- Configured location: `$HYPERION_LOGHOME/essbase/app`

OPMN creates the Essbase console log, located by default in `$ORACLE_INSTANCE/diagnostics/logs/Essbase/<name of Essbase instance>`.

The Essbase Forward Ping log, `EssbasePing.log`, is located by default in `$ORACLE_INSTANCE/diagnostics/logs/OPMN/opmn`.

Lease manager logs are created only in a failover configuration. (See [Section 15.1.4.2, "Protection from Failures and Expected Behaviors."](#)) The lease manager modules within Essbase Agent and Essbase Server write messages to their own logs:

- Default location: `$ORACLE_INSTANCE/diagnostics/logs/Essbase/<ias component name>/essbase`
- Configured location: `$HYPERION_LOGHOME/essbase`

## 15.1.4 Oracle Essbase High Availability Concepts

This section provides conceptual information about using Essbase in a high availability environment.

### 15.1.4.1 Oracle Essbase High Availability Architecture

[Figure 15–2](#) shows Essbase in a highly available Oracle BI EE deployment.

**15.1.4.1.1 Shared Files and Directories** ARBORPATH, which includes configuration files, security files, and all applications and corresponding databases, is on a shared disk. Instance-specific files, binaries, and log files are on a local disk by default but can be shared. For example, you could put log files on a shared disk by changing a configuration setting in `opmn.xml`.

**15.1.4.1.2 Cluster-Wide Configuration Changes** Shared configuration information is in `essbase.cfg`, which is on a shared disk. For example, changes in timeout settings for lease acquisition or renewal are made in `essbase.cfg` on the shared disk.

### 15.1.4.2 Protection from Failures and Expected Behaviors

For failover, you can configure Essbase in a two-node cluster.

Included in a failover cluster setup:

- Two nodes
  - One node is active at all times and one node is passive.
- A shared disk between the nodes
  - The shared disk stores applications, databases, product configuration, and everything that is placed under `$ARBORPATH`.
- Product binaries and others in `$ESSBASEPATH`
  - Each node can have the binaries on a direct attached storage disk.
- Registry and lease databases
  - These databases, which are critical, must be on redundant storage devices.
  - A leasing mechanism ensures that one and only one Essbase Agent and its set of application servers are active. One and only one lease must be active on a resource at any time; the owner of the active lease has ownership rights to the resource. This leasing scheme is implemented using the lease tables in the database. The lease tables are seeded by Repository Configuration Utility (RCU).

To exploit failover functionality, Java API clients must log on using APS Servlet endpoint URL. The URL would specify the cluster name of the failover cluster; for example, `http://WEBHOST1:7777/aps/Essbase?clusterName=cludemo`

Provider Services resolves the cluster name to the host name and port of the active node in the cluster.

Session failover is not supported. If service failover occurs, the client receives an error message and must log on again.

#### 15.1.4.3 Troubleshooting


To troubleshoot Essbase failover, check OPMN and Essbase logs to determine event sequence. For example, the logs might show that OPMN starts Essbase but Essbase does not acquire a lease because of database authentication failure.

### 15.1.5 Oracle Hyperion Provider Services Component Architecture

Oracle Hyperion Provider Services (Provider Services) is a middle-tier data-source provider to Oracle Essbase for Java API, Oracle Hyperion Smart View for Office, Fusion Edition, and XMLA clients and to Oracle Business Intelligence Enterprise Edition. Provider Services supports highly concurrent analytical scenarios and provides scalability and reliability in a distributed Web-enabled enterprise environment.

Figure 15–4 shows the Provider Services single instance architecture.

**Figure 15–4 Provider Services Single Instance Architecture**



Provider Services contains three subcomponents: Essbase Java API, SmartView Provider, and an XMLA (XML for Analysis) provider. These subcomponents service Java API clients, SmartView clients, and XMLA clients, respectively. Provider Services is designed with three different servlets, one for each subcomponent.

Communication between clients and Provider Services is over the HTTP(S) protocol. The JAPI and XMLA clients are designed for Essbase communication. The Smart View client enables users to connect to data sources such as Essbase, Oracle Hyperion Planning, Fusion Edition and Oracle BI EE servers.

#### 15.1.5.1 Oracle Hyperion Provider Services Component Characteristics

The following sections describe Prover Services Component characteristics.

**15.1.5.1.1 State Information** Provider Services session information is stored in memory until the user disconnects or the session times-out.

**15.1.5.1.2 Runtime Processes** Provider Services is a Java-based enterprise application deployed on the WebLogic server application container. At runtime, there is a single Java process, which is the container. Provider Services runs no other external processes.

**15.1.5.1.3 Process Lifecycle** The lifecycle of the Provider Services process is the same as the process lifecycle for any standard Java- based enterprise application deployed on a WebLogic server.

**15.1.5.1.4 Request Flow** A dispatcher such as Oracle HTTP Web Server acts as a gateway, delivering requests to the Provider Services instances that are part of the cluster. Provider Services supports sticky sessions, so that when a user session establishes with a Provider Services instance, the same instance handles all subsequent requests from the client until the user disconnects. If a Provider Services instance fails, the user must disconnect the existing invalid session and log on again, using the same URL (the dispatcher URL) as for the failed session.

**15.1.5.1.5 External Dependencies** Provider Services is dependent on Hyperion Registry. The SmartView Provider component stores the Data Sources (`datasources.xml`) and Smart Slices in the registry database using Registry Java APIs. The Essbase Java API component also uses the Hyperion registry for maintaining domain-related (`domain.db` file) information. All of this information is stored under the Registry's LOGICAL\_WEB\_APP component of Provider Services.

**15.1.5.1.6 Configuration Artifacts** Provider Services has a configuration file, `essbase.properties`, which resides in `$ORACLE_HOME/products/Essbase/aps/bin`. The Provider Services startup script requires the following JVM setting `ESS_ES_HOME` so that it can identify the configuration file at runtime. This setting specifies the `PROVIDER_SERVICES_HOME` folder, in which the `bin/essbase.properties` file resides.

#### **15.1.5.1.7 Deployment Artifacts**

Provider Services deployment artifacts include these files:

- `logging.xml`: Configuration file for managing Provider Services logs
- `essbase.properties`: Provider Services configuration file
- `aps.ear`: Provider Services enterprise application archive

**15.1.5.1.8 Log Files** Provider Services logging is based on ODL logger. The Provider Services startup script requires these JVM settings related to logging:

- `oracle.core.ojdl.logging.config.file`: To specify the ODL `logging.xml` file path for Provider Services.
- `logging.folder`: The location to store the Provider Services server log files.

---

**Note:** In a cluster including several Provider Services instances, a separate unique log location is defined for each instance. If the log location for an instance is under a Provider Services home, the administrator must grant read and write permissions to that folder for all Provider Services instances.

---

## **15.1.6 Oracle Hyperion Provider Services High Availability Concepts**

Provider Services deploys and runs within an application server environment such as Oracle WebLogic Server. To support Provider Services high availability, you can deploy multiple instances of Provider Services, using the WebLogic Server clustering and load balancing capabilities. For details on clustering enterprise applications such as Provider Services, see the WebLogic Server documentation.

### 15.1.6.1 Oracle Hyperion Provider Services High Availability Architecture

Figure 15–2 shows Oracle Hyperion Provider Services in a highly available Oracle BI EE deployment.

**15.1.6.1.1 Hyperion Registry Structure for a Provider Services Cluster** In the preceding figure, LOGICAL\_WEB\_APP represents a cluster consisting of Provider Services instances. Each PROVIDER\_SERVICES\_WEB\_APP component stores the actual server and port of its Provider Services instance; for example:

```
LOGICAL_WEB_APPCluster_URL
|
---- PROVIDER_SERVICES_WEB_APP 1 server1:port
|
---- PROVIDER_SERVICES_WEB_APP 2
|
---- PROVIDER_SERVICES_WEB_APP 2 server2:port_y
```

The LOGICAL\_WEB\_APP component information specifies the Provider Services cluster URL, which is unique.

**15.1.6.1.2 Cluster-Wide Configuration Changes** Provider Services reads the `essbase.properties` file only during startup. If the Provider Services administrator changes this file, the administrator must restart the Provider Services cluster so that all Provider Services instances in the cluster are synchronized with configuration settings.

**15.1.6.1.3 OPMN Essbase Cluster Support** To support Essbase clusters, clients communicate with Provider Services to fetch the active Essbase node in the cluster. API clients (CAPI/JAPI) must specify an Provider Services URL (which contains the OPMN Essbase cluster name) to obtain the active Essbase host instance details. Provider Services then retrieves the active Essbase node details from the Hyperion Registry. The remainder of this section describes the behavior for Provider Services and C/JAPI client use cases.

**15.1.6.1.4 Essbase Database Clustering by Provider Services** Provider Services supports high availability of Essbase databases with active-active clustering. Clustering Essbase cubes (databases) enables load balancing and failover support. Provider Services manages a series of active, duplicate databases that respond to user requests. Users connect to the database by specifying an Essbase cluster; the database that is accessed must be transparent to users. Provider Services facilitates the routing of connections between databases in a cluster, based on availability and precedence rules (using the round-robin technique). The Essbase cluster supports read-only operations on the Essbase database. Write-back operations such as data load and outline editing are not supported.

### 15.1.6.2 Protection from Failures and Expected Behaviors

Provider Services supports sticky sessions: After a user session is established with a Provider Services instance, the same Provider Services instance handles all subsequent requests from the client until the user disconnects. When a Provider Services instance fails (indicated by an invalid session error), the user must disconnect the existing (invalid) session and log on again. The failover to the new Provider Services instance is transparent, so that the user can connect to the same APS URL (the dispatcher's URL)



that was used earlier. A dispatcher could be, for example, Oracle HTTP Server, which acts as a gateway, delivering requests to the Provider Services instances that are part of the cluster.


## 15.1.7 Oracle EPM Workspace Component Architecture

Oracle Enterprise Performance Management (EPM) Workspace, Fusion Edition (Workspace) is the Web user interface that is used to access all Oracle Hyperion and non-Oracle Hyperion content. Oracle Hyperion content includes Reporting and Analysis Framework and Oracle's Hyperion financial applications.

### 15.1.7.1 Workspace Component Characteristics

Figure 15–5 shows the Workspace single instance architecture.

**Figure 15–5 Workspace Single Instance Architecture**



EPM Workspace consists of a single J2EE Web application, which is typically fronted by an HTTP proxy such as Oracle HTTP Server. The HTTP proxy serves static content and proxies other products integrated into EPM Workspace. Content aggregation typically occurs on the client-side. EPM Workspace does not aggregate content within its own process; that is, EPM Workspace does not function as a portal.

**15.1.7.1.1 State Information** EPM Workspace sessions are regular J2EE sessions that are held in memory only. EPM Workspace is stateful in its session use and the sessions are not serializable to any DB.

**15.1.7.1.2 Runtime Processes** EPM Workspace has a normal J2EE runtime. Users typically connect to EPM Workspace, log on to the system, and access another product, such as Oracle Hyperion Financial Reporting. There are no background threads or other forms of server-side services or jobs. All processing is handled in the context of the HTTP request.

EPM Workspace uses a set of servlets, implemented directly or as JSPs, and servlet filters. There are no EJBs or JMS usage. Database access is through a standard JNDI lookup of a JDBC data source.

**15.1.7.1.3 Process Lifecycle** EPM Workspace is managed as a standard J2EE Web application deployed to the WebLogic application server. It is managed, started, and terminated by its application server.

**15.1.7.1.4 Request Flow** Requests are made to EPM Workspace by means of the HTTP proxy, which is the only means for end-users to access EPM Workspace. EPM Workspace does not make requests against other Web applications or databases other than its repository database and the EPM Registry database.

**15.1.7.1.5 External Dependencies** EPM Workspace relies upon its repository database and the Hyperion Registry; both are in the same schema.

Most EPM products rely on EPM Workspace to provide initial authentication and a containing user interface.

**15.1.7.1.6 Configuration Artifacts** The entire EPM Workspace configuration is stored in the Hyperion Registry. Individual user preferences are stored in the EPM Workspace repository.

**15.1.7.1.7 Deployment Artifacts** EPM Workspace has no deployment artifacts other than the JDBC data source.

**15.1.7.1.8 Log Files** Go to the `logs/workspace` subdirectory in the managed server domain home to view `workspace.log` and `Framework.log`.

## 15.1.8 Oracle EPM Workspace High Availability Concepts

This section provides conceptual information about using EPM Workspace in a high availability environment.

### 15.1.8.1 Workspace High Availability Architecture

You can cluster both the HTTP proxy and the J2EE deployment of EPM Workspace with standard WebLogic Server clustering mechanisms and standard use of the application server deployment and data sources. EPM Workspace high availability requires sticky sessions; session replication and failover are not supported.

All clustered versions must be identical. All updates to cluster members must be made at the same time.

The load-balancer hardware or HTTP proxy allocates requests in a round-robin rotation. Requests fail if the repository database or Hyperion Registry is lost. When a request fails, subsequent requests attempt to initialize EPM Workspace and restore its database connectivity.

Database use increases markedly when an EPM Workspace instance is started and initialized. If there are many deployments, the cluster startup may be faster if you stagger startup of the instances.

**15.1.8.1.1 Shared Files and Directories** EPM Workspace has no shared files or directories except for Oracle JRF and EPM common files.

**15.1.8.1.2 Cluster-Wide Configuration Changes** Product binaries and deployment information generated by WebLogic Server are stored in the file system. The configuration is stored in the Hyperion Registry. All Workspace instances share the same configuration via the Hyperion Registry. You can change the configuration through the EPM Workspace administrative user interface or with explicit updates to the registry outside EPM Workspace.

Normally, the file system artifact changes only for a patch or for reconfiguration of the deployment or clustering that requires Oracle WebLogic Server file changes. Configuration changes do not require files system updates. Cluster members see the updated configuration at the next reinitialization. Each cluster member must be directed to reinitialize itself; the Hyperion Registry is not polled.

#### **15.1.8.2 Protection from Failures and Expected Behaviors**

Unfinished requests and sessions that are on the failed node are lost. Subsequent requests from the browser that are routed to a different node are intercepted and require re-authentication to Workspace. There is no provision to migrate any sessions for a single node failure; there is a feature for administrators to send a message to Workspace users at login that could inform them of downtime, etc.

---

**Note:** Database outages affect BI Workspace sessions. With RAC database backends, if a database instance has an outage, BI workspace sessions connected to the database generate an error and may need to restart.

---

#### **15.1.8.3 Troubleshooting**


Go to the `logs/workspace` subdirectory in the managed server domain home to view `Workspace.log` and `Framework.log`.

### **15.1.9 Oracle Hyperion Financial Reporting Component Architecture**

Oracle Hyperion Financial Reporting (Financial Reporting) is a powerful tool for designing and presenting analytic data graphically. You can design traditional financial report formats such as cash management reports, profit and loss statements, and balance sheets. You can also design nontraditional formats for financial or analytic data that include text and graphics.

Use Financial Reporting to create and deliver financial reports to the General Ledger system. Reporting administrators use EPM Workspace with financial reports to create and modify reports, books, and batch definitions and to schedule batch jobs via the reporting scheduler user interface within EPM Workspace. Reporting users run and view reports using the General Ledger user interface; they can launch reports in an IFrame within the General Ledger browser window.

[Figure 15–6](#) shows the Financial Reporting single instance architecture.

**Figure 15–6 Financial Reporting Single Instance Architecture**

### 15.1.9.1 Oracle Hyperion Financial Reporting Component Characteristics

Financial Reporting contains servlets and also contains JSPs for UI and XML generation. It also provides a web service API (JRF based) for use from ESS scheduler implementation for FR (FREssClient.ear). FRWebApp also contains RMI objects which are discovered and distributed via HTTP calls that replace the typical RMI registry, allowing multiple instances per machine with lookup via proxy request to distribute the load across instances.

PrintServer is a stand-alone RMI server (default port 8297) which is registered with the Web application, allowing a Web request for a PDF to locate a print server and generate PDF results.

**15.1.9.1.1 State Information** EPM Workspace browser is either stateless or per session. Connections to Oracle BI Presentation Service cache and restore if they time out. WebLogic pooling handles RDBMS and tries to reconnect in case of failure when making SQL calls. By default, Essbase connections cache for 5 minutes (configurable).

Protocol:

- RDBMS – Managed by JNDI connection configured by WebLogic.
- Oracle BI Presentation Services Web Services URL – Configured in Hyperion Registry.
- EPM Workspace – Web user interface container for Financial Reporting (All interactions with Financial Reporting are in the browser with JavaScript.)

- Essbase/Provider Services – Financial Reporting connects directly to Essbase by way of sockets. Lookup is direct or by way of Provider Services HTTP calls.

#### 15.1.9.1.2 Runtime Processes ■ This section describes Financial Reporting processes at runtime.

- Users launch report or book from the list of available reports. Users can select formats that include HTML, PDF, or Microsoft Office document).
  - Financial Reporting loads a tracking frame, which identifies what report instances are used and catches unloading events for the IFrame to clean up server instances. This frame loads another frame for executing and viewing the report or book.
  - Financial Reporting Web Application loads the report or book design and runs it against Essbase to generate output in the requested format.
    - \* Financial Reporting loads the design from the catalog by means of Oracle BI Presentation Services calls.
    - \* Financial Reporting connects to Essbase either directly or by way of a Provider Services URL.
    - \* Financial Reporting queries Essbase based on the report or book design and generates output in the requested format.

For HTML and Microsoft Office formats, Financial Reporting generates HTML and delivers the associated mime type to the client to allow launching in Microsoft Office when requested.

For PDF format, the Financial Reporting application server makes an RMI call to an available print server to generate PDF content, then returns the PDF file to the browser client.

---

**Note:** Configure and register multiple print servers. If a print server fails or goes offline, Financial Reporting moves to another print server.

---

- EPM Workspace-initiated report or book execution and viewing (HTML, PDF, Microsoft Office)
  - Initiated from a browser
  - XML, HTML, JavaScript, and Bindows – User interface components shown from EPM Workspace, Financial Reporting, and Oracle BI EE
  - Failover not supported; users must restart browser.
- Smart View importing of reports:
  - Initiated from Microsoft Office
  - XML, HTML, JavaScript
  - Failover supported

Report previews appear in an Internet Explorer-based control that shows the report in HTML before returning the Microsoft Office-formatted HTML results.

- Book and batch design and modification (available only from EPM Workspace)
  - Initiated from a browser

- XML, HTML, JavaScript, and Bindows
- Failover not supported; users must restart browser.

Editing book and batch designs uses a complex user interface of bindows components to present and manipulate book and batch contents.

- Financial Reporting opens and saves design objects in the Oracle BI Presentation Services repository by means of web services APIs.
  - Financial Reporting connects to design- specified Essbase cubes through direct Essbase connection or a Provider Services URL (to allow for failover).
  - Financial Reporting uses Bindows user interface components and JavaScript to present the design to Reports Designer to allow for changes to the book or batch design.
- Report Design and Modification (available from Financial Reporting Studio)
    - Initiated from Financial Reporting Studio (Windows 32-bit executable)
    - Visual Basic/Visual C++/Java application
    - Failover not supported, users must restart Financial Reporting Studio
    - Connection and user authentication: Report Designer enters credentials and URL of the Financial Reporting server (normally address of the Oracle HTTP Server proxy). Financial Reporting Studio makes an HTTP call to the Web application to get an RMI stub object; afterward, all communication is by way of RMI.
    - All access to repository and data sources is through remote RMI calls to Financial Reporting Web Application RMI objects, which make calls to repository and data sources. No direct connection is made from Financial Reporting Studio to the repository or Essbase.
    - All access to RDBMS is through remote calls to Financial Reporting Web Application RMI objects, which make the JDBC calls. No direct connection is made from FRStudio to the JDBC datasources.
  - Batch scheduling (Available from EPM Workspace; schedule accessible and modifiable via ESS EM):
    - Initiated from a browser
    - HTML, JavaScript, and Bindows
    - Failover not supported; users must restart browser.

Users open a wizard to select a batch to schedule and job attributes such as point of view, execution time/frequency of execution, and output format.

The Financial Reporting Web Application uses ESS EJB APIs to create and schedule the job according to the user's instructions. ESS RDBMS tables store the job and request.

- Batch Execution (ESS and Financial Reporting backend code)
  - Initiated by ESS
  - Invokes Webservice calls to Financial Reporting
  - Failover supported (Webservice failure during call causes batch result state of error)

- ESS picks up the Financial Reporting scheduled job and calls the Financial Reporting EssClient, which deploys with ESS-dedicated servers in the FSCM domain.
- FREssClient extracts the URL of the target Financial Reporting application server and makes a Webservice call to the Financial Reporting application server to execute the job, passing the XML defined as the job definition.
- Financial Reporting Web Application runs the job according to the parameters defined when the user scheduled the job and returns a status code to FREssClient.
- FREssClient returns the success or failure status to the ESS caller.

---

**Note:** Financial Reporting keeps job results in RDBMS tables. However, ESS determines when a job runs and the FREssClient connects to only one Financial Reporting Web Application. This prevents multiple Financial Reporting servers from processing issues of the same job; ESS manages how it prevents multiple ESS servers from processing the same job.

---

- Batch Results viewer (Available from EPM Workspace)
  - Initiated from a browser
  - HTML, JavaScript, and Bindows
  - Failover not supported; users must restart browser.
  - User gets a list of executed and pending jobs; Financial Reporting Web Application gets the list from ESS in-process EJBs.
  - User can retrieve result content from list of jobs; Financial Reporting Web Application gets results from Financial Reporting Scheduler RDBMS tables.
  - User can delete or replicate jobs, following the same process as for creating and scheduling jobs.

#### 15.1.9.1.3 External Dependencies ■ RDBMS – Hyperion Registry, Annotations Tables (Financial Reporting), Financial Reporting Scheduler Tables (Financial Reporting), ESS Job Tables (ESS)

- Oracle BI Presentation Services Repository (accessed through Web Services and EPM Workspace)
- EPM Workspace (for reporting administrators)
- Essbase/Provider Services
- Essbase-generated cubes
- ESS installed and configured

**15.1.9.1.4 Configuration Artifacts** Hyperion Registry stores configuration information that is not part of the overall system configuration; all Web application instances share Hyperion Registry. MBeans, which you can open and modify with Fusion Middleware Control, exposes configuration information. See the following in Fusion Middleware Control: `com.hyperion.FinancialReporting` and `com.hyperion.Annotations`.

**15.1.9.1.5 Deployment Artifacts** Financial Reporting has these Deployment Artifacts:

- `FRessClient.ear` – Contains the Enterprise Scheduler Service server implementation for Financial Reporting jobs, deployed to the `FinDomain` Enterprise Scheduler Service.
- `HReports.ear` – Primary Financial Reporting Web Application deployed to the `BIDomain`.

Shared Libraries deployed to the `BIDomain` that the Financial Reporting Web Application requires include: `epm-shared-libraries` `epm-fr-libraries` `epm-annotation-libraries` `epm-frweb-libraries` `epm-misc-libraries`.

**15.1.9.1.6 Log Files** Financial Reporting logs are in `$domain` `home/servers/${weblogic.Name}/logs/financialreporting`. Note that `${weblogic.Name}` is the managed server name, for example `bi_server1` or `bi_server2`.

## 15.1.10 Oracle Hyperion Financial Reporting High Availability Concepts

Oracle Hyperion Financial Reporting supports high availability for the following functionality:

- Report Execution and Viewing as HTML or PDF:
  - Prompt
  - POV changes
  - Page dimension drop-down (HTML only)
  - Annotation creation and viewing
  - Expansion
  - Drill-through (Not available in PDF)
- Book execution and viewing as HTML or PDF
- Export to Microsoft Office formats

### 15.1.10.1 Oracle Hyperion Financial Reporting High Availability Architecture

Figure 15–2 shows Essbase in a highly available Oracle BI EE deployment.

**15.1.10.1.1 Shared Files and Directories** There are no shared files or directories. All common properties and documents are either in the relational database or stored in the Oracle BI Presentation Services repository.

**15.1.10.1.2 Cluster-Wide Configuration Changes** Except for `PrintServer`, configuration is cached. Modifications affect only the server where they are made until the next restart of the other server. Print servers are not cached and can be added or removed at any time.

### 15.1.10.2 Protection from Failures and Expected Behaviors

Financial Reporting supports failover and session serialization if you also deploy the following components:

- Oracle Business Intelligence Enterprise Edition/Oracle BI Presentation Services – Financial Reporting uses web service connections to Oracle BI Presentation Services, with failover from the current Oracle BI Presentation Services connection fails to other catalog nodes, assuming these conditions:



- The catalog application server, Presentation Services, and file storage are clustered or replicated.
- The catalog location configured during applications provisioning references the proxy server and not a direct catalog Web application server.
- Essbase – Financial Reporting reports must target Essbase using the Provider Services URL, and Provider Services must be configured with active and passive Essbase servers. When Essbase fails, Financial Reporting connects through the Oracle Hyperion Provider Services URL to the backup server, and Essbase queries are reissued without interruption to the user request.
- RDBMS – The database must support failover. Connections are handled through the WebLogic JDBC configuration and connection pooling.
- Oracle Internet Directory/OAM
- Oracle HTTP Server is expected to provide failover and provide integration point for all components. Financial Reporting configuration should target Oracle BI Presentation Services through Oracle HTTP Server or a load-balanced URL and not directly target the catalog server.
- Oracle Hyperion Financial Reporting, Fusion Edition PrintServer – At least two PrintServer hosts should be deployed on Windows machines and configured with the system.

### 15.1.11 Allocation Manager Component Architecture and Characteristics

Use Allocation Manager to create, validate, deploy, and administer sophisticated calculations that solve Oracle Hyperion Financial Management, Fusion Edition, Oracle Hyperion Planning, Fusion Edition, and Oracle Essbase business problems. You access Allocation Manager within Oracle Enterprise Performance Management Workspace, Fusion Edition.

Allocation Manager is a Java Web application, deployed on WebLogic with multiple instances running in active-active mode. All instances of Allocation Manager are equal and can service any request. Allocation Manager is a stateful application that requires sticky routing. Session failover is not supported. Edits are made only within a single instance and are not saved if the instance fails.

Allocation Manager administration and discovery of external components is done through Hyperion Registry.

### 15.1.12 Oracle BI EE High Availability Configuration Steps

This section describes how to set up a two node highly available configuration for Oracle BI EE and BI Publisher. The architecture targeted for the configuration steps is [Figure 15-2](#).

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

#### 15.1.12.1 Prerequisite Steps Before Setting Up a High Availability Configuration for Oracle BI Enterprise Edition and BI Publisher

The following section describes prerequisite steps before setting up a high availability configuration for Oracle BI EE and Oracle BI Publisher.

**15.1.12.1.1 Database Prerequisites** Oracle BI EE supports the following database versions:

- Oracle Database 10g (10.2.0.4 or later for non-XE database)
- Oracle Database 11g (11.1.0.7 or later for non-XE database)

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

**15.1.12.1.2 VIP and IP Prerequisites** You configure the BI managed servers to listen on different virtual IPs, shown in [Table 15–7](#). This requires the provisioning of the corresponding VIP in the node and related host names in the DNS system in your network. Ensure that the different VIPs are available and are reachable before running the installation.

**Table 15–3 BI EE Virtual Hosts**

| Virtual IP | Virtual IP Maps To | Description                                                                                                                                                                                                                            |
|------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VIP1       | APPHOST1VHN1       | APPHOST1VHN1 is the virtual host name that maps to the listen address for BI_SERVER1 and fails over with server migration of this managed server. It is enabled on the node where BI_SERVER1 process is running (APPHOST1 by default). |
| VIP2       | APPHOST2VHN1       | APPHOST2VHN1 is the virtual host name that maps to the listen address for BI_SERVER2 and fails over with server migration of this managed server. It is enabled on the node where BI_SERVER2 process is running (APPHOST2 by default). |

**15.1.12.1.3 Shared Storage Prerequisites** For proper recovery in case of failure, store both JMS and transaction logs in a location that is accessible to all the nodes that can resume the operations after a failure in a managed server. This requires a shared storage location that can be referenced by multiple nodes. [Table 15–4](#) lists the contents of shared storage.

**Table 15–4 Contents of BI EE Shared Storage**

| Server                    | Type of Data         | Volume in Shared Storage | Directory                                              | Files                                                                                                                      |
|---------------------------|----------------------|--------------------------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| BI_SERVER1 and BI_SERVER2 | Transaction logs     | VOL2                     | ORACLE_BASE/admin/domain_name/cluster_name/tlogs       | Common location (stores decided by WebLogic Server)                                                                        |
| BI_SERVER1 and BI_SERVER2 | JMS stores           | VOL2                     | ORACLE_BASE/admin/domain_name/cluster_name/jms         | Common location, but individual store per server. For example: BipJmsStore for BI_SERVER1 and BipJmsStore2 for BI_SERVER2) |
| BI_SERVER1 and BI_SERVER2 | BI Server Repository | VOL1                     | ORACLE_BASE/admin/domain_name/cluster_name/ClusterRPD  | Common location                                                                                                            |
| BI_SERVER1 and BI_SERVER2 | BI Global Cache      | VOL1                     | ORACLE_BASE/admin/domain_name/cluster_name/GlobalCache | Common location                                                                                                            |

**Table 15–4 (Cont.) Contents of BI EE Shared Storage**

| Server                    | Type of Data                          | Volume in Shared Storage | Directory                                                                                                                                               | Files           |
|---------------------------|---------------------------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| BI_SERVER1 and BI_SERVER2 | BI Presentation Services Repository   | VOL1                     | ORACLE_BASE/admin/domain_name/cluster_name/catalog                                                                                                      | Common location |
| BI_SERVER1 and BI_SERVER2 | BI Publisher Configuration folder     | VOL1                     | ORACLE_BASE/admin/domain_name/cluster_name/bipublisher/config                                                                                           | Common location |
| BI_SERVER1 and BI_SERVER2 | BI Publisher Catalog repository       | VOL1                     | ORACLE_BASE/config/bipublisher/catalog<br><br>This is only required if Oracle BI Publisher - File System is selected for the BI Publisher Catalog type. | Common location |
| BI_SERVER1 and BI_SERVER2 | BI Publisher Scheduler Temp directory | VOL1                     | ORACLE_BASE/admin/domain_name/cluster_name/bipublisher/temp                                                                                             | Common location |

The shared storage can be a NAS or SAN device. The following is an example command based on a NAS over NFS device. Your options may be different from the ones specified in this section.

```
APPHOST1> mount nasfiler:/vol/volX/FMWshared MW_HOME-t nfs -o
rw,bg,hard,nointr,tcp,vers=3,timeo=300,rsize=32768,wsz=32768
```

**15.1.12.1.4 Clock Synchronization** You must synchronize all server clocks that participate in the cluster to within one second difference. To synchronize the clocks, use a single network time server and point each server to it. The procedure to point to the network time server is different from Windows to Linux. Refer to your operating system's manual.

**15.1.12.1.5 Installing and Configuring the Database Repository** This section describes how to install and configure the database repository.

#### Oracle Clusterware

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

#### Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

**Oracle Real Application Clusters**

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

You must install the 11g (11.1.1) Oracle Fusion Middleware Repository into an Oracle Real Application Clusters (Oracle RAC) database before you install the Oracle Fusion Middleware components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. You install RCU in its own, separate Middleware home.

Use the latest version of RCU to install the 11g (11.1.1.) Oracle Fusion Middleware Repository into an Oracle RAC database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

**15.1.12.1.6 Using RCU to Load the Business Intelligence Schemas in the Database** Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and Oracle BI schemas into an Oracle RAC database before you install Oracle Business Intelligence. Follow these steps:

1. Insert the Repository Creation Utility (RCU) DVD, and then start RCU from the bin directory in the RCU home directory.

```
prompt> cd RCU_HOME/bin
prompt> ./rcu
```

2. In the Welcome screen, click **Next**.
3. In the Create Repository screen, select **Create** to load component schemas into a database. Click **Next**.
4. In the Database Connection Details screen, enter connect information for your database:
  - **Database Type:** Select **Oracle Database**.
  - **Host Name:** Specify the name of the node on which the database resides. For the Oracle RAC database, specify the VIP name or one of the node names as the host name, for example:  
BIDBHOST1-VIP
  - **Port:** Specify the listen port number for the database: 1521
  - **Service Name:** Specify the service name of the database:  
biha.mycompany.com
  - **Username:** Specify the name of the user with DBA or SYSDBA privilege: SYS
  - **Password:** Enter the password for the SYS user.
  - **Role:** Select the database user's role from the list: SYSDBA (required by the SYS user).

Click **Next**.

5. In the Select Components screen:
  - Select **Create a new Prefix** and enter a prefix to use for the database schemas, for example, BIHA. You can specify up to six characters as a prefix. Prefixes

are used to create logical groupings of multiple repositories in a database. For more information, see *Oracle Fusion Middleware Repository Creation Utility User's Guide*.

**Tip:** Note the name of this schema because the upcoming steps require this information.

- Select the following components:
  - AS Common Schemas: Metadata Services (automatically selected)
  - Oracle Business Intelligence: Business Intelligence Platform

Click **Next**.

6. In the Schema Passwords screen, enter passwords for the main schema users and click **Next**.

Choose either **Use same passwords for all schemas** or **Specify different passwords for all schemas**, depending on your requirements.

Do not select **Use main schema passwords for auxiliary schemas**. The auxiliary passwords derive from the passwords of the main schema users.

**Tip:** Note the names of the schema passwords because the upcoming steps require this information.

7. In the Map Tablespaces screen, choose the tablespaces for the selected components, and click **Next**.
8. In the Summary screen, click **Create**.
9. In the Completion Summary screen, click **Close**.

**15.1.12.1.7 Configuring Virtual Server Names and Ports for the Load Balancer** This section describes the load balancer prerequisites for deploying an Oracle BI EE high availability environment.

### Load Balancers

Oracle BI EE uses a hardware load balancer when deployed in a high availability configuration with two Oracle HTTP Servers as web tier components.

### Virtual Server Names

bi.mycompany.com is a virtual server name that acts as the access point for all HTTP traffic to the runtime BI components, such as Oracle BI Publisher. Traffic to both SSL and non-SSL ports is configured, and typically non-SSL is redirected to SSL. Clients access this service using the address bi.mycompany.com:443. This virtual server is defined on the load balancer.

**15.1.12.1.8 Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2** This section describes how to install Oracle HTTP Server on WEBHOST1.

1. Verify that the servers meet the following requirements:
  - Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.

- Ensure that port 7777 is not in use by any service on WEBHOST1 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

**On UNIX:**

```
netstat -an | grep "7777"
```

**On Windows:**

```
netstat -an | findstr :7777
```

If the port is in use (if the command returns output identifying the port), you must free the port.

**On UNIX:**

Remove the entries for port 7777 in the `/etc/services` file and restart the services, or restart the computer.

**On Windows:**

Stop the component that is using the port.

- Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
- Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port number for Oracle HTTP Server):

```
The port for Oracle HTTP server
Oracle HTTP Server port = 7777
```

2. Start the Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:

**On UNIX:**

Issue the command: `./runinstaller`

**On Windows:**

Double-click `setup.exe`.

This displays the Specify Inventory Directory screen.

3. On the Specify Inventory Directory screen, enter values for the Oracle Inventory Directory and the Operating System Group Name. For example:

**Specify the Inventory Directory:** `/u01/app/oraInventory`

**Operating System Group Name:** `oinstall`

A dialog box appears with the following message:

"Certain actions need to be performed with root privileges before the install can continue. Please execute the script  
`/u01/app/oraInventory/createCentralInventory.sh` now from another window and then press "Ok" to continue the install. If you do not have the root privileges and wish to continue the install select the "Continue installation with local inventory" option"

Log in as root and run the `"/u01/app/oraInventory/createCentralInventory.sh"`

This sets the required permissions for the Oracle Inventory Directory and then brings up the Welcome screen.

---

**Note:** The Oracle Inventory screen is not shown if an Oracle product was previously installed on the host. If the Oracle Inventory screen does not appear for this installation, verify and see:

- If the `/etc/oraInst.loc` file exists.
  - If the file exists, that the Inventory directory listed is valid.
  - Whether or not the user performing the installation has write permissions for the Inventory directory.
- 

4. On the Welcome screen, click **Next**.
5. On the Select Installation Type screen, select **Install and Configure**, and click **Next**.
6. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.  
Click **Next**.
7. On the Specify Installation Location screen, specify:
  - **Oracle Middleware Home:** `ORACLE_BASE/product/fmw` (grayed out)
  - **Oracle Home Directory:** `Oracle_WT1`Click **Next**.
8. On the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Do not select **Associate Selected Components with Weblogic Domain**.Click **Next**.
9. On the Specify Component Details screen, enter the following values for WEBHOST1:
  - **Instance Home Location:** `ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1`
  - **Instance Name:** `web1`
  - **OHS Component Name:** `ohs1`Click **Next**.
10. On the Configure Ports screen:
  - If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
  - Enter the Oracle HTTP Server port, for example: `7777`Click **Next**.
11. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address and click **Next**.
12. On the Installation Summary screen, ensure that the selections are correct. If they are not, click **Back** and make the necessary fixes. After ensuring that the selections are correct, click **Install**.

13. On the Installation Progress screen on UNIX systems, a dialog appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

14. On the Configuration Progress screen, multiple configuration assistants are launched in succession; this process can be lengthy. Click **Next**.

15. On the Installation Complete screen, click **Finish** to exit.

Repeat the steps for WEBHOST2 and configure your load balancer with a pool containing both the WEBHOST1 and WEBHOST2 addresses.

**15.1.12.1.9 Validating Oracle HTTP Server** To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URLs in the browser:

```
http://WEBHOST1:7777/
http://WEBHOST2:7777/
```

If Oracle HTTP Server is set up correctly, the Oracle FMW Welcome screen appears in the browser.

### 15.1.12.2 Installing Oracle Business Intelligence Enterprise Edition for High Availability

This section describes how to install Oracle Fusion Middleware on all application tier nodes that run Oracle WebLogic Server and Oracle BI EE. Repeat the procedure (described below for APPHOST1) for installing WebLogic Server and Oracle BI EE in APPHOST2. The directory paths you use for binary files and domains when installing new nodes must be identical to those for the first node.

Install the following Oracle Fusion Middleware components:

- Oracle WebLogic Server
- Oracle Business Intelligence

In this section, a software-only install is performed then the `config.sh` configuration script is used to create and scale out the Oracle BI system.

**15.1.12.2.1 Installing Oracle WebLogic Server** Follow these steps to install Oracle WebLogic Server:

1. Start the installer for Oracle WebLogic Server from the installation media.
2. In the Welcome screen, click **Next**.
3. In the Choose Middleware Home Directory screen:
  - Select **Create a new Middleware Home**.
  - For the **Middleware Home Directory**, enter `ORACLE_BASE/product/fmw`

---

**Note:** `ORACLE_BASE` is the base directory under which Oracle products are installed. The recommended value is `/u01/app/oracle`.

---

Click **Next**.



4. In the Register for Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.
5. In the Choose Install Type screen, select **Typical** and click **Next**.
6. In the Choose Product Installation Directories screen, accept the directory `ORACLE_BASE/product/fmw/wlserver_10.3` for WebLogic Server and `ORACLE_BASE/product/fmw/coherence_3.6` for Oracle Coherence and click **Next**.
7. In the Installation Summary screen, click **Next**.  
The Oracle WebLogic Server software is installed.
8. In the Installation Complete screen, clear the Run Quickstart check box and click **Done**.

**15.1.12.2.2 Installing Oracle Fusion Middleware for Business Intelligence Enterprise Edition and Oracle BI Publisher** On the Linux platform, if the `/etc/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

1. Start the installer for Oracle Fusion Middleware Business Intelligence 11g from the installation media:

**On UNIX:**

```
APPHOST1>./runInstaller
```

**On Windows:**

```
APPHOST1> setup.exe
```

2. In the Specify Inventory Directory screen:
  - Enter `HOME/oraInventory`, where `HOME` is the home directory of the user performing the installation (this is the recommended location).
  - Enter the OS group for the user performing the installation.  
Click **Next**.
  - Follow the instructions on screen to execute `/createCentralInventory.sh` as root.  
Click **OK**.
3. In the Welcome screen, click **Next**.
4. In the Select Installation Type screen, select **Software Only Install** and click **Next**.
5. In the Prerequisite Checks screen, verify that all checks complete successfully, and click **Next**.
6. In the Specify Installation Location screen, select the previously installed Middleware Home from the drop-down list. For the Oracle Home directory, enter the directory name (`Oracle_BI1`).  
Click **Next**.
7. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address and click **Next**.

8. In the Summary screen, click **Install**.

The Oracle Fusion Middleware Business Intelligence 11g software is installed.

9. In the Installation Progress screen, click **Next**.
10. In the Complete screen, click **Finish**.

### 15.1.12.3 Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2

The BI domain uses virtual hostnames as the listen addresses for the BI managed servers. You must enable the VIPs, mapping each of these hostnames on the two BI machines (VIP1 on APPHOST1 and VIP2 on APPHOST2), and they must correctly resolve to the virtual hostnames in the network system used by the topology (either by DNS Server or hosts resolution).

### 15.1.12.4 Creating a Domain with the Administration Server and the First BI\_SERVER1 Managed Server

This section describes how to create a domain and the first WebLogic Server BI managed server using the Oracle Business Intelligence Configuration Assistant, Oracle WebLogic Server Administration Console, and Oracle Enterprise Manager. Ensure that the database where you installed the repository is running. For Oracle RAC databases, all the instances must be running.

---

---

**Note:** Oracle strongly recommends that you read the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

Run the Configuration Assistant from the Oracle home directory to create a domain containing the Administration Server and the managed server with BI Publisher:

1. Change the directory to the location of the Configuration Assistant:

```
APPHOST1> cd ORACLE_HOME/bin
```

2. Start the Configuration Assistant:

**On UNIX:**

```
APPHOST1> ./config.sh
```

**On Windows:**

```
APPHOST1> config.cmd
```

3. In the Welcome screen, click **Next**.
4. In the Prerequisite Checks screen, verify that all checks complete successfully, and click **Next**.
5. In the Create or Scale Out BI System screen, select **Create New BI System** and specify the following:
  - **User Name:** weblogic
  - **User Password:** Enter the weblogic user password.
  - **Domain Name:** *bifoundation\_domain*Click **Next**.
6. In the Specify Installation Location screen, enter:

- **Middleware Home:** *ORACLE\_BASE/product/fmw* (grayed out)
- **Oracle Home:** *ORACLE\_BASE/product/fmw/Oracle\_BI1* (grayed out)
- **WebLogic Server Home:** *ORACLE\_BASE/product/fmw/wlserver\_10.3* (grayed out)
- **Domain Home:** *ORACLE\_BASE/product/fmw/user\_projects/domain/bifoundation\_domain*

---

**Note:** The Domain Home must end with the domain name.

---

- **Instance Home:** *ORACLE\_BASE/product/fmw/instance1*
- **Instance Name:** *instance1*

Click **Next**.

7. In the Configure Components screen, select:

- Oracle Business Intelligence
  - Business Intelligence Enterprise Edition
  - Business Intelligence Publisher

Click **Next**.

8. In the BIPLATFORM Schema screen, enter:

- **Database Type:** Oracle Database
- **Connect String:**  
BIDBHOST1:1521:BIDB1^BIDBHOST2:1521:BIDB2@BIHA.MYCOMPANY.COM
- **BIPLATFORM Schema Username:** *BIHA\_BIPLATFORM*
- **BIPLATFORM Schema Password:** Enter the password for the BIPLATFORM schema user.

Click **Next**.

9. In the MDS Schema screen, verify the information and click **Next**.

10. In the Configure Ports screen, select one of the following:

- Auto Port Configuration
- Specify Ports using Configuration File

Click **Next**.

11. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address.  
Click **Next**.

12. In the Summary screen, click **Configure**.

13. In the Configuration Progress screen, verify that all the Configuration Tools have completed successfully and click **Next**.

14. In the Complete screen, click **Finish**.

### 15.1.12.5 Creating boot.properties for the Administration Server on APPHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password.

To create a boot.properties file for the Administration Server on APPHOST1:

1. Go to the following directory:

```
ORACLE_BASE/product/fmw/user_projects/domains/bifoundation_domain/servers/
AdminServer/security
```

2. Enter the following lines in the file, save the file, and close the editor:

```
username=Admin_username
password=Admin_password
```

---

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible for the entries to be encrypted.

---

---

### 15.1.12.6 Starting and Validating the Administration Server on APPHOST1

This section describes procedures for starting and validating the Administration Server on APPHOST1.

**15.1.12.6.1 Starting the Administration Server on APPHOST1** To start the Administration Server on APPHOST1, run the following commands:

```
APPHOST1> cd MW_HOME/user_projects/domains/bifoundation_domain/bin
APPHOST1> ./startWebLogic.sh
```

**15.1.12.6.2 Validating the Administration Server** Verify that the Administration Server is properly configured:

1. In a browser, go to <http://APPHOST1:7001/console> and log in as the administrator.
2. Verify that the BI\_SERVER1 managed server is listed.
3. Verify that the bi\_cluster cluster is listed.
4. Verify that you can access Enterprise Manager at <http://APPHOST1:7001/em>.

### 15.1.12.7 Setting the Listen Address for BI\_SERVER1 Managed Server

To set the managed server listen address:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.

5. Select **bi\_server1** in the **Names** column of the table.

The Setting page for bi\_server1 is displayed.

6. Set the Listen Address to APPHOST1VHN1. Click **Save**.
7. Click **Activate Changes**.
8. Restart the BI\_SERVER1 managed server to make the changes take effect.

- a. In the Summary of Servers screen, select the **Control** tab.
  - b. Select `bi_server1` in the table and then click **Shutdown**.
9. Restart the BI System components. For example:

```
cd ORACLE_BASE/admin/instance1/bin
./opmnctl restartproc
```

**15.1.12.7.1 Updating the Oracle BI Publisher Scheduler Configuration** Follow the steps in this section to update the WebLogic JNDI URL for the Oracle BI Publisher Scheduler.

---

**Note:** See Missing JMS Instances on Oracle BI Publisher Scheduler Diagnostics Page for information about keeping clocks synchronized in Oracle BI Publisher Scheduler.

---

**To update the Oracle BI Publisher Scheduler configuration:**

1. Log in to Oracle BI Publisher at the following URL:  
`http://APPHOST1VHN1:9704/xmlpserver`
2. Click the **Administration** link.
3. Click **Scheduler Configuration** under System Maintenance.
4. Update the **WebLogic JNDI URL** under JMS Configuration, as follows:  
`t3://APPHOST1VHN1:9704`
5. Click **Apply**.
6. Check the Scheduler status from the Scheduler Diagnostics tab.

**15.1.12.8 Disabling Host Name Verification for the BI\_SERVER1 Managed Server**

You must disable host name verification if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you do not configure the server certificates, you receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and validating the topology, and enable it again after you complete the high availability topology configuration.

To disable host name verification:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window. Click **Servers**.
4. Select `bi_server1` in the **Names** column of the table.
5. Open the **SSL** tab.
6. Expand the **Advanced** section of the page.
7. Set **Hostname Verification** to **None**. Click **Save**.
8. Click **Activate Changes**.
9. The changes do not take effect until the `BI_SERVER1` managed server restarts.
  - a. In the Summary of Servers screen, select the **Control** tab.

- b. Select bi\_server1 in the table and then click **Shutdown**.
10. Restart the BI System components, for example:

```
cd ORACLE_BASE/admin/instance1/bin
./opmnctl restartproc
```

### 15.1.12.9 Starting the System in APPHOST1

This section describes how to start Node Manager on APPHOST1 and how to start and validate the BI\_SERVER1 managed server on APPHOST1.

**15.1.12.9.1 Starting Node Manager on APPHOST1** Usually, Node Manager starts automatically when config.sh completes. If Node Manager is not running for some reason, start it on APPHOST1 using these commands:

```
APPHOST1> cd WL_HOME/server/bin
APPHOST1> ./startNodeManager.sh
```

**15.1.12.9.2 Starting and Validating the BI\_SERVER1 Managed Server** To start the BI\_SERVER1 managed server and check that it has the correct configuration:

1. Start the `bi_server1` managed server using Oracle WebLogic Server Administration Console, as follows:
  - a. Expand the **Environment** node in the Domain Structure window.
  - b. Choose **Servers**.
  - c. Click the **Control** tab.
  - d. Select **bi\_server1** and then click **Start**.
2. Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming", wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors
3. When `BI_SERVER1` starts, the following URLs become available:
  - Access `http://APPHOST1VHN1:9704/wsm-pm` to verify the status of Web Services Manager. Click **Validate Policy Manager**. A list of policies and assertion templates available in the data appears.

**Note:** The configuration is incorrect if no policies or assertion templates appear.

- Access <http://APPHOST1VHN1:9704/xmlpserver> to verify the BI Publisher application status.

**15.1.12.9.3 Starting and Validating the BI EE System Components on APPHOST1** You can control Oracle Business Intelligence system components using `opmnctl` commands.

To start the BI EE system components using the OPMNCTL command line:

1. Go to the directory that contains the OPMN command line tool.

---

**Note:** The OPMN command line tool is in the *ORACLE\_INSTANCE/bin* directory.

---

2. Run the `opmnctl` command to start the BI EE system components:

- `opmnctl startall`

This command starts OPMN and all BI EE system components.

- `opmnctl start`

This command starts OPMN only.

- `opmnctl startproc ias-component=componentName`

This command starts a particular system component. For example, where `coreapplication_obips1` is the BI Presentation Services:

```
opmnctl startproc ias-component=coreapplication_obips1
```

3. Check the status of the BI EE system components:

```
opmnctl status
```

#### 15.1.12.10 Configuring Oracle BI EE

This section describes how to configure Oracle BI EE.

**15.1.12.10.1 Setting the Location of the Shared Oracle BI Repository** Perform the following steps in Fusion Middleware Control:

1. Log into Fusion Middleware Control.
2. Expand the **Business Intelligence** node in the *Farm\_domain\_name* window.
3. Click **coreapplication**.
4. Click **Deployment**, then click **Repository**.
5. Click **Lock and Edit Configuration**.
6. Select **Share Repository** and specify the **Shared Location** for the Oracle BI Repository. In a Windows environment, the Enterprise Manager Fusion Middleware Control requires that a UNC path be used to define the shared location of the Repository.
7. Click **Apply**.
8. Click **Activate Changes**.

**15.1.12.10.2 Setting the Shared Global Cache for BI Server** Perform the following steps in Fusion Middleware Control:

1. Log into Fusion Middleware Control.
2. Expand the **Business Intelligence** node in the *Farm\_domain\_name* window.
3. Click **coreapplication**.
4. Click **Capacity Management**, then click **Performance**.
5. Click **Lock and Edit Configuration**.

6. In the **Global Cache** section, specify the shared location for the Oracle BI Server Global Cache and specify 250 MB for the global cache size. In a Windows environment, the Enterprise Manager Fusion Middleware Control requires that a UNC path be used to define the shared location of the Oracle BI Server Global Cache.
7. Click **Apply**.
8. Click **Activate Changes**.

**15.1.12.10.3 Setting the Scheduler Script Path and Default Script Path** if you use server-side scripts with Oracle BI Scheduler, it is recommended that you configure a shared directory for the scripts so that they can be shared by all Oracle BI Scheduler components in a cluster.

Perform these steps only if you are using server-side scripts.

To share Oracle BI Scheduler scripts:

1. Copy the default Oracle BI Scheduler Scripts (for example, `ORACLE_INSTANCE/bifoundation/OracleBISchedulerComponent/coreapplication_obisch1/scripts/common`) and custom Oracle BI Scheduler scripts (for example, `ORACLE_INSTANCE/bifoundation/OracleBISchedulerComponent/coreapplication_obisch1/scripts/scheduler`) to the shared BI Scheduler scripts location.
2. Update the `SchedulerScriptPath` and `DefaultScriptPath` elements of the Oracle BI Scheduler `instanceconfig.xml` file.
  - `SchedulerScriptPath`: Refers to the path where Oracle BI Scheduler-created job scripts are stored. Change this to the path of the shared BI Scheduler scripts location.
  - `DefaultScriptPath`: Specifies the path where user-created job scripts (not agents) are stored. Change this to the path of the shared BI Scheduler scripts location.

You must update this file for each Oracle BI Scheduler component in the deployment

**15.1.12.10.4 Setting the Location of the Shared Oracle BI Presentation Catalog** Each Presentation Services instance loads the Oracle BI Presentation Catalog from the catalog location specified in Fusion Middleware Control.

Perform the following steps:

1. Copy your existing (locally published) Oracle BI Presentation Catalog to the shared location. An example of a locally published catalog is:

```
ORACLE_INSTANCE/bifoundation/OracleBIPresentationServicesComponent/
coreapplication_obipsn/catalog/SampleAppLite
```

You must perform this step before designating the **Catalog Location** from Fusion Middleware Control.

If you plan to use the `SampleAppLite` catalog mentioned as an example in this section as the shared catalog, make sure to copy it from `APPHOST1`.

2. Log into Fusion Middleware Control.
3. Expand the **Business Intelligence** node in the `Farm_domain_name` window.
4. Click **coreapplication**.



5. Click **Deployment**, then click **Repository**.
6. Click **Lock and Edit Configuration**.
7. Specify the **Catalog Location** for the shared Oracle BI Presentation Catalog.  
In a Windows environment, specify a UNC path name.

---

**Note:** In a Windows environment, you must specify shared storage using the Universal Naming Convention (UNC). UNC is a PC format for specifying resource locations on a local area network. UNC uses the format `\\server-name\shared-resource-path-name`.

---

8. Click **Apply**.
9. Click **Activate Changes**.

#### 15.1.12.11 Setting Server Configuration Options

Follow these steps to set server configuration options for Oracle BI Publisher:

1. Copy over the contents of `DOMAIN_HOME/config/bipublisher/repository` to the shared configuration folder location.
2. Log into BI Publisher with Administrator credentials and select the **Administration** tab.
3. Under **System Maintenance**, select **Server Configuration**.
4. In the Path field, enter the path of the shared location for the Configuration Folder.
5. Apply your changes and restart your BI Publisher application.

#### 15.1.12.12 Scaling Out the BI System on APPHOST2

Follow the steps in [Section 15.1.12.2, "Installing Oracle Business Intelligence Enterprise Edition for High Availability"](#) to perform a software-only install of Oracle WebLogic Server and Oracle Business Intelligence on APPHOST2.

Then run the Configuration Assistant from the `ORACLE_HOME` directory to scale out the BI System.

1. Change the directory to the location of the Configuration Assistant:  

```
APPHOST2> cd ORACLE_HOME/bin
```
2. Start the Configuration Assistant:  

```
APPHOST2> ./config.sh
```
3. In the Welcome screen, click **Next**.
4. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.  
Click **Next**.
5. The Create or Scale out BI System screen opens. Select **Scale Out BI System** and then enter the following values:
  - **Host Name:** ADMINHOST
  - **Port:** 7001

- **User name:** weblogic
- **User Password:** Enter the password for the weblogic user.

Click **Next**.

6. In the Scale Out BI System Details screen, enter:
  - **Middleware Home:** *ORACLE\_BASE/product/fmw* (grayed out)
  - **Oracle Home:** *Oracle\_BI1* (grayed out)
  - **WebLogic Server Home:** *ORACLE\_BASE/product/fmw/wlserver\_10.3* (grayed out)
  - **Domain Home:** *ORACLE\_BASE/product/fmw/user\_projects/domain/bifoundation\_domain*
  - **Applications Home:** *ORACLE\_BASE/admin/domain\_name/msserver/applications*
  - **Instance Home:** *ORACLE\_BASE/product/fmw/instance2*
  - **Instance Name:** *instance2* (grayed out)

Click **Next**.

7. In the Configure Ports screen, select one of the following:

- Auto Port Configuration
- Specify Ports using Configuration File

Click **Next**.

8. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address.  
Click **Next**.

9. In the Summary screen, click **Configure**.

10. In the Configuration Progress screen, verify that all the Configuration Tools have completed successfully and click **Next**.

11. In the Complete screen, click **Finish**.

Usually Node Manager starts automatically when `config.sh` completes. If Node Manager is not running, run these commands to start it on APPHOST2:

```
APPHOST2> cd WL_HOME/server/bin
APPHOST2> ./startNodeManager.sh
```

### 15.1.12.13 Scaling Out BI System Components

To scale out BI System Components, complete these steps:

1. Log into Fusion Middleware Control.
2. Expand the **Business Intelligence** node in the *Farm\_domain\_name* window.
3. Click **coreapplication**.
4. Click **Capacity Management**, then click **Scalability**.
5. Click **Lock and Edit Configuration**.
6. For the APPHOST2 instance2 Oracle instance, increment the Oracle Business Intelligence components by 1:

- BI Servers
  - Presentation Services
  - JavaHosts
7. Change the **Port Range From** and **Port Range To** to be the same as the APPHOST1 instance1 Oracle instance.
  8. Click **Apply**. The following message appears:  

Warning: Multiple Presentation Services are configured;  
adjust the catalog location to point to a shared location.
  9. Click **OK**. You configure shared location in a later step.
  10. Click **Activate Changes**. Click **Restart** to apply recent changes.

#### 15.1.12.14 Making Singleton Components Active-Passive

Perform the following steps in Fusion Middleware Control:

1. Log into Fusion Middleware Control.
2. Expand the **Business Intelligence** node in the *Farm\_domain\_name* window.
3. Click **coreapplication**.
4. Click **Capacity Management**, then click **Availability**.
5. Click **Lock and Edit Configuration** to activate the Active/Passive Configuration section of the **Availability** tab.
6. Specify the Passive Host/Instance for BI Scheduler and BI Cluster Controller.
7. Click **Apply**.  

Under **Potential Single Points of Failure**, it should report "No problems - all components have a backup."
8. Click **Activate Changes**.
9. Click **Restart** to apply recent changes.

#### 15.1.12.15 Setting the Listen Address for the BI\_SERVER2 Managed Server

To set the BI\_SERVER2 managed server listen address:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server2** in the **Names** column of the table.  

The Setting page for bi\_server2 opens.
6. Set the Listen Address to APPHOST2VHN1.
7. Click **Save**.
8. Click **Activate Changes**.
9. The changes do not take effect until the BI\_SERVER2 managed server restarts.
  - a. In the Summary of Servers screen, select the **Control** tab.
  - b. Select **bi\_server2** in the table and then click **Shutdown**.

10. Restart the BI System components, for example:

```
cd ORACLE_BASE/admin/instance2/bin
./opmnctl restartproc
```

**15.1.12.15.1 Updating the Oracle BI Publisher Scheduler Configuration on APPHOST1 and APPHOST2** Follow the steps in this section to update the WebLogic JNDI URL for the Oracle BI Publisher Scheduler.

**To update the Oracle BI Publisher Scheduler configuration:**

1. Log into Oracle BI Publisher at the following URLs:

```
http://APPHOST1VHN1:9704/xmlpserver
http://APPHOST2VHN1:9704/xmlpserver
```

2. Click the **Administration** link.
3. Click **Scheduler Configuration** under System Maintenance.
4. Update the **WebLogic JNDI URL** under JMS Configuration, as follows:

```
t3://APPHOST1VHN1:9704,APPHOST2VHN1:9704
```

5. Click **Apply**.
6. Check the Scheduler status from the Scheduler Diagnostics tab.

#### **15.1.12.16 Disabling Host Name Verification for the BI\_SERVER2 Managed Server**

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and validating the topology, and enable it again once the high availability topology configuration is complete.

To disable host name verification:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server2** in the **Names** column of the table.
6. Open the **SSL** tab.
7. Expand the **Advanced** section of the page.
8. Set **Hostname Verification** to **None**.
9. Click **Save**.
10. Click **Activate Changes**.

The changes don't take effect until the BI\_SERVER2 managed server restarts.

#### **15.1.12.17 Configuring Oracle BI for Microsoft Office SSO Properties**


To configure Oracle BI for Microsoft Office Single Sign-On (SSO) properties:

1. Validate the Oracle BI EE Office Server setup by accessing the following URLs:

- <http://APPHOST1VHN1:9704/biooffice/about.jsp>
- <http://APPHOST2VHN1:9704/biooffice/about.jsp>

The About Oracle BI EE Office Server page opens, as shown in [Figure 15–7](#).

**Figure 15–7 About Oracle BI EE Office Server Page**



2. Go to the Oracle BI EE Office Server directory. For example:

`DOMAIN_HOME/servers/managed_server/tmp/_WL_user/biooffice_11.1.1/cvsibb/war/WEB-INF`

If you are not sure how to locate the Oracle BI EE Office Server directory, check the **LogDir** parameter on the About Oracle BI EE Office Server page. The Oracle BI Enterprise Edition Office Server directory is the parent directory of the log directory.

3. On both APPHOST1 and APPHOST2, open `biooffice.xml` for editing and modify the BI Office properties shown in [Table 15–5](#).

**Table 15–5 Oracle BI for Microsoft Office Properties in biooffice.xml**

| Property Name   | Valid Value                                                                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|---------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SawBaseURL      | https://bi.mydomain.com:443/analytics/saw.dll<br>or<br>https://bi.mydomain.com:443/analytics-ws/saw.dll | Load Balancer Virtual Server Name URL for Oracle BI Presentation Services.<br><br><b>Important:</b> If SSO is enabled, enter the URL for the protected analytics servlet that you deployed when configuring Oracle BI for Microsoft Office to integrate with the SSO-enabled Oracle BI Server. Web services requests between the Oracle BI for Microsoft Office Server and Presentation Services use the URL that you specify here. |
| SawUseSSO       | 0 = No (Default)<br>1 = Yes                                                                             | Set this property to 1 if the Oracle Business Intelligence implementation is enabled for SSO.                                                                                                                                                                                                                                                                                                                                       |
| SawWebURLforSSO | https://bi.mydomain.com:443/analytics/saw.dll                                                           | When SSO is enabled, use this property to enter the public URL that allows external users to access Oracle Business Intelligence using SSO from Oracle BI for Microsoft Office.                                                                                                                                                                                                                                                     |

4. Restart the Oracle BI for Microsoft Office application, as follows:
  - a. Log in to the Administration Console.
  - b. Click **Deployments** in the Domain Structure window.
  - c. Select **biooffice(11.1.1)**.
  - d. Click **Stop**.
  - e. After the application stops, click **Start**.
5. Validate that the **SawBaseURL** parameter has been updated on the About Oracle BI EE Office Server page.


**15.1.12.17.1 Validating Oracle BI for Microsoft Office Configuration** Follow these steps to validate configuration for Oracle BI for Microsoft Office:

1. Log in to Oracle BI Presentation Services at:  
https://bi.mydomain.com:443/analytics
2. In the lower left pane, under the Get Started heading, select **Download BI Desktop Tools** and then select **Oracle BI for MS Office**.
3. Install Oracle BI for Microsoft Office by running the Oracle BI for Microsoft Office InstallShield Wizard.
4. Open Microsoft Excel or Microsoft Powerpoint.
5. From the **Oracle BI** menu, select **Preferences**.
6. In the Connections tab, select **New**.
7. Enter values for the following fields:
  - **Server Name:** Enter a name for the connection.
  - **BI Office Server:** Enter the URL for the Oracle BI for Microsoft Office Server.

- **Application Name:** Enter the Application Name that you defined for the Oracle BI for Microsoft Office Server when you deployed the Oracle BI for Microsoft Office Server application to WLS. The default name is **bioffice**.
- **Port:** Enter the Oracle BI for Microsoft Office Server port number.

Figure 15–8 shows the New Connection dialog.


**Figure 15–8 New Connection Dialog for Oracle BI for Microsoft Office**




8. Click **Test Connection** to test the connection between the add-in and the Oracle BI for Microsoft Office Server.

Successful connections receive a Test connection successful message, as shown in Figure 15–9.

**Figure 15–9 Test Connection Successful Message**



9. Log in as an Administrator (for example, weblogic) and validate that you can access the Oracle BI Task Pane, as shown in Figure 15–10.

**Figure 15–10 Oracle BI Task Pane in Microsoft Excel**

### 15.1.12.18 Configuring Oracle BI Publisher

This section describes how to configure Oracle BI Publisher.

**15.1.12.18.1 Setting Oracle BI Presentation Services Options** Follow these steps to configure Oracle BI Presentation Services Integration options:

1. Log into Oracle BI Publisher with Administrator credentials and select the Administration tab.
2. Under **Integration**, select **Oracle BI Presentation Services**.
3. Verify and update the following:
  - **Server Protocol:** http
  - **Server:** bi.mycompany.com
  - **Port URL Suffix:** analytics-ws/saw.dll: 80
4. Click **Apply**.
5. Restart your Oracle BI Publisher application.

**15.1.12.18.2 Setting Scheduler Configuration Options** Follow these steps to configure scheduler configuration options:

1. Log into BI Publisher with Administrator credentials and select the **Administration** tab.
2. Under **System Maintenance**, select **Scheduler Configuration**.
3. Select **Quartz Clustering** under the **Scheduler Selection**.
4. Apply your changes and restart your BI Publisher application.

**15.1.12.18.3 Setting the Oracle BI EE Data Source** The Oracle BI EE data source must point to the clustered BI Servers via the Cluster Controllers.



To set the Oracle BI EE data source in BI Publisher:

1. Log in to BI Publisher with Administrator credentials. Select the Administration tab.
2. Under **Data Sources**, select **JDBC Connection**.
3. Update the Oracle BI EE data source setting by changing the Connection String parameter to the following:

```
<Primary Cluster Controller Port>/PrimaryCCS=<Primary Cluster Controller Host>;PrimaryCCSPort=<Primary Cluster Controller Port>;SecondaryCCS=<Secondary Cluster Controller Host>;SecondaryCCSPort=<Secondary Cluster Controller Port>;
```

For example:

```
jdbc:oraclebi://APPHOST1:9706/PrimaryCCS=APPHOST1;PrimaryCCSPort=9706;SecondaryCCS=APPHOST2;SecondaryCCSPort=9706;
```

4. Select **Use System User**.
5. Click **Test Connection**. You should see Connection established successfully. Click **Apply**.

**15.1.12.18.4 Configuring JMS for Oracle BI Publisher** You must configure the location for all persistence stores to a directory visible from both nodes. Change all persistent stores to use this shared base directory.

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Services** node and then click the **Persistent Stores** node.
3. In the Change Center, click **Lock & Edit**.
4. Click on **BipJmsStore** and enter a directory that is located in the shared storage. This shared storage is accessible from both APPHOST1 and APPHOST2:

```
ORACLE_BASE/admin/domain_name/bi_cluster/jms
```

5. Click **Save** and **Activate Changes**.
6. In the Domain Structure window, expand the **Services** node and then click the **Persistent Stores** node.
7. In the Change Center, click **Lock & Edit**.
8. Click **New**, and then **Create File Store**.
9. Enter a name (for example, BipJmsStore2) and target BI\_SERVER2. Enter a directory that is located in shared storage so that it is accessible from both APPHOST1 and APPHOST2:

```
ORACLE_BASE/admin/domain_name/bi_cluster/jms
```

10. Click **OK** and activate the changes.
11. In the Domain Structure window, expand the **Services** node and then click the **Messaging > JMS Servers** node.
12. In the Change Center, click **Lock & Edit**.
13. Click **New**.
14. Enter a name (for example, BipJmsServer2) and in the **Persistence Store** drop-down list, select **BipJmsStore2** and click **Next**.

15. Select **BI\_SERVER2** as the target.
16. Click **Finish** and **Activate Changes**.
17. In the Domain Structure window, expand the **Services** node and then click the **Messaging > JMS Modules** node.
18. In the Change Center, click **Lock & Edit**.
19. Click **BIPJmsResource** and then click the Subdeployments tab.
20. Target the Subdeployment **BipJmsSubDeployment** to both BipJmsServer1 and BipJmsServer2.
21. Click **Finish** and **Activate Changes**.

To validate the JMS configuration performed for Oracle BI Publisher, perform the steps in [Section 15.1.12.7.1, "Updating the Oracle BI Publisher Scheduler Configuration."](#)

**15.1.12.18.5 Configuring a Default Persistence Store for Transaction Recovery** Each server has a transaction log, which stores information about committed transactions coordinated by the server that may not have been completed. WebLogic Server uses the transaction log when recovering from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to the server.

---

**Note:** Preferably, this location should be a dual-ported SCSI disk or on a Storage Area Network (SAN).

---

To set the location for the default persistence store for BI\_SERVER1:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. In the Domain Structure window, expand the **Environment** node and then click the **Servers** node.
4. Click **BI\_SERVER1** (represented as a hyperlink) in the **Name** column of the table.  
The settings page for BI\_SERVER1 opens with the Configuration tab active.
5. Click the **Services** tab.
6. In the Default Store section, enter the path to the folder where the default persistent stores will store its data files. The path directory structure is as follows:  
`ORACLE_BASE/admin/domain_name/bi_cluster/tlogs`
7. Click **Save** and **Activate Changes**.
8. Repeat these steps for the BI\_SERVER2 server.

---

**Note:** To enable Transaction Recovery Service migration, specify a location on a persistent storage solution that is available to other servers in the cluster. Both APPHOST1 and APPHOST2 must be able to access this directory and the directory must exist before you restart the server.

---

#### 15.1.12.19 Starting the System in APPHOST2

This section describes procedures for starting the system in APPHOST2.

**15.1.12.19.1 Starting Node Manager on APPHOST2** Usually, Node Manager is started automatically when config.sh completes. If Node Manager is not running for some reason, start the Node Manager on APPHOST2 by following the instructions in [Section 15.1.12.9.1, "Starting Node Manager on APPHOST1."](#)

**15.1.12.19.2 Starting and Validating the BI\_SERVER2 Managed Server** To start the BI\_SERVER2 managed server and check that it is configured correctly:

1. Start the bi\_server2 managed server using Oracle WebLogic Server Administration Console, as follows:
  - a. Expand the **Environment** node in the Domain Structure window.
  - b. Choose **Servers**.
  - c. Click the **Control** tab.
  - d. Select **bi\_server2** and then click **Start**.
2. Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming", wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors
3. When BI\_SERVER2 is started, the following URLs become available:
  - Access <http://APPHOST2VHN1:9704/wsm-pm> to verify the status of Web Services Manager. Click **Validate Policy Manager**. A list of policies and assertion templates available in the data opens.

---

**Note:** The configuration is incorrect if no policies or assertion templates appear.

---

- Access <http://APPHOST2VHN1:9704/xmlpserver> to verify the status of the Oracle BI Publisher application.

**15.1.12.19.3 Starting and Validating the Business Intelligence Enterprise Edition System Components** To start the BI Enterprise Edition system components on APPHOST2, repeat the steps in [Section 15.1.12.9.3, "Starting and Validating the BI EE System Components on APPHOST1"](#) on APPHOST2.

### 15.1.12.20 Configuring Oracle HTTP Server for the BI\_SERVERn Managed Servers

To enable Oracle HTTP Server to route to bi\_cluster, which contains the BI\_SERVERn managed servers, you must set the WebLogicCluster parameter to the list of nodes in the cluster:

1. On WEBHOST1 and WEBHOST2, add the following lines to the `ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/config/OHS/ohs1/mod_wl_ohs.conf` file:
 

```
BI Office
<Location /biooffice>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1VHN1:9704,APPHOST2VHN1:9704
 WLProxySSL ON
 WLProxySSLPassThrough ON
</Location>

<Location /bioofficeclient>
```

```
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

WSM-PM
<Location /wsm-pm>
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

BIEE Analytics
<Location /analytics>
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

<Location /mapviewer>
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

<Location /analytics-ws>
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

<Location /bimiddleware>
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>

BI Publisher
<Location /xmlpservlet>
SetHandler weblogic-handler
WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
WLProxySSL ON
WLProxySSLPassThrough ON
</Location>
```

2. Perform the same steps for the Oracle HTTP Server on WEBHOST2.
3. Restart Oracle HTTP Server on WEBHOST1 and WEBHOST2:

```
WEBHOST1> ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/bin/
opmnctlrestartproc ias-component=ohs1
```

```
WEBHOST2> ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/bin/
opmnctlrestartproc ias-component=ohs1
```

**15.1.12.20.1 Validating Access Through Oracle HTTP Server** In the Administration Console, verify that the BI Server status is **Running**. If the status is **Starting** or **Resuming**, wait for the server status to change to **Started**. If the console shows another status (such as **Admin**), check the server output log files for errors.

Verify that you can access these URLs:

- <http://WEBHOST1:7777/wsm-pm>
- <http://WEBHOST2:7777/wsm-pm>
- <http://WEBHOST1:7777/analytics>
- <http://WEBHOST2:7777/analytics>
- <http://WEBHOST1:7777/xmlpserver>
- <http://WEBHOST2:7777/xmlpserver>
- <http://WEBHOST1:7777/bioffice/about.jsp>
- <http://WEBHOST2:7777/bioffice/about.jsp>
- <http://WEBHOST1:7777/mapviewer>
- <http://WEBHOST2:7777/mapviewer>

Verify that you can access these URLs using your load balancing router address:

- <http://bi.mycompany.com:80/analytics>
- <http://bi.mycompany.com:80/xmlpserver>
- <http://bi.mycompany.com:80/wsm-pm>
- <http://bi.mycompany.com:80/bioffice/about.jsp>
- <http://bi.mycompany.com:80/mapviewer>

Follow these instructions to ensure that routing and failover from the HTTP Server to the bi\_cluster is working correctly:

1. While BI\_SERVER2 is running, stop BI\_SERVER1 from Oracle WebLogic Server Administration Console.
2. Access the following URLs and verify the appropriate functionality:
  - [WEBHOST1:7777/wsm-pm](http://WEBHOST1:7777/wsm-pm)
  - [WEBHOST1:7777/analytics](http://WEBHOST1:7777/analytics)
  - [WEBHOST1:7777/xmlpserver](http://WEBHOST1:7777/xmlpserver)
  - [WEBHOST1:7777/bioffice/about.jsp](http://WEBHOST1:7777/bioffice/about.jsp)
3. Start BI\_SERVER1 from Oracle WebLogic Server Administration Console.
4. Stop BI\_SERVER2.
5. Access the URLs in Step 2 above again and verify the appropriate functionality.

#### **15.1.12.21 Setting the Frontend HTTP Host and Port**

You must set the frontend HTTP host and port for the Oracle WebLogic Server cluster. To do this, follow these steps:

1. In the Change Center of the WebLogic Server Administration Console, click **Lock & Edit**.

2. In the left pane, choose **Environment** in the Domain Structure window and then choose **Clusters**. The Summary of Clusters page appears.
3. Select the `bi_cluster` cluster.
4. Select **HTTP**.
5. Set the values for the following:
  - **Frontend Host:** `bi.mycompany.com`
  - **Frontend HTTP Port:** 80
  - **Frontend HTTPS Port:** Leave this field blank

---

**Note:** Make sure this address is correct and available (the load balancer is up). An incorrect value, for example, `http://` in the address or trailing `/` in the hostname, may prevent the BI system from being accessible even when using the virtual IPs to access it.

---

6. Click **Save**.
7. To activate the changes, click **Activate Changes**.
8. Restart the servers to make the Frontend Host directive in the cluster effective.

#### 15.1.12.22 Configuring Server Migration for the BI\_SERVERn Servers

In this high availability topology, you must configure server migration for the `bi_server1` and `bi_server2` Managed Servers:

- Configure the `bi_server1` Managed Server to restart on APPHOST2 should a failure occur.
- Configure the `bi_server2` Managed Server to restart on APPHOST1 should a failure occur.

For this configuration, the `bi_server1` and `bi_server2` servers listen on specific floating IPs that WLS Server Migration fails over.

This section includes the following topics:

- Step 1: [Setting Up a User and Tablespace for the Server Migration Leasing Table](#)
- Step 2: [Creating a Multi Data Source Using the Oracle WebLogic Server Administration Console](#)
- Step 3: [Editing Node Manager's Properties File](#)
- Step 4: [Setting Environment and Superuser Privileges for the wlsifconfig.sh Script](#)
- Step 5: [Configuring Server Migration Targets](#)
- Step 6: [Testing the Server Migration](#)

**15.1.12.22.1 Setting Up a User and Tablespace for the Server Migration Leasing Table** To set up a user and tablespace for the server migration leasing table:

1. Create a tablespace named `leasing`. For example, log on to SQL\*Plus as the `sysdba` user and run the following command:

```
SQL> create tablespace leasing logging datafile 'DB_
HOME/oradata/orcl/leasing.dbf' size 32m autoextend on next 32m maxsize 2048m
extent management local;
```

2. Create a user named `leasing` and assign to it the leasing tablespace:

```
SQL> create user leasing identified by welcome1;
SQL> grant create table to leasing;
SQL> grant create session to leasing;
SQL> alter user leasing default tablespace leasing;
SQL> alter user leasing quota unlimited on LEASING;
```

3. Create the leasing table using the leasing.ddl script:

- Copy the leasing.ddl file located in either the `WL_HOME/server/db/oracle/817` or the `WL_HOME/server/db/oracle/920` directory to your database node.
- Connect to the database as the **leasing** user.
- Run the leasing.ddl script in SQL\*Plus:

```
SQL> @Copy_Location/leasing.ddl;
```

#### 15.1.12.22.2 Creating a Multi Data Source Using the Oracle WebLogic Server Administration Console

This section describes how to create a multi data source.

##### Creating a Multi Data Source

You create a data source to each of the Oracle RAC database instances during the process of setting up the multi data source, both for these data sources and the global leasing multi data source. Note the following when creating a data source:

- Ensure that this is a non-XA data source.
- The names of the multi data sources are in the format of `<MultiDS>-rac0`, `<MultiDS>-rac1`, and so on.
- Use Oracle's Driver (Thin) Version 9.0.1, 9.2.0, 10, 11.
- Data sources do not require support for global transactions. Therefore, do *not* use any type of distributed transaction emulation/participation algorithm for the data source (do not choose the **Supports Global Transactions** option, or the **Logging Last Resource, Emulate Two-Phase Commit**, or **One-Phase Commit** options of the **Supports Global Transactions** option), and specify a service name for your database.
- Target these data sources to the Oracle I/PM cluster.
- Make sure the initial connection pool capacity of the data sources is set to 0 (zero). To do this, select **Services**, then **JDBC**, and then **Datasources**. In the Datasources screen, click the **Datasource Name**, then click the **Connection Pool** tab, and enter 0 (zero) in the **Initial Capacity** field.

To create a multi data source:

- From the Domain Structure window in the Administration Console, expand the **Services** node, then click **Data Sources**. The Summary of JDBC Data Sources page appears.
- In the Change Center, click **Lock & Edit**.
- Click **New**, then click **Multi Data Sources**. The Create a New JDBC Multi Data Source page appears.
- Enter `leasing` as the name.
- Enter `jdbc/leasing` as the JNDI name.

6. Select **Failover** as algorithm (default).
7. Click **Next**.
8. Select `bi_cluster` as the target.
9. Click **Next**.
10. Select **non-XA driver** (the default).
11. Click **Next**.
12. Click **Create New Data Source**.
13. Enter `leasing-rac0` as the name. Enter `jdbc/leasing-rac0` as the JNDI name. Enter `oracle` as the database type. For the driver type, select Oracle Driver (Thin) for RAC server-Instance connection Version 10,11.

---

**Note:** When creating the multi data sources for the leasing table, enter names in the format of `<MultiDS>-rac0`, `<MultiDS>-rac1`, and so on.

---

14. Click **Next**.
15. Deselect **Supports Global Transactions**.
16. Click **Next**.
17. Enter the service name, database name (the RAC Node instance name, for example: `racdb1`, `racdb2`), host port, and password for your leasing schema. Click **Next**.
18. Click **Test Configuration** and verify that the connection works. Click **Next**.
19. Target the data source to `bi_cluster`.
20. Select the data source and add it to the right screen.
21. Click **Create a New Data Source** for the second instance of your Oracle RAC database, target it to the `bi_cluster`, and then repeat the steps for the second instance of your Oracle RAC database.
22. Add the second data source to your multi data source.
23. Click **Activate Changes**.

**15.1.12.22.3 Editing Node Manager's Properties File** This section describes how to edit the Node Manager properties file, `nodemanager.properties`. This needs to be done for the Node Managers in both nodes where server migration is being configured:

```
Interface=eth0
NetMask=255.255.255.0
UseMACBroadcast=true
```

- **Interface:** This property specifies the interface name for the floating IP (`eth0`, for example, on Linux).



---

**Note:** Do not specify the sub interface, such as `eth0:1` or `eth0:2`. This interface is to be used without the `:0`, or `:1`. The Node Manager's scripts traverse the different `:X` enabled IPs to determine which to add or remove. For example, the valid values in Linux environments are `eth0`, `eth1`, or, `eth2`, `eth3`, `ethn`, depending on the number of interfaces configured.

---

---

**Note:** For Windows, the Interface should be set to the Network Interface Name. For example: `Interface="Local Area Connection"`.

---

- **NetMask:** This property specifies the net mask for the interface for the floating IP. The net mask should be the same as the net mask on the interface; `255.255.255.0` is used as an example in this document.
- **UseMACBroadcast:** This property specifies whether or not to use a node's MAC address when sending ARP packets, that is, whether or not to use the `-b` flag in the `arping` command.

Verify in Node Manager's output (shell where Node Manager is started) that these properties are being used, or problems may arise during migration. You should see something like this in Node Manager's output:

```
...
StateCheckInterval=500
Interface=eth0 (Linux) or Interface="Local Area Connection" (Windows)
NetMask=255.255.255.0
...
```

---

**Note:** The steps below are not required if the server properties (start properties) have been properly set and Node Manager can start the servers remotely.

---

1. Set the following property in the `nodemanager.properties` file:
  - **StartScriptEnabled:** Set this property to `true`. This is required for Node Manager to start the managed servers using start scripts.
2. Start Node Manager on `APPHOST1` and `APPHOST2` by running the `startNodeManager.sh` script, which is located in the `WL_HOME/server/bin` directory.

---

**Note:** When running Node Manager from a shared storage installation, multiple nodes are started using the same `nodemanager.properties` file. However, each node may require different `NetMask` or `Interface` properties. In this case, specify individual parameters on a per-node basis using environment variables. For example, to use a different interface (`eth3`) in `APPHOSTn`, use the `Interface` environment variable as follows:  
`APPHOSTn> export JAVA_OPTIONS=-DInterface=eth3` and start Node Manager after the variable has been set in the shell.

---

**15.1.12.22.4 Setting Environment and Superuser Privileges for the wlsifconfig.sh Script** This section describes how to set environment and superuser privileges for the `wlsifconfig.sh` script:

---

**Note:** On Windows, the script is named `wlsifconfig.cmd` and users with the Administrator privilege can run it.

---

1. Ensure that your PATH environment variable includes these files:

**Table 15–6 Files Required for the PATH Environment Variable**

| File                             | Located in this directory                                                            |
|----------------------------------|--------------------------------------------------------------------------------------|
| <code>wlsifconfig.sh</code>      | <code>ORACLE_BASE/admin/domain_name/msserver/domain_name/bin/server_migration</code> |
| <code>wlscontrol.sh</code>       | <code>WL_HOME/common/bin</code>                                                      |
| <code>nodemanager.domains</code> | <code>WL_HOME/common</code>                                                          |

2. Grant sudo configuration for the `wlsifconfig.sh` script.
  - Configure sudo to work without a password prompt.
  - For security reasons, sudo should be restricted to the subset of commands required to run the `wlsifconfig.sh` script. For example, perform these steps to set the environment and superuser privileges for the `wlsifconfig.sh` script:
    - a. Grant sudo privilege to the WebLogic user ('oracle') with no password restriction, and grant execute privilege on the `/sbin/ifconfig` and `/sbin/arping` binaries.
    - b. Make sure the script is executable by the WebLogic user ('oracle'). The following is an example of an entry inside `/etc/sudoers` granting sudo execution privilege for `oracle` and also over `ifconfig` and `arping`:

```
Defaults:oracle !requiretty
oracle ALL=NOPASSWD: /sbin/ifconfig, /sbin/arping
```

---

**Note:** Ask the system administrator for the sudo and system rights as appropriate to this step.

---

**15.1.12.22.5 Configuring Server Migration Targets** This section describes how to configure server migration targets. You first assign all the available nodes for the cluster's members and then specify candidate machines (in order of preference) for each server that is configured with server migration. Follow these steps to configure cluster migration in a migration in a cluster:

1. Log in to the Oracle WebLogic Server Administration Console (`http://Host:Admin_Port/console`). Typically, `Admin_Port` is 7001 by default.
2. In the Domain Structure window, expand **Environment** and select **Clusters**.
3. Click the cluster for which you want to configure migration (`bi_cluster`) in the Name column of the table.
4. Click the **Migration** tab.
5. In the Change Center, click **Lock & Edit**.

6. In the **Available** field, select the machine to which to allow migration and click the right arrow. In this case, select **APPHOST1** and **APPHOST2**.
7. Select the data source to be used for automatic migration. In this case, select the leasing data source.
8. Click **Save**. Click **Activate Changes**.
9. Set the candidate machines for server migration. You must perform this task for all of the managed servers as follows:
  - a. In the Domain Structure window of the Oracle WebLogic Server Administration Console, expand **Environment** and select **Servers**.
 

**Tip:** Click **Customize this table** in the Summary of Servers page and move Current Machine from the Available window to the Chosen window to view the machine on which the server is running. This will be different from the configuration if the server gets migrated automatically.
  - b. Select the server you want to configure migration for.
  - c. Click the **Migration** tab.
  - d. In the **Available** field, located in the Migration Configuration section, select the machines to which to allow migration and click the right arrow. For **bi\_server1**, select **APPHOST2**. For **bi\_server2**, select **APPHOST1**.
  - e. Select **Automatic Server Migration Enabled**. This enables Node Manager to start a failed server on the target node automatically.
  - f. Click **Save**. Click **Activate Changes**.
  - g. Restart the administration server, node managers, and the servers for which server migration has been configured.

**15.1.12.22.6 Testing the Server Migration** The final step is to test the server migration. To verify that server migration is working properly:

**From APPHOST1:**

1. Force stop the **bi\_server1** managed server. Run the command:

```
APPHOST1> kill -9 pid
```

where *pid* specifies the process ID of the managed server. You can identify the pid in the node by running this command:

```
APPHOST1> ps -ef | grep bi_server1
```

---

**Note:** For Windows, the Managed Server can be terminated by using the `taskkill` command. For example:

```
taskkill /f /pid <pid>
```

Where *<pid>* is the process Id of the Managed Server.

To determine the process Id of the Managed Server run the following command and identify the pid of the `bi_servern` Managed Server.

```
MW_HOME\jrockit_160_<version>\bin\jps -l -v
```

---

2. Watch the Node Manager console. You should see a message indicating that `bi_server1`'s floating IP has been disabled.
3. Wait for Node Manager to try a second restart of `bi_server1`. It waits for a fence period of 30 seconds before trying this restart.
4. Once Node Manager restarts the server, stop it again. Node Manager should now log a message indicating that the server will not be restarted again locally.

**From APPHOST2:**

1. Watch the local Node Manager console. After 30 seconds since the last try to restart `bi_server1` on APPHOST1, Node Manager on APPHOST2 should prompt that the floating IP for `bi_server1` is being brought up and that the server is being restarted in this node.
2. Access one of the applications (for example, BI Publisher) using the same IP.

**Verification From the Administration Console**

Migration can also be verified in the Administration Console:

1. Log in to the Administration Console.
2. Click **Domain** on the left console.
3. Click the **Monitoring** tab and then the **Migration** subtype.

The Migration Status table provides information on the status of the migration.

---

**Note:** After a server is migrated, to fail it back to its original node/machine, stop the managed server from the Oracle WebLogic Administration Console and then start it again. The appropriate Node Manager will start the managed server on the machine to which it was originally assigned.

---

### 15.1.12.23 Scaling Up the Oracle BI EE Topology

When you scale up the Oracle BI EE topology, you add additional system components to one of the existing nodes in your Oracle BI EE high availability topology.

This section assumes that you already have a high availability topology that includes at least two nodes, with a managed server and a full set of system components on each node. To scale up the topology, you increase the number of system components running on one of your existing nodes.

Note that it is not necessary to run multiple managed servers on a given node.

To scale up the Oracle BI EE topology:

1. Log in to Fusion Middleware Control.
2. Expand the **Business Intelligence** node in the *Farm\_domain\_name* window.
3. Click **coreapplication**.
4. Click **Capacity Management**, then click **Scalability**.
5. Click **Lock and Edit Configuration**.
6. Use arrow keys to change the number of **BI Servers**, **Presentation Services**, or **JavaHosts**.
7. Click **Apply** then click **Activate Changes**.
8. Click **Overview** then click **Restart**.

#### 15.1.12.24 Scaling Out the Oracle BI EE Topology to a New Node (APPHOST3)

When you scale out the Oracle BI EE topology, you add a new managed server and set of system components to a new node in your topology (APPHOST3). This section assumes that you already have a high availability topology that includes at least two nodes, with a managed server and a full set of system components on each node.

Before performing the steps in this section, check that you meet these requirements:

- There must be existing nodes running Oracle Business Intelligence Managed Servers within the topology.
- The new node (APPHOST3) can access the existing home directories for Oracle WebLogic Server and Oracle Business Intelligence.
- When an ORACLE\_HOME or WL\_HOME is shared by multiple servers in different nodes, it is recommended that you keep the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. The directory paths you use for binary files and domains when installing new nodes must be identical to those for the first node. To update the Middleware home list to add or remove a WL\_HOME, edit the *MW\_HOME/.home* file. See the steps below.
- The new server can use a new individual domain directory or, if the other Managed Servers domain directories reside on shared storage, reuse the domain directories on those servers.

To scale out Oracle Business Intelligence on APPHOST3:

1. On APPHOST3, mount the existing Middleware home, which should include the Oracle Business Intelligence installation and (optionally, if the domain directory for Managed Servers in other nodes resides on shared storage) the domain directory, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To attach *ORACLE\_HOME* in shared storage to the local Oracle Inventory, run the following command:

```
APPHOST3> cd ORACLE_COMMON_HOME/oui/bin/
APPHOST3> ./attachHome.sh -jreLoc ORACLE_BASE/product/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation is in the node) the *MW\_HOME/.home* file and add *ORACLE\_BASE/product/fmw* to it.

3. Enable VIP3 in APPHOST3. See [Section 15.1.12.3, "Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2."](#)
4. Run the Configuration Assistant from one of the shared Oracle homes using the steps in [Section 15.1.12.12, "Scaling Out the BI System on APPHOST2"](#) as a guide.
5. Scale out the system components on APPHOST3 using the steps in [Section 15.1.12.13, "Scaling Out BI System Components"](#) as a guide.
6. Configure the bi\_server3 Managed Server by setting the Listen Address and disabling host name verification, using the steps in [Section 15.1.12.15, "Setting the Listen Address for the BI\\_SERVER2 Managed Server"](#) and [Section 15.1.12.16, "Disabling Host Name Verification for the BI\\_SERVER2 Managed Server"](#) as a guide.
7. Configure JMS for Oracle BI Publisher as [Section 15.1.12.18.4, "Configuring JMS for Oracle BI Publisher"](#) describes.
8. Configure Oracle BI for Microsoft Office on APPHOST3, as described in [Section 15.1.12.17, "Configuring Oracle BI for Microsoft Office SSO Properties."](#)
9. Set the location of the default persistence store for bi\_server 3, as described in [Section 15.1.12.18.5, "Configuring a Default Persistence Store for Transaction Recovery."](#)
10. Configure Oracle HTTP Server for APPHOST3VHN1, using the steps in [Section 15.1.12.20, "Configuring Oracle HTTP Server for the BI\\_SERVERn Managed Servers"](#) as a guide.
11. Start the bi\_server3 Managed Server and the system components running on APPHOST3. See [Section 15.1.12.9.2, "Starting and Validating the BI\\_SERVER1 Managed Server"](#) and [Section 15.1.12.9.3, "Starting and Validating the BI EE System Components on APPHOST1"](#) for details.
12. Configure server migration, as described in [Section 8.9.3.11.5, "Configuring Server Migration Targets"](#) and [Section 15.1.12.22.6, "Testing the Server Migration."](#)
13. To validate the configuration, access the following URLs:
  - <http://APPHOST3VHN1:9704/analytics> to verify the status of bi\_server3.
  - <http://APPHOST3VHN1:9704/wsm-pm> to verify the status of Web Services Manager. Click **Validate Policy Manager**. A list of policies and assertion templates available in the data opens.

**Note:** The configuration is incorrect if no policies or assertion templates appear.
  - <http://APPHOST3VHN1:9704/xmlpserver> to verify the status of the Oracle BI Publisher application.
14. Oracle recommends using host name verification for the communication between Node Manager and the servers in the domain. This requires using certificates for the different addresses communicating with the Administration Server and other servers. See the chapter on setting up Node Manager in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence* for more detail.

## 15.2 High Availability for Oracle Business Intelligence Publisher

Oracle Business Intelligence Publisher (Oracle BI Publisher) offers you the most efficient, scalable reporting solution available for complex, distributed environments. It provides a central architecture for generating and delivering information to employees, customers, and suppliers—both securely and in the right format. Oracle BI Publisher reduces the high costs associated with the development, customization and maintenance of business documents, while increasing the efficiency of reports management


This section describes how to design and deploy a high availability environment for Oracle BI Publisher.

This section includes the following topics:

- [Section 15.2.1, "Oracle BI Publisher Component Architecture"](#)
- [Section 15.2.2, "Oracle BI Publisher High Availability Concepts"](#)
- [Section 15.2.3, "Oracle BI Publisher High Availability Configuration Steps"](#)

### 15.2.1 Oracle BI Publisher Component Architecture

[Figure 15–11](#) shows the Oracle BI Publisher components in a single instance architecture.

**Figure 15–11 Oracle BI Publisher Single Instance Architecture**

### 15.2.1.1 Oracle BI Publisher Component Characteristics

Oracle BI Publisher is a Java EE application that uses a servlet based architecture. JMS is used for scheduled jobs execution; Fusion ESS integrates with BI Publisher by means of web services.

**15.2.1.1.1 State Information** During the execution of Oracle BI Publisher reports, the session information is stored in memory.

BI Publisher is not a stateless application due to the complexity of migrating report execution across servers.

**15.2.1.1.2 Runtime Processes** Oracle BI Publisher is a Java EE component deployed on Oracle WebLogic Server.

- Oracle BI Publisher artifact processing, which includes creating, deleting, and editing these artifacts:
  - Reports
  - Data models



- Layout templates
- Scheduler jobs
- Configuration
- Execution of scheduler reports
- Execution of online reports

**15.2.1.1.3 Process Lifecycle** Oracle BI Publisher depends entirely on the Oracle infrastructure. Use the WebLogic Server Administration Console to deploy, start, stop, and monitor Oracle BI Publisher.

**15.2.1.1.4 Request Flow** JMS is used as a backbone to the Oracle BI Publisher batch report processing system. JMS is used as a messaging board in a transactional fashion. Every report is divided into these stages of execution:

1. Data retrieval
2. Data formatting
3. Report delivery

Each of these stages appears as a job request in the JMS queue. If a node goes down, the other nodes continue to service the queue. However, if a report job is in one of these stages of execution—data retrieval, data formatting, or report delivery—the job is marked as failed and you must resubmit it manually.

**15.2.1.1.5 External Dependencies** Oracle BI Publisher depends on:

- Database repository: for scheduled jobs information.
- Storage of Oracle BI Publisher artifacts: These artifacts include reports, definitions, and layouts. The storage is done in two different ways:
  - In Oracle Business Intelligence Enterprise Edition mode: Oracle BI Publisher uses Oracle BI Presentation Catalog via a short-lived connection over Web services.
  - In standalone mode: Oracle BI Publisher uses a shared file system (such as SAN or NAS).
- DBMS for report data: Oracle BI Publisher queries the DBMS to retrieve the report data that needs to be processed.
- Web services: for data retrieval.

**15.2.1.1.6 Configuration Artifacts** Oracle BI Publisher is a standard Java EE application.

**15.2.1.1.7 Deployment Artifacts** Oracle BI Publisher does not have any special deployment artifacts.

**15.2.1.1.8 Log Files** Oracle BI Publisher is a Java EE application deployed on Oracle WebLogic Server. All log messages are logged in the server log file of the Oracle WebLogic Server that the application is deployed on. The default location of the server log is:

```
WEBLOGIC_SERVER_HOME/user_projects/domains/domainName/servers/serverName/logs/
serverName-diagnostic.log
```


## 15.2.2 Oracle BI Publisher High Availability Concepts

This section provides conceptual information about using Oracle BI Publisher in a high availability environment.

### 15.2.2.1 Oracle BI Publisher High Availability Architecture

Figure 15–12 shows an Oracle BI Publisher high availability architecture.

**Figure 15–12 Oracle BI Publisher High Availability Architecture**



In Figure 15–12, incoming requests are received by the hardware load balancer, which routes them to WEBHOST1 or WEBHOST2 in the web tier. These hosts have Oracle HTTP Server installed.

The Oracle WebLogic Server on which Oracle BI Publisher runs is installed on APPHOST1 and APPHOST2 in the application tier. The Administration Server on APPHOST2 is passive. The Administration Server can be made active on APPHOST2 if APPHOST1 becomes unavailable, as described in Chapter 3, "High Availability for WebLogic Server." Oracle BI Publisher is deployed on the BI\_SERVER1 and BI\_

SERVER2 managed servers on these hosts. These managed servers are configured in an Oracle WebLogic cluster that allows them to work in an active-active manner. Whole server migration is configured for the managed servers.

The BI domain uses virtual hostnames as the listen addresses for the BI managed servers. You must enable VIP mapping for each of these hostnames on the two BI Machines, (VIP1 on APPHOST1 and VIP2 on APPHOST2), and must correctly resolve the virtual hostnames in the network system used by the topology (either by DNS Server or by hosts resolution). [Section 15.1.12.22, "Configuring Server Migration for the BI\\_SERVERn Servers"](#) provides more details on configuring server migration for the BI servers.

In the data tier, shared external storage is set up to store Oracle BI Publisher Repository metadata such as record definitions, report layouts, and data models, as well as the Oracle BI Publisher JMS persistence store.

An Oracle RAC database is used for reports data, and is also recommended for the enterprise database.

Oracle BI Publisher supports an active-active high availability configuration. Each node acts as an independent server that shares a common repository and the scheduler database with the other Oracle BI Publisher nodes.

The Oracle BI Publisher repository (which contains all content, such as reports and data models) is stored in the Oracle BI Presentation Catalog when BI Publisher is configured with Oracle BI EE, and it is stored on a shared file system when Oracle BI Publisher is installed by itself.

The Oracle BI Publisher configuration folder must be on a shared file system.

The shared repository stores the following Oracle BI Publisher artifacts:

- Report definitions
- Report templates
- Data models
- Oracle BI Publisher configuration

The Scheduler database is a shared database that stores the following information:

- Details about the reports (parameters)
- Execution status
- Execution history
- Snapshots of previous report runs

A hardware or software load balancer can be used as a front end to the Oracle BI Publisher nodes in a high availability environment.

In a high availability deployment of Oracle BI Publisher:

- Sticky sessions should be enabled for the load balancer.
- Oracle BI Publisher depends on the WebLogic Server JMS failover mechanism.
- All the Oracle BI Publisher instances in the cluster are independent.
- Artifacts and configurations are stored in the shared repository, and each instance reads the repository independently.

**15.2.2.1.1 Shared Files and Directories** The Oracle BI Publisher shared files are the shared repository and Scheduler database described in the previous section.

**15.2.2.1.2 Cluster-Wide Configuration Changes** The Configuration Folder contains all configuration, security, and data sources that you set up with BI Publisher. When the Configuration Folder is on shared storage, any changes made will be applied to all nodes in the cluster.

### 15.2.2.2 Protection from Failures and Expected Behaviors

Node Manager can be set up to restart a failed Oracle BI Publisher instance.

When an Oracle BI Publisher instance restarts, it does not affect the other running instances.

When a node fails, the behavior is:

- In a deployment where single sign-on is not implemented, the user will be asked to log in again.
- Any transaction that included editing is saved.

### 15.2.2.3 Troubleshooting

Oracle BI Publisher logging is included in the WebLogic Server logs in the default location:

`DOMAIN_HOME/servers/serverName/logs/serverName-diagnostic.log`

Also, the following file includes BI Publisher related messages:

`DOMAIN_HOME/servers/serverName/logs/bipublisher/bipublisher.log`

Configuration information is stored in the shared repository. The configuration files are stored under the Admin directory in the shared repository.

## 15.2.3 Oracle BI Publisher High Availability Configuration Steps

This section describes how to set up a two node highly available configuration for Oracle BI Publisher. The architecture targeted for the configuration steps is shown in [Figure 15-12](#).

---

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

### 15.2.3.1 Preparing the Environment: Prerequisite Steps Before Setting Up an Oracle BI Publisher High Availability Configuration

This section includes the prerequisites for setting up the BI Publisher high availability configuration.

**15.2.3.1.1 Database Prerequisites** Oracle BI Publisher supports the following database versions:

- Oracle Database 10g (10.2.0.4 or later for non-XE database)
- Oracle Database 11g (11.1.0.7 or later for non-XE database)

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

**15.2.3.1.2 VIP and IP Prerequisites** You configure the BI managed servers to listen on different virtual IPs, shown in [Table 15–7](#). This requires the provisioning of the corresponding VIP in the node and related host names in the DNS system in your network. Ensure that the different VIPs are available and are reachable before running the installation.

**Table 15–7 BI Publisher Virtual Hosts**

| Virtual IP | Virtual IP Maps To | Description                                                                                                                                                                                                                             |
|------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VIP1       | APPHOST1VHN1       | APPHOST1VHN1 is the virtual host name that makes to the listen address for BI_SERVER1 and fails over with server migration of this managed server. It is enabled on the node where BI_SERVER1 process is running (APPHOST1 by default). |
| VIP2       | APPHOST2VHN1       | APPHOST2VHN1 is the virtual host name that makes to the listen address for BI_SERVER2 and fails over with server migration of this managed server. It is enabled on the node where BI_SERVER2 process is running (APPHOST2 by default). |

**15.2.3.1.3 Shared Storage Prerequisites** For proper recovery in case of failure, store both JMS and transaction logs in a location that is accessible to all the nodes that can resume the operations after a failure in a managed server. This requires a shared storage location that can be referenced by multiple nodes. [Table 15–8](#) lists the contents of shared storage.

**Table 15–8 Contents of BI Publisher Shared Storage**

| Server                    | Type of Data                          | Volume in Shared Storage | Directory                                             | Files                                                                                                                      |
|---------------------------|---------------------------------------|--------------------------|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| BI_SERVER1 and BI_SERVER2 | Transaction logs                      | VOL2                     | ORACLE_BASE/admin/domain_name/bi_cluster_name/tlogs   | Common location (stores decided by WebLogic Server)                                                                        |
| BI_SERVER1 and BI_SERVER2 | JMS stores                            | VOL2                     | ORACLE_BASE/admin/domain_name/bi_cluster_name/jms     | Common location, but individual store per server. For example: BipJmsStore for BI_SERVER1 and BipJmsStore2 for BI_SERVER2) |
| BI_SERVER1 and BI_SERVER2 | BI Publisher Configuration folder     | VOL1                     | ORACLE_BASE/domain_name/config/bipublisher/repository | Common location                                                                                                            |
| BI_SERVER1 and BI_SERVER2 | BI Publisher Catalog repository       | VOL1                     | ORACLE_BASE/domain_name/config/bipublisher/catalog    | Common location                                                                                                            |
| BI_SERVER1 and BI_SERVER2 | BI Publisher Scheduler Temp directory | VOL1                     | ORACLE_BASE/domain_name/config/bipublisher/temp       | Common location                                                                                                            |

The shared storage can be a NAS or SAN device. The following is an example command based on a NAS device. Your options may be different from the ones specified in this section.

```
APPHOST1> mount nasfiler:/vol/volX/FMWshared MW_HOME-t nfs -o
rw,bg,hard,nointr,tcp,vers=3,timeo=300,rsize=32768,wsiz=32768
```

**15.2.3.1.4 Clock Synchronization** You must synchronize all server clocks that participate in the cluster to within one second difference. To synchronize the clocks, use a single network time server and point each server to it. The procedure to point to the network time server is different from Windows to Linux. Refer to your operating system's manual.

**15.2.3.1.5 Installing and Configuring the Database Repository** This section describes how to install and configure the database repository.

#### **Oracle Clusterware**

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

#### **Automatic Storage Management**

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

#### **Oracle Real Application Clusters**

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

You must install the 11g (11.1.1) Oracle Fusion Middleware Repository into an Oracle Real Application Clusters (Oracle RAC) database before you install the Oracle Fusion Middleware components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. You install RCU in its own, separate Middleware home.

Use the latest version of RCU to install the 11g (11.1.1.) Oracle Fusion Middleware Repository into an Oracle RAC database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

**15.2.3.1.6 Using RCU to Load the Business Intelligence Schemas in the Database** Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and Oracle BI schemas into an Oracle RAC database before you install Oracle BI Publisher. Follow these steps:

1. Insert the Repository Creation Utility (RCU) DVD, and then start RCU from the bin directory in the RCU home directory.

```
prompt> cd RCU_HOME/bin
prompt> ./rcu
```

2. In the Welcome screen, click **Next**.

3. In the Create Repository screen, select **Create** to load component schemas into a database. Click **Next**.
4. In the Database Connection Details screen, enter connect information for your database:
  - **Database Type:** Select **Oracle Database**.
  - **Host Name:** Specify the name of the node on which the database resides. For the Oracle RAC database, specify the VIP name or one of the node names as the host name, for example:  
BIDBHOST1-VIP
  - **Port:** Specify the listen port number for the database: 1521
  - **Service Name:** Specify the service name of the database:  
biha.mycompany.com
  - **Username:** Specify the name of the user with DBA or SYSDBA privilege: SYS
  - **Password:** Enter the password for the SYS user.
  - **Role:** Select the database user's role from the list: SYSDBA (required by the SYS user).

Click **Next**.
5. In the Select Components screen:
  - Select **Create a new Prefix**, and enter a prefix to use for the database schemas, for example, BIHA. You can specify up to six characters as a prefix. Prefixes are used to create logical groupings of multiple repositories in a database. For more information, see *Oracle Fusion Middleware Repository Creation Utility User's Guide*.

**Tip:** Note the name of this schema because the upcoming steps require this information.

  - Select the following components:
    - AS Common Schemas: Metadata Services (automatically selected)
    - Oracle Business Intelligence: Business Intelligence Platform

Click **Next**.
6. In the Schema Passwords screen, enter passwords for the main schema users, and click **Next**.  
 You can choose either **Use same passwords for all schemas** or **Specify different passwords for all schemas**, depending on your requirements.  
 Do not select **Use main schema passwords for auxiliary schemas**. The auxiliary passwords are derived from the passwords of the main schema users.  

**Tip:** Note the names of the schema passwords because the upcoming steps require this information.
7. In the Map Tablespaces screen, choose the tablespaces for the selected components, and click **Next**.
8. In the Summary screen, click **Create**.
9. In the Completion Summary screen, click **Close**.

**15.2.3.1.7 Configuring Virtual Server Names and Ports for the Load Balancer** This section describes the load balancer prerequisites for deploying an Oracle BI Publisher high availability environment.

### Load Balancers

BI Publisher uses a hardware load balancer when deployed in a high availability configuration with two Oracle HTTP Servers as web tier components.

### Virtual Server Names

bi.mycompany.com is a virtual server name that acts as the access point for all HTTP traffic to the runtime BI components, such as Oracle BI Publisher. Traffic to both SSL and non-SSL ports is configured, and typically non-SSL is redirected to SSL. Clients access this service using the address bi.mycompany.com:443. This virtual server is defined on the load balancer.

**15.2.3.1.8 Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2** This section describes how to install Oracle HTTP Server on WEBHOST1.

1. Verify that the servers meet the following requirements:

- Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
- Ensure that port 7777 is not in use by any service on WEBHOST1 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

#### On UNIX:

```
netstat -an | grep "7777"
```

#### On Windows:

```
netstat -an | findstr :7777
```

If the port is in use (if the command returns output identifying the port), you must free the port.

#### On UNIX:

Remove the entries for port 7777 in the `/etc/services` file and restart the services, or restart the computer.

#### On Windows:

Stop the component that is using the port.

- Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
- Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port number for Oracle HTTP Server):

```
The port for Oracle HTTP server
Oracle HTTP Server port = 7777
```

2. Start the Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:



**On UNIX:**

Issue the command: `./runinstaller`

**On Windows:**

Double-click `setup.exe`.

This displays the Specify Inventory Directory screen.

3. On the Specify Inventory Directory screen, enter values for the Oracle Inventory Directory and the Operating System Group Name. For example:

**Specify the Inventory Directory:** `/u01/app/oraInventory`

**Operating System Group Name:** `oinstall`

A dialog box appears with the following message:

"Certain actions need to be performed with root privileges before the install can continue. Please execute the script `/u01/app/oraInventory/createCentralInventory.sh` now from another window and then press "Ok" to continue the install. If you do not have the root privileges and wish to continue the install select the "Continue installation with local inventory" option"

Log in as root and run the "`/u01/app/oraInventory/createCentralInventory.sh`"

This sets the required permissions for the Oracle Inventory Directory and then brings up the Welcome screen.

---

**Note:** The Oracle Inventory screen is not shown if an Oracle product was previously installed on the host. If the Oracle Inventory screen does not appear for this installation, verify:

- If the `/etc/oraInst.loc` file exists.
  - If the file exists, the Inventory directory listed is valid.
  - Whether or not the user performing the installation has write permissions for the Inventory directory.
- 

4. On the Welcome screen, click **Next**.
5. On the Select Installation Type screen, select **Install and Configure**, and click **Next**.
6. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.

Click **Next**.

7. On the Specify Installation Location screen, specify:
  - **Oracle Middleware Home:** `ORACLE_BASE/product/fmw` (grayed out)
  - **Oracle Home Directory:** `Oracle_WT1`

Click **Next**.

8. On the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Do not select **Associate Selected Components with Weblogic Domain**.

Click **Next**.

9. On the Specify Component Details screen, enter the following values for WEBHOST1:

- **Instance Home Location:** `ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1`
- **Instance Name:** `web1`
- **OHS Component Name:** `ohs1`

Click **Next**.

10. On the Configure Ports screen:

- If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
- Enter the Oracle HTTP Server port, for example: `7777`

Click **Next**.

11. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address and click **Next**.
12. On the Installation Summary screen, ensure that the selections are correct. If they are not, click **Back** and make the necessary fixes. After ensuring that the selections are correct, click **Install**.
13. On the Installation Progress screen on UNIX systems, a dialog appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

14. On the Configuration Progress screen, multiple configuration assistants are launched in succession; this process can be lengthy. Click **Next**.
15. On the Installation Complete screen, click **Finish** to exit.

Repeat the steps for WEBHOST2 and configure your load balancer with a pool containing both the WEBHOST1 and WEBHOST2 addresses.

**15.2.3.1.9 Validating Oracle HTTP Server** To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URLs in the browser:

```
http://WEBHOST1:7777/
http://WEBHOST2:7777/
```

If Oracle HTTP Server is set up correctly, the Oracle FMW Welcome screen appears in the browser.

### 15.2.3.2 Installing Oracle Fusion Middleware Home

This section describes the procedure for installing Oracle Fusion Middleware on all nodes in the application tier that run Oracle WebLogic Server and Oracle Fusion Middleware BI EE. Repeat the procedure (described below for APPHOST1) for installing WebLogic Server and BI Publisher in APPHOST2. The directory paths for binary files and domains used when installing new nodes must be exactly the same as those used for first node.

Install the following Oracle Fusion Middleware components:

- Oracle WebLogic Server
- Oracle Fusion Middleware BI EE

**15.2.3.2.1 Installing Oracle WebLogic Server** Follow these steps to install Oracle WebLogic Server:

1. Start the installer for Oracle WebLogic Server from the installation media.
2. In the Welcome screen, click **Next**.
3. In the Choose Middleware Home Directory screen:
  - Select **Create a new Middleware Home**.
  - For the **Middleware Home Directory**, enter `ORACLE_BASE/product/fmw`

---

**Note:** `ORACLE_BASE` is the base directory under which Oracle products are installed. The recommended value is `/u01/app/oracle`.

---

Click **Next**.

4. In the Register for Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.
5. In the Choose Install Type screen, select **Typical** and click **Next**.
6. In the Choose Product Installation Directories screen, accept the directory `ORACLE_BASE/product/fmw/wlserver_10.3` for WebLogic Server and `ORACLE_BASE/product/fmw/coherence_3.6` for Oracle Coherence and click **Next**.
7. In the Installation Summary screen, click **Next**.  
The Oracle WebLogic Server software is installed.
8. In the Installation Complete screen, clear the Run Quickstart check box and click **Done**.

**15.2.3.2.2 Installing Oracle Fusion Middleware for BI Publisher** On the Linux platform, if the `/etc/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

1. Start the installer for Oracle Fusion Middleware Business Intelligence 11g from the installation media:

**On UNIX:**

```
APPHOST1> ./runInstaller
```

**On Windows:**

```
APPHOST1> setup.exe
```

2. In the Specify Inventory Directory screen:
  - Enter `HOME/oraInventory`, where `HOME` is the home directory of the user performing the installation (this is the recommended location).
  - Enter the OS group for the user performing the installation.

Click **Next**.

- Follow the instructions on screen to execute `/createCentralInventory.sh` as root.

Click **OK**.

3. In the Welcome screen, click **Next**.
4. In the Select Installation Type screen, select **Software Only Install** and click **Next**.
5. In the Prerequisite Checks screen, verify that all checks complete successfully, and click **Next**.
6. In the Specify Installation Location screen, select the previously installed Middleware Home from the drop-down list. For the Oracle Home directory, enter the directory name (Oracle\_BI1).

Click **Next**.

7. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address and click **Next**.
8. In the Summary screen, click **Install**.
9. In the Installation Progress screen, click **Next**.
10. In the Complete screen, click **Finish**.

### 15.2.3.3 Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2

The BI domain uses virtual hostnames as the listen addresses for the BI managed servers. You must enable the VIPs, mapping each of these hostnames on the two BI machines (VIP1 on APPHOST1 and VIP2 on APPHOST2), and they must correctly resolve to the virtual hostnames in the network system used by the topology (either by DNS Server or hosts resolution).

### 15.2.3.4 Creating a Domain with the Administration Server and the First BI\_SERVER1 Managed Server

This section describes how to create a domain and the first WebLogic Server BI managed server using the Oracle Business Intelligence Configuration Assistant, Oracle WebLogic Server Administration Console, and Oracle Enterprise Manager. Ensure that the database where you installed the repository is running. For Oracle RAC databases, all the instances must be running.

---

---

**Note:** Oracle strongly recommends that you read the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

Run the Configuration Assistant from the Oracle home directory to create a domain containing the Administration Server and the managed server with BI Publisher:

1. Change the directory to the location of the Configuration Assistant:

```
APPHOST1> cd ORACLE_HOME/bin
```

2. Start the Configuration Assistant:

**On UNIX:**

```
APPHOST1> ./config.sh
```

### On Windows:

```
APPHOST1> config.cmd
```

3. In the Welcome screen, click **Next**.
4. In the Prerequisite Checks screen, verify that all checks complete successfully and click **Next**.
5. The Create or Scale Out BI System screen opens. Select **Create New BI System** and specify the following:
  - **User Name:** weblogic
  - **User Password:** Enter the weblogic user password.
  - **Domain Name:** *bifoundation\_domain*

Click **Next**.

6. In the Specify Installation Location screen, enter:
  - **Middleware Home:** *ORACLE\_BASE/product/fmw* (grayed out)
  - **Oracle Home:** *ORACLE\_BASE/product/fmw/Oracle\_BI1* (grayed out)
  - **WebLogic Server Home:** *ORACLE\_BASE/product/fmw/wlserver\_10.3* (grayed out)
  - **Domain Home:** *ORACLE\_BASE/product/fmw/user\_projects/domain/bifoundation\_domain*

---

**Note:** The Domain Home must end with the domain name.

---

- **Instance Home:** *ORACLE\_BASE/product/fmw/instance1*
- **Instance Name:** *instance1*

Click **Next**.

7. In the Configure Components screen, select **Business Intelligence Publisher**.  
Click **Next**.

8. In the BIPLATFORM Schema screen, enter:
  - **Database Type:** Oracle Database
  - **Connect String:**  
*BIDBHOST1:1521:BIDB1^BIDBHOST2:1521:BIDB2@BIHA.MYCOMPANY.COM*
  - **BIPLATFORM Schema Username:** *BIHA\_BIPLATFORM*
  - **BIPLATFORM Schema Password:** Enter the password for the BIPLATFORM schema user.

Click **Next**.

9. In the Configure Ports screen, select one of the following:
  - Auto Port Configuration
  - Specify Ports using Configuration File

Click **Next**.

10. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.
11. In the Summary screen, click **Configure**.
12. In the Configuration Progress screen, verify that all the Configuration Tools have completed successfully and click **Next**.
13. In the Complete screen, click **Finish**.

#### 15.2.3.5 Creating boot.properties for the Administration Server on APPHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a boot.properties file for the Administration Server on APPHOST1.

1. Go to the following directory:

```
ORACLE_BASE/product/fmw/user_projects/domains/bifoundation_domain/servers/
AdminServer/security
```

2. Enter the following lines in the file, save the file, and close the editor:

```
username=Admin_username
password=Admin_password
```

---

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible for the entries to be encrypted.

---

---

#### 15.2.3.6 Starting and Validating the Administration Server on APPHOST1

This section describes procedures for starting and validating the Administration Server on APPHOST1.

**15.2.3.6.1 Starting the Administration Server on APPHOST1** To start the Administration Server on APPHOST1, run the following commands:

```
APPHOST1> cd MW_HOME/user_projects/domains/bifoundation_domain/bin
APPHOST1> ./startWebLogic.sh
```

**15.2.3.6.2 Validating the Administration Server** verify that the Administration Server is properly configured:

1. In a browser, go to <http://APPHOST1:7001/console>.
2. Log in as the administrator.
3. Verify that the BI\_SERVER1 managed server is listed.
4. Verify that the bi\_cluster cluster is listed.
5. Verify that you can access Enterprise Manager at <http://APPHOST1:7001/em>.

### 15.2.3.7 Setting the Listen Address for BI\_SERVER1 Managed Server

To set the managed server listen address:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server1** in the **Names** column of the table.
6. Set the Listen Address to APPHOST1VHN1.
7. Click **Save**.
8. Save and activate the changes.

The changes will not take effect until the BI\_SERVER1 managed server is restarted.

### 15.2.3.8 Disabling Host Name Verification for the BI\_SERVER1 Managed Server

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you will receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and validating the topology, and enable it again once the high availability topology configuration is complete.

To disable host name verification:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server1** in the **Names** column of the table. The settings page opens.
6. Open the **SSL** tab.
7. Expand the **Advanced** section of the page.
8. Set **Hostname Verification** to **None**.
9. Click **Save**.
10. Save and activate the changes.
11. The change will not take effect until the BI\_SERVER1 managed server is restarted.

### 15.2.3.9 Starting the System in APPHOST1

This section describes how to start Node Manager on APPHOST1 and how to start and validate the BI\_SERVER1 managed server on APPHOST1.

**15.2.3.9.1 Starting Node Manager on APPHOST1** Usually, Node Manager is started automatically when config.sh completes. If Node Manager is not running for some reason, start it on APPHOST1 using these commands:

```
APPHOST1> cd WL_HOME/server/bin
APPHOST1> ./startNodeManager.sh
```

**15.2.3.9.2 Starting and Validating the BI\_SERVER1 Managed Server** To start the BI\_SERVER1 managed server and check that it is configured correctly:

1. Start the bi\_server1 managed server using Oracle WebLogic Server Administration Console, as follows:
  - a. Expand the **Environment** node in the Domain Structure window.
  - b. Choose **Servers**.
  - c. Click the **Control** tab.
  - d. Select **bi\_server1** and then click **Start**.
2. Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming", wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors
3. When BI\_SERVER1 is started, the following URLs become available:
  - Access <http://APPHOST1VHN1:9704/wsm-pm> to verify the status of Web Services Manager. Click **Validate Policy Manager**.

---

**Note:** The configuration is incorrect if no policies or assertion templates appear.

---

- Access <http://APPHOST1VHN1:9704/xmlpserver> to verify the status of the BI Publisher application.

### 15.2.3.10 Prerequisites for Scaling Out the BI System on APPHOST2

This section includes prerequisite tasks to perform before scaling out the BI system on APPHOST2.

**15.2.3.10.1 Setting Server Configuration Options** Follow these steps to configure server configuration options:

1. Copy over the contents of *DOMAIN\_HOME/config/bipublisher/repository* to the shared configuration folder location.
2. Log into BI Publisher with Administrator credentials and select the **Administration** tab.
3. Under **System Maintenance**, select **Server Configuration**.
4. Enter the following fields for the Configuration Folder.
  - **Path:** Enter the path of the shared location for the Configuration Folder.
5. Apply your changes and restart your BI Publisher application.

**15.2.3.10.2 Setting Scheduler Configuration Options** Follow these steps to configure scheduler configuration options:

1. Log into BI Publisher with Administrator credentials and select the **Administration** tab.
2. Under **System Maintenance**, select **Scheduler Configuration**.
3. Select **Quartz Clustering** under the **Scheduler Selection**.



4. Apply your changes and restart your BI Publisher application.

**15.2.3.10.3 Updating the Oracle BI Publisher Scheduler Configuration** Follow the steps in this section to update the WebLogic JNDI URL and the JMS Shared Temp directory for the Oracle BI Publisher Scheduler.

**To update the Oracle BI Publisher Scheduler configuration:**

1. Log in to one of the Oracle BI Publisher instances at the following URLs:

```
http://APPHOST1VHN1:9704/xmlpserver
http://APPHOST2VHN1:9704/xmlpserver
```

2. Click the **Administration** link.
3. Click **Scheduler Configuration** under System Maintenance.
4. Update the **WebLogic JNDI URL** under JMS Configuration, as follows:  
`t3://APPHOST1VHN1:9704,APPHOST2VHN1:9704`
5. Update the Shared Directory by entering a directory that is located in the shared storage. This shared storage is accessible from both APPHOST1 and APPHOST2.
6. Continue with the rest of the steps.
7. Click **Apply**.
8. Check the Scheduler status from the Scheduler Diagnostics tab.

**15.2.3.11 Scaling Out the BI System on APPHOST2**

Run the Configuration Assistant from the ORACLE\_HOME directory to scale out the BI System.

1. Change the directory to the location of the Configuration Assistant:  
`APPHOST2> cd ORACLE_HOME/bin`
2. Start the Configuration Assistant:  
`APPHOST2> ./config.sh`
3. In the Welcome screen, click **Next**.
4. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.  
 Click **Next**.
5. Select **Scale Out BI System** and then enter the following values:
  - **Host Name:** ADMINHOST
  - **Port:** 7001
  - **User name:** weblogic
  - **User Password:** Enter the password for the weblogic user.
 Click **Next**.
6. In the Scale Out BI System Details screen, enter:
  - **Middleware Home:** ORACLE\_BASE/product/fmw (grayed out)
  - **Oracle Home:** Oracle\_BI1 (grayed out)

- **WebLogic Server Home:** *ORACLE\_BASE/product/fmw/wlserver\_10.3* (grayed out)
- **Domain Home:** *ORACLE\_BASE/product/fmw/user\_projects/domain/bifoundation\_domain*

---

**Note:** The Domain Home must end with the domain name.

---

- **Instance Home:** *ORACLE\_BASE/product/fmw/instance2*
- **Instance Name:** instance2 (grayed out)

Click **Next**.

7. In the Configure Ports screen, select one of the following:

- Auto Port Configuration
- Specify Ports using Configuration File

Click **Next**.

8. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.
9. In the Summary screen, click **Configure**.
10. In the Configuration Progress screen, verify that all the Configuration Tools have completed successfully and click **Next**.
11. In the Complete screen, click **Finish**.

#### 15.2.3.12 Setting the Listen Address for the BI\_SERVER2 Managed Server

To set the BI\_SERVER2 managed server listen address:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server2** in the **Names** column of the table.
6. Set the Listen Address to APPHOST2VHN1.
7. Click **Save**.
8. Save and activate the changes.

Changes take effect when the BI\_SERVER2 managed server restarts.

#### 15.2.3.13 Disabling Host Name Verification for the BI\_SERVER2 Managed Server

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you will receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and validating the topology, and enable it again once the high availability topology configuration is complete.

To disable host name verification:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server2** in the **Names** column of the table.
6. Open the **SSL** tab.
7. Expand the **Advanced** section of the page.
8. Set **Hostname Verification** to **None**.
9. Click **Save**.
10. Save and activate the changes.
11. The change does not take effect until you restart the BI\_SERVER2 managed server.

#### 15.2.3.14 Configuring Oracle BI Publisher

This section describes how to configure Oracle BI Publisher.

**15.2.3.14.1 Configuring JMS Persistence Store for BI Publisher** You must configure the location for all persistence stores to a directory visible from both nodes. Change all persistent stores to use this shared base directory.

---

**Note:** See [Considerations for Using File Stores on NFS](#) for important information about JMS messages and transaction logs, as well as releasing locks on file stores.

---

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Services** node and then click the **Persistent Stores** node.
3. In the Change Center, click **Lock & Edit**.
4. Click on **BipJmsStore** and enter a directory that is located in the shared storage. This shared storage is accessible from both APPHOST1 and APPHOST2:

`ORACLE_BASE/admin/domain_name/bi_cluster/jms`

5. Click **Save** and **Activate Changes**.
6. In the Domain Structure window, expand the **Services** node and then click the **Persistent Stores** node.
7. In the Change Center, click **Lock & Edit**.
8. Click **New**, and then **Create File Store**.
9. Enter a name (for example, BipJmsStore2) and target BI\_SERVER2. Enter a directory that is located in shared storage so that it is accessible from both APPHOST1 and APPHOST2:

`ORACLE_BASE/admin/domain_name/bi_cluster/jms`

10. Click **OK** and activate the changes.
11. In the Domain Structure window, expand the **Services** node and then click the **Messaging > JMS Servers** node. The Summary of JMS Servers page opens.

12. In the Change Center, click **Lock & Edit**.
13. Click **New**.
14. Enter a name (for example, BipJmsServer2) and in the **Persistence Store** drop-down list, select **BipJmsStore2** and click **Next**.
15. Select **BI\_SERVER2** as the target.
16. Click **Finish** and **Activate Changes**.
17. In the Domain Structure window, expand the **Services** node and then click the **Messaging > JMS Modules** node.
18. In the Change Center, click **Lock & Edit**.
19. Click **BIPJmsResource** and then click the Subdeployments tab.
20. Target the Subdeployment **BipJmsSubDeployment** to both BipJmsServer1 and BipJmsServer2.
21. Click **Finish** and **Activate Changes**.

To validate the JMS configuration performed for Oracle BI Publisher, perform the steps in [Section 15.2.3.10.3, "Updating the Oracle BI Publisher Scheduler Configuration."](#)

**15.2.3.14.2 Configuring a Default Persistence Store for Transaction Recovery** Each server has a transaction log, which stores information about committed transactions coordinated by the server that may not have been completed. WebLogic Server uses the transaction log when recovering from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to the server.

---

**Note:** Preferably, this location should be a dual-ported SCSI disk or on a Storage Area Network (SAN).

---

To set the location for the default persistence store for BI\_SERVER1:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node and then click the **Servers** node.
3. Click **BI\_SERVER1** (represented as a hyperlink) in the **Name** column of the table.  
The settings page for the BI\_SERVER1 server opens with the Configuration tab active.
4. Click the **Services** tab.
5. In the Change Center, click **Lock & Edit**.
6. In the Default Store section of the page, enter the path to the folder where the default persistent stores will store its data files. The directory structure of the path is as follows:  
  
`ORACLE_BASE/admin/domain_name/bi_cluster/tlogs`
7. Click **Save** and **Activate Changes**.
8. Repeat these steps for the BI\_SERVER2 server.

---

**Note:** To enable migration of the Transaction Recovery Service, specify a location on a persistent storage solution that is available to other servers in the cluster. Both APPHOST1 and APPHOST2 must be able to access this directory. This directory must also exist before you restart the server.

---

### 15.2.3.15 Starting the System in APPHOST2

This section describes procedures for starting the system in APPHOST2.

**15.2.3.15.1 Starting Node Manager on APPHOST2** Usually, Node Manager is started automatically when config.sh completes. If Node Manager is not running for some reason, start the Node Manager on APPHOST2 by following the instructions in [Section 15.2.3.9.1, "Starting Node Manager on APPHOST1."](#)

**15.2.3.15.2 Starting and Validating the BI\_SERVER2 Managed Server** To start the BI\_SERVER2 managed server and check that it is configured correctly:

1. Start the bi\_server2 managed server using Oracle WebLogic Server Administration Console, as follows:
  - a. Expand the **Environment** node in the Domain Structure window.
  - b. Choose **Servers**.
  - c. Click the **Control** tab.
  - d. Select **bi\_server2** and then click **Start**.
2. Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming", wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors
3. When BI\_SERVER2 is started, the following URLs become available:
  - Access `http://APPHOST2VHN1:9704/wsm-pm` to verify the status of Web Services Manager. Click **Validate Policy Manager**.

---

**Note:** The configuration is incorrect if no policies or assertion templates appear.

---

- Access `http://APPHOST2VHN1:9704/xmlpserver` to verify the status of the Oracle BI Publisher application.

### 15.2.3.16 Configuring Oracle HTTP Server for the BI\_SERVERn Managed Servers

To enable Oracle HTTP Server to route to bi\_cluster, which contains the BI\_SERVERn managed servers, you must set the WebLogicCluster parameter to the list of nodes in the cluster:

1. On WEBHOST1 and WEBHOST2, add the following lines to the `ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/config/OHS/ohs1/mod_wl_ohs.conf` file:

```
WSM-PM
<Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
```

```
 WLProxySSL ON
 WLProxySSLPassThrough ON
 </Location>

 # BI Publisher
 <Location /xmlpsserver>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
 WLProxySSL ON
 WLProxySSLPassThrough ON
 </Location>
```

2. Perform the same steps for the Oracle HTTP Server on WEBHOST2.
3. Restart Oracle HTTP Server on WEBHOST1 and WEBHOST2:

```
WEBHOST1> ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/bin/
opmnctlrestartproc ias-component=ohs1
```

```
WEBHOST2> ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/bin/
opmnctlrestartproc ias-component=ohs1
```

#### 15.2.3.17 Validating Access Through Oracle HTTP Server

Verify that the BI Servers status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming," wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors.

Verify that you can access these URLs:

- <http://WEBHOST1:7777/wsm-pm>
- <http://WEBHOST2:7777/wsm-pm>
- <http://WEBHOST1:7777/xmlpsserver>
- <http://WEBHOST2:7777/xmlpsserver>

Verify that you can access these URLs using your load balancing router address:

- <http://bi.mycompany.com:80/xmlpsserver>
- <http://bi.mycompany.com:80/wsm-pm>

Follow these instructions to ensure that routing and failover from the HTTP Server to the bi\_cluster is working correctly:

1. While BI\_SERVER2 is running, stop BI\_SERVER1 from Oracle WebLogic Server Administration Console.
2. Access the following URLs and verify the appropriate functionality:
  - [WEBHOST1:7777/wsm-pm](http://WEBHOST1:7777/wsm-pm)
  - [WEBHOST1:7777/xmlpsserver](http://WEBHOST1:7777/xmlpsserver)
3. Start BI\_SERVER1 from Oracle WebLogic Server Administration Console.
4. Stop BI\_SERVER2.
5. Access the URLs in Step 2 above again and verify the appropriate functionality.

### 15.2.3.18 Configuring Server Migration for the BI\_SERVERn Servers

In this high availability topology, you must configure server migration for the bi\_server1 and bi\_server2 Managed Servers. To do this, follow the instructions in [Section 15.1.12.22, "Configuring Server Migration for the BI\\_SERVERn Servers."](#)

### 15.2.3.19 Scaling Out the Oracle BI Publisher Topology to a New Node (APPHOST3)

When you scale out the Oracle BI Publisher topology, you add a new managed server and set of system components to a new node in your topology (APPHOST3). This section assumes that you already have a high availability topology that includes at least two nodes, with a managed server and a full set of system components on each node.

Before performing the steps in this section, check that you meet these requirements:

- There must be existing nodes running Oracle BI Publisher managed servers within the topology.
- The new node (APPHOST3) can access the existing home directories for Oracle WebLogic Server and Oracle Business Intelligence.
- When an ORACLE\_HOME or WL\_HOME is shared by multiple servers in different nodes, it is recommended that you keep the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. To update the oraInventory in a node and "attach" an installation in a shared storage to it, use `ORACLE_HOME/oui/bin/attachHome.sh`. To update the Middleware home list to add or remove a WL\_HOME, edit the `MW_HOME/.home` file. See the steps below.
- The new server can use a new individual domain directory or, if the other Managed Servers domain directories reside on shared storage, reuse the domain directories on those servers.

To scale out Oracle BI Publisher on APPHOST3:

1. On APPHOST3, mount the existing Middleware home, which should include the Oracle Business Intelligence installation and (optionally, if the domain directory for Managed Servers in other nodes resides on shared storage) the domain directory, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To attach `ORACLE_HOME` in shared storage to the local Oracle Inventory, execute the following command:

```
APPHOST3> cd ORACLE_COMMON_HOME/oui/bin/
APPHOST3> ./attachHome.sh -jreLoc ORACLE_BASE/product/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the `MW_HOME/.home` file and add `ORACLE_BASE/product/fmw` to it.

3. Enable VIP3 in APPHOST3. See [Section 15.2.3.3, "Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2."](#)
4. Run the Configuration Assistant from one of the shared Oracle homes, using the steps in [Section 15.2.3.11, "Scaling Out the BI System on APPHOST2"](#) as a guide.
5. Configure the bi\_server3 Managed Server by setting the Listen Address and disabling host name verification, using the steps in [Section 15.2.3.12, "Setting the Listen Address for the BI\\_SERVER2 Managed Server"](#) and [Section 15.2.3.13, "Disabling Host Name Verification for the BI\\_SERVER2 Managed Server"](#) as a guide.

6. Configure JMS for Oracle BI Publisher, as described in [Section 15.2.3.14.1, "Configuring JMS Persistence Store for BI Publisher."](#)
7. Set the location of the default persistence store for bi\_server 3, as described in [Section 15.2.3.14.2, "Configuring a Default Persistence Store for Transaction Recovery."](#)
8. Configure Oracle HTTP Server for APPHOST3VHN1, using the steps in [Section 15.2.3.16, "Configuring Oracle HTTP Server for the BI\\_SERVERn Managed Servers"](#) as a guide.
9. Start the bi\_server3 Managed Server and the system components running on APPHOST3. See [Section 15.2.3.9.2, "Starting and Validating the BI\\_SERVER1 Managed Server"](#) for details.
10. Configure server migration, as described in [Section 8.9.3.11.5, "Configuring Server Migration Targets"](#) and [Section 15.1.12.22.6, "Testing the Server Migration."](#)
11. To validate the configuration, access the `http://APPHOST3VHN1:9704/xmlpserver` URL to verify the status of the Oracle BI Publisher application.
12. Oracle recommends using host name verification for the communication between Node Manager and the servers in the domain. This requires certificates for the different addresses communicating with the Administration Server and other servers. See the chapter on setting up Node Manager in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence* for further details.

## 15.3 High Availability for Oracle Real-Time Decisions

Oracle Real-Time Decisions (Oracle RTD) is used to provide a ranked list of choices or offers for a request, given a sufficient context for making a decision on the best choices to select from. Oracle RTD always uses a front-end application that is responsible for defining process flow, capturing user interaction, and translating this interaction as input for Oracle RTD decision making. Therefore, Oracle RTD can be seen as more of a decision service rather than a traditional web application.

This section describes how to design and deploy a high availability environment for Oracle RTD.


This section includes the following topics:

- [Section 15.3.1, "Oracle RTD Component Architecture"](#)
- [Section 15.3.2, "Oracle RTD High Availability Concepts"](#)
- [Section 15.3.3, "Oracle RTD High Availability Configuration Steps"](#)

### 15.3.1 Oracle RTD Component Architecture

[Figure 15-13](#) shows the key components in a single instance Oracle RTD architecture:



**Figure 15–13 Oracle RTD Single Instance Architecture**

The components shown in [Figure 15–13](#) include:

- **Decision Server:** Provides run-time services to execute Oracle RTD applications called **Inline Services**.

For more information about **Decision Server** and **Inline Services**, see the "[Decision Server](#)" subsection.

- **Batch Services:** Includes the **Batch Manager** and **Batch Agent**. **Batch Services** is responsible for satisfying requests to perform batch jobs.

For more information about Batch Services, see the "Oracle RTD Batch Framework" chapter in the *Oracle Fusion Middleware Platform Developer's Guide for Oracle Real-Time Decisions*.

- **Decision Center:** Is the web-based UI for business users to view reports about the effectiveness of their analytical models and to tune their hosting Inline Services.

For more information about Decision Center, see the "[Decision Center](#)" subsection.

- **Learning Services:** Read learning records from the Oracle RTD database and periodically generate new predictive models.

For more information about Learning Services, see the "[Learning Services](#)" subsection.

### **Decision Server**

Decision Server provides the necessary run-time services to execute Oracle RTD applications called Inline Services. Inline Services are developed using the Decision Studio development environment, and are deployed to and run in Oracle RTD Decision Server. For more information about Decision Studio, see the "[Decision Studio, RTD Clients and Tools](#)" subsection.

An Inline Service can gather data and analyze characteristics of enterprise business processes on a real-time and continuous basis. It also leverages that data and analysis to provide decision-making capability and feedback to key business processes.

Decision Server stores all its metadata in the Oracle RTD database.

The following elements of an Inline Service play a role in or are affected by a high availability topology (mostly by accessing resources outside of Oracle RTD):

1. **Choices:** Choices represent the offers that will be presented through the Inline Service or the targets of study to be tracked by the self-learning models of Oracle RTD. There are two forms of choices: static and dynamic. Static choices are hard coded using Decision Studio. Dynamic choices are defined in an external data source and loaded during the execution of an Inline Service.
2. **Rules:** Rules are used to target segments of population, decide whether a choice is eligible, or score a particular choice. Rules are constructed through a graphical rules editor. Rules are loaded by Decision Server during the execution of an Inline Service. Oracle RTD has a type of rules called external rules. These rules are stored in an data source external to Oracle RTD.
3. **Entities:** Entities are a logical representation of data used for Oracle RTD modeling and decisioning. The attributes of an entity can be populated via data sources, as incoming parameters from integration points, or derived in real time through custom logic.
4. **Data Sources:** Data sources retrieve data from tables or stored procedures.
5. **Integration Points:** Integration points serve as the touchpoints with outside systems interacting with Oracle RTD. There are two classes of integration points: informants and advisors. Informants receive data from outside systems, whereas advisors receive data and also send recommendations back to outside systems.
6. **Models:** Serve two primary purposes: prediction and reporting. Models are defined to:
  - Predict the likelihood that certain events associated with Choices will occur.
  - Analyze data correlated with those events.

Statistics Collectors are special models that track statistics about entities.

Models are associated at design time with a Choice Group. This defines the list of Choices to which prediction and reports apply.

### Decision Center

Workbench Services, which supports the deployment of Inline Services by Decision Server, also provides the Decision Center web interface. Decision Center displays the structure and decisioning history of Inline Services. Business users use Decision Center to view reports about the effectiveness of their analytical models and to tune their hosting Inline Services.

Decision Center interacts with the Learning Server to query the contents of the Learning Models.

### Learning Services

Learning records are created when an Inline Service session closes. Learning records are batched in memory (configurable) and queued to be stored in the Oracle RTD database. This is done for performance reasons - to minimize the number of database transactions required to store learning records in the Oracle RTD database. By default, Oracle RTD buffers 100 records, with a queue for 50 buffers. A worker thread is responsible for writing learning records to the Oracle RTD database.

Learning Services awakens periodically and:

- Reads learning records from the database
- Generates new predictive models

Each predictive model is associated with a specific Inline Service. Therefore, an Inline Service is only affected when it has a new predictive model. All Inline Services are isolated from updates to another Inline Service's models.

Predictive models are propagated to the Decision Server. Each Decision Server frequently polls the Oracle RTD database for new predictive models. When a new predictive model is found, it is loaded into memory.

### Decision Studio, RTD Clients and Tools

Decision Studio is an Eclipse based development environment for developing Oracle RTD Inline Services.

Decision Studio is also used to deploy and download an Inline Service from an Oracle RTD Server through the Oracle RTD Workbench Services web service.

Oracle RTD provides a number of client components that make it easier to integrate with Oracle RTD Decision Server. For example, Oracle RTD provides a Java client that provides a wrapper for Oracle RTD web service calls, caches default result sets on the client side, and provides default responses when the Oracle RTD server is not responding.

For more information on Oracle RTD client tools, refer to the *Oracle Fusion Middleware Platform Developer's Guide for Oracle Real-Time Decisions*.

### Administration of Oracle RTD

Oracle Enterprise Manager Fusion Middleware Control is used to configure Oracle RTD single instance and cluster-wide attributes.

#### 15.3.1.1 Oracle RTD Component Characteristics

This section describes the Oracle RTD component characteristics.

The Oracle RTD Server and its components are Java or web based components. They run either within the Java or web container. Communication between Decision Studio and Oracle RTD clients with the Oracle RTD Server is through a number of web services.

Oracle RTD manages session information locally and does not rely on a web server for session management. This is because Oracle RTD requires long running sessions to complete a business process vs the typical session associated with a web application. A single Decision Server session might receive requests from several HTTP sessions from different client servers. An Oracle RTD session can be fast, or it can last several hours, days or weeks.

If a server fails, the request will be directed to another server and a new session will be created. For the most part, session information is read from an external data source or calculated. These values will be reread or recalculated on a new session. Oracle RTD uses session information to calculate likelihood and not for storing state of a business transaction (this is delegated to a transactional/operational system). Therefore, the loss of session information does not adversely impact a business process.

**15.3.1.1.1 Component Lifecycle** Oracle RTD is managed through Oracle Enterprise Manager Fusion Middleware Control. The Fusion Middleware Control is used to start, stop, undeploy, deploy, and monitor Oracle RTD services.

**15.3.1.1.2 Process Flow** The process flow for Oracle RTD is:

1. A call is made by an external application to an Oracle RTD integration point.
2. The request is sent to the Oracle RTD's Decision Service web service, which is authenticated by Oracle Web Services Manager before the Decision Service serves the request.
3. The request is then forwarded to the targeted Inline Service and the target informant or advisor within the Inline Service.
4. Informant processing includes executing rules, executing functions, accessing entities, and creating learning records.
5. Advisor processing includes executing decisions, executing rules, executing functions for eligibility (filtering), accessing entities, and returning a list of eligible choices.

**15.3.1.1.3 External Dependencies** Configuration information for Oracle RTD is self-contained and stored in its own database schema. However, an Inline Service can load its entities or external rules from external data sources.

**15.3.1.1.4 Configuration Artifacts** Configuration information is stored in the Oracle RTD database, and it is updated using the Fusion Middleware Control.

Data source configuration information is set using the Oracle WebLogic Server Administration Console.

**15.3.1.1.5 Deployment Artifacts** Oracle RTD Inline Services are created in Decision Studio. Developers use the Decision Studio deployment facility to deploy an Inline Service. The Inline Service is packaged by Oracle RTD Studio and uploaded to an instance of Decision Server. The Decision Server stores all uploaded Inline Services to the Oracle RTD operational tables.

**15.3.1.1.6 Log File Locations** You can view Oracle RTD log files using Oracle Enterprise Manager Fusion Middleware Control. To view the log files from the Oracle RTD home

page, right-click the deployed Oracle RTD application in the navigator pane, and select **Logs > View Log Files**. This displays the Log Messages screen, on which you can view the log files. For more information on viewing Oracle RTD log files, see *Oracle Fusion Middleware Administrator's Guide for Oracle Real-Time Decisions*.


## 15.3.2 Oracle RTD High Availability Concepts

This section provides conceptual information about using Oracle RTD in a high availability deployment.

### 15.3.2.1 Oracle RTD High Availability Architecture

Figure 15-14 shows an Oracle RTD high availability active-active architecture.

**Figure 15-14 Oracle RTD High Availability Architecture**



In Figure 15-14, incoming requests are received by the hardware load balancer, which routes them to WEBHOST1 or WEBHOST2 in the web tier. These hosts have Oracle HTTP Server installed.

Oracle WebLogic Server is installed on APPHOST1 and APPHOST2 in the application tier. The Administration Server can be made active on APPHOST2 if APPHOST1

becomes unavailable, as described in [Chapter 3, "High Availability for WebLogic Server."](#) Oracle RTD is deployed on the BI\_SERVER1 and BI\_SERVER2 managed servers on these hosts. These managed servers are configured in an Oracle WebLogic cluster that allows them to work in an active-active manner.

In the data tier, an Oracle RAC database is used for the Oracle RTD repository, which stores learning models, learning records, prediction models, and Inline Services metadata, and is also recommended for the enterprise database.

---

**Note:** A cluster of Oracle RTD instances front-ended directly by a load balancer (instead of a web tier) is also supported. In this type of deployment, the load balancer has to be configured to ensure sticky session routing using the HTTP headers. See your load balancer documentation for information on configuring it to ensure sticky routing using the HTTP header.

Clients accessing the Oracle RTD cluster in this type of deployment must meet the requirements specified in the "About the Java Smart Client" section of *Oracle Fusion Middleware Platform Developer's Guide for Oracle Real-Time Decisions*. For an example of setting up the client, see the "Using the Java Smart Client" chapter of *Oracle Fusion Middleware Platform Developer's Guide for Oracle Real-Time Decisions*.

---

### Oracle RTD Cluster Considerations

In a clustered installation, instances of RTD components must communicate with other RTD components that can be located in the same host or a remote host. When the components are not in the same host they use an internal Remote Procedure Call (RPC) mechanism to communicate between cluster hosts and convey the response.

In a cluster, Decision Center is typically deployed in the same Oracle RTD instances as Learning Services, because Decision Center needs to communicate with Learning Services. Decision Center is not a singleton, however, so any instances not co-located with the active Learning Services use RPC to communicate with the active Learning Services.

In a cluster, Decision Center is typically deployed in a different Oracle RTD instance than Decision Services, which minimizes the performance and memory impact of Decision Center on Decision Services.

In a cluster, an Inline Service is deployed by Decision Studio to any cluster member, using the same internal WorkBench web service used in a non-clustered installation. The instance to which Decision Studio is connected will load the Inline Service into memory and will save the Inline Service in the database. Other cluster members will notice the new Inline Service in the database, and will load the Inline Service into their own memory.

For performance reasons it is best to deploy Oracle RTD in a cluster with HTTP session affinity enabled. This is because having the context of the Oracle RTD Session is necessary in order to create accurate learning records and making predictions. There is dynamic state in a session that can contribute to learning (for example, the sequence of web pages accessed in this Decision Server session), and that does not come from the database. While this state contributes to learning, it is only in an aggregated manner; there is no personal state, or any state important enough to mirror across servers, even if the state was Serializable. Losing a few sessions because of a server failure will not noticeably impact the integrity of Oracle RTD learning models, which are based on aggregated information.

Each instance of Decision Server in a cluster polls the Oracle RTD operational tables for new or updated Inline Services. If a new or newer version of an Inline Service is found, it is loaded into memory by every member of a cluster.

Oracle RTD uses hardware or software load balancers.

You can inject a Session Key into a web service request. The Session Key is user defined and passed in as a parameter to the Oracle RTD Java client (a wrapper to the Oracle RTD web service). The load balancer must be configured to have session affinity enabled and to retrieve the Session Key from the HTTP header.

Intra-cluster member communication between Oracle RTD instances is performed through its internal intra-cluster RPC technology. The JMS publish/subscribe model is used with a single topic "RTDTopic" created on Oracle RTD deployment (defined in Oracle RTD's Weblogic template). Each Oracle RTD instance subscribes to this single Topic and reads messages that apply to it (through filtering).

An Oracle RTD server may receive a request for which it does not have the request's Session state. In this case, the receiving Decision Server performs a lookup and finds the Decision Server to which this request should be directed. The request will then be forwarded to the correct Decision Server. The forwarding Decision Server will wait on the request to be completed by the receiving Decision Server.

The forwarded request is sent via RPC. Each Oracle RTD server pulls all messages that apply to it from the Oracle RTD JMS Topic.

After the receiving Decision Server completes the request, the call is unwound and the originator of the request forwarding gets the response from the receiving Decision Server. The response is then sent back to the caller.

### Singleton Services

In [Figure 15–14](#), the Oracle RTD Learning Service and Batch Manager are singletons. The Learning Service is responsible for building analytical models from data records written by one or more of the Oracle RTD Decision Service instances. An analytical model provides the likelihood of a specific outcome given a set of input values. For example, an analytical model can determine how likely a customer is to purchase a specific product given a prior history of products purchased. Learning models are built in memory; if the active Learning Service fails while building a model, a new Learning Service is activated and it will build the model from scratch with no loss of data. Batch Manager is responsible for starting and managing batch jobs on behalf of the batch administration client. It communicates with Batch Agent instances on other Oracle RTD instances to distribute the work across the cluster. Both Learning Service and Batch Manager are considered Active-Spare singletons.

### Cluster Coordinator

Oracle RTD has its own cluster coordinator, which performs these tasks:

- It assures that precisely one singleton is active of each type, Learning Service singleton and Batch Manager singleton. Only the cluster coordinator starts or stops a cluster singleton
- The cluster coordinator cleans up resources left open by a member that leaves the cluster (this cleanup logic is the same as the node would normally perform for itself when it shuts down in a controlled manner). The coordinator does this cleanup in case the leaving node died unexpectedly before it had a chance to close its own resources. e.g. All Decision Center sessions hosted by a cluster member will be closed by the coordinator when that cluster member leaves the cluster

On startup, each Oracle RTD cluster member vies to become the cluster coordinator by transactionally testing the database to see if another instance is already the coordinator, and inserts itself into the database as the coordinator if no other coordinator exists. A coordinator's right to be the coordinator expires unless periodically renewed by itself, so if a coordinator dies some other cluster member will successfully become coordinator.

**15.3.2.1.1 Starting and Stopping the Cluster** You can use the Oracle WebLogic Server Administration Console to start and stop single instances in an Oracle RTD cluster.

**15.3.2.1.2 Cluster-Wide Configuration Changes** Configuration information for Oracle RTD is self-contained and stored in its own database schema. It can be updated using Oracle Enterprise Manager Fusion Middleware Control.

When you change the configuration of one instance in an Oracle RTD cluster, the other instances pick up the configuration changes after you click the **Apply** button in Oracle Enterprise Manager Fusion Middleware Control.

### **15.3.2.2 Protection from Failures and Expected Behaviors**

The Oracle RTD topology management layer, running in each cluster member, will detect cluster members that stop updating their heartbeat records in the database. It does not check for failures of individual services within a managed server. So the failure scenarios described in this section are for managed server failures, not individual service failures.

**15.3.2.2.1 Decision Server Failure** If a Decision Server fails:

- The Oracle RTD sessions associated with the Decision Server are lost, along with their stateful information.
- The load balancer will direct requests originally destined to the failed server to another server.
- The server receiving these requests will create a new Oracle RTD session, hydrate session level entities from the appropriate data sources, and satisfy the request.

Learning records are written to the database frequently. They are used for statistical purposes. When a Decision Server fails:

- All non-persisted learning records are lost.
- Oracle RTD does not store any transactional data, such as a specific customer's information or account information, and so on. Therefore, when a large number of learning records have already been learned from and the prediction models have converged, the loss of some data is not material, and it will not significantly affect likelihood calculations and, therefore, the offers returned by an Oracle RTD Informant.

**15.3.2.2.2 Cluster Coordinator Failure** If the failed server is the cluster coordinator:

- It would stop updating the timestamp in the Oracle RTD database.
- Other servers polling the database table would realize that the timestamp has become stale and that there may be a problem with the cluster coordinator.
- All active Oracle RTD instances vie to become the new cluster coordinator.

**15.3.2.2.3 Learning Service Failure** Learning models are built in memory.



If the active Learning Service fails while building a model, a new Learning Service will be activated, and it will build the model from scratch with no loss of data.

If the node with the active Learning Service fails, the cluster coordinator will activate the Learning Service on one of the remaining cluster nodes (to which a Learning Service has been deployed and enabled). The cluster coordinator does this by sending a command to the Oracle RTD instance via an intra-cluster RPC.

**15.3.2.2.4 Decision Center Failure** If a Decision Center fails, a user loses any work in progress that was not saved, for example, rule updates or changes to performance goals that were not saved.

Since multiple Decision Centers run in a cluster, the user will be directed to an active Decision Center.

**15.3.2.2.5 Batch Manager Failure** If the node with the active Batch Manager fails, the cluster coordinator will activate a Batch Manager on one of the other remaining nodes (to which a Batch Manager has been deployed and enabled).

If the cluster's Batch Manager singleton fails, the cluster coordinator will notice the death and start another Batch Manager in a different host. All Batch Agents will be notified of the presence of the new Batch Manager, and the Batch Agents will then register themselves with the new Batch Manager.

If the node running a batch job fails:

- Only the batch jobs on the failed server will stop processing.
- Any batch jobs that were running on the failed node will enter an Interrupted state, and can be manually resumed by issuing a Resume command to the Batch Manager using the Batch Console or other Batch Admin Client.
- When an interrupted job is resumed, it will continue from the point where it last synchronized its state with the Batch Manager. When the Batch Manager is asked to resume a job, the job will resume on any available Batch Agent, selected by the Batch Manager.

### 15.3.3 Oracle RTD High Availability Configuration Steps

This section describes how to set up a two node highly available configuration for Oracle Real-Time Decisions. The architecture targeted for the configuration steps is shown in [Figure 15–14](#).

---

**Note:** Oracle strongly recommends reading the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

#### 15.3.3.1 Prerequisite Steps Before Setting up an Oracle RTD High Availability Configuration

This section includes the prerequisites for setting up the Oracle RTD high availability configuration.

**15.3.3.1.1 Database Prerequisites** Oracle RTD supports the following database versions:

- Oracle Database 10g (10.2.0.4 or later for non-XE database)
- Oracle Database 11g (11.1.0.7 or later for non-XE database)

To determine the database version, execute this query:

```
SQL>select version from sys.product_component_version where
product like 'Oracle%';
```

**15.3.3.1.2 Installing and Configuring the Database Repository** This section describes how to install and configure the database repository.

#### Oracle Clusterware

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Clusterware Installation Guide*.

#### Automatic Storage Management

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

When you run the installer, select the **Configure Automatic Storage Management** option in the **Select Configuration** page to create a separate Automatic Storage Management home.

#### Oracle Real Application Clusters

- For 10g Release 2 (10.2), see *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide*.
- For 11g Release 1 (11.1) and later, see *Oracle Real Application Clusters Installation Guide*.

You must install the 11g (11.1.1) Oracle Fusion Middleware Repository into an Oracle Real Application Clusters (Oracle RAC) database before you install the Oracle Fusion Middleware components. Oracle Fusion Middleware provides a tool, Oracle Fusion Middleware Repository Creation Utility (RCU), to create the component schemas in an existing database. You install RCU in its own, separate Middleware home.

Use the latest version of RCU to install the 11g (11.1.1.) Oracle Fusion Middleware Repository into an Oracle RAC database.

See *Oracle Fusion Middleware Repository Creation Utility User's Guide* for more information about obtaining and running the latest version of RCU.

**15.3.3.1.3 Using RCU to Load the Business Intelligence Schemas into the Database** Install the 11g (11.1.1) Oracle Fusion Middleware Metadata Store and Oracle BI schemas into an Oracle RAC database before you install Oracle RTD. Follow these steps:

1. Insert the Repository Creation Utility (RCU) DVD, and then start RCU from the bin directory in the RCU home directory.

```
prompt> cd RCU_HOME/bin
prompt> ./rcu
```

2. In the Welcome screen, click **Next**.
3. In the Create Repository screen, select **Create** to load component schemas into a database. Click **Next**.
4. In the Database Connection Details screen, enter connect information for your database:

- **Database Type:** Select **Oracle Database**.
- **Host Name:** Specify the name of the node on which the database resides. For the Oracle RAC database, specify the VIP name or one of the node names as the host name, for example:  
BIDBHOST1-VIP
- **Port:** Specify the listen port number for the database: 1521
- **Service Name:** Specify the service name of the database:  
biha.mycompany.com
- **Username:** Specify the name of the user with DBA or SYSDBA privilege: SYS
- **Password:** Enter the password for the SYS user.
- **Role:** Select the database user's role from the list: SYSDBA (required by the SYS user).

Click **Next**.

5. In the Select Components screen:

- Select **Create a new Prefix**, and enter a prefix to use for the database schemas, for example, BIHA. You can specify up to six characters as a prefix. Prefixes are used to create logical groupings of multiple repositories in a database. For more information, see *Oracle Fusion Middleware Repository Creation Utility User's Guide*.

**Tip:** Note the name of this schema because the upcoming steps require this information.

- Select the following components:
  - AS Common Schemas: Metadata Services (automatically selected)
  - Oracle Business Intelligence: Business Intelligence Platform

Click **Next**.

6. In the Schema Passwords screen, enter passwords for the main schema users and click **Next**.

You can choose either **Use same passwords for all schemas** or **Specify different passwords for all schemas**, depending on your requirements.

Do not select **Use main schema passwords for auxiliary schemas**. The auxiliary passwords are derived from the passwords of the main schema users.

**Tip:** Note the names of the schema passwords because the upcoming steps require this information.

7. In the Map Tablespaces screen, choose the tablespaces for the selected components, and click **Next**.
8. In the Summary screen, click **Create**.
9. In the Completion Summary screen, click **Close**.

**15.3.3.1.4 Configuring Virtual Server Names and Ports for the Load Balancer** This section describes the load balancer prerequisites for deploying an Oracle RTD high availability environment.

## Load Balancers

Oracle RTD uses a hardware load balancer when deployed in a high availability configuration with two Oracle HTTP Servers as web tier components.

## Virtual Server Names

bi.mycompany.com is a virtual server name that acts as the access point for all HTTP traffic to the runtime BI components, such as Oracle RTD. Traffic to both SSL and non-SSL ports is configured, and typically non-SSL is redirected to SSL. Clients access this service using the address bi.mycompany.com:443. This virtual server is defined on the load balancer.

**15.3.3.1.5 Installing Oracle HTTP Server on WEBHOST1 and WEBHOST2** This section describes how to install Oracle HTTP Server on WEBHOST1.

1. Verify that the servers meet the following requirements:

- Ensure that the system, patch, kernel and other requirements are met. These are listed in the *Oracle Fusion Middleware Installation Guide for Oracle Web Tier* in the Oracle Fusion Middleware documentation library for the platform and version you are using.
- Ensure that port 7777 is not in use by any service on WEBHOST1 by issuing these commands for the operating system you are using. If a port is not in use, no output is returned from the command.

### On UNIX:

```
netstat -an | grep "7777"
```

### On Windows:

```
netstat -an | findstr :7777
```

If the port is in use (if the command returns output identifying the port), you must free the port.

### On UNIX:

Remove the entries for port 7777 in the `/etc/services` file and restart the services, or restart the computer.

### On Windows:

Stop the component that is using the port.

- Copy the `staticports.ini` file from the `Disk1/stage/Response` directory to a temporary directory.
- Edit the `staticports.ini` file that you copied to the temporary directory to assign the following custom ports (uncomment the line where you specify the port number for Oracle HTTP Server):

```
The port for Oracle HTTP server
Oracle HTTP Server port = 7777
```

2. Start the Oracle Universal Installer for Oracle Fusion Middleware 11g Web Tier Utilities CD installation as follows:

### On UNIX:

Issue the command: `./runinstaller`

### On Windows:

Double-click `setup.exe`.

This displays the Specify Inventory Directory screen.

3. On the Specify Inventory Directory screen, enter values for the Oracle Inventory Directory and the Operating System Group Name. For example:

**Specify the Inventory Directory:** `/u01/app/oraInventory`

**Operating System Group Name:** `oinstall`

A dialog box appears with the following message:

"Certain actions need to be performed with root privileges before the install can continue. Please execute the script  
`/u01/app/oraInventory/createCentralInventory.sh` now from another window  
 and then press "Ok" to continue the install. If you do not have the root privileges  
 and wish to continue the install select the "Continue installation with local  
 inventory" option"

Log in as root and run the "`/u01/app/oraInventory/createCentralInventory.sh`"

This sets the required permissions for the Oracle Inventory Directory and then brings up the Welcome screen.

---

**Note:** The Oracle Inventory screen is not shown if an Oracle product was previously installed on the host. If the Oracle Inventory screen is not displayed for this installation, make sure to check and see:

- If the `/etc/oraInst.loc` file exists.
  - If the file exists, the Inventory directory listed is valid.
  - The user performing the installation has write permissions for the Inventory directory.
- 

4. On the Welcome screen, click **Next**.
5. On the Select Installation Type screen, select **Install and Configure**, and click **Next**.
6. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.

Click **Next**.

7. On the Specify Installation Location screen, specify:
  - **Oracle Middleware Home:** `ORACLE_BASE/product/fmw` (grayed out)
  - **Oracle Home Directory:** `Oracle_WT1`

Click **Next**.

8. On the Configure Components screen:
  - Select **Oracle HTTP Server**.
  - Do not select **Associate Selected Components with Weblogic Domain**.

Click **Next**.

9. On the Specify Component Details screen, enter the following values for `WEBHOST1`:

- **Instance Home Location:** `ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1`
- **Instance Name:** `web1`
- **OHS Component Name:** `ohs1`

Click **Next**.

10. On the Configure Ports screen:

- If you specify a custom port, select **Specify Ports using Configuration File** and then use the Browse function to select the file.
- Enter the Oracle HTTP Server port, for example: `7777`

Click **Next**.

11. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address and click **Next**.

12. On the Installation Summary screen, ensure that the selections are correct. If they are not, click **Back** and make the necessary fixes. After ensuring that the selections are correct, click **Install**.

13. On the Installation Progress screen on UNIX systems, a dialog appears that prompts you to run the `oracleRoot.sh` script. Open a window and run the script, following the prompts in the window.

Click **Next**.

14. On the Configuration Progress screen, multiple configuration assistants are launched in succession; this process can be lengthy. Click **Next**.

15. On the Installation Complete screen, click **Finish** to exit.

Repeat the steps for WEBHOST2 and configure your load balancer with a pool containing both the WEBHOST1 and WEBHOST2 addresses.

**15.3.3.1.6 Validating Oracle HTTP Server** To verify that Oracle HTTP Server is set up correctly, access the root URL context of the server by using the following URLs in the browser:

```
http://WEBHOST1:7777/
http://WEBHOST2:7777/
```

If Oracle HTTP Server is set up correctly, the Oracle FMW Welcome screen appears in the browser.

### 15.3.3.2 Installing Oracle Fusion Middleware Home

This section describes how to install Oracle Fusion Middleware on all nodes in the application tier that run Oracle WebLogic Server and Oracle Fusion Middleware BI EE. Repeat the procedure (described below for APPHOST1) for installing WebLogic Server and Oracle RTD in APPHOST2. The directory paths for binary files and domains used when installing new nodes must be exactly the same as those used for first node.

Install the following Oracle Fusion Middleware components:

- Oracle WebLogic Server
- Oracle Fusion Middleware BI EE

**15.3.3.2.1 Installing Oracle WebLogic Server** Follow these steps to install Oracle WebLogic Server:

1. Start the installer for Oracle WebLogic Server from the installation media.
2. In the Welcome screen, click **Next**.
3. In the Choose Middleware Home Directory screen:
  - Select **Create a new Middleware Home**.
  - For the **Middleware Home Directory**, enter `ORACLE_BASE/product/fmw`

---

**Note:** `ORACLE_BASE` is the base directory under which Oracle products are installed. The recommended value is `/u01/app/oracle`.

---

Click **Next**.

4. In the Register for Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.
5. In the Choose Install Type screen, select **Typical** and click **Next**.
6. In the Choose Product Installation Directories screen, accept the directory `ORACLE_BASE/product/fmw/wlserver_10.3` for WebLogic Server and `ORACLE_BASE/product/fmw/coherence_3.6` for Oracle Coherence and click **Next**.
7. In the Installation Summary screen, click **Next**.  
The Oracle WebLogic Server software is installed.
8. In the Installation Complete screen, clear the Run Quickstart check box and click **Done**.

**15.3.3.2.2 Installing Oracle RTD** On the Linux platform, if the `/etc/oraInst.loc` file exists, check that its contents are correct. Specifically, check that the inventory directory is correct and that you have write permissions for that directory. If the `/etc/oraInst.loc` file does not exist, you can skip this step.

1. Start the installer for Oracle Fusion Middleware Business Intelligence 11g from the installation media:

**On UNIX:**

```
APPHOST1> ./runInstaller
```

**On Windows:**

```
APPHOST1> setup.exe
```

2. In the Specify Inventory Directory screen:
  - Enter `HOME/oraInventory`, where `HOME` is the home directory of the user performing the installation (this is the recommended location).
  - Enter the OS group for the user performing the installation.

Click **Next**.

  - Follow the instructions on screen to execute `/createCentralInventory.sh` as root.

Click **OK**.

3. In the Welcome screen, click **Next**.
4. In the Select Installation Type screen, select **Software Only Install** and click **Next**.
5. In the Prerequisite Checks screen, verify that all checks complete successfully, and click **Next**.
6. In the Specify Installation Location screen, select the previously installed Middleware Home from the drop-down list. For the Oracle Home directory, enter the directory name (Oracle\_BI1).

Click **Next**.

7. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address and click **Next**.
8. In the Summary screen, click **Install**.  
The Oracle Fusion Middleware Business Intelligence 11g software is installed.
9. In the Installation Progress screen, click **Next**.
10. In the Complete screen, click **Finish**.

### 15.3.3.3 Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2

The BI domain uses virtual hostnames as the listen addresses for the BI managed servers. You must enable the VIPs, mapping each of these hostnames on the two BI machines (VIP1 on APPHOST1 and VIP2 on APPHOST2), and they must correctly resolve to the virtual hostnames in the network system used by the topology (either by DNS Server or hosts resolution).

### 15.3.3.4 Creating a Domain with the Administration Server and the First BI\_SERVER1 Managed Server

This section describes how to create a domain and the first WebLogic Server BI managed server using the Oracle Business Intelligence Configuration Assistant, Oracle WebLogic Server Administration Console, and Oracle Enterprise Manager. Ensure that the database where you installed the repository is running. For Oracle RAC databases, all the instances must be running.

---

---

**Note:** Oracle strongly recommends that you read the release notes for any additional installation and deployment considerations prior to starting the setup process.

---

---

Run the Configuration Assistant from the Oracle home directory to create a domain containing the Administration Server and the managed server with Oracle RTD:

1. Change the directory to the location of the Configuration Assistant:

```
APPHOST1> cd ORACLE_HOME/bin
```

2. Start the Configuration Assistant:

**On UNIX:**

```
APPHOST1> ./config.sh
```

**On Windows:**



```
APPHOST1> config.bat
```

3. In the Welcome screen, click **Next**.
4. In the Prerequisite Checks screen, verify that all checks complete successfully, and click **Next**.
5. The Create or Scale Out BI System screen opens. Select **Create New BI System** and specify the following:
  - **User Name:** weblogic
  - **User Password:** Enter the weblogic user password.
  - **Domain Name:** *bifoundation\_domain*

Click **Next**.

6. In the Specify Installation Location screen, enter:
  - **Middleware Home:** *ORACLE\_BASE/product/fmw* (grayed out)
  - **Oracle Home:** *ORACLE\_BASE/product/fmw/Oracle\_BI1* (grayed out)
  - **WebLogic Server Home:** *ORACLE\_BASE/product/fmw/wlserver\_10.3* (grayed out)
  - **Domain Home:** *ORACLE\_BASE/product/fmw/user\_projects/domain/bifoundation\_domain*

---

**Note:** The Domain Home must end with the domain name.

---

- **Instance Home:** *ORACLE\_BASE/product/fmw/instance1*
- **Instance Name:** *instance1*

Click **Next**.

7. In the Configure Components screen, select **Oracle Real-Time Decisions**.  
Click **Next**.

8. In the BIPLATFORM Schema screen, enter:
  - **Database Type:** Oracle Database
  - **Connect String:**  
*BIDBHOST1:1521:BIDB1^BIDBHOST2:1521:BIDB2@BIHA.MYCOMPANY.COM*
  - **BIPLATFORM Schema Username:** *BIHA\_BIPLATFORM*
  - **BIPLATFORM Schema Password:** Enter the password for the BIPLATFORM schema user.

Click **Next**.

9. In the Configure Ports screen, select one of the following:
  - Auto Port Configuration
  - Specify Ports using Configuration File

Click **Next**.

10. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.
11. In the Summary screen, click **Configure**.
12. In the Configuration Progress screen, verify that all the Configuration Tools have completed successfully and click **Next**.
13. In the Complete screen, click **Finish**.

#### 15.3.3.5 Creating boot.properties for the Administration Server on APPHOST1

This is an optional step for enabling the Administration Server to start without prompting you for the administrator username and password. Create a boot.properties file for the Administration Server on APPHOST1.

1. Go to the following directory:

```
ORACLE_BASE/product/fmw/user_projects/domains/bifoundation_domain/servers/
AdminServer/security
```

2. Enter the following lines in the file, save the file, and close the editor:

```
username=Admin_username
password=Admin_password
```

---

---

**Note:** When you start the Administration Server, the username and password entries in the file are encrypted.

For security reasons, minimize the time the entries in the file are left unencrypted. After you edit the file, start the server as soon as possible for the entries to be encrypted.

---

---

#### 15.3.3.6 Starting and Validating the Administration Server on APPHOST1

This section describes procedures for starting and validating the Administration Server on APPHOST1.

**15.3.3.6.1 Starting the Administration Server on APPHOST1** To start the Administration Server on APPHOST1, run the following commands:

```
APPHOST1> cd MW_HOME/user_projects/domains/bifoundation_domain/bin
APPHOST1> ./startWebLogic.sh
```

**15.3.3.6.2 Validating the Administration Server** verify that the Administration Server is properly configured:

1. In a browser, go to <http://APPHOST1:7001/console>.
2. Log in as the administrator.
3. Verify that the BI\_SERVER1 managed server is listed.
4. Verify that the bi\_cluster cluster is listed.
5. Verify that you can access Enterprise Manager at <http://APPHOST1:7001/em>.

#### 15.3.3.7 Setting the Listen Address for BI\_SERVER1 Managed Server

Perform these steps to set the managed server listen address:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server1** in the **Names** column of the table.
6. Set the Listen Address to APPHOST1VHN1.
7. Click **Save**.
8. Save and activate the changes.

The changes will not take effect until the BI\_SERVER1 managed server is restarted.

#### 15.3.3.8 Disabling Host Name Verification for the BI\_SERVER1 Managed Server

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you will receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and validating the topology, and enable it again once the high availability topology configuration is complete.

To disable host name verification:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server1** in the **Names** column of the table. The settings page for the server is displayed.
6. Open the **SSL** tab.
7. Expand the **Advanced** section of the page.
8. Set **Hostname Verification** to **None**.
9. Click **Save**.
10. Save and activate the changes.
11. The change will not take effect until the BI\_SERVER1 managed server is restarted.

#### 15.3.3.9 Starting the System in APPHOST1

This section describes how to start Node Manager on APPHOST1 and how to start and validate the BI\_SERVER1 managed server on APPHOST1.

**15.3.3.9.1 Starting Node Manager on APPHOST1** Usually, Node Manager is started automatically when config.sh completes. If Node Manager is not running for some reason, start it on APPHOST1 using these commands:

```
APPHOST1> cd WL_HOME/server/bin
APPHOST1> ./startNodeManager.sh
```

**15.3.3.9.2 Starting and Validating the BI\_SERVER1 Managed Server** To start the BI\_SERVER1 managed server and check that it is configured correctly:

1. Start the `bi_server1` managed server using Oracle WebLogic Server Administration Console, as follows:
  - a. Expand the **Environment** node in the Domain Structure window.
  - b. Choose **Servers**.
  - c. Click the **Control** tab.
  - d. Select **bi\_server1** and then click **Start**.
2. Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming", wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors
3. When `BI_SERVER1` is started, the following URLs become available:
  - Access `http://APPHOST1:9704/wsm-pm` to verify the status of Web Services Manager. Click **Validate Policy Manager**. A list of policies and assertion templates available in the data is displayed.

---

**Note:** The configuration is incorrect if no policies or assertion templates appear.

---

- Access `http://APPHOST1:9704/ui` to verify the status of the Oracle RTD Decision Center application.

#### 15.3.3.10 Scaling Out the BI System on APPHOST2

Run the Configuration Assistant from the `ORACLE_HOME` directory to scale out the BI System.

1. Change the directory to the location of the Configuration Assistant:

```
APPHOST2> cd ORACLE_HOME/bin
```
2. Start the Configuration Assistant:

```
APPHOST2> ./config.sh
```
3. In the Welcome screen, click **Next**.
4. On the Prerequisite Checks screen, the installer completes the prerequisite check. If any fail, fix them and restart your installation.  
Click **Next**.
5. The Create or Scale out BI System screen opens. Select **Scale Out BI System** and then enter the following values:
  - **Host Name:** ADMINHOST
  - **Port:** 7001
  - **User name:** weblogic
  - **User Password:** Enter the password for the weblogic user.Click **Next**.
6. In the Scale Out BI System Details screen, enter:
  - **Middleware Home:** `ORACLE_BASE/product/fmw` (grayed out)

- **Oracle Home:** Oracle\_BI1 (grayed out)
- **WebLogic Server Home:** *ORACLE\_BASE/product/fmw/wlserver\_10.3* (grayed out)
- **Domain Home:** *ORACLE\_BASE/product/fmw/user\_projects/domain/bifoundation\_domain*

---

**Note:** The Domain Home must end with the domain name.

---

- **Instance Home:** *ORACLE\_BASE/product/fmw/instance2*
- **Instance Name:** instance2 (grayed out)

Click **Next**.

7. In the Configure Ports screen, select one of the following:

- Auto Port Configuration
- Specify Ports using Configuration File

Click **Next**.

8. In the Specify Security Updates screen, choose whether you want to receive security updates from Oracle Support and if you do, enter your email address. Click **Next**.

9. In the Summary screen, click **Configure**.

10. In the Configuration Progress screen, verify that all the Configuration Tools have completed successfully and click **Next**.

11. In the Complete screen, click **Finish**.

#### 15.3.3.11 Setting the Listen Address for the BI\_SERVER2 Managed Server

To set the BI\_SERVER2 managed server listen address:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server2** in the **Names** column of the table.  
The Setting page for bi\_server2 is displayed.
6. Set the Listen Address to APPHOST2VHN1.
7. Click **Save**.
8. Save and activate the changes.

The changes do not take effect until the BI\_SERVER2 managed server restarts.

#### 15.3.3.12 Disabling Host Name Verification for the BI\_SERVER2 Managed Server

This step is required if you have not set up the appropriate certificates to authenticate the different nodes with the Administration Server. If you have not configured the server certificates, you will receive errors when managing the different WebLogic Servers. To avoid these errors, disable host name verification while setting up and

validating the topology, and enable it again once the high availability topology configuration is complete.

To disable host name verification:

1. Log into Oracle WebLogic Server Administration Console.
2. In the Change Center, click **Lock & Edit**.
3. Expand the **Environment** node in the Domain Structure window.
4. Click **Servers**.
5. Select **bi\_server2** in the **Names** column of the table. The settings page for the server is displayed.
6. Open the **SSL** tab.
7. Expand the **Advanced** section of the page.
8. Set **Hostname Verification** to **None**.
9. Click **Save**.
10. Save and activate the changes.
11. The change will not take effect until the Administration Server is restarted (ensure that Node Manager is up and running):
  - a. In the Summary of Servers screen, select the **Control** tab.
  - b. Select **AdminServer(admin)** (represented as a hyperlink) in the table and then click **Shutdown**.
  - c. Start the Administration Server again from the command line:

```
APPHOST1> cd ORACLE_COMMON_HOME/bin
APPHOST1> ./wlst.sh
```

Once in the WLST shell, execute the following command (make sure Node Manager is up and running):

```
wls:/offline> nmConnect ('Admin_User','Admin_Password', 'APPHOST1', '9556',
'domain_name', 'ORACLE_BASE/product/fmw/user_projects/domain/
bifoundation_domain')

wls:/nm/domain_name> nmStart ('AdminServer')
```

### 15.3.3.13 Configuring Oracle RTD

This section includes information about configuring Oracle RTD.

**15.3.3.13.1 Configuring RTD Cluster-Specific Properties** Perform the following steps in Oracle Enterprise Manager Fusion Middleware Control:

1. Log into Fusion Middleware Control.
2. Expand the **Application Deployments** node in the *Farm\_domain\_name* window.
3. Click **OracleRTD(11.1.1)(bi\_cluster)**.
4. Click on any node under it. For example, **OracleRTD(11.1.1)(bi\_server1)**.
5. On the right pane, click **Application Deployment**, and then select **System MBean Browser**.

6. On the System MBean Browser pane, expand **Application Defined MBeans**.
7. For each server under Oracle RTD, navigate to the MBean and set the attribute shown in [Table 15–9](#).

**Table 15–9 Attribute Value to Set for Oracle RTD MBeans**

| MBean                          | Attribute              | BI_SERVER1 Value        |
|--------------------------------|------------------------|-------------------------|
| SDClusterPropertyManager->Misc | DecisionServiceAddress | http://bi.mycompany.com |

8. Click **Apply**.

### 15.3.3.14 Starting the System in APPHOST2

This section describes procedures for starting the system in APPHOST2.

**15.3.3.14.1 Configuring a Default Persistence Store for Transaction Recovery** Each server has a transaction log, which stores information about committed transactions coordinated by the server that may not have been completed. WebLogic Server uses the transaction log when recovering from system crashes or network failures. To leverage the migration capability of the Transaction Recovery Service for the servers within a cluster, store the transaction log in a location accessible to the server.

---

**Note:** Preferably, this location should be a dual-ported SCSI disk or on a Storage Area Network (SAN).

---

To set the location for the default persistence store for BI\_SERVER1:

1. Log into the Oracle WebLogic Server Administration Console.
2. In the Domain Structure window, expand the **Environment** node and then click the **Servers** node.
3. Click **BI\_SERVER1** (represented as a hyperlink) in the **Name** column of the table.  
The settings page for the BI\_SERVER1 server opens with the Configuration tab active.
4. Click the **Services** tab.
5. In the Change Center, click **Lock & Edit**.
6. In the Default Store section of the page, enter the path to the folder where the default persistent stores will store its data files. The directory structure of the path is as follows:  
`ORACLE_BASE/admin/domain_name/bi_cluster/tlogs`
7. Click **Save** and **Activate Changes**.
8. Repeat these steps for the BI\_SERVER2 server.

---

**Note:** To enable migration of the Transaction Recovery Service, specify a location on a persistent storage solution that is available to other servers in the cluster. Both APPHOST1 and APPHOST2 must be able to access this directory. This directory must also exist before you restart the server.

---

**15.3.3.14.2 Starting Node Manager on APPHOST2** Usually, Node Manager is started automatically when config.sh completes. If Node Manager is not running for some reason, start the Node Manager on APPHOST2 by repeating the steps in [Section 15.3.3.9.1, "Starting Node Manager on APPHOST1."](#)

**15.3.3.14.3 Starting and Validating the BI\_SERVER2 Managed Server** To start the BI\_SERVER2 managed server and check that it is configured correctly:

1. Start the bi\_server2 managed server using Oracle WebLogic Server Administration Console, as follows:
  - a. Expand the **Environment** node in the Domain Structure window.
  - b. Choose **Servers**.
  - c. Click the **Control** tab.
  - d. Select **bi\_server2** and then click **Start**.
2. Verify that the server status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming", wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors
3. When BI\_SERVER2 is started, the following URLs become available:
  - Access <http://APPHOST2:9704/wsm-pm> to verify the status of Web Services Manager. Click **Validate Policy Manager**. A list of policies and assertion templates available in the data is displayed.

---

**Note:** The configuration is incorrect if no policies or assertion templates appear.

---

- Access <http://APPHOST2:9704/ui> to verify the status of the Oracle RTD Decision Center application.

### 15.3.3.15 Configuring Oracle HTTP Server for the BI\_SERVERn Managed Servers

To enable Oracle HTTP Server to route to bi\_cluster, which contains the BI\_SERVERn managed servers, you must set the WebLogicCluster parameter to the list of nodes in the cluster:

1. On WEBHOST1 and WEBHOST2, add the following lines to the `ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/config/OHS/ohs1/mod_wl_ohs.conf` file:

```
<Location /rtis>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1:9704, APPHOST2:9704
 WLPProxySSL ON
 WLPProxySSLPassThrough ON
</Location>

<Location /schema>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1:9704, APPHOST2:9704
 WLPProxySSL ON
 WLPProxySSLPassThrough ON
</Location>

<Location /ws>
```



```

 SetHandler weblogic-handler
 WebLogicCluster APPHOST1VHN1:9704, APPHOST2VHN1:9704
 WLProxySSL ON
 WLProxySSLPassThrough ON
 </Location>

 # WSM-PM
 <Location /wsm-pm>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1:9704, APPHOST2:9704
 WLProxySSL ON
 WLProxySSLPassThrough ON
 </Location>

 # RTD
 <Location /ui>
 SetHandler weblogic-handler
 WebLogicCluster APPHOST1:9704, APPHOST2:9704
 WLProxySSL ON
 WLProxySSLPassThrough ON
 </Location>

```

2. Perform the same steps for the Oracle HTTP Server on WEBHOST2.
3. Restart Oracle HTTP Server on WEBHOST1 and WEBHOST2:

```

WEBHOST1> ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/bin/
opmnctlrestartproc ias-component=ohs1

```

```

WEBHOST2> ORACLE_BASE/product/fmw/Oracle_WT1/instances/web1/bin/
opmnctlrestartproc ias-component=ohs1

```

### 15.3.3.16 Validating Access Through Oracle HTTP Server

Verify that the BI Servers status is reported as "Running" in the Administration Console. If the server is shown as "Starting" or "Resuming," wait for the server status to change to "Started". If another status is reported (such as "Admin" or "Failed"), check the server output log files for errors.

Verify that you can access these URLs:

- <http://WEBHOST1:7777/wsm-pm>
- <http://WEBHOST2:7777/wsm-pm>
- <http://WEBHOST1:7777/ui>
- <http://WEBHOST2:7777/ui>

Verify that you can access these URLs using your load balancing router address:

- <http://bi.mycompany.com:80/ui>
- <http://bi.mycompany.com:80/wsm-pm>

Follow these instructions to ensure that routing and failover from the HTTP Server to the bi\_cluster is working correctly:

1. While BI\_SERVER2 is running, stop BI\_SERVER1 from Oracle WebLogic Server Administration Console.
2. Access the following URLs and verify the appropriate functionality:
  - [WEBHOST1:7777/wsm-pm](http://WEBHOST1:7777/wsm-pm)

- WEBHOST1:7777/ui
- 3. Start BI\_SERVER1 from Oracle WebLogic Server Administration Console.
- 4. Stop BI\_SERVER2.
- 5. Access the URLs in Step 2 above again and verify the appropriate functionality.

#### 15.3.3.17 Configuring Server Migration for the BI\_SERVERn Servers

In this high availability topology, you must configure server migration for the bi\_server1 and bi\_server2 Managed Servers. To do this, follow the instructions in [Section 15.1.12.22, "Configuring Server Migration for the BI\\_SERVERn Servers."](#)

#### 15.3.3.18 Scaling Out the Oracle RTD Topology to a New Node (APPHOST3)

When you scale out the Oracle RTD topology, you add a new managed server and set of system components to a new node in your topology (APPHOST3). This section assumes that you already have a high availability topology that includes at least two nodes, with a managed server and a full set of system components on each node.

Before performing the steps in this section, check that you meet these requirements:

- There must be existing nodes running Oracle RTD managed servers within the topology.
- The new node (APPHOST3) can access the existing home directories for Oracle WebLogic Server and Oracle Business Intelligence.
- When an ORACLE\_HOME or WL\_HOME is shared by multiple servers in different nodes, it is recommended that you keep the Oracle Inventory and Middleware home list in those nodes updated for consistency in the installations and application of patches. To update the oraInventory in a node and "attach" an installation in a shared storage to it, use `ORACLE_HOME/oui/bin/attachHome.sh`. To update the Middleware home list to add or remove a WL\_HOME, edit the `MW_HOME/.home` file. See the steps below.
- The new server can use a new individual domain directory or, if the other Managed Servers domain directories reside on shared storage, reuse the domain directories on those servers.

To scale out Oracle RTD on APPHOST3:

1. On APPHOST3, mount the existing Middleware home, which should include the Oracle Business Intelligence installation and (optionally, if the domain directory for Managed Servers in other nodes resides on shared storage) the domain directory, and ensure that the new node has access to this directory, just like the rest of the nodes in the domain.
2. To attach `ORACLE_HOME` in shared storage to the local Oracle Inventory, run the following command:

```
APPHOST3> cd ORACLE_COMMON_HOME/oui/bin/
APPHOST3> ./attachHome.sh -jreLoc ORACLE_BASE/product/fmw/jrockit_160_<version>
```

To update the Middleware home list, create (or edit, if another WebLogic installation exists in the node) the `MW_HOME/.home` file and add `ORACLE_BASE/product/fmw` to it.

3. Enable VIP3 in APPHOST3. See [Section 15.2.3.3, "Enabling VIP1 in APPHOST1 and VIP2 in APPHOST2."](#)

4. Run the Configuration Assistant from one of the shared Oracle homes, using the steps in [Section 15.3.3.10, "Scaling Out the BI System on APPHOST2"](#) as a guide.
5. Configure the bi\_server3 Managed Server by setting the Listen Address and disabling host name verification, using the steps in [Section 15.3.3.11, "Setting the Listen Address for the BI\\_SERVER2 Managed Server"](#) and [Section 15.3.3.12, "Disabling Host Name Verification for the BI\\_SERVER2 Managed Server"](#) as a guide.
6. Perform additional Oracle Real-Time Decisions configuration steps for the bi\_server3 Managed Server on APPHOST3, using the steps in [Section 15.3.3.13, "Configuring Oracle RTD"](#) as a guide.
7. Set the location of the default persistence store for bi\_server 3, as described in [Section 15.3.3.14.1, "Configuring a Default Persistence Store for Transaction Recovery."](#)
8. Configure Oracle HTTP Server for APPHOST3VHN1, using the steps in [Section 15.3.3.15, "Configuring Oracle HTTP Server for the BI\\_SERVERn Managed Servers"](#) as a guide.
9. Start the bi\_server3 Managed Server and the system components running on APPHOST3. See [Section 15.3.3.9.2, "Starting and Validating the BI\\_SERVER1 Managed Server"](#) for details.
10. To validate the configuration, access the `http://APPHOST3VHN1:9704/ui` URL to verify the status of the Oracle RTD application.
11. Oracle recommends using host name verification for the communication between Node Manager and the servers in the domain. This requires the use of certificates for the different addresses communicating with the Administration Server and other servers. See the chapter on setting up Node Manager in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence* for further details.



---

## Using HA Power Tools

You can configure the Java Object Cache using the HA Power Tools tab in the Oracle WebLogic Server Administration Console.

This chapter contains the following sections:

- [Section 16.1, "Enabling HA Power Tools in the Oracle WebLogic Administration Console"](#)
- [Section 16.2, "Configuring Java Object Cache for a Cluster Using HA Power Tools"](#)
- [Section 16.3, "Configuring Java Object Cache for Managed Servers Using HA Power Tools"](#)

### 16.1 Enabling HA Power Tools in the Oracle WebLogic Administration Console

To use HA Power Tools, you must enable them using the following procedure:

1. Copy the following two .war files located in the *MW\_HOME/oracle\_common/hapowertools* directory to the *DOMAIN\_HOME/console-ext* directory.
  - powertools-core.war
  - powertools-configurejoc.war
2. Restart the Administration Server.
3. Click **Domain** in the left navigation tree to see the **HA Power Tools** tab.

### 16.2 Configuring Java Object Cache for a Cluster Using HA Power Tools

To configure Java Object Cache for a cluster using HA Power Tools:

1. Log in to Oracle WebLogic Server Administration Console, and click **Domain** in left navigation tree.
2. Click the **HA Power Tools** tab.
3. Select the **Configure JOC for Cluster** radio button.
4. Select the cluster name from the drop-down list.
5. Enter the **Discover Port**.
6. Enter the Hosts (separate host names by comma).
7. Click **Configure JOC**.
8. Restart the cluster.

9. Verify JOC configuration using the CacheWatcher utility. See [Appendix F, "Configuring the Distributed Java Object Cache."](#)

## 16.3 Configuring Java Object Cache for Managed Servers Using HA Power Tools

To configure Java Object Cache for Managed Servers using HA Power Tools:

1. Log in to Oracle WebLogic Server Administration Console, and click **Domain** in left navigation tree.
2. Click the **HA Power Tools** tab
3. Select the **Configure JOC for Managed Server(s)** radio button.
4. Enter the Managed Server names and Discover Port.
5. Enter the Hosts (separate host names by comma).
6. Click **Configure JOC**.
7. Restart the cluster.
8. Verify JOC configuration using the Cache Watcher utility. See [Appendix F, "Configuring the Distributed Java Object Cache."](#)

Some of the managed servers can be excluded from the JOC configuration by specifying the server names (comma separated) in the **Exclude Server** section. The Distribution Mode for all the specified servers are set to **false**.

Use the **Disable Distribution Mode** option to disable Distribution Mode for all the affected managed servers. This option is applicable for configuring JOC for a cluster as well as for configuring JOC for Managed Servers.

---

## Setting Up Auditing with an Oracle RAC Database Store

With Oracle Fusion Middleware 11g, you have the option of setting up the Oracle Fusion Middleware Audit Framework service. Auditing provides a measure of accountability and answers "who has done what and when" types of questions.

The Oracle Fusion Middleware Audit Framework is designed to provide a centralized audit framework for middleware products. The framework provides audit service for the following:

- **Middleware Platform** - This includes Java components such as Oracle Platform Security Services (OPSS) and Oracle Web Services. These are components that are leveraged by applications deployed in the middleware. Indirectly, all the deployed applications leveraging these Java components will benefit from the audit framework auditing events that are happening at the platform level.
- **JavaEE applications** - The objective is to provide a framework for JavaEE applications, starting with Oracle's own Java components. JavaEE applications will be able to create application-specific audit events. In the current release, the Java EE components using the Oracle Fusion Middleware Audit Framework are internal Oracle components.
- **System components** - For system components in the middleware that are managed by Oracle Process Manager and Notification Server (OPMN), the audit framework also provides an end-to-end service similar to that for Java components.

See the "Introduction to Oracle Fusion Middleware Audit Framework" chapter in the *Oracle Fusion Middleware Application Security Guide* for more introductory information about Oracle Fusion Middleware Audit Framework.

Out of the box, the Audit Framework uses the file system to store audit records. In a production environment, however, Oracle recommends that you use a database audit store to provide scalability and high availability for the audit framework. In high availability configurations such as the configurations described in this chapter, Oracle recommends that you use an Oracle Real Application Clusters (Oracle RAC) database as the database audit store.

The "Configuring and Managing Auditing" chapter in the *Oracle Fusion Middleware Application Security Guide* includes the steps for configuring auditing. The "Managing the Audit Store" section in that chapter includes steps for setting up a database as the audit data store.

When you set up the Oracle Fusion Middleware Audit Framework with an Oracle RAC database audit store, you must manually configure the following:

- Data sources and multi data sources for the audit data source using WebLogic Server
- The JDBC string for the OPMN loader in the opmn.xml file

The following sections provide additional information specific to configuring auditing when an Oracle RAC database is used as the audit data store.

## A.1 Using WebLogic Server to Configure Audit Data Sources and Multi Data Sources

To set up the audit data source and multi data sources for an Oracle RAC database, follow the instructions in the "Managing the Audit Store" section of the *Oracle Fusion Middleware Application Security Guide*. Use the information in the "Set Up Audit Data Sources" section to set up the audit data sources and the information in the "Multiple Data Sources" section to configure an Oracle RAC database as the audit data store.

Use the information in the "Set Up Audit Data Sources" section to set up the audit data sources. To use an Oracle RAC database as the audit data store, you must create two individual data sources pointing to each individual Oracle RAC instance where the audit schemas are installed. The following settings are required:

- The connection URL should be in the following format:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=PROTOCOL=TCP) (HOST=host-vip)
(PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=dbservice) (INSTANCE_NAME=inst1))
```

Note that the service name and instance name are required, in addition to the host and port.

- The driver used is oracle.jdbc.OracleDriver
- The following property should be set:

```
<property>
<name>oracle.net.CONNECT_TIMEOUT</name>
<value>10000</value>
</property>
```

- The following settings are required for the individual data sources:
  - initial-capacity: **0**
  - connection-creation-retry-frequency-seconds: **10**
  - test-frequency-seconds: **300**
  - test-connections-on-reserve: **true**
  - test-table-name: **SQL SELECT 1 FROM DUAL**
  - seconds-to-trust-an-idle-pool-connection: **0**
  - global-transactions-protocol: **None**

Use the information in the "Multiple Data Sources" section to configure an Oracle RAC database as the audit data store. Create a multi data source with JNDI name jdbc/AuditDB. This multi data source should point to the individual data sources you created.

The following settings are required for the multi data source:

- test-frequency-seconds: **5**
- algorithm-type: **Load-Balancing**



- `data-source-list`: point to a list of comma separated child data sources ("**JDBC Data Source-0,JDBC Data Source-1**"). This list is the same set of data sources that you created for each individual node of the Oracle RAC database.

## A.2 Configuring the JDBC String for the Audit Loader

If you have an audit store configured, Oracle Process Manager and Notification Server (OPMN) manages several system components running in Oracle WebLogic Server. For these components, OPMN pushes the audit events to the database audit store.

The "Configure a Database Audit Store for System Components" section in the *Oracle Fusion Middleware Application Security Guide* describes how to set up the OPMN startup audit loader.

During the setup of the OPMN startup audit loader, you must modify the `rmd-definitions` element in the `opmn.xml` file. By default, the `rmd-definitions` element includes a JDBC string for a single instance database in this format:

```
jdbc:oracle:thin:@host:port:sid
```

When you are using an Oracle RAC database as the audit data store, you must use a JDBC string for an Oracle RAC database in the following format in the `rmd-definitions` element:

```
jdbc:oracle:thin@(DESCRIPTION=(ADDRESS_LIST=(LOAD_BALANCE=on) (ADDRESS=(PROTOCOL=tcp) (HOST=node1-vip) (PORT=1521)) (ADDRESS=(PROTOCOL=tcp) (HOST=node2-vip) (PORT=1521)) (CONNECT_DATA=SERVICE_NAME=service-name.mycompany.com)))
```

If you also need to configure the Oracle RAC database audit store for Java components, refer to the instructions in the "Configure a Database Audit Store for Java Components" section in the *Oracle Fusion Middleware Application Security Guide*.



## Recommended Multi Data Sources

The following are examples of multi data source configuration files. Multi pool DS JDBC Multi Data Source-0 is made up of two data sources, JDBC Data Source-0, and JDBC Data Source-1. The property settings in bold text are recommended settings:

### B.1 JDBC Multi Data Source-0

```
<?xml version='1.0' encoding='UTF-8'?>
<jdbc-data-source xmlns="http://www.bea.com/ns/weblogic/jdbc-data-source"
 xmlns:sec="http://www.bea.com/ns/weblogic/90/security"
 xmlns:wls="http://www.bea.com/ns/weblogic/90/security/wls"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="http://www.bea.com/ns/weblogic/jdbc-data-source
http://www.bea.com/ns/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd">
 <name>JDBC Multi Data Source-0</name>
 <jdbc-connection-pool-params>
 <test-frequency-seconds>5</test-frequency-seconds>
 </jdbc-connection-pool-params>
 <jdbc-data-source-params>
 <jndi-name>jdbc/OracleDS</jndi-name>
 <algorithm-type>Load-Balancing</algorithm-type>
 <data-source-list>JDBC Data Source-0,JDBC Data Source-1</data-source-list>
 <failover-request-if-busy>>false</failover-request-if-busy>
 </jdbc-data-source-params>
</jdbc-data-source>
```

### B.2 JDBC Data Source-0 (non-XA)

```
<?xml version='1.0' encoding='UTF-8'?>
<jdbc-data-source xmlns="http://www.bea.com/ns/weblogic/jdbc-data-source"
 xmlns:sec="http://www.bea.com/ns/weblogic/90/security"
 xmlns:wls="http://www.bea.com/ns/weblogic/90/security/wls"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="http://www.bea.com/ns/weblogic/jdbc-data-source
http://www.bea.com/ns/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd">
 <name>JDBC Data Source-0</name>
 <jdbc-driver-params>
 <url>jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=host-vip)
(PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=dbservice) (INSTANCE_NAME=inst1)))</url>
 <driver-name>oracle.jdbc.OracleDriver</driver-name>
 <properties>
 <property>
 <name>user</name>
 <value>username</value>
 </property>
 </properties>
 </jdbc-driver-params>
```

```

 <property>
 <name>oracle.net.CONNECT_TIMEOUT</name>
 <value>10000</value>
 </property>
 </properties>
 <password-encrypted>{3DES}LMWCj6T00uI=</password-encrypted>
</jdbc-driver-params>
<jdbc-connection-pool-params>
 <initial-capacity>0</initial-capacity>
 <max-capacity>20</max-capacity>
 <capacity-increment>1</capacity-increment>
 <shrink-frequency-seconds>900</shrink-frequency-seconds>
 <highest-num-waiters>2147483647</highest-num-waiters>

 <connection-creation-retry-frequency-seconds>10</connection-creation-retry-frequen
cy-seconds>
 <connection-reserve-timeout-seconds>10</connection-reserve-timeout-seconds>
 <test-frequency-seconds>300</test-frequency-seconds>
 <test-connections-on-reserve>true</test-connections-on-reserve>
 <ignore-in-use-connections-enabled>true</ignore-in-use-connections-enabled>
 <inactive-connection-timeout-seconds>0</inactive-connection-timeout-seconds>
 <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
 <login-delay-seconds>0</login-delay-seconds>
 <statement-cache-size>10</statement-cache-size>
 <statement-cache-type>LRU</statement-cache-type>
 <remove-infected-connections>true</remove-infected-connections>

 <seconds-to-trust-an-idle-pool-connection>0</seconds-to-trust-an-idle-pool-connect
ion>
 <statement-timeout>-1</statement-timeout>
 <pinned-to-thread>false</pinned-to-thread>
</jdbc-connection-pool-params>
<jdbc-data-source-params>
 <jndi-name>jdbc/OracleDS0</jndi-name>
 <global-transactions-protocol>None</global-transactions-protocol>
</jdbc-data-source-params>
</jdbc-data-source>

```

## B.3 JDBC Data Source-0 (XA)

```

<?xml version='1.0' encoding='UTF-8'?>
<jdbc-data-source xmlns="http://www.bea.com/ns/weblogic/jdbc-data-source"
 xmlns:sec="http://www.bea.com/ns/weblogic/90/security"
 xmlns:wls="http://www.bea.com/ns/weblogic/90/security/wls"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="http://www.bea.com/ns/weblogic/jdbc-data-source
http://www.bea.com/ns/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd">
 <name>JDBC Data Source-0</name>
 <jdbc-driver-params>
 <url>
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=host-vip)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=dbservice)(INSTANCE_NAME=inst1)))</url>
 <driver-name>oracle.jdbc.xa.client.OracleXADataSource</driver-name>
 </properties>
 <property>
 <name>user</name>
 <value>username</value>
 </property>
 <property>
 <name>oracle.net.CONNECT_TIMEOUT</name>

```

```

 <value>10000</value>
 </property>
</properties>
<password-encrypted>{3DES}LMWCj6T00uI=</password-encrypted>
<use-xa-data-source-interface>true</use-xa-data-source-interface>
</jdbc-driver-params>
<jdbc-connection-pool-params>
 <initial-capacity>0</initial-capacity>
<connection-creation-retry-frequency-seconds>10</connection-creation-retry-frequen
cy-seconds>
 <max-capacity>15</max-capacity>
 <capacity-increment>1</capacity-increment>
 <shrink-frequency-seconds>900</shrink-frequency-seconds>
 <highest-num-waiters>2147483647</highest-num-waiters>
 <connection-reserve-timeout-seconds>10</connection-reserve-timeout-seconds>
 <test-frequency-seconds>300</test-frequency-seconds>
 <test-connections-on-reserve>true</test-connections-on-reserve>
 <ignore-in-use-connections-enabled>true</ignore-in-use-connections-enabled>
 <inactive-connection-timeout-seconds>0</inactive-connection-timeout-seconds>
 <test-table-name>SQL SELECT 1 FROM DUAL</test-table-name>
 <login-delay-seconds>0</login-delay-seconds>
 <statement-cache-size>10</statement-cache-size>
 <statement-cache-type>LRU</statement-cache-type>
 <remove-infected-connections>true</remove-infected-connections>
<seconds-to-trust-an-idle-pool-connection>0</seconds-to-trust-an-idle-pool-connect
ion>
 <statement-timeout>-1</statement-timeout>
 <jdbc-xa-debug-level>10</jdbc-xa-debug-level>
 <pinned-to-thread>false</pinned-to-thread>
</jdbc-connection-pool-params>
<jdbc-data-source-params>
 <jndi-name>jdbc/OracleDS1</jndi-name>
 <global-transactions-protocol>TwoPhaseCommit</global-transactions-protocol>
</jdbc-data-source-params>
<jdbc-xa-params>
 <keep-xa-conn-till-tx-complete>true</keep-xa-conn-till-tx-complete>
 <need-tx-ctx-on-close>false</need-tx-ctx-on-close>
 <xa-end-only-once>false</xa-end-only-once>
 <keep-logical-conn-open-on-release>false</keep-logical-conn-open-on-release>
 <resource-health-monitoring>true</resource-health-monitoring>
 <recover-only-once>false</recover-only-once>
 <xa-set-transaction-timeout>false</xa-set-transaction-timeout>
 <xa-transaction-timeout>0</xa-transaction-timeout>
 <rollback-local-tx-upon-conn-close>false</rollback-local-tx-upon-conn-close>
 <xa-retry-duration-seconds>300</xa-retry-duration-seconds>
 <xa-retry-interval-seconds>60</xa-retry-interval-seconds>
</jdbc-xa-params>
</jdbc-data-source>

```

The only difference between JDBC Data Source-1 for XA and non-XA is they point to a different instance of Oracle RAC (host:port).



---

## Whole Server Migration for Windows

This appendix describes the procedure for enabling a virtual IP address for Windows XP systems.

### C.1 Using Windows Control Panel

This procedure assumes the system is Windows XP, configured for static IP addresses.

1. Go to **Control Panel**.
2. Select **Network and Internet Connections**.
3. Select **Network Connections**.
4. Select **Local Area Connection**.
5. Choose **Internet Protocol (TCP / IP)** and click the **Properties** button.
6. In the next window, click the **Advanced** button at the bottom of the dialog.
7. Click the **Add** button next to the displayed IP addresses to add an IP address with a corresponding sub-net mask. (this works only if the system is configured for static IPs)

### C.2 Using the netsh Command Line

The command line method is to use netsh. For example:

```
netsh interface ip show address <interface>
netsh interface ip add address <interface> <address> <netmask>
netsh interface ip delete address <interface> addr=<address>
```

The following commands are available:

Commands in this context:

**Table C-1 Application URLs**

| Application | URL                                          |
|-------------|----------------------------------------------|
| ?           | Displays a list of commands.                 |
| add         | Adds a configuration entry to a table        |
| delete      | Deletes a configuration entry from a table.  |
| dump        | Displays a configuration script.             |
| help        | Displays a list of commands.                 |
| ip          | Changes to the 'netsh interface ip' context. |

**Table C-1 (Cont.) Application URLs**

| <b>Application</b> | <b>URL</b>                                          |
|--------------------|-----------------------------------------------------|
| ipv6               | Changes to the `netsh interface ipv6' context.      |
| portproxy          | Changes to the `netsh interface portproxy' context. |
| reset              | Resets information.                                 |
| set                | Sets configuration information                      |
| show               | Displays information.                               |

The following sub-contexts are available:

`ip ipv6 portproxy`

To view Help for a command, type the command, followed by a space, and then ?.



---

## Component Workbooks

This appendix includes workbooks that you can use when performing the high availability configurations for the components documented in this manual.

It includes the following topics:

- [Section D.1, "Oracle SOA Suite Workbook"](#)
- [Section D.2, "Oracle Identity Management Workbook"](#)
- [Section D.3, "Oracle WebCenter Portal Workbook"](#)
- [Section D.4, "Oracle Portal, Forms, Reports, and Discoverer Workbook"](#)
- [Section D.5, "Oracle Enterprise Content Management Workbook"](#)
- [Section D.6, "Oracle Data Integrator Workbook"](#)
- [Section D.7, "Oracle Business Intelligence Platform Workbook"](#)

### D.1 Oracle SOA Suite Workbook

This section is a workbook for Oracle SOA Suite.

Using the information in [Section 5.13, "Configuring High Availability for Oracle SOA Service Infrastructure and Component Service Engines"](#) and [Section 5.15, "Configuring High Availability for Oracle BAM"](#) as an example, you can fill out the equivalent names that you plan to use in your environment when configuring high availability for the Oracle SOA Suite.

#### D.1.1 Workbook Tables for Oracle SOA Suite

Use [Table D–1](#) for application URLs.

**Table D–1 Application URLs**

| Application                                    | URL |
|------------------------------------------------|-----|
| SOA-INFRA Application                          |     |
| Oracle Business Process<br>Web Application URL |     |
| SOA Worklist Application                       |     |
| WSM-PM Validation<br>Application               |     |
| UMS Preferences<br>Application                 |     |

**Table D–1 (Cont.) Application URLs**

| Application        | URL |
|--------------------|-----|
| BPM Composer       |     |
| BPM Workspace      |     |
| B2B UI Application |     |
| BAM Application    |     |
| Admin Console      |     |
| Web Logic Console  |     |
| FMW Control        |     |

Use [Table D–2](#) for Virtual IP Addresses.

**Table D–2 Virtual IP Addresses**

| Purpose               | IP Address | DNS Name |
|-----------------------|------------|----------|
| Administration Server |            |          |
| WLS_SOA1              |            |          |
| WLS_SOA2              |            |          |
| WLS_BAM1              |            |          |
| WLS_BAM2              |            |          |

Use [Table D–3](#) for file locations - Application Tier.

**Table D–3 File Locations - Application Tier**

| Environment Variable or Artifact | Location | Shared |
|----------------------------------|----------|--------|
| ORACLE_BASE                      |          | No     |
| MW_HOME                          |          | Yes    |
| ORACLE_HOME                      |          | Yes    |
| WL_HOME                          |          | Yes    |
| DOMAIN_HOME                      |          | No     |
| TLOG Directory                   |          | Yes    |
| JMS Persistence Store            |          | Yes    |

Use [Table D–4](#) for file locations - Web Tier.

**Table D–4 File Locations - Web Tier**

| Environment Variable or Artifact | Location | Shared |
|----------------------------------|----------|--------|
| ORACLE_BASE                      |          | No     |
| MW_HOME                          |          | Yes    |
| ORACLE_HOME                      |          | Yes    |
| ORACLE_INSTANCE                  |          | No     |

Use [Table D-5](#) for Identity Management Details.

**Table D-5 Identity Management Details**

| Identity Management Artifact | Value |
|------------------------------|-------|
| LDAP Host Name               |       |
| LDAP Host                    |       |
| LDAP SSL Port                |       |
| LDAP user DN                 |       |
| LDAP Password                |       |
| Access Manager Host          |       |
| Access Manager Port          |       |
| Application Profile Password |       |

Use [Table D-6](#) for Database Information

**Table D-6 Database Information - Metadata Repository**

| SOA Database Artifact                     | Value |
|-------------------------------------------|-------|
| Database Hosts (VIPs if using Oracle RAC) |       |
| Listener Port                             |       |
| Database Service Name                     |       |
| RCU Prefix                                |       |
| Main Schema User/Password                 |       |
| Auxiliary Schema User/Password            |       |
| MDS User/Password                         |       |
| User Messaging User/Password              |       |
| BAM User/Password                         |       |
| SOA Infrastructure User/Password          |       |
| SYS Password                              |       |

Use [Table D-7](#) for Load Balancing Configuration

**Table D-7 Load Balancer Configuration**

| Purpose             | Virtual Name | Port | Externally Visible | SSL | Destination Hosts | Destination Ports | SSL |
|---------------------|--------------|------|--------------------|-----|-------------------|-------------------|-----|
| External Site       |              |      |                    |     |                   |                   |     |
| Internal Site       |              |      |                    |     |                   |                   |     |
| Administration Site |              |      |                    |     |                   |                   |     |

Use [Table D-8](#) for Port Information

**Table D–8 Port Information**

| Component                      | Host(s)       | Port |
|--------------------------------|---------------|------|
| SSL Port on the Load Balancer  | Load Balancer |      |
| HTTP Port on Load Balancer     | Load Balancer |      |
| OHS HTTP Listening             |               |      |
| OHS Admin                      |               |      |
| OPMN                           |               |      |
| WebLogic Administration Server |               |      |
| Single Sign-on Listening       |               |      |
| WLS_SOA1 listen address        |               |      |
| WLS_SOA2 listen address        |               |      |
| WLS_BAM1 listen address        |               |      |
| WLS_BAM2 listen address        |               |      |
| Node Manager SOAHOST1          |               |      |
| Node Manager SOAHOST2          |               |      |
| Node Manager BAMHOST1          |               |      |
| Node Manager BAMHOST2          |               |      |

## D.2 Oracle Identity Management Workbook

This section is a workbook for Oracle Identity Management.

Using the information in [Chapter 8, "Configuring High Availability for Identity Management Components"](#) an example, you can fill out the equivalent names that you plan to use in your environment when configuring high availability for the Oracle Identity Management components.

### D.2.1 Workbook Tables for Oracle Identity Management

Use the following tables to record the names you plan to use in your Oracle Identity Management high availability configuration.

Enter the application URLs for your configuration in [Table D–9](#).

**Table D–9 Application URLs**

| Application                                         | URL |
|-----------------------------------------------------|-----|
| Oracle WebLogic Administration Console              |     |
| Oracle Enterprise Manager Fusion Middleware Control |     |
| Oracle Access Manager Console                       |     |

Enter the virtual IP address for your configuration in [Table D–10](#).

**Table D–10 Virtual IP Addresses**

| Purpose                                  | IP Address | DNS Name |
|------------------------------------------|------------|----------|
| Oracle WebLogic<br>Administration Server |            |          |

Enter the following generic file locations for your configuration in [Table D–11](#).

**Table D–11 Generic File Locations**

| Type        | Location | Shared |
|-------------|----------|--------|
| ORACLE_BASE |          | No     |
| MW_HOME     |          | No     |

Enter the file locations for IDMHOST $n$  for your configuration in [Table D–12](#).

**Table D–12 File Locations for IDMHOST $n$** 

| Environment Artifact                     | Directory Location | Shared |
|------------------------------------------|--------------------|--------|
| ORACLE_HOME                              |                    | No     |
| Administration Server<br>ORACLE_INSTANCE |                    | No     |
| WLS_ODS1 ORACLE_<br>INSTANCE             |                    | No     |
| WLS_ODS2 ORACLE_<br>INSTANCE             |                    | No     |
| WL_HOME                                  |                    | No     |
| DOMAIN_HOME                              |                    | No     |

Enter the file locations for OIDHOST $n$  for your configuration in [Table D–13](#).

**Table D–13 File Locations for OIDHOST $n$** 

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| ORACLE_HOME          |                    | No     |
| ORACLE_INSTANCE      |                    | No     |

Enter the file locations for OVDHOST $n$  for your configuration in [Table D–14](#).

**Table D–14 File Locations for OVDHOST $n$** 

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| ORACLE_HOME          |                    | No     |
| ORACLE_INSTANCE      |                    | No     |

Enter the file locations for OAMHOST $n$  for your configuration in [Table D–15](#).

**Table D–15 File Locations for OAMHOST $n$** 

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| IDM_ORACLE_HOME      |                    | No     |

**Table D–15 (Cont.) File Locations for OAMHOST $n$** 

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| IAM_ORACLE_HOME      |                    | No     |
| SOA_ORACLE_HOME      |                    | No     |
| ORACLE_INSTANCE      |                    | No     |

Enter the file locations for OIMHOST $n$  for your configuration in [Table D–16](#).

**Table D–16 File Locations for OIMHOST $n$** 

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| IAM_ORACLE_HOME      |                    | No     |
| ORACLE_HOME          |                    | No     |
| ORACLE_INSTANCE      |                    | No     |

Enter the file locations for OAAMHOST $n$  for your configuration in [Table D–17](#).

**Table D–17 File Locations for OAAMHOST $n$** 

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| IDM_ORACLE_HOME      |                    | No     |
| IAM_ORACLE_HOME      |                    | No     |
| SOA_ORACLE_HOME      |                    | No     |
| ORACLE_INSTANCE      |                    | No     |

Enter the file locations for OIFHOST $n$  for your configuration in [Table D–18](#).

**Table D–18 File Locations for OIFHOST $n$** 

| Environment Artifact                     | Directory Location | Shared |
|------------------------------------------|--------------------|--------|
| ORACLE_HOME                              |                    | No     |
| Administration Server<br>ORACLE_INSTANCE |                    | No     |
| WLS_OIF1 ORACLE_<br>INSTANCE             |                    | No     |
| WLS_OIF2 ORACLE_<br>INSTANCE             |                    | No     |
| WL_HOME                                  |                    | No     |
| DOMAIN_HOME                              |                    | No     |

Enter the file locations for the web tier for your configuration in [Table D–19](#).

**Table D–19 File Locations for the Web Tier**

| Environment Artifact | Directory Location | Shared |
|----------------------|--------------------|--------|
| ORACLE_HOME          |                    | No     |
| ORACLE_INSTANCE      |                    | No     |

Enter Oracle Identity Management details for your configuration in [Table D-20](#).

**Table D-20 Identity Management Artifacts**

| Identity Management Artifact             | Value |
|------------------------------------------|-------|
| Single Sign-On URL                       |       |
| Oracle Internet Directory Host Name      |       |
| Oracle Internet Directory Non-SSL Port   |       |
| Oracle Internet Directory SSL Enabled    |       |
| Oracle Internet Directory SSL Port       |       |
| Oracle Internet Directory Security Realm |       |
| Oracle Virtual Directory Host Name       |       |
| Oracle Virtual Directory Port            |       |
| Oracle Virtual Directory SSL Enabled     |       |
| Oracle Virtual Directory SSL Port        |       |

Enter Authentication LDAP Artifacts details for the Oracle Identity Federation configuration in [Table D-21](#).

**Table D-21 Authentication LDAP Artifacts for Oracle Identity Federation**

| Authentication LDAP Artifact | Value |
|------------------------------|-------|
| LDAP Type                    |       |
| LDAP URL                     |       |
| LDAP Bind DN                 |       |
| LDAP Bind DN Password        |       |
| User Credential ID Attribute |       |
| User Unique ID Attribute     |       |
| Person Object Class          |       |
| Base DN                      |       |

Enter the Oracle Identity Federation Artifacts when using an LDAP User Data Store in your configuration in [Table D-22](#).

**Table D-22 LDAP User Data Store Artifacts for Oracle identity Federation**

| LDAP User Data Store Artifact | Value |
|-------------------------------|-------|
| LDAP Type                     |       |

**Table D–22 (Cont.) LDAP User Data Store Artifacts for Oracle identity Federation**

| <b>LDAP User Data Store Artifact</b> | <b>Value</b> |
|--------------------------------------|--------------|
| LDAP URL                             |              |
| LDAP Bind DN                         |              |
| LDAP Bind DN Password                |              |
| User Description Attribute           |              |
| User ID Attribute                    |              |
| Person Object Class                  |              |
| Base DN                              |              |

Enter the Oracle Identity Federation Artifacts when using an LDAP Federation Data Store in your configuration in [Table D–23](#).

**Table D–23 LDAP Federation Data Store Artifacts for Oracle Identity Federation**

| <b>Federation Data Store Artifact</b> | <b>Value</b> |
|---------------------------------------|--------------|
| LDAP Type                             |              |
| LDAP URL                              |              |
| LDAP Bind DN                          |              |
| LDAP Bind DN Password                 |              |
| User Federation Record Context        |              |
| LDAP Container Object Class           |              |

Enter the Oracle Identity Federation Artifacts when using an RDBMS User Data Store in your configuration in [Table D–24](#).

**Table D–24 RDBMS User Data Store Artifacts for Oracle Identity Federation**

| <b>RDBMS User Data Store Artifact</b> | <b>Value</b> |
|---------------------------------------|--------------|
| Hostname                              |              |
| Username                              |              |
| Password                              |              |
| Login Table                           |              |
| User ID Attribute                     |              |
| User Description Attribute            |              |

Enter the Oracle Identity Federation Artifacts when using an RDBMS Federation Data Store in your configuration in [Table D–25](#).



**Table D–25 RDBMS Federation Data Store Artifacts for Oracle Identity Federation**

| <b>RDBMS Federation Data Store Artifacts</b> | <b>Value</b> |
|----------------------------------------------|--------------|
| Hostname                                     |              |
| Username                                     |              |
| Password                                     |              |

Enter RDBMS Transient Data Store Artifacts for the Oracle Identity Federation configuration in [Table D–26](#).

**Table D–26 RDBMS Transient Data Store Artifacts for Oracle Identity Federation**

| <b>Transient Data Store Artifact</b> | <b>Value</b> |
|--------------------------------------|--------------|
| Hostname                             |              |
| Username                             |              |
| Password                             |              |

Enter database information for the metadata repository for your configuration in [Table D–27](#).

**Table D–27 Database Information for the Metadata Repository**

| <b>Database Details</b>                   | <b>Value</b> |
|-------------------------------------------|--------------|
| Database Hosts (VIPs if using Oracle RAC) |              |
| Listener Port                             |              |
| Database Service Name                     |              |
| RCU Prefix                                |              |
| Main Schema Name/Password                 |              |
| Auxiliary Schema Name/Password            |              |
| SYS Password                              |              |

Enter load balancer configuration information for your configuration in [Table D–28](#).

**Table D–28 Load Balancer Configuration**

| <b>Purpose</b>            | <b>Virtual Name</b> | <b>Port</b> | <b>Externally Visible</b> | <b>SSL</b> | <b>Destination Hosts</b> | <b>Destination Ports</b> | <b>SSL</b> |
|---------------------------|---------------------|-------------|---------------------------|------------|--------------------------|--------------------------|------------|
| Oracle Internet Directory |                     |             |                           |            |                          |                          |            |
| Oracle Virtual Directory  |                     |             |                           |            |                          |                          |            |
| Administration Server     |                     |             |                           |            |                          |                          |            |
| Oracle Identity Manager   |                     |             |                           |            |                          |                          |            |
| Single Sign-On            |                     |             |                           |            |                          |                          |            |

Enter port information for your configuration in [Table D–29](#).

**Table D–29 Port Information**

| Component                                              | Host(s) | Port |
|--------------------------------------------------------|---------|------|
| Oracle Internet Directory                              |         |      |
| Oracle Internet Directory (SSL)                        |         |      |
| Oracle Virtual Directory                               |         |      |
| Oracle Virtual Directory (SSL)                         |         |      |
| WebLogic Server Console                                |         |      |
| Oracle Enterprise Manager Fusion<br>Middleware Control |         |      |
| Oracle Directory Services Manager                      |         |      |
| Oracle Access Manager Server                           |         |      |
| Oracle Identity Manager Server                         |         |      |
| Oracle Identity Manager                                |         |      |
| Oracle SOA                                             |         |      |
| Oracle Adaptive Access Manager<br>Server               |         |      |
| Oracle Adaptive Access Manager<br>Admin Server         |         |      |
| Oracle HTTP Server                                     |         |      |
| Oracle HTTP Server (SSL)                               |         |      |
| Oracle HTTP Server Admin                               |         |      |
| OPMN                                                   |         |      |
| Node Manager                                           |         |      |

## D.3 Oracle WebCenter Portal Workbook

This section is a workbook for [Section 6–2, "WebCenter Portal's Managed Servers and Applications."](#)

Using the information in [Section 6.4, "Configuring High Availability for WebCenter Portal"](#) as an example, you can fill out the equivalent names that you plan to use in your environment when configuring high availability for the Oracle WebCenter Portal components.

### D.3.1 Workbook Tables for Oracle WebCenter Portal

Use the following tables to record the names you plan to use in your Oracle WebCenter Portal high availability configuration.

Use [Table D–30](#) for application URLs.

**Table D–30 Application URLs**

| Application      | URL |
|------------------|-----|
| WebCenter Spaces |     |
| Application      |     |

**Table D–30 (Cont.) Application URLs**

| Application                    | URL |
|--------------------------------|-----|
| Portal Tools Application       |     |
| WSM Policy Manager Application |     |
| WebCenter Discussions Forum    |     |
| WebCenter Analytics            |     |
| WebCenter Activity Graph       |     |
| WebCenter Personalization      |     |
| Admin Console                  |     |
| WebLogic Console               |     |
| FMW Control                    |     |

Use [Table D–31](#) for Virtual IP Addresses.

**Table D–31 Virtual IP Addresses**

| Purpose               | IP Address | DNS Name |
|-----------------------|------------|----------|
| Administration Server |            |          |

Use [Table D–32](#) for File Locations - App Tier.

**Table D–32 File Locations - App Tier**

| Type              | Location | Shared |
|-------------------|----------|--------|
| ORACLE_BASE       |          | No     |
| MW_HOME           |          | Yes    |
| ORACLE_HOME       |          | Yes    |
| WL_HOME           |          | Yes    |
| DOMAIN_HOME       |          | No     |
| ADMIN_DOMAIN_HOME |          | Yes    |

Use [Table D–33](#) for File Locations - Web Tier.

**Table D–33 File Locations - Web Tier**

| Application       | URL | Shared |
|-------------------|-----|--------|
| ORACLE_BASE       |     | No     |
| MW_HOME           |     | Yes    |
| ORACLE_HOME       |     | Yes    |
| WL_HOME           |     | Yes    |
| DOMAIN_HOME       |     | No     |
| ADMIN_DOMAIN_HOME |     | Yes    |

Use [Table D–34](#) for Identity Management Details.

**Table D–34 Identity Management Details**


---

LDAP Host Name

LDAP Port

LDAP SSL Enabled

LDAP user DN

LDAP Password

Access Manager Host

Access Manager Port

Application Profile Password

---

Use [Table D–35](#) for Database Information - Metadata Repository.

**Table D–35 Database Information - Metadata Repository**


---

Database Hosts (VIPs if using Oracle RAC)

Listener Port

Database Service Name

RCU Prefix

Main Schema User/Password

Auxiliary Schema User/Password

MDS User/Password

WebCenter Spaces User/Password

Portlets User/Password

Discussion Forum User/Password

Content Server User/Password

SYS Password

---

Use [Table D–36](#) for Load Balancer Configuration

**Table D–36 Load Balancer Configuration**

| Purpose                 | Virtual Name | Port | Externally Visible | SSL | Destination Hosts | Destination Ports |
|-------------------------|--------------|------|--------------------|-----|-------------------|-------------------|
| WC Application          |              |      |                    |     |                   |                   |
| Admin                   |              |      |                    |     |                   |                   |
| WC Internal Application |              |      |                    |     |                   |                   |

---

Use [Table D–37](#) for Port Information

**Table D–37 Port Information**

| Comment                        | Host(s)       | Port |
|--------------------------------|---------------|------|
| SSL Port on the Load Balancer  | Load Balancer |      |
| HTTP Port on Load Balancer     | Load Balancer |      |
| OHS HTTP Listening             |               |      |
| OHS Admin                      |               |      |
| WebLogic Administration Server |               |      |
| OPMN                           |               |      |
| Single Sign on Listening       |               |      |
| WLS_WSM1                       |               |      |
| WLS_WSM2                       |               |      |
| WC_SPACES1                     |               |      |
| WC_SPACES2                     |               |      |
| WC_PORTLETS1                   |               |      |
| WC_PORTLETS2                   |               |      |
| WC_COLLABORATION1              |               |      |
| WC_COLLABORATION2              |               |      |
| WC_UTILITIES1                  |               |      |
| WC_UTILITIES2                  |               |      |
| Node Manager                   |               |      |

## D.4 Oracle Portal, Forms, Reports, and Discoverer Workbook

This section is a workbook for [Section 14.6, "Configuring Oracle Portal, Forms, Reports, and Discoverer for High Availability."](#)

Using the information in [Section 14.6, "Configuring Oracle Portal, Forms, Reports, and Discoverer for High Availability"](#) as an example, you can fill out the equivalent names that you plan to use in your environment when configuring high availability for the Oracle Portal, Forms, Reports, and Discoverer.

### D.4.1 Workbook Tables for Oracle Portal, Forms, Reports, and Discoverer

Use [Table D–38](#) for application URLs.

**Table D–38 Application URLs**

| Application       | URL |
|-------------------|-----|
| Portal            |     |
| Forms             |     |
| Reports           |     |
| Discoverer        |     |
| Web Logic Console |     |
| FMW Control       |     |

Use [Table D-39](#) for Virtual IP Addresses.

**Table D-39 Virtual IP Addresses**

| Purpose               | IP Address | DNS Name |
|-----------------------|------------|----------|
| Administration Server |            |          |

Use [Table D-40](#) for file locations.

**Table D-40 File Locations**

| Type            | Location | Shared |
|-----------------|----------|--------|
| ORACLE_BASE     |          | No     |
| MW_HOME         |          | No     |
| ORACLE_HOME     |          | No     |
| ORACLE_INSTANCE |          | No     |
| WL_HOME         |          | No     |
| DOMAIN_HOME     |          | No     |
| Reports Cache   |          | Yes    |

Use [Table D-41](#) for Identity Management Details.

**Table D-41 Database Information - Metadata Repository**

|                    |
|--------------------|
| Single Sign-On URL |
| OID Host Name      |
| OID Port           |
| OID SSL Enabled    |

Use [Table D-42](#) for Database Information.

**Table D-42 Database Information - Metadata Repository**

|                                           |
|-------------------------------------------|
| Database Hosts (VIPS if using Oracle RAC) |
| Listener Port                             |
| Database Service Name                     |
| RCU Prefix                                |
| Portal User Name/Password                 |
| Portlet User Name/Password                |
| Discoverer User Name/Password             |
| SYS Password                              |

Use [Table D-43](#) for Load Balancing Configuration.

**Table D–43 Load Balancer Configuration**

| <b>Purpose</b>          | <b>Virtual Name</b> | <b>Port</b> | <b>Externally Visible</b> | <b>SSL</b> | <b>Destination Hosts</b> | <b>Destination Ports</b> |
|-------------------------|---------------------|-------------|---------------------------|------------|--------------------------|--------------------------|
| Main Site               |                     |             |                           |            |                          |                          |
| Internal Site           |                     |             |                           |            |                          |                          |
| Web Cache Invalidations |                     |             |                           |            |                          |                          |

Use [Table D–44](#) for Port Information.

**Table D–44 Port Information**

| <b>Comment</b>                 | <b>Host(s)</b> | <b>Port</b> |
|--------------------------------|----------------|-------------|
| SSL Port on the Load Balancer  | Load Balancer  |             |
| HTTP Port on Load Balancer     | Load Balancer  |             |
| Web Cache HTTP                 |                |             |
| Web Cache Invalidation         |                |             |
| Web Cache Administration       |                |             |
| Web Cache Statistics           |                |             |
| OHS HTTP Listening             |                |             |
| OHS HTTPS Listening            |                |             |
| OHS Admin                      |                |             |
| WebLogic Administration Server |                |             |
| OPMN                           |                |             |
| Single Sign on Listening       |                |             |
| WebLogic Managed Server        |                |             |
| WebLogic Managed Server        |                |             |
| OID HTTP                       |                |             |
| WLS_PORTAL                     |                |             |
| WLS_PORTAL1                    |                |             |
| WLS_FORMS                      |                |             |
| WLS_FORMS1                     |                |             |
| WLS_REPORTS                    |                |             |
| WLS_REPORTS1                   |                |             |
| WLS_DISCO                      |                |             |
| WLS_DISCO1                     |                |             |
| Node Manager                   |                |             |

## D.5 Oracle Enterprise Content Management Workbook

This section is a workbook for [Section 11.3, "Oracle ECM High Availability Configuration Steps."](#)

Using the information in [Section 11.3, "Oracle ECM High Availability Configuration Steps"](#) as an example, you can fill out the equivalent names that you plan to use in your environment when configuring high availability for the Oracle WebCenter Content.

## D.5.1 Workbook Tables for Oracle WebCenter Content

Use [Table D-45](#) for application URLs.

**Table D-45 Application URLs**

| Application               | URL |
|---------------------------|-----|
| UCM Application           |     |
| I/PM Application          |     |
| URM Application           |     |
| IBR Location              |     |
| WebLogic Console          |     |
| Fusion Middleware Control |     |

Use [Table D-46](#) for Virtual IP Addresses.

**Table D-46 Virtual IP Addresses**

| Purpose               | IP Address | DNS Name |
|-----------------------|------------|----------|
| Administration Server |            |          |
| WLS_IPM1              |            |          |
| WLS_IPM2              |            |          |

Use [Table D-47](#) for File Locations in the application tier.

**Table D-47 File Locations - Application Tier**

| Environment Variable or Artifact | Location | Shared |
|----------------------------------|----------|--------|
| ORACLE_BASE                      |          | No     |
| MW_HOME                          |          | Yes    |
| ORACLE_HOME                      |          | Yes    |
| WL_HOME                          |          | Yes    |
| DOMAIN_HOME                      |          | No     |
| UCM Config Directory             |          | Yes    |
| URM Config Directory             |          | Yes    |
| I/PM Config Directory            |          | Yes    |
| IBR Config Directory             |          | No     |

Use [Table D-48](#) for File Locations in the web tier.



**Table D–48 File Locations - Web Tier**

| Environment Variable or Artifact | Location | Shared |
|----------------------------------|----------|--------|
| ORACLE_BASE                      |          | No     |
| MW_HOME                          |          | Yes    |
| ORACLE_HOME                      |          | Yes    |

Use [Table D–49](#) for database information.

**Table D–49 Database Information - Metadata Repository**

| ECM Database Artifact                     | Value |
|-------------------------------------------|-------|
| Database Hosts (VIPS if using Oracle RAC) |       |
| Listener Port                             |       |
| Database Service Name                     |       |
| RCU Prefix                                |       |
| UCM Schema User/Password                  |       |
| IPM Schema User/Password                  |       |
| URM User/Password                         |       |

Use [Table D–50](#) for load balancing configuration.

**Table D–50 Load Balancer Configuration**

| Purpose             | Virtual Name | Port | Externally Visible | SSL | Destination Hosts | Destination Ports | SSL |
|---------------------|--------------|------|--------------------|-----|-------------------|-------------------|-----|
| External Site       |              |      |                    |     |                   |                   |     |
| Internal Site       |              |      |                    |     |                   |                   |     |
| Administration Site |              |      |                    |     |                   |                   |     |

Use [Table D–51](#) for port information.

**Table D–51 Port Information**

| Component                         | Host(s)       | Port |
|-----------------------------------|---------------|------|
| SSL Port on the Load Balancer     | Load Balancer |      |
| HTTP Port on the Load Balancer    | Load Balancer |      |
| Oracle HTTP Server HTTP Listening |               |      |
| Oracle HTTP Server Admin          |               |      |
| WebLogic Administration Server    |               |      |
| WLS_IPM1 listen address           |               |      |
| WLS_IPM2 listen address           |               |      |
| Node Manager ECMHOST1             |               |      |
| Node Manager ECMHOST2             |               |      |
| Node Manager IBRHOST1             |               |      |

**Table D–51 (Cont.) Port Information**

| Component             | Host(s) | Port |
|-----------------------|---------|------|
| Node Manager IBRHOST2 |         |      |

## D.6 Oracle Data Integrator Workbook

This section is a workbook for [Section 7.4, "Configuring High Availability for Oracle Data Integrator."](#)

Using the information in [Section 7.4, "Configuring High Availability for Oracle Data Integrator"](#) as an example, you can fill out the equivalent names that you plan to use in your environment when configuring high availability for Oracle Data Integrator.

### D.6.1 Workbook Tables for Oracle Data Integrator

Use [Table D–52](#) for application URLs.

**Table D–52 Application URLs**

| Application                                         | URL |
|-----------------------------------------------------|-----|
| Oracle Data Integrator Agent                        |     |
| Oracle Data Integrator Explorer                     |     |
| Oracle Data Integrator Web Services                 |     |
| WebLogic Console                                    |     |
| Oracle Enterprise Manager Fusion Middleware Control |     |

Use [Table D–53](#) for Virtual IP Addresses.

**Table D–53 Virtual IP Addresses**

| Purpose               | IP Address | DNS Name |
|-----------------------|------------|----------|
| Administration Server |            |          |

Use [Table D–54](#) for file locations in the application tier.

**Table D–54 File Locations - Application Tier**

| Environment Variable or Artifact | Location | Shared |
|----------------------------------|----------|--------|
| ORACLE_BASE                      |          | No     |
| MW_HOME                          |          | Yes    |
| ORACLE_HOME                      |          | Yes    |
| WL_HOME                          |          | Yes    |
| DOMAIN_HOME                      |          | No     |

Use [Table D–55](#) for file locations in the web tier.

**Table D–55 File Locations - Web Tier**

| Environment Variable or Artifact | Location | Shared |
|----------------------------------|----------|--------|
| ORACLE_BASE                      |          | No     |
| MW_HOME                          |          | Yes    |
| ORACLE_HOME                      |          | Yes    |

Use [Table D–56](#) for database information.

**Table D–56 Database Information - Metadata Repository**

| Oracle Data Integrator Database Artifact  | Value |
|-------------------------------------------|-------|
| Database Hosts (VIPs if using Oracle RAC) |       |
| Listener Ports                            |       |
| Database Service Name                     |       |
| RCU Prefix                                |       |
| ODI Master Schema User/Password           |       |
| ODI Work Schema User/Password             |       |

Use [Table D–57](#) for load balancing configuration.

**Table D–57 Load Balancer Configuration**

| Purpose             | Virtual Name | Port | Externally Visible | SSL | Destination Hosts | Destination Ports | SSL |
|---------------------|--------------|------|--------------------|-----|-------------------|-------------------|-----|
| External Site       |              |      |                    |     |                   |                   |     |
| Internal Site       |              |      |                    |     |                   |                   |     |
| Administration Site |              |      |                    |     |                   |                   |     |

Use [Table D–58](#) for port information.

**Table D–58 Port Information**

| Component                         | Host(s)       | Port |
|-----------------------------------|---------------|------|
| SSL Port on the Load Balancer     | Load Balancer |      |
| HTTP Port on the Load Balancer    | Load Balancer |      |
| Oracle HTTP Server HTTP Listening |               |      |
| Oracle HTTP Server Admin          |               |      |
| WebLogic Administrator Server     |               |      |
| odi_server1 listen address        |               |      |
| odi_server2 listen address        |               |      |
| Node Manager APPHOST2             |               |      |
| Node Manager APPHOST2             |               |      |

## D.7 Oracle Business Intelligence Platform Workbook

This section is a workbook for Oracle BI Platform.

Use the information in [Section 15.1.12, "Oracle BI EE High Availability Configuration Steps"](#) as an example; you can fill out the equivalent names that you plan to use in your environment when configuring high availability for the Oracle BI EE components.

### D.7.1 Workbook Tables for Oracle Business Intelligence Platform

Use the following tables to record the names you plan to use in your BI Platform high availability configuration.

Use [Table D-59](#) for application URLs.

**Table D-59 Application URLs**

| Application                     | URL |
|---------------------------------|-----|
| BI EE Application               |     |
| BI Publisher Application        |     |
| Real-Time Decisions Application |     |
| WMS-PM Validation Application   |     |
| Administration Console          |     |
| WebLogic Console                |     |
| Fusion Middleware Control       |     |

Use [Table D-60](#) for Virtual IP Addresses.

**Table D-60 Virtual IP Addresses**

| Purpose               | IP Address | DNS Name |
|-----------------------|------------|----------|
| Administration Server |            |          |
| BI_SERVER1            |            |          |
| BI_SERVER2            |            |          |

Use [Table D-61](#) for file locations - Application Tier.

**Table D-61 File Locations - Application Tier**

| Environment Variable or Artifact      | Location | Shared |
|---------------------------------------|----------|--------|
| ORACLE_BASE                           |          | No     |
| MW_HOME                               |          | Yes    |
| ORACLE_HOME                           |          | Yes    |
| WL_HOME                               |          | Yes    |
| WLS Administration Server DOMAIN_HOME |          | Yes    |
| BI_SERVER1 DOMAIN_HOME                |          | No     |

**Table D–61 (Cont.) File Locations - Application Tier**

| <b>Environment Variable or Artifact</b>   | <b>Location</b> | <b>Shared</b> |
|-------------------------------------------|-----------------|---------------|
| BI_SERVER2 DOMAIN_HOME                    |                 | No            |
| TLOG Directory                            |                 | Yes           |
| JMS Persistence Store                     |                 | Yes           |
| WLS Administration Server ORACLE_INSTANCE |                 | No            |
| BI_SERVER1 ORACLE_INSTANCE                |                 | No            |
| BI_SERVER2 ORACLE_INSTANCE                |                 | No            |
| Global Cache                              |                 | Yes           |
| BI Presentation Catalog                   |                 | Yes           |
| BI Repository Publishing Directory        |                 | Yes           |
| BI Publisher Catalog                      |                 | Yes           |
| BI Publisher Configuration                |                 | Yes           |
| BI Publisher Temp Directory               |                 | Yes           |

Use [Table D–62](#) for file locations - Web Tier.

**Table D–62 File Locations - Web Tier**

| <b>Environment Variable or Artifact</b> | <b>Location</b> | <b>Shared</b> |
|-----------------------------------------|-----------------|---------------|
| ORACLE_BASE                             |                 | No            |
| MW_HOME                                 |                 | Yes           |
| ORACLE_HOME                             |                 | Yes           |
| ORACLE_INSTANCE                         |                 | No            |

Use [Table D–63](#) for database information.

**Table D–63 Database Information**

| <b>BI Database Artifact</b>               | <b>Value</b> |
|-------------------------------------------|--------------|
| Database Hosts (VIPs if using Oracle RAC) |              |
| Listener Port                             |              |
| Database Service Name                     |              |
| RCU Prefix                                |              |
| BIPLATFORM Schema User/Password           |              |
| MDS User/Password                         |              |
| SYS Password                              |              |

Use [Table D–64](#) for Load Balancing Configuration.

**Table D–64 Load Balancing Configuration**

| <b>Purpose</b>      | <b>Virtual Name Port</b> | <b>Externally Visible SSL</b> | <b>Destination Hosts</b> | <b>Destination Ports</b> | <b>SSL</b> |
|---------------------|--------------------------|-------------------------------|--------------------------|--------------------------|------------|
| External Site       |                          |                               |                          |                          |            |
| Internal Site       |                          |                               |                          |                          |            |
| Administration Site |                          |                               |                          |                          |            |

Use [Table D–65](#) for Port Information.

**Table D–65 Port Information**

| <b>Component</b>                  | <b>Host(s)</b> | <b>Port</b> |
|-----------------------------------|----------------|-------------|
| SSL Port on the Load Balancer     | Load Balancer  |             |
| HTTP Port on the Load Balancer    | Load Balancer  |             |
| Oracle HTTP Server HTTP Listening |                |             |
| Oracle HTTP Server Admin OPMN     |                |             |
| WebLogic Administrator Server     |                |             |
| BI_SERVER1 listen address         |                |             |
| BI_SERVER2 listen address         |                |             |
| Node Manager APPHOST2             |                |             |
| Node Manager APPHOST2             |                |             |

---

## ascrsctl Online Help

The detailed usage of ascrsctl can be obtained by following the instructions generated from the command `ascrsctl help`. This appendix provides the full content of the help pages to serve as a complete offline reference.

### E.1 start

#### TOPIC

**start** - start an ASCRS resource

#### COMMAND

**ascrsctl start -name** <string> [**-type** <string>] [**-node** <string>]

#### DESCRIPTION

This ASCRS command is used to start a resource already created with ASCRS.

Mandatory arguments:

**-name, -n**

This argument specifies the resource name.

Optional arguments:

**-type, -t**

This argument specifies the resource type if the resource name is in short form. It is not needed if the name is in canonical form.

**-node**

This argument specifies the cluster node for starting the resource. If it is not specified, the node will be chosen by CRS based on the placement policy for this resource.

#### EXAMPLE(S)

**ascrsctl start -n mydisk -t disk**

**ascrsctl start -n ora.myvip.cfcvip -node hostA.mycompany.com**

## E.2 stop

### TOPIC

**stop** - stop an ASCRS resource

### COMMAND

```
ascrsctl stop -name <string> [-type <string>] [-force] [-noprompt]
```

### DESCRIPTION

This ASCRS command is used to stop a resource created with ASCRS.

Mandatory arguments:

**-name, -n**

This argument specifies the resource name.

Non-mandatory arguments:

**-type, -t**

This argument specifies the resource type if the resource name is in short form. It is not needed if the name is in canonical form.

**-force, -f**

This argument shuts down the named resource and takes it offline from CRS management. This option guarantees to take offline the CRS monitoring of the resource, but does not guarantee to shut down the resource if it is already in an unmanageable state.

**-noprompt, -np**

When this argument is specified, the user is not prompted for confirmation, and the resource and its dependents are taken offline.

### EXAMPLE(S)

```
ascrsctl stop -n mydisk -t disk
ascrsctl stop -n ora.myvip.cfcvip -f -np
```

## E.3 status

### TOPIC

**status** - check the status of ASCRS resources

### COMMAND

```
ascrsctl status [-name <string>] [-type <string>] [-long]
```

### DESCRIPTION

This ASCRS command is used to check the status of one or all resources created with ASCRS. The status of a resource includes its current running state, its basic CRS profile information, and its relationship to the other ASCRS resources.



**-name, -n**

This argument specifies the resource name. If not specified, check all resources.

**-type, -t**

This argument specifies the type of resources to be checked. It is not needed if the name is in canonical form.

**-long, -l**

When this argument is specified, the status information is displayed in a detailed format.

**EXAMPLE(S)**

```
ascrsctl status
ascrsctl status -name ora.mydisk.cfcdisk
ascrsctl status -l
```

## E.4 switch

**TOPIC**

**switch** - switchover an ASCRS resource to another cluster node

**COMMAND**

```
ascrsctl switch -name <string> [-type <string>]
 [-node <string>] [-noprompt]
```

**DESCRIPTION**

This ASCRS command is used to switch over an ASCRS resource that is currently in online state to another node in the cluster. All resources this resource depends upon are also switched over.

Mandatory arguments:

**-name, -n**

This argument specifies the resource name.

Non-mandatory arguments:

**-type, -t**

This argument specifies the resource type if the resource name is in short form. It is not needed if the name is in canonical form.

**-node**

This argument specifies the target cluster node of this resource. If it is not specified, the target node will be chosen by CRS based on the placement policy for this resource.

**-noprompt, -np**

When this argument is specified, the user is not prompted for confirmation.

**EXAMPLE(S)**

```
ascrsctl switch -n mydisk -t disk hostB.mycompany.com
ascrsctl switch -n ora.myvip.cfcvip -np
```

## E.5 delete

### TOPIC

**delete** - delete an ASCRS resource

### COMMAND

**ascrsctl delete -name** <string> [**-type** <string>] [**-noprompt**]

### DESCRIPTION

This ASCRS command is used to delete a resource created with ASCRS. Once a resource is successfully deleted, the resource is no longer managed by CRS.

An ASCRS resource cannot be deleted if there are still one or more other resources depending on it, or if it is not in offline state.

Mandatory argument:

**-name, -n**

This argument specifies the resource name.

Non-mandatory arguments:

**-type, -t**

This argument specifies the resource type if the resource name is in short form. It is not needed if the name is in canonical form.

**-noprompt, -np**

When specified, the user is not prompted for confirmation.

### EXAMPLE(S)

```
ascrsctl delete -n mydisk -np
ascrsctl delete -n ora.myvip.cfcvip
```

## E.6 create/disk

### TOPIC

**create/disk** - create **disk** ASCRS resource

### COMMAND

**ascrsctl create -name** <string> **-type disk -path** <string>  
**-mountCommand** <string> **-umountCommand** <string>  
[options]

## DESCRIPTION

This ASCRS command is used to create (or register) a shared disk resource in CRS. To successfully create a disk resource, a signature file needs to be created on the root of the shared disk. See the *Oracle Fusion Middleware Administrator's Guide* for details.

These are mandatory arguments:

**-name, -n**

This argument specifies the resource name to be created.

**-type, -t**

This argument specifies the resource type. Its value must be **disk**.

**path**

This argument specifies the mount point of the shared disk.

**-mountCommand, -mc**

This argument specifies a platform-specific command to be invoked when mounting the shared disk. Command "nop" takes no action. On Windows, if the shared disk is NTFS, use the command "diskmgr online <disk number>", where <disk number> can be identified in the Microsoft Disk Manager window.

**-umountCommand, -umc**

This argument specifies a platform-specific command to be invoked when unmounting the shared disk. Command "nop" takes no action. On Windows, if the shared disk is NTFS, use the command "diskmgr offline <disk number>", where <disk number> can be identified in the Microsoft Disk Manager window.

These are non-mandatory options:

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. If it is not specified, all the nodes are included.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties of this resource, for example, as=1,rt=400.

The property names and their value ranges are shown in [Table E-1](#), where, except for as, st, and ra, all the other numbers are in seconds. These properties can be configured through the ASCRS configuration file. If a property is neither configured nor specified with this option, the default is assumed.

**Table E-1 Resource Values for Create Commands**

| Parameter | Min | Max  | Default | Purpose           |
|-----------|-----|------|---------|-------------------|
| as        | 0   | 1    | 1       | Auto Start        |
| ci        | 5   | 6000 | 600     | Check Interval    |
| fd        | 5   | 600  | 50      | Failover Delay    |
| fi        | 5   | 6000 | 50      | Failure Interval  |
| ft        | 0   | 20   | 5       | Failure Threshold |
| ra        | 0   | 20   | 3       | Restart Attempts  |

**Table E-1 (Cont.) Resource Values for Create Commands**

| Parameter | Min | Max  | Default | Purpose         |
|-----------|-----|------|---------|-----------------|
| st        | 20  | 3600 | 30      | Script Timeouts |
| rt        | 20  | 3600 | 30      | Start Timeout   |
| pt        | 20  | 3600 | 30      | Stop Timeout    |

**-policy**

This argument specifies what resource parameter values are used. These policies and values are available in the ASCRS configuration.

The valid values are **normal** or **fast**. If it not specified, **normal** is assumed. However, the **-resourceParams** values always take precedence over the policy.

**EXAMPLE(S)**

UNIX:

```
ascrsctl create -n dbhome -t disk -path /cfcd1
-mc "/bin/mount /dev/sda1 /cfcd1" -p fd=30
ascrsctl create -n dbhome -t disk -path /cfcd1
-mc "/bin/mount /dev/sda1 /cfcd1"
-umc "/bin/umount /dev/sda1"
```

Windows:

```
ascrsctl create -n asdisk -t disk -path c:\oracle\asdisk
-mc "diskmgr online 2" -mc "diskmgr offline 2"
-p fd=30
```

## E.7 update/disk

**TOPIC**

**update/disk** - update **disk** ASCRS resource

**COMMAND**

```
ascrsctl update -name <string> [-type disk] [-path <string>]
[-mountCommand <string>] [-umountCommand <string>]
[options]
```

**DESCRIPTION**

This ASCRS command is used to update a **disk** resource created with ASCRS.

Mandatory argument:

**-name, -n**

This argument specifies the resource name to be updated. If it is a fully qualified name, the **-type** option can be omitted.

These are non-mandatory options:

**-type, -t**

This argument specifies the resource type. Its value must be **disk**.

**-path**

This argument specifies the mount point of the shared disk.

**-mountCommand, -mc**

This argument specifies a platform-specific fully qualified command or script name to be executed for mounting the shared disk. On Windows, if the shared disk is NTFS, use the command "diskmgr online <disk number>", where <disk number> can be identified in the Microsoft Disk Manager window.

**-umountCommand, -umc**

This argument specifies a platform-specific fully qualified command or script name to be executed for unmounting the shared disk. On Windows, if the shared disk is NTFS, use the command "diskmgr offline <disk number>", where <disk number> can be identified in the Microsoft Disk Manager window.

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. The special value **default** includes all nodes.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1, rt=400**.

The property names and their value ranges are shown in [Table E-2](#), where, except for **as**, **st**, and **ra**, all the other numbers are in seconds.

**Table E-2 Resource Values for Update Commands**

| Parameter | Min | Max  | Purpose           |
|-----------|-----|------|-------------------|
| as        | 0   | 1    | Auto Start        |
| ci        | 5   | 6000 | Check Interval    |
| fd        | 5   | 600  | Failover Delay    |
| fi        | 5   | 6000 | Failure Interval  |
| ft        | 0   | 20   | Failure Threshold |
| ra        | 0   | 20   | Restart Attempts  |
| st        | 20  | 3600 | Script Timeout    |
| rt        | 20  | 3600 | Start Timeout     |
| pt        | 20  | 3600 | Stop Timeout      |

**EXAMPLE(S)**

UNIX:

```
ascsctl update -n mydisk -t disk -umfc "/bin/umount -l /sharedisk"
```

Windows:

```
ascsctl update -n mydisk -t disk -mc "diskmgr online 1"
 -umc "diskmgr offline 1"
```

## E.8 create/vip

### TOPIC

**create/vip** - create **vip** ASCRS resource

### COMMAND

```
ascrsctl create -name <string> -type vip -ipAddr <ip> -netmask <string>
 -interface <string> [options]
```

### DESCRIPTION

This ASCRS command is used to create (or register) a virtual IP resource in CRS

These are mandatory arguments:

**-name, -n**

This argument specifies the resource name to be created.

**-type, -t**

This argument specifies the resource type. Its value must be **vip**.

**-ipAddr, -ip**

This argument specifies the IP address of the virtual IP or its hostname.

**-netmask, -nm**

This argument specifies the network mask for the above virtual IP.

**-interface, -if**

This argument specifies the network interface(s) on which the IP should be enabled.

On UNIX, the value can be one or more interface names such as `eth0` or `"eth0|eth1"`.

On Windows, the value can be one or more network connection names, such as `"Public network1|Public network2"`.

These are non-mandatory options:

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. If it is not specified, all the nodes are included.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, `as=1,rt=400`.

The property names and their value ranges are listed in [Table E-1](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds. These properties can be configured through the ASCRS configuration file. If a property is neither configured nor specified with this option, the default is assumed.

**-policy**

This argument specifies what resource parameter values are used. These policies and values are available in the ASCRS configuration file.

The valid values are **normal** or **fast**. If it not specified, **normal** is assumed. However, the **-resourceParams** values always take precedence over the policy.

### EXAMPLE(S)

UNIX:

```
ascrsctl create -n myvip -t vip -ip 192.168.1.10 -nm 255.255.255.0
 -if eth1 -p ci=5
```

Windows:

```
ascrsctl create -n myvip -t vip -ip 192.168.1.10 -nm 255.255.255.0
 -if "Public network" -p ci=5
```

## E.9 update/vip

### TOPIC

**update/vip** - update **vip** ASCRS resource

### COMMAND

```
ascrsctl update -name <string> [-type vip] [-ipAddr <string>
 [-netmask <string>] [-interface <string>] [options]
```

### DESCRIPTION

This ASCRS command is used to update a vip resource created with ASCRS.

Mandatory argument:

**-name, -n**

This argument specifies the resource name to be updated. If it is a fully qualified name, the **-type** option can be omitted.

These are non-mandatory options:

**-type, -t**

This argument specifies the resource type. Its value must be **vip**.

**-ipAddr, -ip**

This argument specifies the IP address of the virtual IP or its hostname.

**-netmask, -nm**

This argument specifies the network mask for the above virtual IP.

**-interface, -if**

This argument specifies the network interface(s) on which the IP should be enabled.

On UNIX, the value can be one or more interface names such as eth0 or "eth0 | eth1".

On Windows, the value can be one or more network connection names, such as "Public network1 | Public network2".

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. This value is a space-separated or comma-separated list of node names in the cluster. The special value **default** includes all the nodes.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-2](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds.

**EXAMPLE(S)**

UNIX:

```
ascrsctl update -n ora.myvip.cfcvip -ip 192.168.1.10
ascrsctl update -n ora.myvip.cfcvip -if eth1 -p ci=3
```

Windows:

```
ascrsctl update -n ora.myvip.cfcvip -if Public -p ci=3
```

## E.10 create/dblsnr

**TOPIC**

**create/dblsnr** - create **dblsnr** ASCRS resource

**COMMAND**

```
ascrsctl create -name <string> -type dblsnr -listenerName <string>
 -listenerOracleHome <string>
 -vip <string> -disk <string>
 [-tnsAdmin <string>] [options]
```

**DESCRIPTION**

This ASCRS command is used to create (or register) an Oracle database listener resource in CRS.

These are mandatory arguments:

**-name, -n**

This argument specifies the resource name to be created.

**-type, -t**

This argument specifies the resource type. Its value must be **dblsnr**.

**-listenerName, -ln**

This argument specifies the database listener name.

**-listenerOracleHome, -lsnroh, -loh**

This argument specifies the Oracle Home of the database that owns this listener.

**-vip**

This argument specifies the vip resource this listener runs on.

**-disk**

This argument specifies the disk resource the Oracle Home resides on.

The other non-mandatory options:



**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. This value is a space-separated or comma-separated list of node names in the cluster. The special value **default** includes all the nodes. If it is not specified, all the nodes are included.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-1](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds. These properties can be configured through the ASCRS configuration file. If a property is neither configured nor specified with this option, the default is assumed.

**-policy**

This argument specifies what resource parameter values are used. These policies and values are available in the ASCRS configuration file.

The valid values are **normal** or **fast**. If it not specified, **normal** is assumed. However, the **-resourceParams** values always take precedence over the policy.

**-tnsAdmin, -ta**

This argument specifies the location of the listener configuration if it is not in the default location within the Oracle Home.

**EXAMPLE(S)**

UNIX:

```
ascrsctl create -name mydblsnr -type dblsnr -listenerName orcl
 -listenerOracleHome /cfcdbl
 -vip 192.168.1.10 -disk ohdisk
```

Windows:

```
ascrsctl create -name mydblsnr -type dblsnr -listenerName orcl
 -listenerOracleHome c:\oraasshare\cfcdbl
 -vip myvip -disk ohdisk
```

## E.11 update/dblsnr

**TOPIC**

**update/dblsnr** - update **dblsnr** ASCRS resource

**COMMAND**

```
ascrsctl update -name <string> [-type dblsnr]
 [-listenerName <string>]
 [-listenerOracleHome <string>]
 [-vip <string>] [-disk <string>]
 [-tnsAdmin <string>] [options]
```

**DESCRIPTION**

This ASCRS command is used to update a **dblsnr** resource created with ASCRS.

Mandatory argument:

**-name, -n**

This argument specifies the resource name to be updated. If it is a fully qualified name, the **-type** option can be omitted.

These are non-mandatory options:

**-type, -t**

This argument specifies the resource type. Its value must be **dblsnr**.

**-listenerName, -ln**

This argument specifies the database listener name.

**-listenerOracleHome, -lsnroh, -loh**

This argument specifies the Oracle Home of the database that owns this listener.

**-vip**

This argument specifies the vip resource this listener runs on.

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. The special value **default** includes all the nodes.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-2](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds.

**-tnsAdmin, -ta**

This argument specifies the new location of the listener configuration file. The special value "nil" sets this location to its default.

**-disk**

This argument specifies the disk resource the Oracle Home resides on.

### EXAMPLE(S)

```
ascrsctl update -n mydblsnr -t dblsnr -vip newvip
ascrsctl update -n mydblsnr -t dblsnr -disk newdisk -p st=30,pt=40,rt=40
```

## E.12 create/db

### TOPIC

**create/db** - create **db** ASCRS resource

### COMMAND

For database instance:

```
ascrsctl create -name <string> -type db -oraHome <string>
 -oraSID <string> -disk <string> [<string> ...]
 -lsnr <string> [-pfile <string>]
```

```
[-componentID dbinstance] [options]
```

For database console:

```
ascrsctl create -name <string> -type db -oraHome <string>
-oraSID <string> -disk <string> -vip <string>
[-componentID dbconsole] [options]
```

For job scheduler (Windows only):

```
ascrsctl create -name <string> -type db -oraHome <string>
-oraSID <string> -disk <string>
[-componentID jobscheduler] [options]
```

For VSS writer (Windows only):

```
ascrsctl create -name <string> -type db -oraHome <string>
-oraSID <string> -disk <string>
[-componentID vsswriter] [options]
```

## DESCRIPTION

This ASCRS command is used to create (or register) an Oracle database resource in CRS. Depending on the specified component ID, the database resource represents either a core database instance, a dbconsole service, or, if the platform is Windows, an Oracle job scheduler service or an Oracle Volume Shadow Service (VSS).

These are mandatory arguments for all database resources:

**-name, -n**

This argument specifies the resource name to be created.

**-type, -t**

This argument specifies the resource type. Its value must be **db**.

**-oraHome, -oh**

This argument specifies the database Oracle Home location.

**-oraSID, -sid**

This argument specifies the Oracle SID name.

**-disk**

This argument specifies the shared disks hosting the Oracle Home and data files. For a database instance component, the value for this parameter is a space-separated or comma-separated list of ASCRS disk resource names; for other components, its value is the disk resource hosting the Oracle home.

The other non-mandatory options:

**-componentID, -c**

This argument identifies the database component to be managed. It assumes **dbinstance**, **dbconsole**, **jobscheduler** (Windows only), or **vsswriter** (Windows only). If this is not specified, **dbinstance** is assumed.

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. If it is not specified, all the nodes are included.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-1](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds. These properties can be configured through the ASCRS configuration file. If a property is neither configured nor specified with this option, the default is assumed.

**-policy**

This argument specifies what resource parameter values are used. These policies and values are available in the ASCRS configuration file.

The valid values are **normal** or **fast**. If it not specified, **normal** is assumed. However, the **-resourceParams** values always take precedence over the policy.

Component-specific options:

For database instance component:

**-lsnr**

This argument specifies the listener resource used by this database. It is mandatory.

**-pfile, -pf**

This argument specifies the database pfile for starting the database. It is only valid for the dbinstance component.

For database console component:

**-vip**

This argument specifies the virtual IP resource used by this database console resource. It is mandatory.

**EXAMPLE(S)**

UNIX:

```
ascrsctl create -n mydb -t db -oh /cfedb1 -sid orcl
 -disk ohdisk datafiledisk -lsnr mydblsnr
```

Windows:

```
ascrsctl create -n mydb -t db -oh c:\oraasshare\cfedb1 -sid orcl
 -disk ohdisk datafiledisk -lsnr mydblsnr
```

## E.13 update/db

**TOPIC**

**update/db** - update **db** ASCRS resource

**COMMAND**

```
ascrsctl update -name <string> -type db [-oraHome <string>]
 [-oraSID <string>] [-disk <string> [<string> ...]]
 [-lsnr <string>] [-pfile <string>]
 [-vip <string>] [options]
```

**DESCRIPTION**

This ASCRS command is used to update an ASCRS **db** resource registered in CRS.

Mandatory argument:

**-name, -n**

This argument specifies the resource name to be updated.

Non-mandatory arguments:

**-type, -t**

This argument specifies the resource type. Its value must be **db**. If the resource name is in canonical form, the **-type** option can be omitted.

**-oraHome, -oh**

This argument specifies the database Oracle Home location.

**-oraSID, -sid**

This argument specifies the Oracle SID name

**-disk**

For a database instance component, the value for this parameter is a space-specified or comma-specified list of ASCRS disk resource names; for other components, its value is the disk resource hosting the Oracle home.

**lsnr**

This argument specifies the listener resource. It is only applicable for the dbinstance component.

**-vip**

This argument specifies the virtual IP resource used by a database console, so it is applicable for a database console resource only.

The other non-mandatory options:

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. The special value **default** includes all the nodes.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-2](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds.

**-pfile, -pf**

This argument specifies the database pfile for starting the database. The special value "nil" sets this file to its default. It is only valid for the dbinstance component.

**EXAMPLE(S)**

```
ascrsctl update -n ora.mydb.cfcdB -disk ohdisk -lsnr newlsnr
 -p st=60,rt=60,pt=60
```

## E.14 create/as

### TOPIC

**create/as** - create **as** ASCRS resource

### COMMAND

```
ascrsctl create -name <string> -type as -componentHome <string>
 [-componentIDs <string> [<string> ...]]
 -vip <string> -disk <string> [<string> ...]
 [-db <string> [<string> ...]]
 [-as <string> [<string> ...]]
 [options]
```

### DESCRIPTION

This ASCRS command is used to create (or register) a middleware resource in CRS.

These are mandatory arguments:

#### **-name, -n**

This argument specifies the resource name to be created.

#### **-type, -t**

This argument specifies the resource type. Its value must be **as**.

#### **-componentHome, -ch**

This argument specifies the location of the WebLogic domain that contains the targeted Administration Server, managed servers, or the OPMN managed servers.

#### **-vip**

This argument specifies the virtual IP resource the middleware servers depend upon.

#### **-disk**

This argument specifies the shared disks hosting the WebLogic software, the WebLogic server domain directory, or, if it is for OPMN managed components, the instance home, and other shared disks this resource directly depends upon. If a shared disk is used for multiple purposes, it needs to be specified only once. The value for this parameter is a space-separated or comma-separated list of ASCRS disk resource names.

Non-mandatory options:

#### **-componentIDs, -ci**

This argument specifies the names of the servers to be managed in the WebLogic domain or the OPMN instance. If this argument is missing or if the value is **default**, all the servers are included.

#### **db**

This argument specifies the database resources this **as** resource directly depends upon.

#### **as**

This argument specifies the **as** resources this resource directly depends upon.

#### **-m**

This argument specifies the health monitors that the WebLogic servers should use. If this option is missing or its value is **default**, the TCP ping monitor is used for all the managed servers. To use user-defined monitors, this option should be followed by a list of monitor assignments such as '**-m** AdminServer=mon1 wlsapp=mon2', where AdminServer and wlsapp are valid server names, while mon1 and mon2 are valid monitor names defined in *CRS\_HOME/ascrs/config/mconfig.xml*.

**-clusterNodes, -nodes, -cn**

This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. If it is not specified, all the nodes are included.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-1](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds. These properties can be configured through the ASCRS configuration file. If a property is neither configured nor specified with this option, the default is assumed.

**-policy**

This argument specifies what resource parameter values are used. These policies and values are available in the ASCRS configuration file.

The valid values are **normal** or **fast**. If it is not specified, **normal** is assumed. However, the **-resourceParams** values always take precedence over the policy.

## EXAMPLE(S)

UNIX:

```
ascrsctl create -n idm.weblogic -t as
 -ch /sharedisk/fmw/user_projects/domains/IDMDomain
 -vip idmvip -disk idmdisk
```

```
ascrsctl create -n idm.weblogic -t as
 -ch /sharedisk/fmw/user_projects/domains/IDMDomain
 -vip idmvip -disk idmdisk -p st=800,rt=800,pt=800
```

```
ascrsctl create -n idm.opmn -t as
 -ch /sharedisk/fmw/asinst_1
 -vip idmvip -disk idmdisk -p st=800,rt=800,pt=800
```

Windows:

```
ascrsctl create -n adminserver -t as
 -ch c:\oracle\asdisk\fmw\user_projects\domains\adminserverDomain
 -vip adminservervip -disk adminserverdisk
```

## E.15 update/as

### TOPIC

**update/as** - update **as** ASCRS resource

### COMMAND

```
ascrsctl update -name <string> [-type as] [-componentHome <string>]
```

```
[-componentIDs <string> [<string> ...]
[-vip <string>] [-disk <string> [<string> ...]
[-db <string> [<string> ...]
[-as <string> [<string> ...]
[options]
```

## DESCRIPTION

This ASCRS command is used to update an **as** resource created with ASCRS. On Windows, OPMN resources are not supported.

Mandatory argument:

**-name, -n**

This argument specifies the resource name to be updated.

Non-mandatory arguments:

**-type, -t**

This argument specifies the resource type. Its value must be **as**. If the resource name is in canonical form, the **-type** option can be omitted.

**-componentHome, -ch**

This argument specifies the location of the WebLogic domain or OPMN instance home.

**-componentIDs, -ci**

This argument specifies the names of the servers managed in the WebLogic domain or the OPMN instance. Value 'default' means all the server names are included.

**-vip**

This argument specifies the virtual IP resource this component depends upon.

**-disk**

This argument specifies the shared **disk** resources this **as** resource directly depends upon. The value is a space-separated or comma-separated list of disk resource names.

**-db**

This argument specifies the database resources this **as** resource directly depends upon. Value 'nil' removes all these dependencies.

**-as**

This argument specifies the **as** resources this resource directly depends upon. Value 'nil' removes all these dependencies.

**-m**

This argument specifies the health monitors that WebLogic servers should use. If its value is **default**, the TCP ping monitor is used for all the managed servers. To fine-tune the monitor assignments, this option should be followed by a list of monitor assignments such as '**-m** AdminServer=mon1 wlsapp=mon2', where AdminServer and wlsapp are valid server names, while mon1 and mon2 are either valid monitor names defined in CRS\_HOME/ascrs/config/mconfig.xml or 'default'.

**-clusterNodes, -nodes, -cn**



This argument specifies the valid nodes of the cluster that can host this resource. The value is a space-separated or comma-separated list of node names in the cluster. The special value **default** includes all the nodes.

**-resourceParams, -params, -p**

This argument specifies space-separated or comma-separated "name=value" pairs to set the CRS properties for this resource, for example, **as=1,rt=400**.

The property names and their value ranges are listed in [Table E-2](#), where, except for **as**, **st** and **ra**, all the other numbers are in seconds.

**EXAMPLE(S)**

```
ascrsctl update -n myas -t as -vip newvip -disk instdisk wldisk
ascrsctl update -n ora.myas.cfcds -p st=800,rt=800,pt=800
```



---

# Configuring the Distributed Java Object Cache

The Java Object Cache (JOC) should be configured among all the servers running Oracle Web Services Manager and among all the servers running the Oracle WebCenter Spaces application. This local cache is provided to increase the performance of Oracle WSM and the Spaces application.

## F.1 Configuring the Java Object Cache

Configure the distributed Java Object Cache using the `MW_HOME/oracle_common/bin/configure-joc.py` script, a Python script that you can use to configure JOC in the managed servers. The script runs in WLST online mode and expects the Administration Server to be up and running.

When configuring JOC ports for Oracle products, Oracle recommends using ports in the 9988 to 9998 range.

---

**Note:** After configuring the Java Object Cache using the `wlst` commands or `configure-joc.py` script, you should restart all affected managed servers for the configurations to take effect.

---

### Usage

This section demonstrates how to configure the distributed Java Object Cache for Oracle Web Services Manager or for the Spaces application. For other applications such as ADF or WebCenter Custom Apps, substitute the relevant host, port, and cluster names where appropriate.

1. Connect to the Administration Server using the command-line Oracle WebLogic Scripting Tool (WLST).

```
MW_HOME/oracle_common/common/bin/wlst.sh
$ connect()
```

Enter the Oracle WebLogic Administration user name and password.

2. After connecting to the Administration Server using `wlst`, start the script using the `execfile` command, for example:

```
wls:/mydomain/serverConfig>execfile('MW_HOME/oracle_
common/bin/configure-joc.py')
```

3. Configure JOC for all the managed servers for a given cluster.

Enter 'y' when the script prompts whether you want to specify a cluster name, and also specify the cluster name and discover port, when prompted. This discovers all the managed servers for the given cluster and configures the JOC. The discover port is common for the entire JOC configuration across the cluster.

For Oracle Web Services Manager:

```
Do you want to specify a cluster name (y/n) <y>
Enter Cluster Name : WSM_Cluster
Enter Discover Port : 9991
```

For the Spaces application:

```
Do you want to specify a cluster name (y/n) <y>
Enter Cluster Name : Spaces_Cluster
Enter Discover Port : 9988
```

Here is a walkthrough for using `configure-joc.py` for HA environments:

For Oracle Web Services Manager:

```
execfile('MW_HOME/oracle_common/bin/configure-joc.py')
.
Enter Hostnames (eg host1,host2) : SOAHOST1, SOAHOST2
.
Do you want to specify a cluster name (y/n) <y>y
.
Enter Cluster Name : WSM_Cluster
.
Enter Discover Port : 9991
.
Enter Distribute Mode (true|false) <true> : true
.
Do you want to exclude any server(s) from JOC configuration (y/n) <n> n
```

For the Spaces application:

```
execfile('MW_HOME/oracle_common/bin/configure-joc.py')
.
Enter Hostnames (eg host1,host2) : APPHOST1, APPHOST2
.
Do you want to specify a cluster name (y/n) <y>y
.
Enter Cluster Name : Spaces_Cluster
.
Enter Discover Port : 9988
.
Enter Distribute Mode (true|false) <true> : true
.
Do you want to exclude any server(s) from JOC configuration (y/n) <n> n
```

The script can also be used to perform the following JOC configurations:

- Configure JOC for all specified managed servers.

Enter 'n' when the script prompts whether you want to specify a cluster name and also specify the managed server and discover port, when prompted.

For Oracle Web Services Manager:

```
Do you want to specify a cluster name (y/n) <y>n
Enter Managed Server and Discover Port (eg WLS_WSM1:9991, WLS_WSM1:9991) : WLS_
WSM1:9991,WLS_WSM2:9991
```

For the Spaces application:

```
Do you want to specify a cluster name (y/n) <y>n
Enter Managed Server and Discover Port (eg WC_SPACES1:9988, WC_SPACES2:9988) :
WC_SPACES1:9988,WC_SPACES2:9988
```

- Exclude JOC configuration for some managed servers.

The script allows you to specify the list of managed servers for which the JOC configuration "DistributeMode" will be set to 'false'. Enter 'y' when the script prompts whether you want to exclude any servers from JOC configuration, and enter the managed server names to be excluded, when prompted.

For Oracle Web Services Manager:

```
Do you want to exclude any server(s) from JOC configuration (y/n) <n>y
Exclude Managed Server List (eg Server1,Server2) : WLS_WSM1,WLS_WSM3
```

For the Spaces application:

```
Do you want to exclude any server(s) from JOC configuration (y/n) <n>y
Exclude Managed Server List (eg Server1,Server2) : WC_SPACES1,WC_SPACES3
```

- Disable the distribution mode for all managed servers.

The script allows you to disable the distribution to all the managed servers for a specified cluster. Specify 'false' when the script prompts for the distribution mode. By default, the distribution mode is set to 'true'.

You can also configure the Java Object Cache using the HA Power Tools tab in the Oracle WebLogic Administration Console; see [Chapter 16, "Using HA Power Tools."](#)

Verify the Java Object Cache configuration using the CacheWatcher utility. See [Section F.2, "Running CacheWatcher."](#)

## F.2 Running CacheWatcher

You can use the CacheWatcher utility to verify your Java Object Cache distributed configuration. CacheWatcher itself becomes a member of the Cache cluster.

For this example, the managed servers in the cluster are WLS\_Server1 and WLS\_Server2:

- WLS\_Server1 is running on APPHOST1.
- Start CacheWatcher on APPHOST2. WLS\_Server2 should be down.

To run CacheWatcher:

1. For ORACLE\_HOME, use the location of the Oracle Web Center installation. For example, MW\_HOME/Oracle\_WC1.
2. For DOMAIN\_HOME, use the full path to the location of your domain, under MW\_HOME/user\_projects/domains.
3. Server\_Name refers to a managed server name. For example, WLS\_Server2. Cachewatcher can use the same javacache.xml file that the Managed Server would use.
4. On APPHOST2, execute the following:

```
"java -classpath ORACLE_HOME/modules/oracle.javacache_
11.1.1/cache.jar:ORACLE_HOME/modules/oracle.odl_11.1.1/ojdl.jar
oracle.ias.cache.CacheUtil
watch -config=DOMAIN_HOME/config/fmwconfig/servers/<Server_Name>/javacache.xml
```

5. At the CacheWatcher's prompt, type `lc` and then press **Enter**.

A group view appears.

The group view indicates the number of cache members. The following is an example of a CacheWatcher run:

```
java -classpath MW_HOME/modules/oracle.javacache_11.1.1/
cache.jar:MW_HOME/modules/oracle.odl_11.1.1/ojdl.jar
oracle.ias.cache.CacheUtil watch -config=DOMAIN_HOME
/config/fmwconfig/servers/WLS_Server_Server/javacache.xml "
cache> lc
VID: 2
Size: 2
Column: 0123456789
CL: JJ
GRP0: 11
...
Member Table:
#1. J57161:145.87.9.15:86AAC2E:11B28D57BE4:-7FFE, 0
#2. J35578:145.87.9.14:-9E92850:11B28D5F7E0:-7FFF, 1
```

After starting CacheWatcher on each host, find two members from the view.

---

---

**Note:** CacheWatcher itself is considered a distributed cache instance. To terminate the CacheWatcher shell, type `exit` and press **Enter**.

---

---

---

## Configuring Distributed Notifications for MDS

The MDS database repository has the ability to use Java Object Cache for caching MetadataObjects and their contents. Each node using MDS has its own Java Object Cache instance and will cache MetadataObjects as they are read from the repository by that node, invalidating them from the local cache when local updates to documents comprising that MetadataObject occur.

MDS does not use Distributed Java Object Cache for replicating MetadataObjects. MDS will detect changes made to the repository by other nodes whenever a local change is flushed to that repository and invalidate the relevant local cache entries accordingly. In addition, each cluster member polls the repository every 30 seconds (this is the default interval, which can be configured) to learn if any documents have changed and will invalidate local cache entries in the same way.

However, the MDS database repository also has the capability of sending cluster-wide notifications when a document in that repository is changed by a cluster member. When this feature is enabled, all the cluster members are notified of changes to documents when the change occurs, instead of waiting for a local commit to be performed or for the repository poll to be run. Therefore, in a high availability environment, it is recommended for performance reasons to enable distributed notifications. Configuring the distributed Java Object Cache on every cluster member enables distributed notifications. When a cluster member receives notification of a change to a document, it invalidates its current local version of that document and uses the updated document instead.

For this feature to work, you must configure Java Object Cache on each node, ensure that the JNDI name that each node is using to refer to the repository is the same on every node, and polling must be enabled (it is enabled by default, but it must not be disabled).

For instructions on configuring the distributed Java Object Cache on the nodes in a cluster, see [Section F.1, "Configuring the Java Object Cache."](#)

