

Oracle® Fusion Middleware

Disaster Recovery Guide

11g Release 1 (11.1.1)

E15250-04

November 2011

Oracle Fusion Middleware Disaster Recovery Guide, 11g Release 1 (11.1.1)

E15250-04

Copyright © 2009, 2011, Oracle and/or its affiliates. All rights reserved.

Primary Author: Showvik Roychowdhuri

Contributors: Bharath Reddy, Dinu Teofilovici, Fermin Castro, Kevin Clugage, Pradeep Bhat, Shailesh Dwivedi, Susan Kornberg

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	ix
Audience	ix
Documentation Accessibility	ix
Related Documents	ix
Conventions	x
1 Disaster Recovery Introduction	
1.1 Disaster Recovery Overview	1-1
1.1.1 Problem Description and Common Solutions	1-1
1.1.2 Terminology	1-2
1.2 Disaster Recovery for Oracle Fusion Middleware Components	1-4
1.2.1 Oracle Fusion Middleware Disaster Recovery Architecture Overview	1-4
1.2.2 Components Described in this Document	1-8
2 Recommendations for Fusion Middleware Components	
2.1 Recommendations for Oracle WebLogic Server	2-1
2.1.1 Recommendations for Oracle WebLogic Server JMS and T-Logs	2-2
2.1.2 Recommendations for Oracle Platform Security Services	2-4
2.2 Recommendations for Oracle ADF	2-5
2.3 Recommendations for Oracle WebCenter Portal	2-6
2.3.1 Recommendations for Oracle WebCenter Portal: Spaces	2-7
2.3.2 Recommendations for Oracle WebCenter Portal's Portlet Producers	2-8
2.3.3 Recommendations for Oracle WebCenter Portal's Discussion Server	2-8
2.3.4 Recommendations for Oracle WebCenter Content Server	2-9
2.3.5 Recommendations for Oracle WebCenter Portal's Analytics Collector	2-9
2.3.6 Recommendations for Oracle WebCenter Portal Activity Graph Engines	2-10
2.4 Recommendations for Oracle SOA Suite	2-10
2.4.1 Recommendations for Oracle SOA Service Infrastructure	2-12
2.4.2 Recommendations for Oracle BPEL Process Manager	2-13
2.4.3 Recommendations for Oracle Mediator	2-14
2.4.4 Recommendations for Oracle Human Workflow	2-15
2.4.5 Recommendations for Oracle B2B	2-15
2.4.6 Recommendations for Oracle Web Services Manager	2-16
2.4.7 Recommendations for Oracle User Messaging Service	2-17
2.4.8 Recommendations for Oracle JCA Adapters	2-18

2.4.9	Recommendations for Oracle Business Activity Monitoring.....	2-19
2.4.10	Recommendations for Oracle Business Process Management.....	2-20
2.5	Recommendations for Oracle Identity Management	2-21
2.5.1	Recommendations for Oracle Internet Directory.....	2-22
2.5.2	Recommendations for Oracle Virtual Directory	2-23
2.5.3	Recommendations for Oracle Directory Integration Platform.....	2-23
2.5.4	Recommendations for Oracle Identity Federation.....	2-24
2.5.5	Recommendations for Oracle Directory Services Manager.....	2-25
2.5.6	Recommendations for Oracle Access Manager.....	2-25
2.5.7	Recommendations for Oracle Adaptive Access Manager	2-26
2.5.8	Recommendations for Oracle Identity Manager.....	2-27
2.5.9	Recommendations for Oracle Identity Navigator.....	2-28
2.6	Recommendations for Oracle Portal, Forms, Reports, and Discoverer	2-29
2.6.1	Recommendations for Oracle Portal.....	2-30
2.6.2	Recommendations for Oracle Forms	2-31
2.6.3	Recommendations for Oracle Reports.....	2-32
2.6.4	Recommendations for Oracle Business Intelligence Discoverer.....	2-32
2.7	Recommendations for Oracle Web Tier Components.....	2-33
2.7.1	Recommendations for Oracle HTTP Server.....	2-33
2.7.2	Recommendations for Oracle Web Cache.....	2-35
2.8	Recommendations for Oracle WebCenter Content.....	2-36
2.8.1	Recommendations for Oracle WebCenter Content	2-37
2.8.2	Recommendations for Oracle WebCenter Content: Inbound Refinery	2-38
2.8.3	Recommendations for Oracle WebCenter Content: Imaging.....	2-38
2.8.4	Recommendations for Oracle WebCenter Content: Information Rights.....	2-39
2.8.5	Recommendations for Oracle WebCenter Content: Records	2-40
2.9	Recommendations for Oracle Business Intelligence.....	2-41
2.9.1	Recommendations for Oracle Business Intelligence Enterprise Edition (EE).....	2-42
2.9.2	Recommendations for Oracle Business Intelligence Publisher	2-43
2.9.3	Recommendations for Oracle Real-Time Decisions	2-43

3 Design Considerations

3.1	Network Considerations.....	3-4
3.1.1	Planning Host Names	3-5
3.1.1.1	Host Name Resolution.....	3-14
3.1.1.2	Resolving Host Names Locally.....	3-15
3.1.1.3	Resolving Host Names Using Separate DNS Servers	3-16
3.1.1.4	Resolving Host Names Using a Global DNS Server	3-17
3.1.1.5	Testing the Host Name Resolution	3-18
3.1.2	Virtual IP and Virtual Hostname Considerations.....	3-18
3.1.3	Load Balancers Considerations	3-26
3.1.4	Virtual Server Consideration	3-27
3.1.5	Wide Area DNS Operations.....	3-31
3.1.5.1	Using a Global Load Balancer	3-31
3.1.5.2	Manually Changing DNS Names.....	3-31
3.2	Storage Considerations	3-32
3.2.1	Oracle Fusion Middleware Artifacts.....	3-32

3.2.2	Oracle Home and Oracle Inventory	3-33
3.2.3	Storage Replication	3-33
3.2.4	File-Based Persistent Store.....	3-34
3.3	Database Considerations	3-34
3.3.1	Making TNSNAMES.ORA Entries for Databases.....	3-35
3.3.2	Manually Forcing Database Synchronization with Oracle Data Guard.....	3-35
3.3.3	Setting Up Database Host Name Aliases	3-36
3.4	Starting Points	3-37
3.4.1	Starting with an Existing Site.....	3-37
3.4.1.1	Migrating an Existing Production Site to Shared Storage	3-37
3.4.2	Starting with New Sites	3-38
3.5	Topology Considerations.....	3-39
3.5.1	Design Considerations for a Symmetric Topology.....	3-39
3.5.2	Design Considerations for an Asymmetric Topology.....	3-39

4 Setting Up and Managing Disaster Recovery Sites

4.1	Setting Up the Site.....	4-1
4.1.1	Directory Structure and Volume Design.....	4-2
4.1.1.1	Directory Structure Recommendations for Oracle SOA Suite.....	4-2
4.1.1.1.1	Volume Design for Oracle SOA Suite	4-5
4.1.1.1.2	Consistency Group Recommendations for Oracle SOA Suite.....	4-11
4.1.1.2	Directory Structure Recommendations for Oracle WebCenter Portal.....	4-12
4.1.1.2.1	Volume Design for Oracle WebCenter Portal	4-12
4.1.1.2.2	Consistency Group Recommendations for Oracle WebCenter Portal	4-17
4.1.1.3	Directory Structure Recommendations for Oracle Identity Management.....	4-18
4.1.1.3.1	Volume Design for Oracle Identity Management	4-19
4.1.1.3.2	Consistency Group Recommendations for Oracle Identity Management.	4-29
4.1.1.4	Directory Structure Recommendations for Oracle Portal, Forms, Reports, and Discoverer 4-31	
4.1.1.4.1	Volume Design for Oracle Portal, Forms, Reports, and Discoverer.....	4-36
4.1.1.4.2	Consistency Group Recommendations for Oracle Portal, Forms, Reports, and Discoverer 4-37	
4.1.1.5	Directory Structure Recommendations for Oracle WebCenter Content.....	4-38
4.1.1.5.1	Volume Design for Oracle WebCenter Content	4-39
4.1.1.5.2	Consistency Group Recommendations for Oracle WebCenter Content....	4-42
4.1.1.6	Directory Structure Recommendations for Oracle Business Intelligence	4-43
4.1.1.6.1	Volume Design for Oracle Business Intelligence.....	4-44
4.1.1.6.2	Consistency Group Recommendations for Oracle Business Intelligence ..	4-48
4.1.2	Storage Replication.....	4-49
4.1.3	Database	4-50
4.1.3.1	Setting Up Oracle Data Guard.....	4-50
4.1.3.1.1	Prerequisites and Assumptions.....	4-50
4.1.3.1.2	Oracle Data Guard Environment Description	4-50
4.1.3.1.3	Gather Files and Perform Backup.....	4-51
4.1.3.1.4	Configure Oracle Net Services on the Standby Site	4-52
4.1.3.1.5	Create Instances and Database on the Standby Site.....	4-53
4.1.3.1.6	Test Database Switchover and Switchback	4-57

4.1.4	Node Manager.....	4-58
4.1.4.1	Generate Self Signed Certificates	4-58
4.1.4.2	Create an Identity KeyStore	4-59
4.1.4.3	Create Trust KeyStore	4-59
4.1.4.4	Configure Node Manager for Custom KeyStores.....	4-60
4.2	Creating a Production Site.....	4-61
4.2.1	Creating the Production Site for the Oracle SOA Suite Topology.....	4-61
4.2.2	Creating the Production Site for the Oracle Identity Management Topology	4-62
4.2.3	Creating the Production Site for the Oracle WebCenter Portal Topology	4-62
4.2.4	Creating the Production Site for the Oracle WebCenter Content Topology.....	4-63
4.2.5	Creating the Production Site for the Oracle Business Intelligence Topology	4-64
4.2.6	Creating the Production Site for the Oracle Portal, Forms, Reports, and Discoverer Topology 4-64	
4.2.7	Validating the Production Site Setup.....	4-65
4.3	Creating a Standby Site	4-66
4.3.1	Creating the Standby Site	4-66
4.3.1.1	Database Setup.....	4-66
4.3.1.2	Middle Tier Setup	4-67
4.3.2	Validating the Standby Site Setup	4-67
4.4	Creating an Asymmetric Standby Site.....	4-67
4.4.1	Creating the Asymmetric Standby Site	4-68
4.4.1.1	Creating an Asymmetric Standby Site with Fewer Hosts and Instances	4-71
4.4.2	Validating the Asymmetric Standby Site Setup	4-74
4.5	Performing Site Operations and Administration	4-74
4.5.1	Synchronizing the Sites.....	4-74
4.5.2	Performing a Switchover	4-75
4.5.3	Performing a Switchback.....	4-75
4.5.4	Performing a Failover.....	4-76
4.5.5	Performing Periodic Testing of the Standby Site.	4-77
4.5.6	Using Peer to Peer File Copy for Testing	4-79
4.5.6.1	Using rsync and Oracle Data Guard for Oracle Fusion Middleware Disaster Recovery Topologies 4-80	
4.5.6.1.1	Using rsync for Oracle Fusion Middleware Middle Tier Components.....	4-80
4.5.6.1.2	Performing Failover and Switchover Operations.....	4-81
4.6	Patching an Oracle Fusion Middleware Disaster Recovery Site.....	4-82

5 Troubleshooting Disaster Recovery

5.1	Troubleshooting Oracle Fusion Middleware Disaster Recovery Topologies	5-1
5.1.1	Verify Host Name Resolution at the Production and Standby Sites.....	5-1
5.1.2	Resolving Issues with Components in a Disaster Recovery Topology.....	5-1
5.1.3	Resolving Issues with Components Deployed on Shared Storage.....	5-2
5.2	Need More Help?.....	5-2

A Managing Oracle Inventory

A.1	Updating Oracle Inventory	A-1
-----	---------------------------------	-----

A.2 Updating the Windows Registry A-1

Index

Preface

This preface contains these sections:

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

Audience

This document is intended for administrators, developers, and others whose role is to deploy and manage the Oracle Fusion Middleware Disaster Recovery solution using storage replication technology.

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 Documents

For more information, see the following documents in the Oracle Fusion Middleware documentation set:

- *Oracle Fusion Middleware High Availability Guide*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence*

- *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.

Disaster Recovery Introduction

This chapter provides an introduction to the Oracle Fusion Middleware Disaster Recovery solution.

It contains the following topics:

- [Disaster Recovery Overview](#)
- [Disaster Recovery for Oracle Fusion Middleware Components](#)

1.1 Disaster Recovery Overview

This section provides an overview of Oracle Fusion Middleware Disaster Recovery.

It contains the following topics:

- [Problem Description and Common Solutions](#)
- [Terminology](#)

1.1.1 Problem Description and Common Solutions

Providing Maximum Availability Architecture is one of the key requirements for any Oracle Fusion Middleware enterprise deployment. Oracle Fusion Middleware includes an extensive set of high availability features such as: process death detection and restart, server clustering, server migration, clusterware integration, GridLink, load balancing, failover, backup and recovery, rolling upgrades, and rolling configuration changes, which protect an Enterprise Deployment from unplanned down time and minimize planned downtime.

Additionally, enterprise deployments need protection from unforeseen disasters and natural calamities. One protection solution involves setting up a standby site at a geographically different location than the production site. The standby site may have equal or fewer services and resources compared to the production site. Application data, metadata, configuration data, and security data are replicated to the standby site on a periodic basis. The standby site is normally in a passive mode; it is started when the production site is not available. This deployment model is sometimes referred to as an active/passive model. This model is normally adopted when the two sites are connected over a WAN and network latency does not allow clustering across the two sites.

A core strategy for and a key feature of Oracle Fusion Middleware is hot-pluggability. Built for the heterogeneous enterprise, Oracle Fusion Middleware consists of modular component software that runs on a range of popular platforms and interoperates with middleware technologies and business applications from other software vendors such as IBM, Microsoft, and SAP. For instance, Oracle Fusion Middleware products and

technologies such as ADF, Oracle BPEL Process Manager, Oracle Enterprise Service Bus, Oracle Web Services Manager, Adapters, Oracle Access Manager, Oracle Identity Manager, Rules, Oracle TopLink, and Oracle Business Intelligence Publisher can run on non-Oracle containers such as IBM Websphere and JBoss, in addition to running on the Oracle WebLogic Server container.

The Oracle Fusion Middleware Disaster Recovery solution uses storage replication technology for disaster protection of Oracle Fusion Middleware middle tier components. It supports hot-pluggable deployments, and it is compatible with third party vendor recommended solutions.

Disaster protection for Oracle databases that are included in your Oracle Fusion Middleware is provided through Oracle Data Guard.

This document describes how to deploy the Oracle Fusion Middleware Disaster Recovery solution for enterprise deployments on Linux and UNIX operating systems, making use of storage replication technology and Oracle Data Guard technology.

1.1.2 Terminology

This section defines the following Disaster Recovery terminology:

- **asymmetric topology:** An Oracle Fusion Middleware Disaster Recovery configuration that is different across tiers on the production site and standby site. For example, an asymmetric topology can include a standby site with fewer hosts and instances than the production site. [Section 4.4, "Creating an Asymmetric Standby Site"](#) describes how to create asymmetric topologies.
- **disaster:** A sudden, unplanned catastrophic event that causes unacceptable damage or loss. A disaster is an event that compromises an organization's ability to provide critical functions, processes, or services for some unacceptable period of time and causes the organization to invoke its recovery plans.
- **Disaster Recovery:** The ability to safeguard against natural or unplanned outages at a production site by having a recovery strategy for applications and data to a geographically separate standby site.
- **alias host name:** This guide differentiates between the terms alias host name and physical host name.

The alias host name is an alternate way to access the system besides its real network name. Typically, it resolves to the same IP address as the network name of the system. This can be defined in the name resolution system such as DNS, or locally in the local hosts file on each system. Multiple alias host names can be defined for a given system.

See also the **physical host name** definition later in this section.

- **physical host name:** The physical host name is the host name of the system as returned by the `gethostname()` call or the `hostname` command. Typically, the physical host name is also the network name used by clients to access the system. In this case, an IP address is associated with this name in the DNS (or the given name resolution mechanism in use) and this IP is enabled on one of the network interfaces to the system.

A given system typically has one physical host name. It can also have one or more additional network names, corresponding to IP addresses enabled on its network interfaces, that are used by clients to access it over the network. Further, each network name can be aliased with one or more alias host names.

See also the **alias host name** definition earlier in this section.

- **virtual host name:** Virtual host name is a network addressable host name that maps to one or more physical machines via a load balancer or a hardware cluster. For load balancers, the name "virtual server name" is used interchangeably with virtual host name in this book. A load balancer can hold a virtual host name on behalf of a set of servers, and clients communicate indirectly with the machines using the virtual host name. A virtual host name in a hardware cluster is a network host name assigned to a cluster virtual IP. Because the cluster virtual IP is not permanently attached to any particular node of a cluster, the virtual host name is not permanently attached to any particular node either.

Note: Whenever the term "virtual host name" is used in this document, it is assumed to be associated with a virtual IP address. In cases where just the IP address is needed or used, it will be explicitly stated.

- **virtual IP:** Also, cluster virtual IP and load balancer virtual IP. Generally, a virtual IP can be assigned to a hardware cluster or load balancer. 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 which are members of the cluster. A virtual IP can be assigned to a server load balancer or a hardware cluster.

A hardware cluster uses a cluster virtual IP to present to the outside world the entry point into the cluster (it can also be set up on a standalone machine). The hardware cluster's software manages the movement of this IP address between the two physical nodes of the cluster while clients connect to this IP address without the need to know which physical node this IP address is currently active on. In a typical two-node hardware cluster configuration, each machine has its own physical IP address and physical host name, while there could be several cluster IP addresses. These cluster IP addresses float or migrate between the two nodes. The node with current ownership of a cluster IP address is active for that address.

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.

- **production site setup:** The process of creating the production site. To create the production site using the procedure described in this manual, you must plan and create physical host names and alias host names, create mount points and symbolic links (if applicable) on the hosts to the Oracle home directories on the shared storage where the Oracle Fusion Middleware instances will be installed, install the binaries and instances, and deploy the applications. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes. See [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
- **site failover:** The process of making the current standby site the new production site after the production site becomes unexpectedly unavailable (for example, due to a disaster at the production site). This book also uses the term "failover" to refer to a site failover.
- **site switchback:** The process of reverting the current production site and the current standby site to their original roles. Switchbacks are planned operations done after the switchover operation has been completed. A switchback restores the original roles of each site: the current standby site becomes the production site

and the current production site becomes the standby site. This book also uses the term "switchback" to refer to a site switchback.

- **site switchover:** The process of reversing the roles of the production site and standby site. Switchovers are planned operations done for periodic validation or to perform planned maintenance on the current production site. During a switchover, the current standby site becomes the new production site, and the current production site becomes the new standby site. This book also uses the term "switchover" to refer to a site switchover.
- **site synchronization:** The process of applying changes made to the production site at the standby site. For example, when a new application is deployed at the production site, you should perform a synchronization so that the same application will be deployed at the standby site, also.
- **standby site setup:** The process of creating the standby site. To create the standby site using the procedure described in this manual, you must plan and create physical host names and alias host names, and create mount points and symbolic links (if applicable) to the Oracle home directories on the standby shared storage. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes. See [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
- **symmetric topology:** An Oracle Fusion Middleware Disaster Recovery configuration that is completely identical across tiers on the production site and standby site. In a symmetric topology, the production site and standby site have the identical number of hosts, load balancers, instances, and applications. The same ports are used for both sites. The systems are configured identically and the applications access the same data. This manual describes how to set up a symmetric Oracle Fusion Middleware Disaster Recovery topology for an enterprise configuration.
- **topology:** The production site and standby site hardware and software components that comprise an Oracle Fusion Middleware Disaster Recovery solution.

1.2 Disaster Recovery for Oracle Fusion Middleware Components

This section provides an introduction to setting up Disaster Recovery for a common Oracle Fusion Middleware enterprise deployment.

It contains the following topics:

- [Oracle Fusion Middleware Disaster Recovery Architecture Overview](#)
- [Components Described in this Document](#)

1.2.1 Oracle Fusion Middleware Disaster Recovery Architecture Overview

This section describes the deployment architecture for Oracle Fusion Middleware components.

The product binaries and configuration for Oracle Fusion Middleware components and applications gets deployed in Oracle home directories on the middle tier. Additionally, most of the products also have metadata or run-time data stored in a database repository.

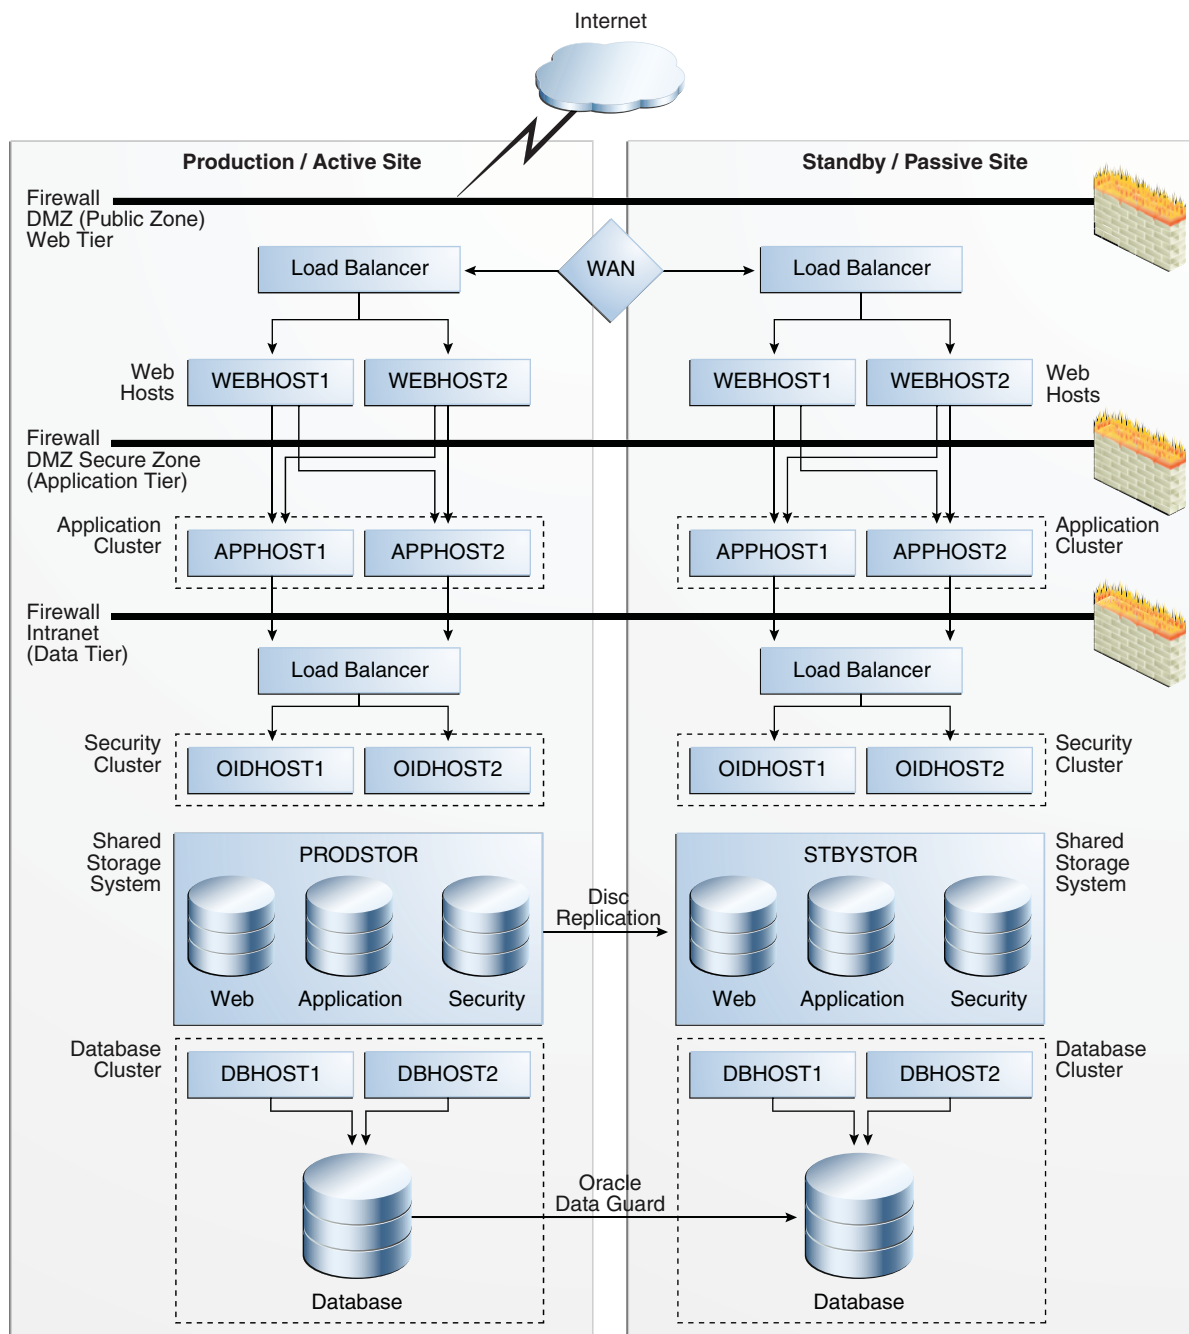
Therefore, the Oracle Fusion Middleware Disaster Recovery solution keeps middle tier file system data and middle tier data stored in databases at the production site synchronized with the standby site.

The Oracle Fusion Middleware Disaster Recovery solution supports these methods of providing data protection for Oracle Fusion Middleware data and database content:

- Oracle Fusion Middleware product binaries, configuration, and metadata files
Use storage replication technologies.
- Database content
Use Oracle Data Guard for Oracle databases (and vendor-recommended solutions for third party databases).

[Figure 1-1](#) shows an overview of an Oracle Fusion Middleware Disaster Recovery topology:

Figure 1–1 Production and Standby Site for Oracle Fusion Middleware Disaster Recovery Topology



Some of the key aspects of the solution in [Figure 1–1](#) are:

- The solution has two sites. The current production site is running and active, while the second site is serving as a standby site and is in passive mode.
- Hosts on each site have mount points defined for accessing the shared storage system for the site.
- On both sites, the Oracle Fusion Middleware components are deployed on the site’s shared storage system. This involves creating all the Oracle home directories, which include product binaries and configuration data for middleware

components, in volumes on the production site's shared storage and then installing the components into the Oracle home directories on the shared storage. In [Figure 1-1](#), a separate volume is created in the shared storage for each Oracle Fusion Middleware host cluster (note the Web, Application, and Security volumes created for the Web Cluster, Application Cluster, and Security Cluster in each site's shared storage system).

- Mount points must be created on the shared storage for the production site. The Oracle Fusion Middleware software for the production site will be installed into Oracle home directories using the mount points on the production site shared storage. Symbolic links may also need to be set up on the production site hosts to the Oracle Fusion Middleware home directories on the shared storage at the production site. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes. See [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
- Mount points must be created on the shared storage for the standby site. Symbolic links also need to be set up on the standby site hosts to the Oracle Fusion Middleware home directories on the shared storage at the standby site. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes. See [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links. The mount points and symbolic links for the standby site hosts must be identical to those set up for the equivalent production site hosts.
- Storage replication technology is used to copy the middle tier file systems and other data from the production site's shared storage to the standby site's shared storage.
- After storage replication is enabled, application deployment, configuration, metadata, data, and product binary information is replicated from the production site to the standby site.
- It is not necessary to perform any Oracle software installations at the standby site hosts. When the production site storage is replicated at the standby site storage, the equivalent Oracle home directories and data are written to the standby site storage.
- Schedule incremental replications at a specified interval. The recommended interval is once a day for the production deployment, where the middle tier configuration does not change very often. Additionally, you should force a manual synchronization whenever you make a change to the middle tier configuration at the production site (for example, if you deploy a new application at the production site). Some Oracle Fusion Middleware components generate data on the file system, which may require more frequent replication based on recovery point objectives. Please refer to [Chapter 2, "Recommendations for Fusion Middleware Components"](#) for detailed Disaster Recovery recommendations for Oracle Fusion Middleware components.
- Before forcing a manual synchronization, you should take a snapshot of the site to capture its current state. This ensures that the snapshot gets replicated to the standby site storage and can be used to roll back the standby site to a previous synchronization state, if desired. Recovery to the point of the previously successful replication (for which a snapshot was created) is possible when a replication fails.
- Oracle Data Guard is used to replicate all Oracle database repositories, including Oracle Fusion Middleware repositories and custom application databases. For information about using Oracle Data Guard to provide disaster protection for Oracle databases, see [Section 3.3, "Database Considerations."](#)

- If your Oracle Fusion Middleware Disaster Recovery topology includes any third party databases, use the vendor-recommended solution for those databases.
- User requests are initially routed to the production site.
- When there is a failure or planned outage of the production site, you perform the following steps to enable the standby site to assume the production role in the topology:
 1. Stop the replication from the production site to the standby site (when a failure occurs, replication may have already been stopped due to the failure).
 2. Perform a failover or switchover of the Oracle databases using Oracle Data Guard.
 3. Start the services and applications on the standby site.
 4. Use a global load balancer to re-route user requests to the standby site. At this point, the standby site has assumed the production role.

1.2.2 Components Described in this Document

The Oracle Fusion Middleware Disaster Recovery solution supports components from various Oracle product suites, including:

- Oracle WebLogic Server
See [Section 2.1, "Recommendations for Oracle WebLogic Server"](#) for Disaster Recovery recommendations for Oracle WebLogic Server components.
- Oracle ADF
See [Section 2.2, "Recommendations for Oracle ADF"](#) for Disaster Recovery recommendations for Oracle Application Development Framework (Oracle ADF).
- Oracle WebCenter Portal components:
 - Oracle WebCenter Portal: Spaces
 - Oracle WebCenter Portal's Portlet Producers
 - Oracle WebCenter Portal's Discussion Server
 - Oracle WebCenter Content Server
 - Oracle WebCenter Portal Pagelet Producer
 - Oracle WebCenter Portal Activity Graph Engines
 - Oracle WebCenter Portal's Personalization
 - Oracle WebCenter Portal's Analytics Collector
 - Oracle WebCenter Portal Services ProducerSee [Section 2.3, "Recommendations for Oracle WebCenter Portal"](#) for Disaster Recovery recommendations for Oracle WebCenter Portal components.
- Oracle SOA Suite components:
 - Oracle SOA Service Infrastructure
 - Oracle BPEL Process Manager
 - Oracle Mediator
 - Oracle Human Workflow
 - Oracle B2B

- Oracle Web Services Manager
- Oracle User Messaging Service
- Oracle JCA Adapters
- Oracle Business Activity Monitoring
- Oracle Business Process Management

See [Section 2.4, "Recommendations for Oracle SOA Suite"](#) for Disaster Recovery recommendations for Oracle SOA Suite components.

- Oracle Identity Management components:
 - Oracle Internet Directory
 - Oracle Virtual Directory
 - Oracle Directory Integration Platform
 - Oracle Identity Federation
 - Oracle Directory Services Manager
 - Oracle Access Manager
 - Oracle Adaptive Access Manager
 - Oracle Identity Manager
 - Oracle Identity Navigator

See [Section 2.5, "Recommendations for Oracle Identity Management"](#) for Disaster Recovery recommendations for Oracle Identity Management components.

- Oracle Portal, Forms, Reports, and Business Intelligence Discoverer components:
 - Oracle Portal
 - Oracle Forms
 - Oracle Reports
 - Oracle Business Intelligence Discoverer (Discoverer)

See [Section 2.6, "Recommendations for Oracle Portal, Forms, Reports, and Discoverer"](#) for Disaster Recovery recommendations for these components.

- Oracle Web Tier components:
 - Oracle HTTP Server
 - Oracle Web Cache

See [Section 2.7, "Recommendations for Oracle Web Tier Components"](#) for Disaster Recovery recommendations for Oracle Web Tier components.

- Oracle WebCenter Content:
 - Oracle WebCenter Content
 - Oracle WebCenter Content: Inbound Refinery
 - Oracle WebCenter Content: Imaging
 - Oracle WebCenter Content: Information Rights
 - Oracle WebCenter Content: Records

See [Section 2.8, "Recommendations for Oracle WebCenter Content"](#) for Disaster Recovery recommendations for Oracle WebCenter Content components.

- Oracle Business Intelligence:
 - Oracle Business Intelligence Enterprise Edition (EE)
 - Oracle Business Intelligence Publisher
 - Oracle Real-Time Decisions

See [Section 2.9, "Recommendations for Oracle Business Intelligence"](#) for Disaster Recovery recommendations for Oracle Enterprise Content Management components.

Recommendations for Fusion Middleware Components

This chapter describes the disaster protection requirements for Oracle Fusion Middleware components in different Oracle product suites and also provides recommendations for synchronizing those components. As mentioned previously, use storage replication to synchronize middle tier content, and use Oracle Data Guard to synchronize data in Oracle database repositories or custom application databases included in your Oracle Fusion Middleware Disaster Recovery topology.

This chapter provides Disaster Recovery recommendations for components in the following Oracle product suites:

- [Recommendations for Oracle WebLogic Server](#)
- [Recommendations for Oracle ADF](#)
- [Recommendations for Oracle WebCenter Portal](#)
- [Recommendations for Oracle SOA Suite](#)
- [Recommendations for Oracle Identity Management](#)
- [Recommendations for Oracle Portal, Forms, Reports, and Discoverer](#)
- [Recommendations for Oracle Web Tier Components](#)
- [Recommendations for Oracle WebCenter Content](#)
- [Recommendations for Oracle Business Intelligence](#)

Common Artifacts Across Oracle Product Suites

Certain artifacts like the Oracle Inventory, the `beahomelist` file, the `oratab` file and `oraInst.loc` file are common across all Oracle product deployments. These artifacts change very rarely and need not be a part of the regular storage replication and synchronization activity. It is recommended to have the Oracle Inventory, the `beahomelist` file, the `oratab` file, and the `oraInst.loc` file on the local disk of the machines. These artifacts should be manually updated upon creation, as well as upon applying patch updates. If required by your environment, these artifacts can also be on shared storage.

2.1 Recommendations for Oracle WebLogic Server

Oracle WebLogic Server is a scalable, enterprise-ready Java Platform, Enterprise Edition (Java EE) application server. The Oracle WebLogic Server infrastructure supports the deployment of many types of distributed applications and is an ideal foundation for building applications based on Service Oriented Architectures (SOA).

Common Artifacts and Considerations for Oracle WebLogic Server

The following artifacts and considerations apply to all the WebLogic Server components along with the component-specific recommendations.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries.

Domain home: The domain home contains the configuration data and the applications for the WebLogic domain.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

If your environment requires whole server migration to be configured, it is recommended to use the virtual host name as the listen address of the Managed Servers that are configured for whole server migration. To avoid manually updating the listen address after a disaster recovery operation, make sure that the host name can be resolved on both the primary and standby sites.

The load balancer virtual hosts used for accessing the WebLogic Server applications should be configured on both the production and standby sites

The rest of this section describes Disaster Recovery recommendations for the following Oracle WebLogic Server components:

- [Recommendations for Oracle WebLogic Server JMS and T-Logs](#)
- [Recommendations for Oracle Platform Security Services](#)

2.1.1 Recommendations for Oracle WebLogic Server JMS and T-Logs

This section describes various Oracle WebLogic Server JMS and transaction log (T-log) artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

File-based persistent stores: The file store location for the JMS/T-log when using a file based persistent store.

Artifacts in the Database

The schema containing the JMS messages, when using database based persistent stores. The schema containing Logging Last Resource (LLR) transaction log records for WebLogic applications that leverage the JDBC LLR option.

When automatic whole server migration is configured, the required leasing table is in the database.

Special Considerations

- Messages are lost if they were en-queued after the system restore point time but never processed. Message duplicates are generated for messages enqueued before the restore point time, but dequeued and acknowledged or committed (processed) after this time.

- If the persistent store is a custom store that is dedicated to JMS use, then you can delete the entire store.
- Restoring different parts of the system to different points in time can lead to inconsistent data. This can occur when the message store, transaction log, or application database are synchronized differently. For example, a message may reference a database row that does not exist, or vice-versa. This may delete unprocessed messages in addition to duplicate messages.
- If the store is not dedicated to JMS use, use the Oracle WebLogic Server JMS message management administrative tooling. This tooling can perform import, export, move, and delete operations from the Administration Console, MBeans, and WLST.
- When applications use both queues and topics, make sure to manipulate both the queue and topic subscriptions.

Synchronization Recommendations

- If JMS data is critical, it is recommended to synchronize transaction log data and JMS data in real time using synchronous replication. Note that using synchronous replication may have performance implications.
- If data consistency between tiers is important, ensure that the database and application tiers are replicated at the same time. This helps ensure that the different tiers recover to the same exact point in time.
- Use Oracle Data Guard to replicate the primary site and standby site when using database based persistent stores.
- When using a storage device that does not support block-level snapshot capabilities, shut down the JMS server to take a consistent backup. This is to ensure that the persistence store is not being written to while the copy operation is being performed. In a clustered environment, you can do so by shutting down one server at a time, backing it up and restarting it. You also can create a script to perform these operations using WLST.

Recovery Recommendations

Recover the database schemas containing the persistent store to the most recent point in time, the Administration Server, and the Managed Servers in the WebLogic Server domain.

Also, follow the recovery recommendations below for avoiding duplicate messages.

Avoiding Duplicate Messages

Use the following procedure before recovery to drain messages in the JMS queue after persistent-store recovery to avoid processing duplicate messages:

Note: Do not drain and discard messages without first being certain that the messages contain no data that must be preserved. The recovered messages may include unprocessed messages with important application data, in addition to duplicate messages that have already been processed.

1. Log into the Oracle WebLogic Server Administration Console.
2. Before recovery, configure JMS server to pause Production, Insertion, and consumption operations at boot-time to ensure that no new messages are

produced or inserted into the destination or consumed from the destination before you drain stale messages. To do this:

- a. Expand **Services**, then **Messaging**, and then **JMS Servers**.
- b. On the Summary of JMS Servers page, click the JMS server you want to configure for message pausing.
- c. On the Configuration: General page, click **Advanced** to define the message pausing options. Select **Insertion Paused At Startup**, **Production Paused At Startup**, and **Consumption Paused At Startup**.
- d. Click **Save**.
- e. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

Use the following procedure after recovery:

1. After recovering the persistent store, start the Managed Servers.
2. Drain the stale messages from JMS destinations by following these steps:
 - a. Expand **Services**, then **Messaging**, and then **JMS Modules**.
 - b. Select a JMS module, then select a destination.
 - c. Select **Monitoring**, then **Show Messages**.
 - d. Click **Delete All**.

Resume operations by following these steps:

1. Expand **Services**, then **Messaging**, and then **JMS Servers**.
2. On the Summary of JMS Servers page, click the JMS server you want to configure for message pausing.
3. On the Configuration: General page, click **Advanced**. Select **Insertion Paused At Startup**, **Production Paused At Startup**, and **Consumption Paused At Startup**.
4. Click **Save**.
5. To activate these changes, in the Change Center of the Administration Console, click **Activate Changes**.

2.1.2 Recommendations for Oracle Platform Security Services

This section describes various Oracle Platform Security Services artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Not applicable because Oracle Platform Security Services does not have any database dependencies.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in

Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

Recover the Administration Server and the Managed Servers in the WebLogic Server domain.

2.2 Recommendations for Oracle ADF

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

This section describes various Oracle ADF artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries.

Oracle_Common_Home: 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: The domain home contains the configuration data and the applications for the Oracle ADF domain.

Custom Applications Directory: This is the stage location for various customer applications and their related libraries.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

The load balancer virtual hosts used for accessing the applications should be configured on both the production and standby sites.

Artifacts in the Database

Oracle ADF applications may use the MDS repository to persist the application state and configuration data. Persisting data depends on how the application is coded.

Most Oracle ADF applications do not use the MDS repository to store application data, but instead use a separate data store (usually a database) to store application data.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying composites, applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Servers running the ADF or WebCenter applications.

2.3 Recommendations for Oracle WebCenter Portal

Oracle 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.

Common Artifacts and Considerations for Oracle WebCenter Portal

The artifacts and considerations below apply to all the Oracle WebCenter Portal products along with the product-specific considerations.

Artifacts on the File System

`MW_HOME`: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries and an Oracle home containing the Oracle WebCenter Portal binaries.

`Oracle_Common_Home`: 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: The domain home contains the configuration data for the Oracle WebCenter Portal domain.

Artifacts in the Database

The Oracle WebCenter Portal metadata stores data for some of the Oracle WebCenter Portal services and is part of Oracle WebCenter Portal databases, and the MDS repository stores WebCenter metadata and configuration information.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

The load balancer virtual hosts required for accessing the Oracle WebCenter Portal products should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for the Oracle database containing the Oracle WebCenter Portal schema and the metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The database containing the Oracle WebCenter Portal schemas must be recovered to the most recent point in time, along with the Managed Servers running the Oracle WebCenter Portal applications.

The rest of this section describes Disaster Recovery recommendations for the following Oracle WebCenter Portal products:

- [Recommendations for Oracle WebCenter Portal: Spaces](#)
- [Recommendations for Oracle WebCenter Portal's Portlet Producers](#)
- [Recommendations for Oracle WebCenter Portal's Discussion Server](#)
- [Recommendations for Oracle WebCenter Content Server](#)
- [Recommendations for Oracle WebCenter Portal's Analytics Collector](#)
- [Recommendations for Oracle WebCenter Portal Activity Graph Engines](#)

2.3.1 Recommendations for Oracle WebCenter Portal: Spaces

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.

This section describes various Spaces artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The WEBCENTER schema stores data for some of the Oracle WebCenter Portal services and the MDS repository stores WebCenter Portal metadata and configuration information.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for the Oracle database containing the WEBCENTER schema and the metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database containing the WebCenter Portal schema and MDS repository must be recovered to the most recent point in time, along with the Oracle WebCenter Portal domain.

2.3.2 Recommendations for Oracle WebCenter Portal's Portlet Producers

Oracle WebCenter Portal supports deployment and execution of both standards-based portlets (JSR 168, WSRP 1.0 and 2.0), and traditional Oracle PDK-Java based portlets. Oracle WebCenter Portal provides several out-of-the-box producers, such as OmniPortlet, Web Clipping, Rich Text Portlet, and WSRP Tools.

This section describes various Portlets Producer artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The PORTLET schema stores user customizations and is part of the Oracle WebCenter Portal schemas, and the MDS repository stores portlet metadata and configuration information.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for the Oracle database containing the PORTLET schema and the metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database containing the PORTLET schema and MDS repository must be recovered to the most recent point in time, along with the Oracle WebCenter Portal domain.

2.3.3 Recommendations for Oracle WebCenter Portal's Discussion Server

Oracle WebCenter Portal's Discussion Server provides the ability to integrate discussion forums and announcements into your applications.

This section describes various Discussion Server artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The DISCUSSIONS schema stores metadata and data and is part of the Oracle WebCenter Portal schemas.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for the Oracle database containing the DISCUSSIONS schema and the metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database containing the Discussion Server schema must be recovered to the most recent point in time, along with the Oracle WebCenter Portal domain.

2.3.4 Recommendations for Oracle WebCenter Content Server

Oracle WebCenter Content Server provides the ability to integrate wikis and blogs into WebCenter Portal applications. It also supports features that enable application users to create their own wikis and blogs.

This section describes various Oracle WebCenter Wiki and Blog Server artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The OCSERVER schema stores metadata and data and is part of the Oracle WebCenter Portal schemas.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for the Oracle database containing the OCSERVER schema.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database containing the OCSERVER schema must be recovered to the most recent point in time, along with the Oracle WebCenter Portal domain.

2.3.5 Recommendations for Oracle WebCenter Portal's Analytics Collector

Oracle WebCenter Portal's Analytics Collector allows portal managers and business owners to track and analyze portal usage.

This section describes the Oracle WebCenter Portal's Analytics data that must be backed up and restored.

Configuration Files

Configuration information is stored in the Analytics schema, ACTIVITIES.

Database Repository Dependencies

ACTIVITIES and MDS schema

Backup Recommendations

Back up the Oracle home, the domain home, and the database containing the ACTIVITIES and MDS schemas.

Recovery Recommendations

Recover the Oracle home and the domain home.

Recover the database to the most recent point in time, if needed.

2.3.6 Recommendations for Oracle WebCenter Portal Activity Graph Engines

Oracle WebCenter Portal Activity is a WebCenter Portal service that provides suggestions of people, items, and spaces that users may be interested in interacting with.

The engine used by the Activity Graph service to provide a central repository for actions that are collected by enterprise applications. The data stored in the activity graph is analyzed to calculate ranks for nodes, predict new actions, and make recommendations.

This section describes the Oracle WebCenter Portal Activity Graph data that must be backed up and restored.

Configuration Files

Configuration information is stored in the ACTIVITIES schema.

Database Repository Dependencies

ACTIVITIES schema

Backup Recommendations

Back up the Oracle home, the domain home, and the database containing the ACTIVITIES schema.

Recovery Recommendations

Recover the Oracle home and the domain home.

Recover the database to the most recent point in time, if needed.

2.4 Recommendations for Oracle SOA Suite

Oracle SOA Suite is a middleware component of Oracle Fusion Middleware. Oracle SOA Suite provides a complete set of service infrastructure components for designing, deploying, and managing SOA composite applications. Oracle SOA Suite enables services to be created, managed, and orchestrated into SOA composite applications. Composites enable you to easily assemble multiple technology components into one SOA composite application. SOA composite applications consist of:

- **Service components:** Service components are the basic building blocks of SOA composite applications. Service components implement a part of the overall business logic of the SOA composite application. Oracle BPEL Process Manager, Oracle Mediator, Oracle Human Workflow and Business Rules are examples of service components.
- **Binding components:** Binding components connect SOA composite applications to external services, applications, and technologies. Binding components are organized into two groups:

- Services: Provide the outside world with an entry point to the SOA composite application. The WSDL file of the service advertises its capabilities to external applications. The service bindings define how a SOA composite service can be invoked (for example, through SOAP).
- References: Enable messages to be sent from the SOA composite application to external services (for example, the same functionality that partner links provide for BPEL processes, but at the higher SOA composite application level).

Note: In Oracle SOA Suite release 11.1.1.1, the soa-infra and service engine configuration files were stored in local or shared storage files as part of the domain configuration.

Starting in Oracle SOA Suite 11.1.1.2, those files were moved into the metadata repository. Thus, soa-infra and service-engine configuration changes are now immediately propagated across a cluster.

In Oracle SOA Suite 11.1.1.3, you can deploy Oracle BPM Suite on top of a SOA Suite installation.

The Disaster Recovery recommendations for Oracle SOA Suite assume that you are using Oracle SOA Suite 11.1.1.2.

Oracle SOA Suite artifacts are stored on the local or shared file system as well as in the metadata repositories. Composite artifacts are stored in the metadata repository, and binaries and domain-related configuration files are stored on a local or shared file system.

Common Artifacts and Considerations for All SOA Suite Components

The artifacts and considerations below apply to all the SOA Suite components, along with the component-specific considerations.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries and an Oracle home containing the Oracle SOA Suite binaries.

Oracle_Common_Home: 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: The domain home contains the configuration data and SOA composites for the SOA domain.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation. See *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for instructions updating an IP address to a virtual host name.

The load balancer virtual hosts required for accessing the SOA Suite components should be configured on both the production and standby sites.

Artifacts in the Database

Oracle SOA Suite schemas, Service Infrastructure and Service Engine configurations, and composite definitions are stored in the Oracle SOA Suite database and metadata repository.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes, deploying composites, and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The database must be recovered to the most recent point in time to ensure that the latest composite definitions and in-flight instances are restored.

In-flight instances require the matching composite definition to continue processing. For this reason, the metadata repository (where composite definitions are stored) and Oracle SOA Suite database (where process state is maintained) must be recovered to the same point in time.

In case of redeployed composites, a database recovery ensures consistency between the dehydrated in-flight processes and their corresponding definition since the process definition is stored in database repository where dehydrated instances are stored.

This section describes Disaster Recovery recommendations for the following Oracle SOA Suite components:

- [Recommendations for Oracle SOA Service Infrastructure](#)
- [Recommendations for Oracle BPEL Process Manager](#)
- [Recommendations for Oracle Mediator](#)
- [Recommendations for Oracle Human Workflow](#)
- [Recommendations for Oracle B2B](#)
- [Recommendations for Oracle Web Services Manager](#)
- [Recommendations for Oracle User Messaging Service](#)
- [Recommendations for Oracle JCA Adapters](#)
- [Recommendations for Oracle Business Activity Monitoring](#)
- [Recommendations for Oracle Business Process Management](#)

2.4.1 Recommendations for Oracle SOA Service Infrastructure

Oracle SOA Service Infrastructure is a Java EE application that provides the foundation services for running Oracle Fusion Middleware SOA Suite. This Java EE

application is a run-time engine that is automatically deployed when Oracle Fusion Middleware SOA Suite is installed. You deploy composites (the basic artifacts in a Service 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 run-time operations.

This section describes various Oracle SOA Service Infrastructure artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Composite definition and configuration files are stored in the MDS repository. The composite instance state persistence is stored in the SOA Service Infrastructure database.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes, deploying composites, and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time to ensure that the latest composite definitions and in-flight instances are restored.

2.4.2 Recommendations for Oracle BPEL Process Manager

The Oracle BPEL Process engine is the service engine running in SOA Service Infrastructure that allows 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.

This section describes various Oracle BPEL Process Manager artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Process definition and configuration files are stored in the MDS repository. The BPEL process state persistence is stored in the Oracle SOA Suite database.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time to ensure that the latest process definitions and in-flight instances are restored. Idempotent Oracle BPEL Process Manager processes are recommended, since no cleanup is required after performing a Disaster Recovery operation. If non-idempotent processes Oracle BPEL Process Manager processes are used, then processes must be cleaned up from the dehydration store after a Disaster Recovery operation is performed, especially when a process is in flight.

2.4.3 Recommendations for Oracle Mediator

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 Mediator service engine runs in-place with the SOA Service Infrastructure Java EE application.

This section describes various Oracle Mediator artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The Mediator service engine stores messages in the database for asynchronous routing for parallel routing rules (note that sequential routing rules do not persist their messages into the database as part of the execution). The Mediator component instance state and audit details are also stored in the database.

The metadata repository stores the Mediator component definition as part of the composite definition.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Administration Server and the Managed Server running the soa-infra application.

2.4.4 Recommendations for Oracle Human Workflow

Oracle Human Workflow is a service engine running in the Oracle SOA Service Infrastructure that allows 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 several services that handle various aspects of human interaction with a business process.

This section describes various Oracle Human Workflow artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Human workflow instance data and other worklist data such as vacation rules, group rules, flex field mappings, view definitions are stored in the database.

The metadata repository is used to store shared human workflow service definitions and schemas used by SOA composites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running the soa-infra application. Oracle Human Workflow's engine uses Oracle User Messaging Service to send and receive notifications. See [Section 2.4.7, "Recommendations for Oracle User Messaging Service"](#) for details about Oracle User Messaging Service.

2.4.5 Recommendations for Oracle B2B

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. Oracle B2B is implemented as a binding component within the SOA Service Infrastructure.

This section describes various Oracle B2B artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

JMS Store: The volume containing the file-based JMS persistent store. [Table 2-1](#) shows the JMS queues and topics used internally by Oracle B2B.

Table 2–1 JMS Queues and Topics Used by Oracle B2B

JMS Artifact Name	Type	JNDI Name
dist_B2BEventQueue_auto	Distributed queue	jms/b2b/B2BEventQueue
dist_B2B_IN_QUEUE_auto	Distributed queue	jms/b2b/B2B_IN_QUEUE
dist_B2B_OUT_QUEUE_auto	Distributed queue	jms/b2b/B2B_OUT_QUEUE
dist_B2BBroadcastTopic_auto	Distributed topic	jms/b2b/B2BBroadcastTopic

Artifacts in the Database

Oracle B2B message and message state persistence are stored in the Oracle SOA Suite database along with the partners, documents, and channels definitions. The metadata repository is used for storing Oracle B2B metadata.

Special Considerations

The external FTP servers and email servers should be available on the standby site if these adapters are used.

Synchronization Recommendations

For information about Oracle B2B JMS queue synchronization and recovery, refer to [Section 2.1.1, "Recommendations for Oracle WebLogic Server JMS and T-Logs."](#)

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running the soa-infra application. Oracle B2B stores state information within JMS queues and the SOA run-time database, so recovering the database and the Managed Server will ensure that the application runs normally.

2.4.6 Recommendations for Oracle Web Services Manager

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. Oracle Web Services Manager is made up of the Policy Manager and the Agent.

The Policy Manager reads and writes security and management policies, including predefined and custom policies from the MDS repository. Policy Manager is a stateless Java EE application. It exposes its capabilities through stateless session beans. Although the Policy Manager does not cache any data, the underlying MDS infrastructure does.

The Agent is responsible for policy enforcement, execution and gathering of run-time statistics. The agent is available on all Oracle Fusion Middleware Managed Servers and is configured on the same server as the application which it protects. The agent consists of two pieces: the Policy Access Point (PAP) and the Policy Interceptor

This section describes various Oracle Web Services Manager artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The MDS repository is used for storing the policies.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running the soa-infra application. All policies are stored in the MDS repository, so recovering the database and the Managed Server will ensure that the application runs normally.

2.4.7 Recommendations for Oracle User Messaging Service

Oracle User Messaging Service (Oracle UMS) enables two way communications between users and deployed applications. It has support for a variety of channels, such as email, IM, SMS, and text-to-voice messages. Oracle User Messaging Service is integrated with Oracle Fusion Middleware components such as Oracle BPEL PM, Oracle Human Workflow, Oracle BAM and Oracle WebCenter Portal. It is typically deployed along with Oracle SOA Service Infrastructure. Oracle User Messaging Service is made up of UMS Server, UMS Drivers and UMS Client applications.

This section describes various Oracle User Messaging Service artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

JMS Store: The volume containing the file based JMS persistent store. [Table 2–2](#) shows the JMS resources used internally by Oracle User Messaging Service.

Table 2–2 JMS Resources Used by Oracle User Messaging Service

JMS Artifact Name	Type	JNDI Name
OraSDPMAAppDefRcvQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMAAppDefRcvQ1
OraSDPMDriverDefSndQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMDriverDefSndQ1
OraSDPMEngineCmdQ_auto	Distributed queue	OraSDPM/Queues/OraSDPMEngineCmdQ

Table 2–2 (Cont.) JMS Resources Used by Oracle User Messaging Service

JMS Artifact Name	Type	JNDI Name
OraSDPMEngineRcvQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMEngineRcvQ1
OraSDPMEngineSndQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMEngineSndQ1
OraSDPMWSRcvQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMWSRcvQ1

Artifacts in the Database

Oracle User Messaging Service depends on an external database repository to maintain message and configuration state.

Special Considerations

Oracle User Messaging Service 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.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying additional Oracle User Messaging Service drivers, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running the usermessagingserver application. Oracle User Messaging Service maintains message and configuration state in an external database repository along with persisting messages in JMS queues, so recovering the database and the Managed Server ensures that the application functions without any issues. For recommendations on synchronizing JMS data, refer to the "Synchronization Recommendations" subsection in [Section 2.1.1, "Recommendations for Oracle WebLogic Server JMS and T-Logs."](#)

2.4.8 Recommendations for Oracle JCA Adapters

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 broad categories of Oracle JCA Adapters are:

- Oracle Technology Adapters
- Legacy Adapters
- Packaged-Application Adapters
- Oracle Adapter for Oracle Applications

See *Oracle Fusion Middleware User's Guide for Technology Adapters* for additional information about the types of Oracle JCA Adapters.

This section describes various Oracle JCA Adapter artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

Certain adapters by their nature use local or shared-storage files, for example:

- JMS adapters utilizing WebLogic JMS with file-based persistence store: The persistence store must be synchronized with the standby site to resume processing after failover.
- Inbound and outbound files from either File or FTP adapters: The relevant files must be synchronized with the standby site to resume processing after failover.

Adapter configuration is maintained in the `weblogic-ra.xml` deployment descriptor for the ear JCA resource (RAR). The file location of each `weblogic-ra.xml` is determined by the administrator when the file is created, and must be replicated to the standby site.

Artifacts in the Database

Adapter artifacts are generated at design time as part of the composite project. These artifacts are stored along with the rest of the composite definition in the metadata repository.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes (that is, adapter configuration changes) and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running the JCA adapters and the Administration Server.

2.4.9 Recommendations for Oracle Business Activity Monitoring

Oracle Business Activity Monitoring (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 BAM provides the necessary tools and run-time services for creating dashboards that display real-time data inflow and define rules to send alerts under specified conditions.

This section describes various Oracle BAM artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle BAM data and report metadata is stored in the Oracle BAM database that contains Oracle BAM schemas.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database containing the BAM schema and the metadata repository.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running Oracle BAM.

2.4.10 Recommendations for Oracle Business Process Management

The Oracle Business Process Management (BPM) Suite provides an integrated environment for developing, administering, and using business applications centered around business processes. It provides a seamless integration of all stages of the application development life cycle from design-time and implementation to runtime and application management.

The Oracle BPM Suite is layered on the Oracle SOA Suite and shares many of the same product components, including:

- Oracle Business Rules
- Human workflow
- Oracle adapter framework for integration
- SOA Composite Architecture

This section describes various Oracle BPM artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

BPM JMS Persistent Store (BPMJMSFileStore_auto): The file-based JMS persistent store. The persistence store must be synchronized with the standby site to resume processing after failover.

Artifacts in the Database

Process definition, deployed applications, and configuration files are stored in the Metadata Service (MDS) repository. Oracle BPM also uses a separate MDS partition to share projects and project templates between process analysts and process developers.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making domain-related configuration changes and applying patches.

Oracle Data Guard should be configured for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, the standby database is also updated to be up to the same point in time. This is recommended if a snapshot Standby database is used.

Recovery Recommendations

The database must be recovered to the most recent point in time, along with the Managed Server running the soa-infra application.

2.5 Recommendations for Oracle Identity Management

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. These products also enable you 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.

Common Artifacts and Considerations for all Oracle Identity Management Components

The artifacts and considerations below apply to all the Oracle Identity Management components along with the component-specific considerations.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries and an Oracle home containing the Oracle Identity Management binaries.

Oracle_Common_Home: 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: The domain home contains the Administration Server and Managed Server configuration data and Identity Management applications for the domain.

Oracle Instance: The Oracle instance contains the configuration data for non-J2EE Identity Management applications like Oracle Internet Directory and Oracle Virtual Directory. It also has OPMN configuration and Enterprise Manager Agent configuration data.

Artifacts in the Database

The Identity Management schemas are in the Identity Management database.

Network Artifacts

Oracle recommends using the host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

The load balancer virtual hosts used to access the Identity Management components must be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The database containing the Identity Management schemas must be recovered to the most recent point in time, along with the Identity Management component in question.

The rest of this section describes Disaster Recovery recommendations for the following Oracle Identity Management components:

- [Recommendations for Oracle Internet Directory](#)
- [Recommendations for Oracle Virtual Directory](#)
- [Recommendations for Oracle Directory Integration Platform](#)
- [Recommendations for Oracle Identity Federation](#)
- [Recommendations for Oracle Directory Services Manager](#)
- [Recommendations for Oracle Access Manager](#)
- [Recommendations for Oracle Adaptive Access Manager](#)
- [Recommendations for Oracle Identity Manager](#)
- [Recommendations for Oracle Identity Navigator](#)

2.5.1 Recommendations for Oracle Internet Directory

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.

This section describes various Oracle Internet Directory artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The ODS and ODSSM schemas used by Oracle Internet Directory are part of the Identity Management database.

Special Considerations

The load balancer virtual hosts required for the Oracle Internet Directory should be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Oracle Internet Directory must be recovered with the ODS and ODSSM schemas to the most recent point in time.

2.5.2 Recommendations for Oracle Virtual Directory

Oracle Virtual Directory is an LDAP Version 3-enabled service that provides an abstracted view of one or more enterprise data sources. 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.

This section describes various Oracle Virtual Directory artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle Virtual Directory does not have any database dependencies.

Special Considerations

The load balancer virtual hosts required for Oracle Virtual Directory should be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Recovery Recommendations

Recover the Oracle Virtual Directory instance.

2.5.3 Recommendations for Oracle Directory Integration Platform

Oracle Directory Integration Platform is a J2EE application that enables you to synchronize data between other directories or databases 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 meta-directory solutions.

This section describes various Oracle Directory Integration Platform artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The ODS and ODSSM schemas are part of the Identity Management database. Quartz jobs are in the ODSSM schema.

Synchronization Recommendations

The application tier and data tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Recovery Recommendations

Oracle Internet Directory must be recovered with the ODS and ODSSM schemas to the most recent point in time, along with the Managed Server running the Oracle Directory Integration Platform application, and the associated Oracle Internet Directory instances.

2.5.4 Recommendations for Oracle Identity Federation

Oracle Identity Federation enables companies to provide services and share identities across their respective security domains, while providing protection from unauthorized access.

This section describes various Oracle Identity Federation artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle Identity Federation uses the OIF schema, which is part of the Identity Management database. When an RDBMS user store, configuration store, Federation store, message data store, or session store is configured for Oracle Identity Federation, an external database contains those stores.

Special Considerations

Load balancer virtual hosts for Oracle Identity Federation should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories and the data stores.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The databases containing the Oracle Identity Federation schemas and the data stores must be recovered to the most recent point in time, along with the Managed Server running the Oracle Identity Federation application.

2.5.5 Recommendations for Oracle Directory Services Manager

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.

This section describes various Oracle Directory Services Manager artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Not applicable because the Oracle Directory Services Manager application does not have any database dependencies.

Special Considerations

Load balancer virtual hosts for Oracle Directory Services Manager should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the Managed Servers running the Oracle Directory Services Manager application, along with the Administration Server.

2.5.6 Recommendations for Oracle Access Manager

Oracle Access Manager provides a full range of identity administration and security functions that include Web single sign-on; user self-service and self-registration; sophisticated workflow functionality; auditing and access reporting; policy management; dynamic group management; and delegated administration.

This section describes various Oracle Access Manager artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle Access Manager uses the OAM, IAU (Audit Services) schemas, and the ODS schemas (for Oracle Internet Directory) which are part of the Oracle Identity Management database.

LDAP Store

Oracle Access Manager stores configuration and user information in the Oracle Internet Directory.

Special Considerations

The Oracle Access Manager Console is deployed to the Administration Server in the domain. Load balancer virtual hosts for Oracle Access Manager should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories and the data stores.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the Managed Server running the Oracle Access Manager application, the Administration Server and the associated Oracle Internet Directory instances. The OAM, IAU and the ODS schemas must be recovered to the most recent point in time.

2.5.7 Recommendations for Oracle Adaptive Access Manager

Oracle Adaptive Access Manager is the Oracle Identity Management solution for Web-access real-time fraud detection and multifactor online authentication security for the enterprise. Oracle Adaptive Access Manager is designed to support complex, heterogeneous enterprise environments.

This section describes various Oracle Adaptive Access Manager artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle Adaptive Access Manager uses the OAAM, OAAM PARTN, MDS, IAU (Audit Services), and the ODS (Oracle Internet Directory) schemas which are part of the Oracle Identity Management database.

For Oracle Adaptive Access Manager with partition schema support, select the Identity Management - Oracle Adaptive Access Manager (Partition Supp...) schema. By default, the AS Common Schemas - Metadata Services schema is also selected.

LDAP Store

Oracle Adaptive Access Manager stores configuration and user information in the Oracle Internet Directory.

Special Considerations

Load balancer virtual hosts for Oracle Adaptive Access Manager should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories and the data stores.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the managed server running the Oracle Adaptive Access Manager application, and the associated Oracle Internet Directory instances. The OAAM, OAAM_PARTN, MDS, IAU (and the ODS schemas must be recovered to the most recent point in time.

2.5.8 Recommendations 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; and improves regulatory compliance by providing granular reports that identify which users have access to which applications. Oracle Identity Manager is available as a stand-alone product or as part of Oracle's Identity and Access Management Suite.

Oracle Identity Manager uses Oracle SOA for workflow, you must also follow the Oracle SOA disaster recovery recommendations. For more information, see [Recommendations for Oracle SOA Suite](#).

This section describes various Oracle Identity Manager artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

JMS Store: The volume containing the file-based JMS persistent store.

Artifacts in the Database

Oracle Identity Manager uses OIM, SOAINFRA, ORASPDM, and MDS schemas, which are part of the Oracle Identity Management database.

LDAP Store

Oracle Identity Manager does not have any dependency on an external LDAP store when used in the standalone mode. Oracle Identity Manager synchronizes users with and external LDAP store when LDAP Sync is enabled or when integrated with Oracle Access Manager or Oracle Identity Federation.

Special Considerations

Load balancer virtual hosts for Oracle Identity Manager should be configured on both the production and standby sites.

The Oracle Identity Management and SOA Managed Servers are configured to listen on a floating IP addresses, this is required for Server Migration. Ensure that the floating IP addresses are configured with the same Virtual Names on both the production and the standby sites.

The connectors in Oracle Identity Manager are file-based. They are used to provision or reconcile records from different enterprise applications. Ensure that the connectors are available for Oracle Identity Manager and the applications.

Oracle Identity Management is also dependent on the JMS persistence store. For more information, see [Recommendations for Oracle WebLogic Server JMS and T-Logs](#).

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories and the data stores.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

For JMS persistence store, see [Recommendations for Oracle WebLogic Server JMS and T-Logs](#).

Recovery Recommendations

Recover the managed server running the Oracle Identity Manager application, and the associated Oracle Internet Directory instances. The OIM, SOAINFRA, ORASPD, and MDS schemas must be recovered to the most recent point in time. Oracle Identity Management is dependent on the ODS schema when LDAP sync is enabled, in such cases make sure to recover the ODS to the most recent point in time as well

2.5.9 Recommendations for Oracle Identity Navigator

Oracle Identity Navigator is an administrative portal designed to act as an application console for all of the Oracle Identity Management products. It does not replace the individual product consoles. You access Oracle Identity Navigator through a browser and use it to access consoles for Oracle Access Manager, Oracle Adaptive Access Manager, Oracle Identity Manager, Directory Services (ODSM), and other Oracle Identity Management services. You configure Oracle Identity Navigator to connect to these consoles either by specifying the URLs directly, or by employing the product discovery feature. The Oracle Identity Management Navigator simplifies the management of all of the Oracle Identity consoles from one site.

This section describes various Oracle Identity Navigator artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle Identity Navigator does not use any schema.

LDAP Store

Oracle Identity Navigator stores configuration information in an LDAP store.

Special Considerations

Load balancer virtual hosts for Oracle Identity Navigator should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the administration server running the Oracle Identity Navigator application.

2.6 Recommendations for Oracle Portal, Forms, Reports, and Discoverer

This section describes artifacts and Disaster Recovery recommendations for the following suite of Oracle components:

- Oracle Portal
- Oracle Forms
- Oracle Reports
- Oracle Business Intelligence Discoverer (Discoverer)

Common Artifacts and Considerations for Oracle Portal, Forms, Reports, and Discoverer

The artifacts and considerations in this section are common to Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries and an Oracle home containing the binaries for Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer.

Oracle_Common_Home: 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: The domain home contains the Administration Server and Managed Server configuration data and the Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer applications for the domain.

Oracle instance: The Oracle instance contains the configuration data for the Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer components. It also contains configuration data for OPMN and the EM Agent.

Artifacts in the Database

Database metadata repositories containing the Oracle Portal, Oracle Reports, and Discoverer component schemas and any user-configured databases.

There is no Oracle Forms component schema in the database metadata repository (RCU), but there will likely be customer data in user-configured databases that is being accessed through Oracle Forms.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

The load balancer virtual hosts used to access Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer must be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The databases containing the Oracle Portal, Oracle Reports, and Discoverer component schemas must be recovered to the most recent point in time. The Managed Servers running the Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer component applications and the Oracle instances must be restored.

The rest of this section describes Disaster Recovery recommendations for the following components:

- [Recommendations for Oracle Portal](#)
- [Recommendations for Oracle Forms](#)
- [Recommendations for Oracle Reports](#)
- [Recommendations for Oracle Business Intelligence Discoverer](#)

2.6.1 Recommendations for Oracle Portal

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.

This section describes various Oracle Portal artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The Portal, Portal_Demo, Portal_App, Portal_Public, and Portal_Approval schemas are part of the Oracle Portal metadata repository.

Special Considerations

Load balancer virtual hosts for accessing the portal should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The databases containing the Oracle Portal schemas must be recovered to the most recent point in time. The Managed Server running the Oracle Portal application and the Oracle instance must be restored.

2.6.2 Recommendations for Oracle Forms

Oracle Forms is Oracle's long-established technology to design and build enterprise applications quickly and efficiently.

This section describes artifacts that are unique to Oracle Forms and provides recommendations for disaster recovery.

Artifacts in the Database

Any user configured databases for the Oracle Forms applications.

Special Considerations

Load balancer virtual hosts for accessing Oracle Forms should be configured on both the production and standby sites

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for the user configured application databases.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The Managed Server running the Oracle Forms application and the Oracle instance must be restored. If there are any user configured databases, they must be recovered to the most recent point in time.

2.6.3 Recommendations for Oracle Reports

Oracle Reports 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 to publish in both Web-based and non-Web-based environments.

This section describes various Oracle Reports artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

Reports home: This is a user defined location that contains the Report definition files. This may also be under the Oracle instance or Oracle home.

Artifacts in the Database

Oracle Reports can be configured to store job-related information, such as scheduled job data, past job data, or job status data in a database. This is a user-configured database.

Special Considerations

Load balancer virtual hosts for accessing Oracle Reports should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

If the Reports Server is configured to store configuration or user information in other components like Oracle Portal and Oracle Internet Directory, make sure to synchronize these components between the production and standby sites as well.

Oracle Data Guard should be configured for Oracle Reports databases.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The Managed Server running the Oracle Reports application and the Oracle instance must be restored. If there are any user configured databases, they must be recovered to the most recent point in time. Also recover any associated Oracle Portal and Oracle Internet Directory instances.

2.6.4 Recommendations for Oracle Business Intelligence Discoverer

Oracle Business Intelligence Discoverer (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.

This section describes various Discoverer artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The `DISCOVERER` and `DISCOVERER_PS` schemas are part of the Discoverer metadata repository.

Special Considerations

Load balancer virtual hosts for accessing Discoverer should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The database containing the Discoverer schemas must be restored to the most recent point in time.

The Managed Server running the Discoverer application and the Oracle instance must be restored.

2.7 Recommendations for Oracle Web Tier Components

The web tier of a J2EE application server is responsible for interacting with the end users, such as Web browsers primarily in the form of HTTP requests and responses. It is the outermost tier in the HTTP stack, closest to the end user.

This section describes Disaster Recovery recommendations for the following Oracle Web Tier components:

- [Recommendations for Oracle HTTP Server](#)
- [Recommendations for Oracle Web Cache](#)

2.7.1 Recommendations for Oracle HTTP Server

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.

Oracle HTTP Server is based on the Apache 2.2.x 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.

This section describes various Oracle HTTP Server artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

Oracle Home: The Oracle home consists of the Oracle HTTP Server binaries.

Oracle_Common_Home: The Oracle home that contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF).

Oracle Instance: The Oracle instance contains the configuration and diagnostic data for Oracle HTTP Server. It also contains configuration data for OPMN and the Enterprise Manager Agent.

Static Content Volume: This volume contains the static content served by the Oracle HTTP Server instance. This volume that contains static HTML content is optional; Oracle Fusion Middleware can operate normally without it.

Artifacts in the Database

Not applicable because Oracle HTTP Server does not have any database dependencies.

Special Considerations

These are the special considerations for Oracle HTTP Server:

- Load balancer virtual hosts for accessing the Oracle HTTP Server should be configured on both the production and standby sites.
- This manual directs you to install 11g Oracle Fusion Middleware instances (including Oracle HTTP Server instances) on shared storage. When an Oracle instance for Oracle HTTP Server 11g is configured on shared storage (for example, NAS storage, NFS storage, or SAN storage) that does not provide reliable file locking, Oracle HTTP Server may experience performance problems.

Some shared storage systems do not provide the reliable file locking that 11g Oracle HTTP Server requires. In these cases, the `LockFile` directive in the `httpd.conf` file must be changed to point at a local file system.

See the Apache documentation at the following URL for more information about the `LockFile` directive:

http://httpd.apache.org/docs/2.2/mod/mpm_common.html#acceptmutex

If any 11g Oracle HTTP Server instance is installed on shared storage and is experiencing performance problems, perform these steps for the Oracle HTTP Server instance to point the `LockFile` directive at a local file system:

1. By default, the `LockFile` directive is in the following format, configured under both the `prefork` and `worker` MPM configuration blocks in the `httpd.conf` file:

```
${ORACLE_INSTANCE}/diagnostics/logs/${COMPONENT_TYPE}/${COMPONENT_NAME}/http_lock
```

2. Edit the `ORACLE_INSTANCE/config/OHS/<ohs_name>/httpd.conf` file using the appropriate method.
3. Change the `LockFile` directive under the correct MPM configuration to point to a local file system:

```
LockFile /<local_disk>/<path>/http_lock
```

4. Restart the Oracle HTTP Server.

After performing these steps, verify that the `http_lock` file exists in the directory specified by the `LockFile` directive.

Synchronization Recommendations

The Oracle HTTP Server instance must be manually synchronized with the standby site after making configuration changes.

Recovery Recommendations

Restore the Oracle HTTP Server instance and related configuration files on the standby site.

2.7.2 Recommendations for Oracle Web Cache

Oracle Web Cache is a content-aware server accelerator, or a reverse proxy, for the Web tier. It improves the performance, scalability, and availability of Web sites that run on any Web server or application server, such as Oracle HTTP Server and Oracle WebLogic 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 Fusion Middleware by storing frequently accessed URLs in memory.

This section describes various Oracle Web Cache artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

Oracle Home: The Oracle directory home consists of the Oracle Web Cache binaries.

Oracle_Common_Home: The Oracle home that contains the binary and library files required for the Oracle Enterprise Manager Fusion Middleware Control and Java Required Files (JRF).

Oracle Instance: The Oracle instance contains the configuration data for Oracle Web Cache. It also contains configuration data for OPMN and the Enterprise Manager Agent.

Artifacts in the Database

Not applicable because Oracle Web Cache does not have any database dependencies.

Special Considerations

Load balancer virtual hosts for accessing Oracle Web Cache should be configured on both the production and standby sites.

Synchronization Recommendations

The Oracle Web Cache instance must be manually synchronized with the standby site after making configuration changes.

Recovery Recommendations

Restore the Oracle Web Cache instance and related configuration files on the standby site.

2.8 Recommendations for Oracle WebCenter Content

Oracle WebCenter Content is an integrated suite of products designed for managing content. This Oracle WebCenter Content platform 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.

Common Artifacts and Considerations for all Oracle WebCenter Content Components

The artifacts and considerations below apply to all the Oracle WebCenter Content components along with the component-specific considerations.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries and an Oracle home containing the Oracle WebCenter Content binaries.

Oracle_Common_Home: 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: The domain home contains the Administration Server and Managed Server configuration data and Oracle WebCenter Content applications for the domain.

Oracle Instance: The Oracle instance contains the configuration data for non-J2EE Oracle WebCenter Content applications like OPMN configuration and Enterprise Manager Agent configuration data.

Artifacts in the Database

The Oracle WebCenter Content schemas are in the Oracle WebCenter Content database.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

The load balancer virtual hosts used to access the Oracle Enterprise Content Management components must be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby

database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The database containing the Oracle WebCenter Content schemas must be recovered to the most recent point in time, along with the Identity Management component in question.

The rest of this section describes Disaster Recovery recommendations for the following Oracle Enterprise Content Management components:

- [Recommendations for Oracle WebCenter Content](#)
- [Recommendations for Oracle WebCenter Content: Inbound Refinery](#)
- [Recommendations for Oracle WebCenter Content: Imaging](#)
- [Recommendations for Oracle WebCenter Content: Information Rights](#)
- [Recommendations for Oracle WebCenter Content: Records](#)

2.8.1 Recommendations for Oracle WebCenter Content

Oracle WebCenter Content provides a unified application for several different kinds of content management.

Oracle WebCenter Content is an enterprise content management platform that enables you to leverage document management, Web content management, digital asset management, and records retention functionality to build and complement 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 for centralized management. Through user-friendly interfaces, roles-based authentication and security models, Oracle WebCenter Content empowers users throughout the enterprise to view, collaborate on or retire content, ensuring that all accessible distributed or published information is secure, accurate and up-to-date.

This section describes various Oracle WebCenter Content artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

Oracle Secure Files or File-based persistent stores.

Artifacts in the Database

The OCS schema is part of the Oracle WebCenter Content database.

Special Considerations

The load balancer virtual hosts required for the Oracle WebCenter Content should be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

For File-based persistent stores, you must synchronize the File-based persistent stores on the standby site.

Recovery Recommendations

Oracle WebCenter Content must be recovered with the OCS and MDS schemas to the most recent point in time.

2.8.2 Recommendations for Oracle WebCenter Content: Inbound Refinery

Oracle WebCenter Content: Inbound Refinery is a conversion server that manages file conversions for electronic assets such as documents, digital images, and motion video. In addition to conversion, Oracle WebCenter Content: Inbound Refinery provides thumbnail functionality for documents and images, storyboarding for video, and 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. You can use Oracle WebCenter Content: Inbound Refinery to convert content items stored in Oracle Content Server.

This section describes various Oracle WebCenter Content: Inbound Refinery artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

Oracle WebCenter Content: Inbound Refinery does not have any database dependencies.

Special Considerations

The load balancer virtual hosts required for Oracle WebCenter Content: Inbound Refinery should be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Recovery Recommendations

Recover the Oracle WebCenter Content: Inbound Refinery instance.

2.8.3 Recommendations for Oracle WebCenter Content: Imaging

Oracle WebCenter Content: Imaging provides organizations with a scalable solution upon which to develop process-oriented imaging applications and image-enablement solutions for enterprise applications. Oracle WebCenter Content: Imaging enables image capture via Oracle Document Capture and Oracle Distributed Document Capture, annotation and markup of images, workflow support for routing and approval automation, and support for high-volume applications for millions of items. With Oracle WebCenter Content: Imaging, organizations can quickly integrate their content and processes directly with Oracle enterprise applications, such as Oracle E-Business Suite, PeopleSoft Enterprise, and JD Edwards EnterpriseOne. Users benefit

by having a single source for all transaction-based content, eliminating the need for double entry.

This section describes various Oracle WebCenter Content: Imaging artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

The IPM Input Agent pulls files from a common file share. This IPM file share must be synchronized with the standby site.

Artifacts in the Database

The IPM schema is part of the Oracle WebCenter Content: Imaging database. Oracle WebCenter Content: Imaging also requires OCS schema, to use Oracle WebCenter Content as the Oracle WebCenter Content: Imaging repository.

Synchronization Recommendations

The application tier and data tier must be manually synchronized with the standby site after making configuration changes and applying patches.

The IPMJMSStore which is the JMS store must be synchronized with a file-based persistent store. For more information, see [Section 2.1.1, "Recommendations for Oracle WebLogic Server JMS and T-Logs"](#).

Recovery Recommendations

Oracle WebCenter Content: Imaging must be recovered with the IPM schema to the most recent point in time, along with the Managed Server running the Oracle WebCenter Content: Imaging application, and the associated instances.

2.8.4 Recommendations for Oracle WebCenter Content: Information Rights

Oracle WebCenter Content: Information Rights distributes rights management between centralized servers and desktop agents.

Oracle WebCenter Content: Information Rights enables documents or emails to be automatically or manually sealed at any stage in their lifecycle, using sealing tools integrated into the Microsoft Windows desktop, authoring applications, email clients, and content management and collaborative repositories. Sealing wraps documents and emails within a layer of strong encryption and digital signatures, together with indelible links back to network-hosted servers (operated by the organization to which the information belongs) that store the decryption keys and associated access rights.

Sealed documents and emails can be distributed by any existing means, including email, the Web, and file sharing.

This section describes various Oracle WebCenter Content: Information Rights artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The ORAIRM schema is part of the Oracle WebCenter Content: Information Rights database. In addition to the database, when you install Oracle WebCenter Content: Information Rights a key store is created that is used by the Oracle IRM J2EE application. It is recommended that the key store is located in a directory under the domain home, so when the domain is restored the key store is automatically available. If the key store is not stored under the domain, then you must manually copy the key store to the replicated system.

Special Considerations

Load balancer virtual hosts for Oracle Identity Federation should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories and the data stores.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

The databases containing the Oracle WebCenter Content: Information Rights schemas and the data stores must be recovered to the most recent point in time, along with the Managed Server running the Oracle WebCenter Content: Information Rights application.

2.8.5 Recommendations for Oracle WebCenter Content: Records

Oracle WebCenter Content: Records effectively manages content items on a retention schedule, which determines the life cycle of that content item.

The focus of records management tends to be the preservation of content for historical, legal, or archival purposes while also performing retention management functions. The focus of retention management tends to be the scheduled elimination of content in which the costs of retaining content outweighs the value of keeping it.

Oracle WebCenter Content: Records combines both record and retention management into one software system. You can use Oracle WebCenter Content: Records to track and to preserve content as needed, or to dispose of content when it is no longer required.

This section describes various Oracle WebCenter Content: Records artifacts and provides recommendations for disaster recovery.

Artifacts in the Database

The URMSERVER and MDS schemas are part of the Oracle WebCenter Content: Records database.

Special Considerations

Load balancer virtual hosts for Oracle WebCenter Content: Records should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs

automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the Managed Servers running the Oracle WebCenter Content: Records application, along with the Administration Server.

2.9 Recommendations for Oracle Business Intelligence

Oracle Business Intelligence (BI) is a portfolio of technology and applications that provides the industry's first integrated, end-to-end Enterprise Performance Management System, including BI foundation and tools as well as financial performance management applications, operational BI applications, and data warehousing.

Common Artifacts and Considerations for all Oracle Business Intelligence Components

The artifacts and considerations below apply to all the Oracle Business Intelligence components along with the component-specific considerations.

Artifacts on the File System

MW_HOME: The Middleware home consists of a WebLogic home that has the WebLogic Server binaries and an Oracle home containing the Oracle Business Intelligence binaries.

Oracle_Common_Home: 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: The domain home contains the Administration Server and Managed Server configuration data and Oracle Business Intelligence applications for the domain.

Oracle Instance: The Oracle instance contains configuration files, log files, and temporary files for one or more Oracle system components (such as Oracle BI Server, Oracle BI Presentation Services, Oracle HTTP Server, and so on).

Artifacts in the Database

The Oracle Business Intelligence schemas are in the Oracle Business Intelligence database.

Network Artifacts

Oracle recommends using the virtual host name as the listen address for both the Oracle WebLogic Administration Server and Managed Server. As long as this host name can be resolved on both the production and standby site, there is no need to update this value after a Disaster Recovery operation.

The load balancer virtual hosts used to access the Oracle Business Intelligence components must be configured on both the production and standby sites.

Synchronization Recommendations

The directory tier must be manually synchronized with the standby site after making configuration changes, deploying applications, and applying patches.

Oracle Data Guard should be configured for Oracle database metadata repositories.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database. For more information, see [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) and "Applying Redo Data to Physical Standby Databases" in the *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

The database containing the Oracle Business Intelligence schemas must be recovered to the most recent point in time, along with the Identity Management component in question.

The rest of this section describes Disaster Recovery recommendations for the following Oracle Business Intelligence components:

- [Recommendations for Oracle Business Intelligence Enterprise Edition \(EE\)](#)
- [Recommendations for Oracle Business Intelligence Publisher](#)
- [Recommendations for Oracle Real-Time Decisions](#)

2.9.1 Recommendations for Oracle Business Intelligence Enterprise Edition (EE)

Oracle Business Intelligence Enterprise Edition (Oracle BI EE) is a comprehensive set of enterprise business intelligence tools and infrastructure, including a scalable and efficient query and analysis server, an ad-hoc query and analysis tool, interactive dashboards, proactive intelligence and alerts, real-time predictive intelligence, and an enterprise reporting engine. Oracle Business Intelligence Enterprise Edition is designed to bring greater business visibility and insight to a wide variety of users.

Artifacts in the Database

BIPLATFORM schema is part of the Oracle Business Intelligence Enterprise Edition (EE) database.

Special Considerations

Load balancer virtual hosts for Oracle Business Intelligence Enterprise Edition (EE) should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the Middleware home, the domain, and the Oracle instance containing the Oracle Business Intelligence Enterprise Edition (EE) components. On Windows,

import Oracle Business Intelligence Enterprise Edition (EE) Registry entries. Recover the database to the most recent point in time, if needed.

2.9.2 Recommendations for Oracle Business Intelligence Publisher

Oracle Business Intelligence Publisher (BI Publisher, formerly XML Publisher) allows you to create highly formatted reports that are suitable for printing. BI Publisher reports are built on top of BI Publisher data models. A BI Publisher data model can consist of data sets from a wide range of sources, such as subject areas from the BI Server or analyses, SQL queries against relational data bases, MDX queries against Subbase or other OLAP sources, LDAP, Web Services, Microsoft Excel, HTTP feeds, or XML files. BI Publisher supports a wide range of layout types, so you can create the full range of documents that your organization might need. Within Oracle BI EE, you can view, create, edit, and schedule BI Publisher reports and then include them in dashboard pages.

Artifacts in the Database

BIPLATFORM schema is a part of the Oracle Business Intelligence Publisher database.

Special Considerations

Load balancer virtual hosts for Oracle Business Intelligence Publisher should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the Managed Server containing the Oracle Business Intelligence Publisher component.

2.9.3 Recommendations for Oracle Real-Time Decisions

Oracle Real-Time Decisions (Oracle RTD) enables you to develop adaptive enterprise software solutions. These adaptive solutions continuously learn from business transactions while they execute and automatically optimize each transaction, in real time, by way of close loop business rules and predictive models.

Artifacts in the Database

Analytic models and the RTD schema are part of the Oracle Real-Time Decisions database.

Special Considerations

Load balancer virtual hosts for Oracle Real-Time Decisions should be configured on both the production and standby sites.

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after making configuration changes and applying patches.

It is recommended that the standby database be synchronized when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.

Recovery Recommendations

Recover the Managed Server containing the Oracle Real-Time Decisions component.

Design Considerations

This chapter describes design considerations for an Oracle Fusion Middleware Disaster Recovery solution for an enterprise deployment.

It contains the following topics:


- [Network Considerations](#)
- [Storage Considerations](#)
- [Database Considerations](#)
- [Starting Points](#)
- [Topology Considerations](#)

This chapter describes detailed instructions for setting up an Oracle Fusion Middleware 11g Disaster Recovery production site and standby site for the Linux and UNIX operating systems. It primarily uses the Oracle SOA Suite enterprise deployment shown in [Figure 3-1](#) in the examples of how to set up the Oracle Fusion Middleware 11g Disaster Recovery solution for an enterprise deployment. After you understand how to set up Disaster Recovery for the Oracle SOA Suite enterprise topology, you can use the information for other 11g enterprise deployments in this chapter to set up Disaster Recovery for those deployments as well.

Note: This chapter describes an Oracle Fusion Middleware 11g Disaster Recovery symmetric topology that uses the Oracle SOA Suite enterprise deployment shown in [Figure 3-1](#) at *both* the production site and the standby site. [Figure 3-1](#) shows the deployment for only one site; the high level of detail shown for this deployment precludes showing the deployment for both sites in a single figure.

[Figure 1-1](#) shows a Disaster Recovery symmetric production site and standby site in a single figure.

Figure 3–1 *Deployment Used at Production and Standby Sites for Oracle Fusion Middleware Disaster Recovery*



[Figure 3–1](#) shows the mySOACompany with Oracle Access Manager enterprise deployment from the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*. See the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for detailed information on installing and configuring this Oracle SOA Suite enterprise deployment.

The Oracle Fusion Middleware Disaster Recovery topology that you design must be symmetric for the following at the production site and standby site.

- **Directory names and paths**

Every file that exists at a production site host must exist in the same directory and path at the standby site peer host.

Thus, Oracle home names and directory paths must be the same at the production site and standby site.
- **Port numbers**

Port numbers are used by listeners and for the routing of requests. Port numbers are stored in the configuration and have to be the same at the production site hosts and their standby site peer hosts.

[Section 3.4.1, "Starting with an Existing Site"](#) describes how to check for port conflicts between production site and standby site hosts.
- **Security**

The same user accounts must exist at both the production site and standby site. Also, the file system, SSL, and Single Sign-On must be configured identically at the production site and standby site. For example, if the production site uses SSL, the standby site must also use SSL that is configured in exactly the same way as the production site.
- **Load balancers and virtual server names**

A front-end load balancer should be set up with virtual server names for the production site, and an identical front-end load balancer should be set up with the same virtual server names for the standby site.
- **Software**

The same versions of software must be used on the production site and standby site. Also, the operating system patch level must be the same at both sites, and patches to Oracle or third party software must be made to both the production site and standby site.

3.1 Network Considerations

This section describes the following network considerations:

- [Planning Host Names](#)
- [Virtual IP and Virtual Hostname Considerations](#)
- [Load Balancers Considerations](#)
- [Virtual Server Consideration](#)
- [Wide Area DNS Operations](#)

3.1.1 Planning Host Names

In a Disaster Recovery topology, the production site host names must be resolvable to the IP addresses of the corresponding peer systems at the standby site. Therefore, it is important to plan the host names for the production site and standby site. After failover from a primary site to a standby site, the alias host name for the middle tier host on the standby site become active. You do not need to reconfigure hostname for the host on the standby site, if you setup alias for the standby site.

Creating aliases for physical host names is required only when using a single Global DNS Server to resolve host names.

This section describes how to plan physical host names and alias host names for the middle tier hosts that use the Oracle Fusion Middleware instances at the production site and standby site. It uses the Oracle SOA Suite enterprise deployment shown in [Figure 3–1](#) for the host name examples. The host name examples in this section assume that a symmetric Disaster Recovery site is being set up, where the production site and standby site have the same number of hosts. Each host at the production site and standby site has a peer host at the other site. The peer hosts are configured the same, for example, using the same ports as their counterparts at the other site.

When configuring each component, use hostname-based configuration instead of IP-based configuration, unless the component requires you to use IP-based configuration. For example, if you are configuring the listen address of an Oracle Fusion Middleware component to a specific IP address such as 123.1.2.113, use the host name SOAHOST1.MYCOMPANY.COM, which resolves to 123.1.2.113.

The following subsections show how to set up host names at the Disaster Recovery production site and standby site for the following enterprise deployments:

- [Host Names for the Oracle SOA Suite Production Site and Standby Site Hosts](#)
- [Host Names for the Oracle WebCenter Portal Production Site and Standby Site Hosts](#)
- [Host Names for the Oracle Identity Management Production Site and Standby Site Hosts](#)
- [Host Names for the Oracle Portal, Forms, Reports, and Discoverer Production Site and Standby Site Hosts](#)
- [Host Names for the Oracle WebCenter Content Production Site and Standby Site Hosts](#)
- [Host Names for the Oracle Business Intelligence Production Site and Standby Site Hosts](#)

Note: In this book's examples, IP addresses for hosts at the initial production site have the format 123.1.x.x and IP addresses for hosts at the initial standby site have the format 123.2.x.x.

Host Names for the Oracle SOA Suite Production Site and Standby Site Hosts

[Table 3–1](#) shows the IP addresses and physical host names that will be used for the Oracle SOA Suite EDG deployment production site hosts. [Figure 3–1](#) shows the configuration for the Oracle SOA Suite EDG deployment at the production site.

Table 3–1 IP Addresses and Physical Host Names for SOA Suite Production Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.1.2.111	WEBHOST1	None
123.1.2.112	WEBHOST2	None
123.1.2.113	SOAHOST1	None
123.1.2.114	SOAHOST2	None

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

[Table 3–2](#) shows the IP addresses, physical host names, and alias host names that will be used for the Oracle SOA Suite EDG deployment standby site hosts. [Figure 3–2](#) shows the physical host names used for the Oracle SOA Suite EDG deployment at the standby site. The alias host names shown in [Table 3–2](#) should be defined for the SOA Oracle Suite standby site hosts in [Figure 3–2](#).

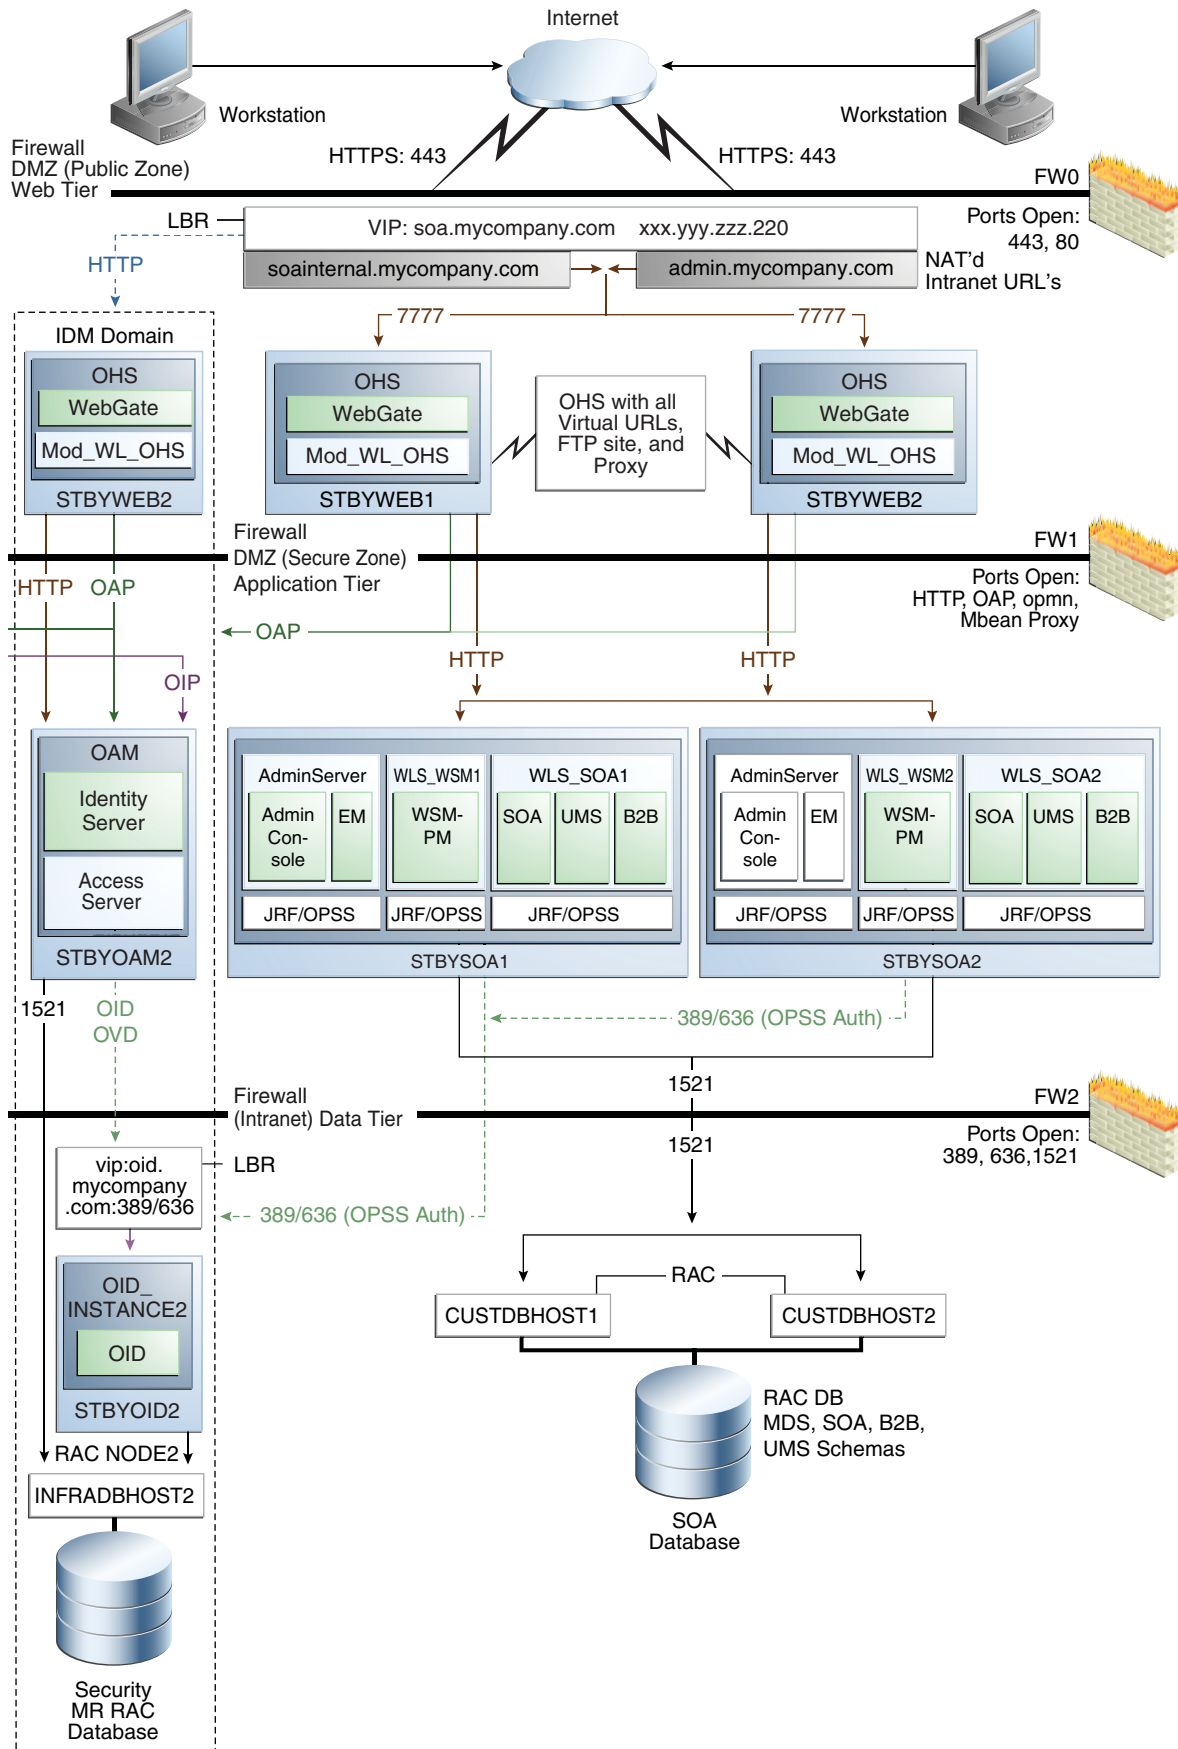
Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in [Table 3–2](#). See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) for more information about using separate DNS servers to resolve host names.

Table 3–2 IP Addresses, Physical Host Names, and Alias Host Names for SOA Suite Standby Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.2.2.111	STBYWEB1	WEBHOST1
123.2.2.112	STBYWEB2	WEBHOST2
123.2.2.113	STBYSOA1	SOAHOST1
123.2.2.114	STBYSOA2	SOAHOST2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

Figure 3-2 *Physical Host Names Used at Oracle SOA Suite Deployment Standby Site*



Host Names for the Oracle WebCenter Portal Production Site and Standby Site Hosts

Table 3–3 shows the IP addresses and physical host names that will be used for the Oracle WebCenter Portal EDG deployment production site hosts. Figure 4–4 shows the configuration for the Oracle WebCenter Portal EDG deployment at the production site.

Table 3–3 IP Addresses and Physical Host Names for Oracle WebCenter Portal Production Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.1.2.111	WEBHOST1	None
123.1.2.112	WEBHOST2	None
123.1.2.113	SOAHOST1	None
123.1.2.114	SOAHOST2	None
123.1.2.115	WCHOST1	None
123.1.2.116	WCHOST2	None

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Table 3–4 shows the IP addresses, physical host names, and alias host names that will be used for the Oracle WebCenter Portal EDG deployment standby site hosts. Figure 4–4 shows the configuration for the Oracle WebCenter Portal EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in Table 3–4. See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" for more information about using separate DNS servers to resolve host names.

Table 3–4 IP Addresses, Physical Host Names, and Alias Host Names for Oracle WebCenter Portal Standby Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.2.2.111	STBYWEB1	WEBHOST1
123.2.2.112	STBYWEB2	WEBHOST2
123.2.2.113	STBYSOA1	SOAHOST1
123.2.2.114	STBYSOA2	SOAHOST2
123.2.2.115	STBYWC1	WCHOST1
123.2.2.116	STBYWC2	WCHOST2

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Host Names for the Oracle Identity Management Production Site and Standby Site Hosts

Table 3–5 shows the IP addresses and physical host names that will be used for the Oracle Identity Management EDG deployment production site hosts. Figure 4–6 shows the configuration for the Oracle Identity Management EDG deployment at the production site.

Table 3–5 IP Addresses and Physical Host Names for Identity Management Production Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.1.2.111	WEBHOST1	None
123.1.2.112	WEBHOST2	None
123.1.2.118	IDMHOST1	None
123.1.2.119	IDMHOST2	None
123.1.2.122	OIDHOST1	None
123.1.2.123	OIDHOST2	None
123.1.2.124	OVDHOST1	None
123.1.2.125	OVDHOST2	None
123.1.2.126	OIFHOST1	None
123.1.2.127	OIFHOST2	None
123.1.2.128	OAMHOST1	None
123.1.2.129	OAMHOST2	None
123.1.2.130	OAAMHOST1	None
123.1.2.131	OAAMHOST2	None
123.1.2.132	OIMHOST1	None
123.1.2.133	OIMHOST2	None

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Table 3–6 shows the IP addresses, physical host names, and alias host names that will be used for the Oracle Identity Management EDG deployment standby site hosts. Figure 4–6 shows the configuration for the Oracle Identity Management EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in Table 3–6. See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" for more information about using separate DNS servers to resolve host names.

Table 3–6 IP Addresses, Physical Host Names, and Alias Host Names for Identity Management Standby Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.2.2.111	STBYWEB1	WEBHOST1
123.2.2.112	STBYWEB2	WEBHOST2
123.2.2.117	STBYADM	OAMADMINHOST
123.2.2.118	STBYIDM1	IDMHOST1
123.2.2.119	STBYIDM2	IDMHOST2
123.2.2.122	STBYOID1	OIDHOST1
123.2.2.123	STBYOID2	OIDHOST2
123.2.2.124	STBYOVD1	OVDHOST1
123.2.2.125	STBYOVD2	OVDHOST2
123.2.2.126	STBYOIF1	OIFHOST1
123.2.2.127	STBYOIF2	OIFHOST2
123.2.2.128	STBYOAM1	OAAMHOST1
123.2.2.129	STBYOAM2	OAAMHOST2
123.2.2.130	STBYOAAM1	OAAMHOST1
123.2.2.131	STBYOAAM2	OAAMHOST2
123.2.2.132	STBYOIM1	OIMHOST1
123.2.2.133	STBYOIM2	OIMHOST2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

The Administration Server, the Oracle Identity Manager Managed Servers and the SOA Managed Servers require a floating IP addresses to be provisioned on each site ([Table 3–7](#)). Ensure that you provision the floating IP addresses with the same virtual host names on production site and the standby site.

Table 3–7 Floating IP Addresses

Physical Host Name	Virtual Host Name	Floating IP
AdminServer	ADMINVHN	122.1.2.134
OIMHOST1	OIMVHN1	122.1.2.135
OIMHOST2	OIMVHN2	122.1.2.136
SOAHOST1	SOAVHN1	122.1.2.137
SOAHOST2	SOAVHN2	122.1.2.138

Note: For more information, see the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management*.

Host Names for the Oracle Portal, Forms, Reports, and Discoverer Production Site and Standby Site Hosts

Table 3–8 shows the IP addresses and physical host names that will be used for the Oracle Portal, Forms, Reports, and Discoverer enterprise deployment production site hosts. Figure 4–9 shows the configuration for the Oracle Portal enterprise deployment at the production site and Figure 4–10 shows the configuration for the Oracle Forms, Reports, and Discoverer enterprise deployment at the production site.

Table 3–8 IP Addresses and Physical Host Names for Oracle Portal, Forms, Reports, and Discoverer Production Site Hosts

IP Address	Physical Host Names ¹	Alias Host Names
123.1.2.111	WEBHOST1	None
123.1.2.112	WEBHOST2	None
123.1.2.126	APPHOST1	None
123.1.2.127	APPHOST2	None

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Table 3–9 shows the IP addresses, physical host names, and alias host names that will be used for the Oracle Portal, Forms, Reports, and Discoverer enterprise deployment standby site hosts. Figure 4–9 shows the configuration for the Oracle Portal enterprise deployment at the production site and Figure 4–10 shows the configuration for the Oracle Forms, Reports, and Discoverer enterprise deployment at the production site.

Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in Table 3–9. See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" for more information about using separate DNS servers to resolve host names.

Table 3–9 IP Addresses, Physical Host Names, and Alias Host Names for Oracle Portal, Forms, Reports, and Discoverer Standby Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.2.2.111	STBYWEB1	WEBHOST1
123.2.2.112	STBYWEB2	WEBHOST1
123.2.2.126	STBYAPP1	APPHOST1
123.2.2.127	STBYAPP2	APPHOST2

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

The alias host names in Table 3–2, Table 3–4, Table 3–6, and Table 3–8 are resolved locally at the standby site to the correct IP address. Section 3.1.1.1, "Host Name Resolution" describes two ways to configure host name resolution in an Oracle Fusion Middleware Disaster Recovery topology.

Host Names for the Oracle WebCenter Content Production Site and Standby Site Hosts

Table 3–10 shows the IP addresses and physical host names that will be used for the Oracle WebCenter Content EDG deployment production site hosts. Figure 4–11 shows the configuration for the Oracle WebCenter Content EDG deployment at the production site.

Table 3–10 IP Addresses and Physical Host Names for Oracle WebCenter Content Production Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.1.2.111	WEBHOST1	None
123.1.2.112	WEBHOST2	None
123.1.2.113	WCCHOST1	None
123.1.2.114	WCCHOST1	None

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Table 3–11 shows the IP addresses, physical host names, and alias host names that will be used for the Oracle WebCenter Content EDG deployment standby site hosts. Figure 4–11 shows the configuration for the Oracle WebCenter Content EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in Table 3–11. See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" for more information about using separate DNS servers to resolve host names.

Table 3–11 IP Addresses, Physical Host Names, and Alias Host Names for Oracle WebCenter Content Standby Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.2.2.111	STBYWEB1	WEBHOST1
123.2.2.112	STBYWEB2	WEBHOST2
123.2.2.113	STBYECM1	WCCHOST1
123.2.2.114	STBYECM2	WCCHOST2

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Host Names for the Oracle Business Intelligence Production Site and Standby Site Hosts

Table 3–12 shows the IP addresses and physical host names that will be used for the Oracle Business Intelligence EDG deployment production site hosts. Figure 4–11 shows the configuration for the Oracle Business Intelligence EDG deployment at the production site.

Table 3–12 IP Addresses and Physical Host Names for Oracle Business Intelligence Production Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.1.2.111	WEBHOST1	None
123.1.2.112	WEBHOST2	None
123.1.2.113	BIHOST1	None
123.1.2.114	BIHOST2	None

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

[Table 3–13](#) shows the IP addresses, physical host names, and alias host names that will be used for the Oracle Business Intelligence EDG deployment standby site hosts. [Figure 4–11](#) shows the configuration for the Oracle Business Intelligence EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in [Table 3–11](#). See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) for more information about using separate DNS servers to resolve host names.

Table 3–13 IP Addresses, Physical Host Names, and Alias Host Names for Oracle Business Intelligence Standby Site Hosts

IP Address	Physical Host Name ¹	Alias Host Name
123.2.2.111	STBYWEB1	WEBHOST1
123.2.2.112	STBYWEB2	WEBHOST2
123.2.2.113	STBYBI1	BIHOST1
123.2.2.114	STBYBI2	BIHOST2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

3.1.1.1 Host Name Resolution

Host name resolution is the process of resolving a host name to the proper IP address for communication. Host name resolution can be configured in one of the following ways:

- Resolving host names locally
 - Local host name resolution uses the host name to IP address mapping that is specified by the `/etc/hosts` file on each host.
 - See [Section 3.1.1.2, "Resolving Host Names Locally"](#) for more information about using the `/etc/hosts` file to implement local host name file resolution.
- Resolving host names using DNS

A DNS Server is a dedicated server or a service that provides DNS name resolution in an IP network.

See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for more information about two methods of implementing DNS server host name resolution.

You must determine the method of host name resolution you will use for your Oracle Fusion Middleware Disaster Recovery topology when you are planning the deployment of the topology. Most site administrators use a combination of these resolution methods in a precedence order to manage host names.

The Oracle Fusion Middleware hosts and the shared storage system for each site must be able to communicate with each other.

Host Name Resolution Precedence

To determine the host name resolution method used by a particular host, search for the value of the `hosts` parameter in the `/etc/nsswitch.conf` file on the host.

As shown in [Example 3–1](#), make the `files` entry the first entry for the `hosts` parameter if you want to resolve host names locally on the host. When `files` is the first entry for the `hosts` parameter, entries in the host's `/etc/hosts` file will be used first to resolve host names:

Example 3–1 Specifying the Use of Local Host Name Resolution

```
hosts:  files  dns  nis
```

As shown in [Example 3–2](#), make the `dns` entry the first entry for the `hosts` parameter if you want to resolve host names using DNS on the host. When `dns` is the first entry for the `hosts` parameter, DNS server entries will be used first to resolve host names:

Example 3–2 Specifying the Use of DNS Host Name Resolution

```
hosts:  dns   files  nis
```

For simplicity and consistency, it is recommended that all the hosts within a site (production site or standby site) use the same host name resolution method (resolving host names locally or resolving host names using separate DNS servers or a global DNS server).

The recommendations in the following sections are high-level recommendations that you can adapt to meet the host name resolution standards used by your enterprise.

3.1.1.2 Resolving Host Names Locally

Local host name resolution uses the host name to IP mapping defined in the `/etc/hosts` file of a host. When you use this method to resolve host names for your Disaster Recovery topology, the following guidelines apply:

1. Ensure that the `hosts` parameter in the `/etc/nsswitch.conf` file on all the production site and standby site hosts looks like this:

```
hosts:  files  dns  nis
```

2. The `/etc/hosts` file entries on the hosts of the production site should have their physical host names mapped to their IP addresses. For the sake of simplicity and ease of maintenance, it is recommended to have the same entries on all the hosts of

the production site. [Example 3–3](#) shows the `/etc/hosts` file for the production site of a SOA Enterprise Deployment topology:

Example 3–3 Making `/etc/hosts` File Entries for a Production Site Host

```
127.0.0.1      localhost.localdomain  localhost
123.1.2.111   WEBHOST1.MYCOMPANY.COM WEBHOST1
123.1.2.112   WEBHOST2.MYCOMPANY.COM WEBHOST2
123.1.2.113   SOAHOST1.MYCOMPANY.COM SOAHOST1
123.1.2.114   SOAHOST2.MYCOMPANY.COM SOAHOST2
```

3. The `/etc/hosts` file entries on the hosts of the standby site should have their physical host names mapped to their IP addresses along with the physical host names of their corresponding peer on the production site defined as the alias host names. For the sake of simplicity and ease of maintenance, it is recommended to have the same entries on all the hosts of the standby site. [Example 3–4](#) shows the `/etc/hosts` file for the production site of a SOA Enterprise Deployment topology:

Example 3–4 Making `/etc/hosts` File Entries for a Standby Site Host

```
127.0.0.1      localhost.localdomain  localhost
123.2.2.111   STBYWEB1.MYCOMPANY.COM WEBHOST1
123.2.2.112   STBYWEB2.MYCOMPANY.COM WEBHOST2
123.2.2.113   STBYSOA1.MYCOMPANY.COM SOAHOST1
123.2.2.114   STBYSOA2.MYCOMPANY.COM SOAHOST2
```

4. After setting up host name resolution using `/etc/host` file entries, use the `ping` command to test host name resolution. For a system configured with static IP addressing and the `/etc/hosts` file entries shown in [Example 3–3](#), a `ping webhost1` command on the production site would return the correct IP address (123.1.2.111) and also indicate that the host name is fully qualified.
5. Similarly, for a system configured with static IP addressing and the `/etc/hosts` file entries shown in [Example 3–4](#), a `ping webhost1` command on the standby site will return the correct IP address (123.2.2.111) and it will also show that the name `WEBHOST1` is also associated with that IP address.

3.1.1.3 Resolving Host Names Using Separate DNS Servers

This manual uses the term "separate DNS servers" to refer to a Disaster Recovery topology where the production site and the standby site have their own DNS servers. When you use separate DNS servers to resolve host names for your Disaster Recovery topology, the following guidelines apply:

1. Ensure that the `hosts` parameter in the `/etc/nsswitch.conf` file on all the production site and standby site hosts looks like this:

```
hosts:  dns  files  nis
```
2. The DNS servers on the production site and standby site must not be aware of each other and must contain entries for host names used within their own site.
3. The DNS server entries on the production site should have the physical host names mapped to their IP addresses. [Example 3–5](#) shows the DNS server entries for the production site of a SOA Enterprise Deployment topology:

Example 3–5 DNS Entries for a Production Site Host in a Separate DNS Servers Configuration

```

WEBHOST1.MYCOMPANY.COM    IN    A    123.1.2.111
WEBHOST2.MYCOMPANY.COM    IN    A    123.1.2.112
SOAHOST1.MYCOMPANY.COM    IN    A    123.1.2.113
SOAHOST2.MYCOMPANY.COM    IN    A    123.1.2.114

```

4. The DNS server entries on the standby site should have the physical host names of the production site mapped to their IP addresses. [Example 3–6](#) shows the DNS server entries for the standby site of a SOA Enterprise Deployment topology:

Example 3–6 DNS Entries for a Standby Site Host in a Separate DNS Servers Configuration

```

WEBHOST1.MYCOMPANY.COM    IN    A    123.2.2.111
WEBHOST2.MYCOMPANY.COM    IN    A    123.2.2.112
SOAHOST1.MYCOMPANY.COM    IN    A    123.2.2.113
SOAHOST2.MYCOMPANY.COM    IN    A    123.2.2.114

```

5. Make sure there are no entries in the `/etc/hosts` file for any host at the production site or standby site.
6. Test the host name resolution using the `ping` command. For a system configured with the production site DNS entries shown in [Example 3–5](#), a `ping webhost1` command on the production site would return the correct IP address (123.1.2.111) and also indicate that the host name is fully qualified.
7. Similarly, for a system configured with the standby site DNS entries shown in [Example 3–6](#), a `ping webhost1` command on the standby site will return the correct IP address (123.2.2.111) and it will also indicate that the host name is fully qualified.

3.1.1.4 Resolving Host Names Using a Global DNS Server

This manual uses the term "global DNS server" to refer to a Disaster Recovery topology where a single DNS server is used for both the production site and the standby site. When you use a global DNS server to resolve host names for your Disaster Recovery topology, the following guidelines apply:

1. When using a global DNS server, for the sake of simplicity, a combination of local host name resolution and DNS host name resolution is recommended.
2. In this example, it is assumed that the production site uses DNS host name resolution and the standby site uses local host name resolution.
3. The global DNS server should have the entries for both the production and standby site hosts. [Example 3–7](#) shows the entries for a SOA Enterprise Deployment topology:

Example 3–7 DNS Entries for Production Site and Standby Site Hosts When Using a Global DNS Server Configuration

```

WEBHOST1.MYCOMPANY.COM    IN    A    123.1.2.111
WEBHOST2.MYCOMPANY.COM    IN    A    123.1.2.112
SOAHOST1.MYCOMPANY.COM    IN    A    123.1.2.113
SOAHOST2.MYCOMPANY.COM    IN    A    123.1.2.114
STBYWEB1.MYCOMPANY.COM    IN    A    123.2.2.111
STBYWEB2.MYCOMPANY.COM    IN    A    123.2.2.112
STBYSOA1.MYCOMPANY.COM    IN    A    123.2.2.113
STBYSOA2.MYCOMPANY.COM    IN    A    123.2.2.114

```

4. Ensure that the `hosts` parameter in the `/etc/nsswitch.conf` file on all the production site hosts looks like this:

```
hosts:  dns  files  nis
```

5. Ensure that the `hosts` parameter in the `/etc/nsswitch.conf` file on all the standby site hosts looks like this:

```
hosts:  files  dns  nis
```

6. The `/etc/hosts` file entries on the hosts of the standby site should have their physical host names mapped to their IP addresses along with the physical host names of their corresponding peer on the production site defined as the alias host names. For the sake of simplicity and ease of maintenance, it is recommended to have the same entries on all the hosts of the standby site. [Example 3–8](#) shows the `/etc/hosts` file for the production site of a SOA Enterprise Deployment topology:

Example 3–8 Standby Site `/etc/hosts` File Entries When Using a Global DNS Server Configuration

```
127.0.0.1    localhost.localdomain  localhost
123.2.2.111  STBYWEB1.MYCOMPANY.COM  WEBHOST1
123.2.2.112  STBYWEB2.MYCOMPANY.COM  WEBHOST2
123.2.2.113  STBYSOA1.MYCOMPANY.COM  SOAHOST1
123.2.2.114  STBYSOA2.MYCOMPANY.COM  SOAHOST2
```

7. Test the host name resolution using the `ping` command. A `ping webhost1` command on the production site would return the correct IP address (123.1.2.111) and also indicate that the host name is fully qualified.
8. Similarly, a `ping webhost1` command on the standby site would return the correct IP address (123.2.2.111) and also indicate that the host name is fully qualified.

3.1.1.5 Testing the Host Name Resolution

Validate that you have assigned host names properly by connecting to each host at the production site and using the `ping` command to ensure that the host can locate the other hosts at the production site.

Then, connect to each host at the standby site and use the `ping` command to ensure that the host can locate the other hosts at the standby site.

3.1.2 Virtual IP and Virtual Hostname Considerations

Virtual IP addresses and hostnames are required to enable the Oracle WebLogic Administration Server to continue servicing requests when the machine hosting the Oracle WebLogic Administration Server fails. Virtual IP addresses enable Managed Servers in your domain to participate in server migration. Virtual servers should be provisioned in the application tier so that they can be bound to a network interface on any host in the application tier.

In a disaster recovery topology, the production site virtual IP hostnames must be resolvable to the IP addresses of the corresponding peer systems at the standby site. Therefore, it is important to plan the host names for the production site and the standby site. After failover from a primary site to a standby site, the alias host name for the middle tier host on the standby site becomes active. You do not need to reconfigure a hostname for the host on the standby site if you setup aliases for the standby site.

This section describes how to plan virtual IP host names and alias host names for the middle tier hosts that use the Oracle Fusion Middleware instances at the production site and the standby site. This is required when you have a single corporate DNS.

It uses the Oracle SOA Suite enterprise deployment shown in [Figure 3–1](#) for the host name examples. The host name examples in this section assume that a symmetric disaster recovery site is being set up, where the production site and standby site have the same number of hosts. Each host at the production site and the standby site has a peer host at the other site. The peer hosts are configured the same, for example, using the same ports as their counterparts at the other site.

The following subsections show how to set up virtual IP addresses and host names at the Disaster Recovery production site and standby site for the following enterprise deployments:

- [Virtual IP Addresses and Virtual Host Names for the Oracle SOA Suite Production Site and Standby Site Hosts](#)
- [Virtual IP Addresses and Virtual Host Names for the Oracle WebCenter Portal Production Site and Standby Site Hosts](#)
- [Virtual IP Addresses and Virtual Host Names for the Oracle Identity Management Production Site and Standby Site Hosts](#)
- [Virtual IP Addresses and Virtual Host Names for the Oracle Portal, Forms, Reports, and Discoverer Production Site and Standby Site Hosts](#)
- [Virtual IP Addresses and Virtual Host Names for the Oracle WebCenter Content Production Site and Standby Site Hosts](#)
- [Virtual IP Addresses and Virtual Host Names for the Oracle Business Intelligence Production Site and Standby Site Hosts](#)

Virtual IP Addresses and Virtual Host Names for the Oracle SOA Suite Production Site and Standby Site Hosts

[Table 3–14](#) shows the virtual IP addresses and virtual host names that will be used for the Oracle SOA Suite EDG deployment production site hosts. [Figure 3–1](#) shows the configuration for the Oracle SOA Suite EDG deployment at the production site.

Table 3–14 *Virtual IP Addresses and Virtual Host Names for the SOA Suite Production Site Hosts*

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.1.2.115	ADMINVHN	None
123.1.2.116	SOAVHN1	None
123.1.2.117	SOAVHN2	None

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

[Table 3–15](#) shows the virtual IP addresses, virtual host names, and alias host names that will be used for the Oracle SOA Suite EDG deployment standby site hosts. [Figure 3–2](#) shows the physical host names used for the Oracle SOA Suite EDG deployment at the standby site. The alias host names shown in [Table 3–15](#) should be defined for the SOA Oracle Suite standby site hosts in [Figure 3–2](#).

Note: If you use separate DNS servers to resolve host names, then you can use the same virtual IP addresses and virtual host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in [Table 3–2](#). See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) for more information about using separate DNS servers to resolve host names.

Table 3–15 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for SOA Suite Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.115	STBYADMINVHN	ADMINVHN
123.2.2.116	STBYSOAVHN1	SOAVHN1
123.2.2.117	STBYSOAVHN2	SOAVHN2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

Virtual IP Addresses and Virtual Host Names for the Oracle WebCenter Portal Production Site and Standby Site Hosts

[Table 3–16](#) shows the virtual IP addresses and virtual host names that will be used for the Oracle WebCenter Portal EDG deployment production site hosts. [Figure 4–4](#) shows the configuration for the Oracle WebCenter Portal EDG deployment at the production site.

Table 3–16 Virtual IP Addresses and Virtual Host Names for Oracle WebCenter Portal Production Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.1.2.111	ADMINVHN	None
123.1.2.113	SOAVHN1	None
123.1.2.114	SOAVHN2	None
123.1.2.115	WCVHN1	None
123.1.2.116	WCVHN2	None

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

[Table 3–17](#) shows the virtual IP addresses, virtual host names, and alias host names that will be used for the Oracle WebCenter Portal EDG deployment standby site hosts. [Figure 4–4](#) shows the configuration for the Oracle WebCenter Portal EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same virtual IP addresses and virtual host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in [Table 3–4](#). See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) for more information about using separate DNS servers to resolve host names.

Table 3–17 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for WebCenter Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.111	STBYADMINVHN	ADMINVHN
123.2.2.113	STBYSOAVHN1	SOAVHN1
123.2.2.114	STBYSOAVHN2	SOAVHN2
123.2.2.115	STBYWCVHN1	WCVHN1
123.2.2.116	STBYWCVHN2	WCVHN2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

Virtual IP Addresses and Virtual Host Names for the Oracle Identity Management Production Site and Standby Site Hosts

[Table 3–18](#) shows the virtual IP addresses and virtual host names that will be used for the Oracle Identity Management EDG deployment production site hosts. [Figure 4–6](#) shows the configuration for the Oracle Identity Management EDG deployment at the production site.

Table 3–18 Virtual IP Addresses and Virtual Host Names for Identity Management Production Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.1.2.111	ADMINVHN	None
123.1.2.118	IDMHVHN1	None
123.1.2.119	IDMVHN2	None
123.1.2.122	OIDVHN1	None
123.1.2.123	OIDVHN2	None
123.1.2.124	OVDVHN1	None
123.1.2.125	OVDVHN2	None
123.1.2.126	OIFVHN1	None
123.1.2.127	OIFVHN2	None
123.1.2.128	OAMVHN1	None
123.1.2.129	OAMVHN2	None
123.1.2.130	OAAMVHN1	None
123.1.2.131	OAAMVHN2	None
123.1.2.132	OIMVHN1	None

Table 3–18 (Cont.) Virtual IP Addresses and Virtual Host Names for Identity Management Production Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.1.2.133	OIMVHN2	None

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

[Table 3–19](#) shows the virtual IP addresses, virtual host names, and alias host names that will be used for the Oracle Identity Management EDG deployment standby site hosts. [Figure 4–6](#) shows the configuration for the Oracle Identity Management EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same physical host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in [Table 3–6](#). See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) for more information about using separate DNS servers to resolve host names.

Table 3–19 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Identity Management Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.111	STBYADMINVHN	ADMINVHN
123.2.2.118	STBYIDMVHN1	IDMVHN1
123.2.2.119	STBYIDMVHN2	IDMVHN2
123.2.2.122	STBYOIDVHN1	OIDVHN1
123.2.2.123	STBYOIDVHN2	OIDVHN2
123.2.2.124	STBYOVDVHN1	OVDVHN1
123.2.2.125	STBYOVDVHN2	OVDVHN2
123.2.2.126	STBYOIFVHN1	OIFVHN1
123.2.2.127	STBYOIFVHN2	OIFVHN2
123.2.2.128	STBYOAMVHN1	OAAMVHN1
123.2.2.129	STBYOAMVHN2	OAAMVHN2
123.2.2.130	STBYOAAMVHN1	OAAMVHN1
123.2.2.131	STBYOAAMVHN2	OAAMVHN2
123.2.2.132	STBYOIMVHN1	OIMVHN1
123.2.2.133	STBYOIMVHN2	OIMVHN2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

The Administration Server, the Oracle Identity Manager Managed Servers and the SOA Managed Servers require a floating IP addresses to be provisioned on each site

(Table 3–7). Ensure that you provision the floating IP addresses with the same virtual host names on production site and the standby site.

Table 3–20 Floating IP Addresses

Physical Host Name	Virtual Host Name	Floating IP
AdminServer	ADMINVHN	123.1.2.134
OIMHOST1	OIMVHN1	123.1.2.135
OIMHOST2	OIMVHN2	123.1.2.136
SOAHOST1	SOAVHN1	123.1.2.137
SOAHOST2	SOAVHN2	123.1.2.138

Note: For more information, see the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management*.

Virtual IP Addresses and Virtual Host Names for the Oracle Portal, Forms, Reports, and Discoverer Production Site and Standby Site Hosts

Table 3–21 shows the virtual IP addresses and virtual host names that will be used for the Oracle Portal, Forms, Reports, and Discoverer enterprise deployment production site hosts. Figure 4–9 shows the configuration for the Oracle Portal enterprise deployment at the production site and Figure 4–10 shows the configuration for the Oracle Forms, Reports, and Discoverer enterprise deployment at the production site.

Table 3–21 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle Portal, Forms, Reports, and Discoverer Production Site Hosts

Virtual IP Address	Virtual Host Names ¹	Alias Host Names
123.1.2.111	ADMINVHN	None
123.1.2.126	APPVHN1	None
123.1.2.127	APPVHN2	None

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Table 3–22 shows the Virtual IP addresses, virtual host names, and alias host names that will be used for the Oracle Portal, Forms, Reports, and Discoverer enterprise deployment standby site hosts. Figure 4–9 shows the configuration for the Oracle Portal enterprise deployment at the production site and Figure 4–10 shows the configuration for the Oracle Forms, Reports, and Discoverer enterprise deployment at the production site.

Note: If you use separate DNS servers to resolve host names, then you can use the same virtual IP addresses and virtual host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in Table 3–9. See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" for more information about using separate DNS servers to resolve host names.

Table 3–22 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle Portal, Forms, Reports, and Discoverer Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.111	STBYADMINVHN	ADMINVHN
123.2.2.126	STBYAPPVHN1	APPVHN1
123.2.2.127	STBYAPPVHN2	APPVHN2

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

The alias host names in [Table 3–2](#), [Table 3–4](#), [Table 3–6](#), and [Table 3–8](#) are resolved locally at the standby site to the correct IP address. [Section 3.1.1.1, "Host Name Resolution"](#) describes two ways to configure host name resolution in an Oracle Fusion Middleware Disaster Recovery topology.

Virtual IP Addresses and Virtual Host Names for the Oracle WebCenter Content Production Site and Standby Site Hosts

[Table 3–23](#) shows the virtual IP addresses and virtual host names that will be used for the Oracle WebCenter Content EDG deployment production site hosts. [Figure 4–11](#) shows the configuration for the Oracle WebCenter Content EDG deployment at the production site.

Table 3–23 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle Enterprise Content Management Production Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.1.2.111	ADMINVHN	None
123.1.2.112	ECMVHN1	None
123.1.2.114	ECMVHN2	None

¹ See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) and [Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server"](#) for information on defining physical host names.

[Table 3–24](#) shows the virtual IP addresses, Virtual host names, and alias host names that will be used for the Oracle WebCenter Content EDG deployment standby site hosts. [Figure 4–11](#) shows the configuration for the Oracle WebCenter Content EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same virtual IP addresses and virtual host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in [Table 3–11](#). See [Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers"](#) for more information about using separate DNS servers to resolve host names.

Table 3–24 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle WebCenter Content Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.111	STBYADMINVHN	ADMINVHN

Table 3–24 (Cont.) Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle WebCenter Content Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.112	STBYECMVHN1	ECMVHN1
123.2.2.113	STBYECMVHN2	ECMVHN2

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Virtual IP Addresses and Virtual Host Names for the Oracle Business Intelligence Production Site and Standby Site Hosts

Table 3–25 shows the virtual IP addresses and virtual host names that will be used for the Oracle Business Intelligence EDG deployment production site hosts. Figure 4–14 shows the configuration for the Oracle Business Intelligence EDG deployment at the production site.

Table 3–25 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle Business Intelligence Production Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.1.2.115	ADMINVHN	None
123.1.2.116	BIVHN1	None
123.1.2.117	BIVHN2	None

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

Table 3–26 shows the virtual IP addresses, virtual host names, and alias host names that will be used for the Oracle Business Intelligence EDG deployment standby site hosts. Figure 4–11 shows the configuration for the Oracle Business Intelligence EDG deployment at the standby site.

Note: If you use separate DNS servers to resolve host names, then you can use the same virtual IP addresses and virtual host names for the production site hosts and standby site hosts, and you do not need to define the alias host names on the standby site hosts that are recommended in Table 3–11. See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" for more information about using separate DNS servers to resolve host names.

Table 3–26 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for Oracle Business Intelligence Standby Site Hosts

Virtual IP Address	Virtual Host Name ¹	Alias Host Name
123.2.2.115	STBYADMINVHN	ADMINVHN
123.2.2.116	STBYBIVHN1	BIVHN1
123.2.2.117	STBYBIVHN2	BIVHN2

¹ See Section 3.1.1.3, "Resolving Host Names Using Separate DNS Servers" and Section 3.1.1.4, "Resolving Host Names Using a Global DNS Server" for information on defining physical host names.

3.1.3 Load Balancers Considerations

Oracle Fusion Middleware components require a hardware load balancer when deployed in high availability topologies. It is recommended that the hardware load balancer 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.
- Monitoring of ports (HTTP and 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, for Oracle Internet Directory clusters, the load balancer must 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, and process failure detection: The load balancer must be able to detect service and node failures (through notification or some other means) and stop directing non-Oracle Net traffic to the failed node. If your load balancer can automatically detect failures, you should use this feature.
- 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: This feature is recommended, but not required.
- For the Identity Management configuration with Oracle Access Manager, configure the virtual servers in the load balancer for the directory tier with a high value for the connection timeout for TCP connections. This value should be more than the maximum expected time over which no traffic is expected between the Oracle Access Manager and the directory tier.
- 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 to preserve the Client IP Address.

3.1.4 Virtual Server Consideration

The virtual servers and associated ports must be configured on the load balancer for different types of network traffic and monitoring. These should be configured to the appropriate real hosts and ports for the services running. Also, the load balancer should be configured to monitor the real host and ports for availability so that the traffic to these is stopped as soon as possible when a service is down. This will ensure that incoming traffic on a given virtual host is not directed to an unavailable service in the other tiers.

It is recommended to use two load balancers when dealing with external and internal traffic. In such a topology, one load balancer is set up for external HTTP traffic and the other load balancer is set up for internal LDAP traffic. A deployment may choose to have a single load balancer device due to a variety of reasons. While this is supported, the deployment should consider the security implications of doing this and if found appropriate, open up the relevant firewall ports to allow traffic across the various DMZs. It is worth noting that in either case, it is highly recommended to deploy a given load balancer device in fault tolerant mode.

Some of the virtual servers defined in the load balancer are used for inter-component communication. These virtual servers are used for internal traffic and are defined in the internal DNS of a company. It is recommended to create aliases for these virtual servers when using a single Global DNS Server to resolve host names.

Creating aliases is not required when using separate DNS servers to resolve host names.

The virtual servers required for the various Oracle Fusion Middleware products are described in the following tables.

Table 3–27 Virtual Servers for Oracle SOA Suite Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle SOA	External	soa.mycompany.com	None
Oracle SOA	Internal	soainternal.mycompany.com	None
Administration Consoles	Internal	admin.mycompany.com	None

Table 3–28 Virtual Servers for Oracle SOA Suite Standby Site

Components	Access	Virtual Server Name	Alias Virtual Server Name
Oracle SOA	External	soa.mycompany.com	None
Oracle SOA	Internal	stbysoainternal.mycompany.com	soainternal.mycompany.com
Administration Consoles	Internal	admin.mycompany.com	None

Table 3–29 Virtual Servers for Oracle WebCenter Portal Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle WebCenter Portal	External	wc.mycompany.com	None
Oracle WebCenter Portal	Internal	wcinternal.mycompany.com	None

Table 3–29 (Cont.) Virtual Servers for Oracle WebCenter Portal Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle SOA Internal	Internal	soainternal.mycompany.com ¹	None
Administration Consoles	Internal	admin.mycompany.com	None

¹ Required when extending with SOA domain.

Table 3–30 Virtual Servers for Oracle WebCenter Portal Suite Standby Site

Components	Access	Virtual Server Name	Alias Virtual Server Name
Oracle WebCenter Portal	External	wc.mycompany.com	None
Oracle WebCenter Portal	Internal	stbywcinternal.mycompany.com	wcinternal.mycompany.com
Oracle SOA Internal	Internal	soainternal.mycompany.com ¹	soainternal.mycompany.com
Administration Consoles	Internal	admin.mycompany.com	None

¹ Required when extending with SOA domain.

Table 3–31 Virtual Servers for Oracle Identity Management Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle Internet Directory	External	oid.mycompany.com	None
Oracle Internet Directory	Internal	oidinternal.mycompany.com	None
Oracle Virtual Directory	External	ovd.mycompany.com	None
Oracle Virtual Directory	Internal	ovdinternal.mycompany.com	None
Oracle Identity Federation	External	oif.mycompany.com	None
Oracle Identity Federation	Internal	oifinternal.mycompany.com	None
Oracle Identity Manager	External	oim.mycompany.com	None
Oracle Identity Manager	Internal	oiminternal.mycompany.com	None
Single Sign-On	External	sso.mycompany.com	None
Single Sign-On	Internal	ssointernal.mycompany.com	None
Administration Consoles	Internal	admin.mycompany.com	None

Table 3–32 Virtual Servers for Oracle Identity Management Standby Site

Components	Access	Virtual Server Name	Alias Name
Oracle Internet Directory	External	oid.mycompany.com	None
Oracle Internet Directory	Internal	stbyoidinternal.mycompany.com	oidinternal.mycompany.com
Oracle Virtual Directory	External	ovd.mycompany.com	None

Table 3–32 (Cont.) Virtual Servers for Oracle Identity Management Standby Site

Components	Access	Virtual Server Name	Alias Name
Oracle Virtual Directory	Internal	stbyovdinternal.mycompany.com	ovdinternal.mycompany.com
Oracle Identity Federation	External	oif.mycompany.com	None
Oracle Identity Federation	Internal	stbyoifinternal.mycompany.com	oifinternal.mycompany.com
Oracle Identity Manager	External	oim.mycompany.com	None
Oracle Identity Manager	Internal	stbyoiminternal.mycompany.com	oiminternal.mycompany.com
Single Sign-On	External	sso.mycompany.com	None
Single Sign-On	Internal	stbyssointernal.mycompany.com	ssointernal.mycompany.com
Administration Consoles	Internal	admin.mycompany.com	None

Table 3–33 Virtual Servers for Oracle Portal, Forms, Reports, and Discoverer Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle Portal	External	portal.mycompany.com	None
Oracle Portal	Internal	portalinternal.mycompany.com	None
Oracle Forms and Oracle Reports	External	forms.mycompany.com	None
Oracle Forms and Oracle Reports	Internal	formsinternal.mycompany.com	None
Discoverer	External	disco.mycompany.com	None
Discoverer	Internal	discointernal.mycompany.com	None
Administration Consoles	Internal	admin.mycompany.com	None

Table 3–34 Virtual Servers for Oracle Portal, Reports, Forms, and Discoverer Standby Site

Components	Access	Virtual Server Name	Alias Virtual Server Name
Oracle Portal	External	portal.mycompany.com	None
Oracle Portal	Internal	stbyportalinternal.mycompany.com	portalinternal.mycompany.com
Oracle Forms and Oracle Reports	External	forms.mycompany.com	None
Oracle Forms and Oracle Reports	Internal	stbyformsinternal.mycompany.com	formsinternal.mycompany.com

Table 3–34 (Cont.) Virtual Servers for Oracle Portal, Reports, Forms, and Discoverer Standby Site

Components	Access	Virtual Server Name	Alias Virtual Server Name
Discoverer	External	disco.mycompany.com	None
Discoverer	Internal	stbydiscointernal.mycompany.com	discointernal.mycompany.com
Administration Consoles	Internal	admin.mycompany.com	None

Table 3–35 Virtual Servers for Oracle WebCenter Content Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle Enterprise Content Management	External	ecm.mycompany.com	None
Oracle Enterprise Content Management	Internal	ecminternal.mycompany.com	None
Oracle SOA Internal	Internal	soainternal.mycompany.com ¹	None
Administration Consoles	Internal	admin.mycompany.com	None

¹ Required when extending with SOA domain.

Table 3–36 Virtual Servers for Oracle WebCenter Content Standby Site

Components	Access	Virtual Server Name	Alias Virtual Server Name
Oracle Enterprise Content Management	External	ecm.mycompany.com	None
Oracle Enterprise Content Management	Internal	stbyecminternal.mycompany.com	ecminternal.mycompany.com
Oracle SOA Internal	Internal	stbysoainternal.mycompany.com ¹	soainternal.mycompany.com
Administration Consoles	Internal	admin.mycompany.com	None

¹ Required when extending with Oracle SOA domain.

Table 3–37 Virtual Servers for Oracle Business Intelligence Production Site

Components	Access	Virtual Server Name	Alias Name
Oracle Business Intelligence	External	bi.mycompany.com	None
Oracle Business Intelligence	Internal	biinternal.mycompany.com	None
Oracle Access Manager Internal	Internal	oaminternal.mycompany.com ¹	None
Administration Consoles	Internal	admin.mycompany.com	None

¹ Required when extending with Oracle Access Manager domain.

Table 3–38 Virtual Servers for Oracle Business Intelligence Standby Site

Components	Access	Virtual Server Name	Alias Virtual Server Name
Oracle Business Intelligence	External	bi.mycompany.com	None
Oracle Business Intelligence	Internal	stbybiinternal.mycompany.com	biinternal.mycompany.com
Oracle Access Manager Internal	Internal	stbyoaminternal.mycompany.com ¹	oaminternal.mycompany.com
Administration Consoles	Internal	admin.mycompany.com	None

¹ Required when extending with Oracle Access Manager domain.

3.1.5 Wide Area DNS Operations

When a site switchover or failover is performed, client requests must be redirected transparently to the new site that is playing the production role. To direct client requests to the entry point of a production site, use DNS resolution. To accomplish this redirection, the wide area DNS that resolves requests to the production site has to be switched over to the standby site. The DNS switchover can be accomplished by either using a global load balancer or manually changing DNS names.

Note: A hardware load balancer is assumed to be front-ending each site. Check for supported load balancers at:

<http://support.oracle.com>

The following topics are described in this section:

- [Using a Global Load Balancer](#)
- [Manually Changing DNS Names](#)

3.1.5.1 Using a Global Load Balancer

When a global load balancer is deployed in front of the production and standby sites, it provides fault detection services and performance-based routing redirection for the two sites. Additionally, the load balancer can provide authoritative DNS name server equivalent capabilities.

During normal operations, the global load balancer can be configured with the production site's load balancer name-to-IP mapping. When a DNS switchover is required, this mapping in the global load balancer is changed to map to the standby site's load balancer IP. This allows requests to be directed to the standby site, which now has the production role.

This method of DNS switchover works for both site switchover and failover. One advantage of using a global load balancer is that the time for a new name-to-IP mapping to take effect can be almost immediate. The downside is that an additional investment must be made for the global load balancer.

3.1.5.2 Manually Changing DNS Names

This method of DNS switchover involves the manual change of the name-to-IP mapping that is originally mapped to the IP address of the production site's load

balancer. The mapping is changed to map to the IP address of the standby site's load balancer. Follow these instructions to perform the switchover:

1. Make a note of the current Time to Live (TTL) value of the production site's load balancer mapping. This mapping is in the DNS cache and it will remain there until the TTL expires. As an example, let's assume that the TTL is 3600 seconds.
2. Modify the TTL value to a short interval (for example, 60 seconds).
3. Wait one interval of the original TTL. This is the original TTL of 3600 seconds from Step 1.
4. Ensure that the standby site is switched over to receive requests.
5. Modify the DNS mapping to resolve to the standby site's load balancer, giving it the appropriate TTL value for normal operation (for example, 3600 seconds).

This method of DNS switchover works for switchover or failover operations. The TTL value set in Step 2 should be a reasonable time period where client requests cannot be fulfilled. The modification of the TTL is effectively modifying the caching semantics of the address resolution from a long period of time to a short period. Due to the shortened caching period, an increase in DNS requests can be observed.

3.2 Storage Considerations

This section provides recommendations for designing storage for the Disaster Recovery solution for your enterprise deployment.

3.2.1 Oracle Fusion Middleware Artifacts

The Oracle Fusion Middleware components in a given environment are usually interdependent on each other, so it is important to have the components in the topology be in sync. This is an important consideration for designing volumes and consistency groups. Some of the artifacts are static while others are dynamic.

Static Artifacts

Static artifacts are files and directories that do not change frequently. These include:

- **MW_HOME:** The Oracle Middleware home usually consists of an Oracle home and an Oracle WebLogic Server home.
- **Oracle Inventory:** The `oraInst.loc` and `oratab` files, which are located in the `/etc` directory.
- **BEA Home List:** On UNIX, this is located at `user_home/boa/beahomelist`.

Dynamic or Run-Time Artifacts

Dynamic or run-time artifacts are files that change frequently. Run-time artifacts include:

- **Domain Home:** Domain directories of the Administration Server and the Managed Servers.
- **Oracle Instances:** Oracle Instance home directories.
- **Application artifacts,** such as `.ear` or `.war` files.
- **Database artifacts** such as the MDS repository.
- **Database metadata repositories** used by Oracle Fusion Middleware.

- Persistent stores, such as JMS Providers and transaction logs.

3.2.2 Oracle Home and Oracle Inventory

Oracle Fusion Middleware allows creating multiple Managed Servers from one single binary installation. This allows the installation of binaries in a single location on a shared storage and the reuse of this installation by the servers in different nodes. However, for maximum availability, Oracle recommends using redundant binary installations.

When an `ORACLE_HOME` or a `WL_HOME` is shared by multiple servers in different nodes, it is recommended to 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 `user_home/bea/beahomelist` file. This would be required for any nodes installed additionally to the ones shown in this topology.

3.2.3 Storage Replication

This section provides guidelines on creating volumes on the shared storage. Depending on the capabilities of the storage replication technology available with your preferred storage device you may need to create mount points, directories and symbolic links on each of the nodes within a tier.

If your storage device's storage replication technology guarantees consistent replication across multiple volumes:

- Create one volume per server running on that tier. For example, on the application tier, you can create one volume for the WebLogic Administration Server and another volume for the Managed Servers.
- Create one consistency group for each tier with the volumes for that tier as its members.
- Note that if a volume is mounted by two systems simultaneously, a clustered file system may be required for this, depending on the storage subsystem. However, there is no known case of a single file or directory tree being concurrently accessed by Oracle processes on different systems. NFS is a clustered file system, so no additional clustered file system software is required if you are using NFS-attached storage.

If your storage device's storage replication technology does not guarantee consistent replication across multiple volumes:

- Create a volume for each tier. For example, you can create one volume for the application tier, one for the web tier, and so on.
- Create a separate directory for each node in that tier. For example, you can create a directory for `SOAHOST1` under the application tier volume; create a directory for `WEBHOST1` under the web tier volume, and so on.
- Create a mount point directory on each node to the directory on the volume.
- Create a symbolic link to the mount point directory. A symbolic link should be created so that the same directory structure can be used across the nodes in a tier.
- Note that if a volume is mounted by two systems simultaneously, a clustered file system may be required for this, depending on the storage subsystem. However,

there is no known case of a single file or directory tree being concurrently accessed by Oracle processes on different systems. NFS is a clustered file system, so no additional clustered file system software is required if you are using NFS-attached storage.

Note: Before you set up the shared storage for your Disaster Recovery sites, read the high availability chapter in the *Oracle Fusion Middleware Release Notes* to learn of any known shared storage-based deployment issues in high availability environments.

The release notes for Oracle Fusion Middleware can be found at this URL:

<http://www.oracle.com/technology/documentation/middleware.html>

3.2.4 File-Based Persistent Store

The WebLogic Server application servers are usually clustered for high-availability. For the local site high availability of the Oracle SOA Suite topology, a file-based persistent store is used for the Java Message Services (JMS) and Transaction Logs (TLogs). This file-based persistent store must reside on shared storage that is accessible by all members of the cluster.

A SAN storage system should use either a host based clustered or shared file system technology such as the Oracle Clustered File System (OCFS2). OCFS2 is a symmetric shared disk cluster file system which allows each node to read and write both metadata and data directly to the SAN.

Additional clustered file systems are not required when using NAS storage systems.

3.3 Database Considerations

This section provides the recommendations and considerations for setting up Oracle databases that will be used in the Oracle Fusion Middleware Disaster Recovery topology.

1. Oracle recommends creating Real Application Cluster databases on both the production site and standby site as required by your topology.
2. Oracle Data Guard is the recommended disaster protection technology for the databases running the metadata repositories. You can also use Oracle Active Data Guard or Oracle GoldenGate.

Note: You can only use Oracle GoldenGate in an active-passive configuration.

For more information, refer to the following:

- Oracle Active Data Guard at:
<http://www.oracle.com/in/products/database/options/active-data-guard/index.html>
- Oracle GoldenGate at:

<http://www.oracle.com/technetwork/middleware/goldengate/overview/index.html>

3. The Oracle Data Guard configuration used should be decided based on the data loss requirements of the database as well as the network considerations such as the available bandwidth and latency when compared to the redo generation. Make sure that this is determined correctly before setting up the Oracle Data Guard configuration.

Please refer to *Oracle Data Guard Concepts and Administration* as well as related Maximum Availability Architecture collateral at the following URL for more information:

<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

4. Ensure that your network is configured for low latency with sufficient bandwidth, since synchronous redo transmission can cause an impact on response time and throughput.
5. The `LOG_ARCHIVE_DEST_n` parameter on standby site databases should have the `SYNC` or `ASYN` attributes. `ASYN` is the default attribute if no attributes are specified.
6. The standby site database should be in the Managed Recovery mode. This ensures that the standby site databases are in a constant state of media recovery. The managed recovery mode is enables for shorter failover times.
7. The `tnsnames.ora` file on the production site and the standby site must have entries for databases on both the production and standby sites.
8. It is strongly recommended to force Data Guard to perform manual database synchronization whenever middle tier synchronization is performed. This is especially true for components that store configuration data in the metadata repositories.
9. It is strongly recommended to set up aliases for the database host names on both the production and standby sites. This enables seamless switchovers, switchbacks and failovers.

3.3.1 Making TNSNAMES.ORA Entries for Databases

Because Oracle Data Guard is used to synchronize production and standby databases, the production database and standby database must be able to reference each other.

Oracle Data Guard uses `tnsnames.ora` file entries to direct requests to the production and standby databases, so entries for production and standby databases must be made to the `tnsnames.ora` file. See *Oracle Data Guard Concepts and Administration* in the Oracle Database documentation set for more information about using `tnsnames.ora` files with Oracle Data Guard.

3.3.2 Manually Forcing Database Synchronization with Oracle Data Guard

For Oracle Fusion Middleware components that store middle tier configuration data in Oracle database repositories, use Oracle Data Guard to manually force a database synchronization whenever a middle tier synchronization is performed. Use the `SQL alter system archive log all` statement to switch the logs, which forces the synchronization of the production site and standby site databases.

[Example 3-9](#) shows the SQL statement to use to force the synchronization of a production site database and standby site database.

Example 3–9 Manually Forcing an Oracle Data Guard Database Synchronization

```
ALTER SYSTEM ARCHIVE LOG ALL;
```

3.3.3 Setting Up Database Host Name Aliases

Optionally, you can set up database host name aliases for the databases at your production site and standby site. The alias must be defined in DNS or in the `/etc/hosts` file on each node running a database instance.

In a Disaster Recovery environment, the site that actively accepts connections is the production site. At the completion of a successful failover or switchover operation, the standby site becomes the new production site.

This section includes an example of defining an alias for database hosts named `custdbhost1` and `stbycustdbhost1`. [Table 3–39](#) shows the database host names and the connect strings for the databases before the alias is defined.

Table 3–39 Database Host Names and Connect Strings

Site	Database Host Name	Database Connect String
Production	<code>custdbhost1.us.oracle.com</code>	<code>custdbhost1.us.oracle.com:1521:orcl</code>
Standby	<code>stbycustdbhost1.us.oracle.com</code>	<code>stbycustdbhost1.us.oracle.com:1521:orcl</code>

In this example, all database connect strings on the production site take the form "`custdbhost1.us.oracle.com:1521:orcl`." After a failover or switchover operation, this connect string must be changed to "`stbycustdbhost1.us.oracle.com:1521:orcl`." However, by creating an alias of "`proddb1`" for the database host name as shown in [Table 3–40](#), you can avoid manually changing the connect strings, which enables seamless failovers and switchovers:

Table 3–40 Specifying an Alias for a Database Host

Site	Database Host Name	Alias	Database Connect String
Production	<code>custdbhost1.us.oracle.com</code>	<code>proddb1.us.oracle.com</code>	<code>proddb1.us.oracle.com:1521:orcl</code>
Standby	<code>stbycustdbhost1.us.oracle.com</code>	<code>proddb1.us.oracle.com</code>	<code>proddb1.us.oracle.com:1521:orcl</code>

In this example, the production site database host name and the standby site database host name are aliased to "`proddb1.us.oracle.com`" and the connect strings on the production site and the standby site can take the form "`proddb1.us.oracle.com:1521:orcl`". On failover and switchover operations, the connect string does not need to change, thus enabling a seamless failover and switchover.

The format for specifying aliases in `/etc/hosts` file entries is:

```
<IP> <ALIAS WITH DOMAIN> <ALIAS> <HOST NAME WITH DOMAIN> <HOST NAME>
```

In this example, you create a database host name alias of `proddb1` for host `custdbhost1` at the production site and for host `stbycustdbhost1` at the standby site. The hosts file entry should specify the fully qualified database host name alias with the `<ALIAS WITH DOMAIN>` parameter, the short database host name alias with the `<ALIAS>` parameter, the fully qualified host name with the `<HOST NAME WITH DOMAIN>` parameter, and the short host name with the `<HOST NAME>` parameter.

So, in the `/etc/hosts` files at the production site, make sure the entry for host `custdbhost1` looks like this:

```
152.68.196.213 proddb1.us.oracle.com proddb1 custdbhost1.us.oracle.com custdbhost1
```

And, in the `/etc/hosts` files at the standby site, make sure the entry for host `stbycustdbhost1` looks like this:

```
140.87.25.40 proddb1.us.oracle.com proddb1 stbycustdbhost1.us.oracle.com
stbycustdbhost1
```

3.4 Starting Points

Before setting up the standby site, the administrator must evaluate the starting point of the project. The starting point for designing an Oracle Fusion Middleware Disaster Recovery topology is usually one of the following:

- The production site is already created and the standby site is being planned and created.

[Section 3.4.1, "Starting with an Existing Site"](#) describes how to design the Oracle Fusion Middleware Disaster Recovery standby site when you have an existing production site.

- There is no existing production site or standby site. Both need to be designed and created.

[Section 3.4.2, "Starting with New Sites"](#) describes how to design a new Oracle Fusion Middleware Disaster Recovery production site and standby site when you do not have an existing production site or standby site.

- Some hosts or components may exist at a current production site, but new hosts or components need to be added at that site or at a standby site to set up a functioning Oracle Fusion Middleware Disaster Recovery topology.

Use the pertinent information in this chapter to design and implement an Oracle Fusion Middleware Disaster Recovery topology.

3.4.1 Starting with an Existing Site

When the administrator's starting point is an existing production site, the configuration data and the Oracle binaries for the production site already exist on the file system. Also, the host names, ports, and user accounts are already defined. When a production site exists, the administrator can choose to:

- Design a symmetric standby site. See [Section 3.5.1, "Design Considerations for a Symmetric Topology."](#)
- Design an asymmetric standby site. See [Section 3.5.2, "Design Considerations for an Asymmetric Topology."](#)
- Migrate the production site to shared storage, if not already on shared storage, and then create either a symmetric standby site or asymmetric standby site. See [Section 3.4.1.1, "Migrating an Existing Production Site to Shared Storage."](#)

3.4.1.1 Migrating an Existing Production Site to Shared Storage

The Oracle Fusion Middleware Disaster Recovery solution relies on shared storage to implement storage replication for disaster protection of the Oracle Fusion Middleware middle tier configuration. When a production site has already been created, it is likely

that the Oracle home directories for the Oracle Fusion Middleware instances that comprise the site are not located on the shared storage. If this is the case, then these homes have to be migrated completely to the shared storage to implement the Oracle Fusion Middleware Disaster Recovery solution.

Follow these guidelines for migrating the production site from the local disk to shared storage:

1. All backups performed must be offline backups. For more information, see "Types of Backups" and "Recommended Backup Strategy" in *Oracle Fusion Middleware Administrator's Guide*.
2. The backups must be performed as the root user and the permissions must be preserved. See the "Overview of the Backup Strategies" section in *Oracle Fusion Middleware Administrator's Guide*.
3. This is a one-time operation, so it is recommended to recover the entire domain.
4. The directory structure on the shared storage must be set up as described in [Section 4.1.1, "Directory Structure and Volume Design."](#)
5. For Oracle SOA Suite, see "Backup and Recovery Recommendations for Oracle SOA Suite" in *Oracle Fusion Middleware Administrator's Guide*.
6. For Oracle WebCenter Portal, see "Backup and Recovery Recommendations for Oracle WebCenter Portal" in *Oracle Fusion Middleware Administrator's Guide*.
7. For Oracle Identity Management, see "Backup and Recovery Recommendations for Oracle Identity Management" in *Oracle Fusion Middleware Administrator's Guide*.
8. For Oracle WebLogic Server, see "Backup and Recovery Recommendations for Oracle JRF Installations" in *Oracle Fusion Middleware Administrator's Guide*.
9. For the Web Tier, see "Backup and Recovery Recommendations for Web Tier Installations" in *Oracle Fusion Middleware Administrator's Guide*.
10. For the Oracle Business Intelligence, see "Backup and Recovery Recommendations for Oracle Business Intelligence Installations" in *Oracle Fusion Middleware Administrator's Guide*.
11. For Oracle Portal, Oracle Forms, Oracle Reports, and Discoverer backup and recovery recommendations, see "Backup and Recovery Recommendations for Oracle Portal, Oracle Forms Services, and Oracle Reports Installations" in *Oracle Fusion Middleware Administrator's Guide*.

3.4.2 Starting with New Sites

This section presents the logic to implementing a new production site for an Oracle Fusion Middleware Disaster Recovery topology. It describes the planning and setup of the production site by pre-planning host names, configuring the hosts to resolve the alias host names and physical host names, and ensuring that storage replication is set up to copy the configuration based on these names to the standby site. When you design the production site, you should also plan the standby site, which can be a symmetric standby site or an asymmetric standby site.

When you are designing a new production site (not using a pre-existing production site), you will use Oracle Universal Installer to install software on the production site, and parameters such as alias host names and software paths must be carefully designed to ensure that they are the same for both sites.

The flexibility you have when you create a new Oracle Fusion Middleware Disaster Recovery production site and standby site includes:

1. You can design your Oracle Fusion Middleware Disaster Recovery solution so that each host at the production site and at the standby site has the desired alias host name and physical host name. Host name planning was discussed in [Section 3.1.1, "Planning Host Names."](#)
2. When you design and create your production site from scratch, you can choose the Oracle home name and Oracle home directory for each Fusion Middleware installation.

Designing and creating your site from scratch is easier than trying to modify an existing site to meet the design requirements described in this chapter.

3. You can assign ports for the Fusion Middleware installations for the production site hosts that will not conflict with the ports that will be used at the standby site hosts.

This is easier than having to check for and resolve port conflicts between an existing production site and standby site.

3.5 Topology Considerations

This section describes design considerations for:

- A symmetric topology
- An asymmetric topology

3.5.1 Design Considerations for a Symmetric Topology

A symmetric topology is an Oracle Fusion Middleware Disaster Recovery configuration that is completely identical across tiers on the production site and standby site. In a symmetric topology, the production site and standby site have the identical number of hosts, load balancers, instances, and applications. The same ports are used for both sites. The systems are configured identically and the applications access the same data. This manual describes how to set up a symmetric Oracle Fusion Middleware Disaster Recovery topology for an enterprise configuration.

3.5.2 Design Considerations for an Asymmetric Topology

An asymmetric topology is an Oracle Fusion Middleware Disaster Recovery configuration that is different across tiers on the production site and standby site. In an asymmetric topology, the standby site can use less hardware (for example, the production site could include four hosts with four Fusion Middleware instances while the standby site includes two hosts with four Fusion Middleware instances. Or, in a different asymmetric topology, the standby site can use fewer Fusion Middleware instances (for example, the production site could include four Fusion Middleware instances while the standby site includes two Fusion Middleware instances). Another asymmetric topology might include a different configuration for a database (for example, using a Real Application Clusters database at the production site and a single instance database at the standby site).

Setting Up and Managing Disaster Recovery Sites

This chapter uses the Oracle SOA Suite enterprise deployment and the Oracle Identity Management enterprise deployment topologies as examples to illustrate the steps required to set up the production site and standby site.

It includes the following topics:

- [Setting Up the Site](#)
- [Creating a Production Site](#)
- [Creating a Standby Site](#)
- [Creating an Asymmetric Standby Site](#)
- [Performing Site Operations and Administration](#)
- [Patching an Oracle Fusion Middleware Disaster Recovery Site](#)

4.1 Setting Up the Site

This section provides the steps to create the production site. The Oracle SOA enterprise deployment topology and the Oracle Identity Management Enterprise deployment topology are used as examples.

It includes the following topics:

- [Directory Structure and Volume Design](#)
- [Storage Replication](#)
- [Database](#)
- [Node Manager](#)

Ensure that you have performed the following prerequisites before you start creating the production site:

- Set up the host name aliases for the middle tier hosts, which was described in [Section 3.1.1, "Planning Host Names."](#)
- Create the required volumes on the shared storage on the production site, which is described in [Section 4.1.1, "Directory Structure and Volume Design."](#)
- Create the mount points and the symbolic links (if required). Refer to [Section 3.2.3, "Storage Replication"](#) to determine whether you must create symbolic links for the production site.

- The Oracle Data Guard configuration used should be decided based on the data loss requirements of the database as well as the network considerations such as the available bandwidth and latency when compared to the redo generation. Ensure that this is determined correctly before setting up the Oracle Data Guard configuration.

Please refer to *Oracle Data Guard Concepts and Administration* as well as related Maximum Availability Architecture collateral at the following URL for more information:

<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

4.1.1 Directory Structure and Volume Design

The following section details the directory structure recommended by Oracle. The end user is free to choose other directory layouts, but the model adopted here enables maximum availability, providing the best isolation of components and symmetry in the configuration, and facilitating backup and disaster recovery.

This list describes directories and directory environment variables:

- `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 Oracle WebLogic Server.
- `ORACLE_HOME`: This environment variable and related directory path refers to the location where a product suite (such as Oracle SOA Suite, Oracle WebCenter Portal, or Oracle Identity Management) is installed.
- `DOMAIN` directory: This directory path refers to the location where the Oracle WebLogic Domain information (configuration artifacts) is stored. Different WebLogic Servers can use different domain directories even when in the same node as described below.
- `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.

For more information, see the following:

- [Directory Structure Recommendations for Oracle SOA Suite](#)
- [Directory Structure Recommendations for Oracle WebCenter Portal](#)
- [Directory Structure Recommendations for Oracle Identity Management](#)
- [Directory Structure Recommendations for Oracle Portal, Forms, Reports, and Discoverer](#)
- [Directory Structure Recommendations for Oracle WebCenter Content](#)
- [Directory Structure Recommendations for Oracle Business Intelligence](#)


4.1.1.1 Directory Structure Recommendations for Oracle SOA Suite

Oracle Fusion Middleware 11g allows creating multiple SOA Managed Servers from one single binary installation. This allows the installation of binaries in a single

location on a shared storage and the reuse of this installation by the servers in different nodes. However, for maximum availability, Oracle recommends using redundant binary installations. In this model, two MW HOMEs (each of which has a `WL_HOME` and an `ORACLE_HOME` for each product suite) are installed in a shared storage. Additional servers (when scaling out or up) of the same type can use either one of these two locations without requiring more installations. Ideally, users should use two different volumes for redundant binary location, thus isolating as much as possible the failures in each volume. For additional protection, Oracle recommends using storage replication for these volumes. If multiple volumes are not available, Oracle recommends using mount points to simulate the same mount location in a different directory in the shared storage. Although this does not guarantee the protection that multiple volumes provide, it does allow protection from user deletions and individual file corruption.

Oracle also recommends separating the domain directory used by the Administration Server from the domain directory used by Managed Servers. This allows a symmetric configuration for the domain directories used by Managed Servers, and isolates the failover of the Administration Server. The domain directory for the Administration Server must reside in a shared storage to allow failover to another node with the same configuration. It is also recommended to have the Managed Servers' domain directories on a shared storage, even though having them on the local file system is also supported. This is especially true when designing a production site with the disaster recovery site in mind. [Figure 4-1](#) represents the directory structure layout for Oracle SOA Suite.




Figure 4–1 Directory Structure for SOA



Detailed information about setting up this directory structure is included in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* and in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal*.

Table 4–1 explains what the color-coded elements in Figure 4–1 mean. The directory structure in Figure 4–1 does not show other required internal directories such as oracle_common and jrockit.

Table 4–1 Directory Structure Elements

Element	Explanation
	The Administration Server domain directories, the Managed Server domain, applications, deployment plans, file adapter control directory, JMS and TX logs, and the entire MW_HOME are on a shared disk.
	Fixed name.
	Installation-dependent name.

4.1.1.1.1 Volume Design for Oracle SOA Suite [Figure 4-2](#) and [Figure 4-3](#) shows an Oracle SOA Suite topology diagram. The volume design described in this section is for this Oracle SOA Suite topology. Detailed instructions for installing and configuring this topology are provided in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

Figure 4–2 *MySOACompany Topology with Oracle Access Manager*


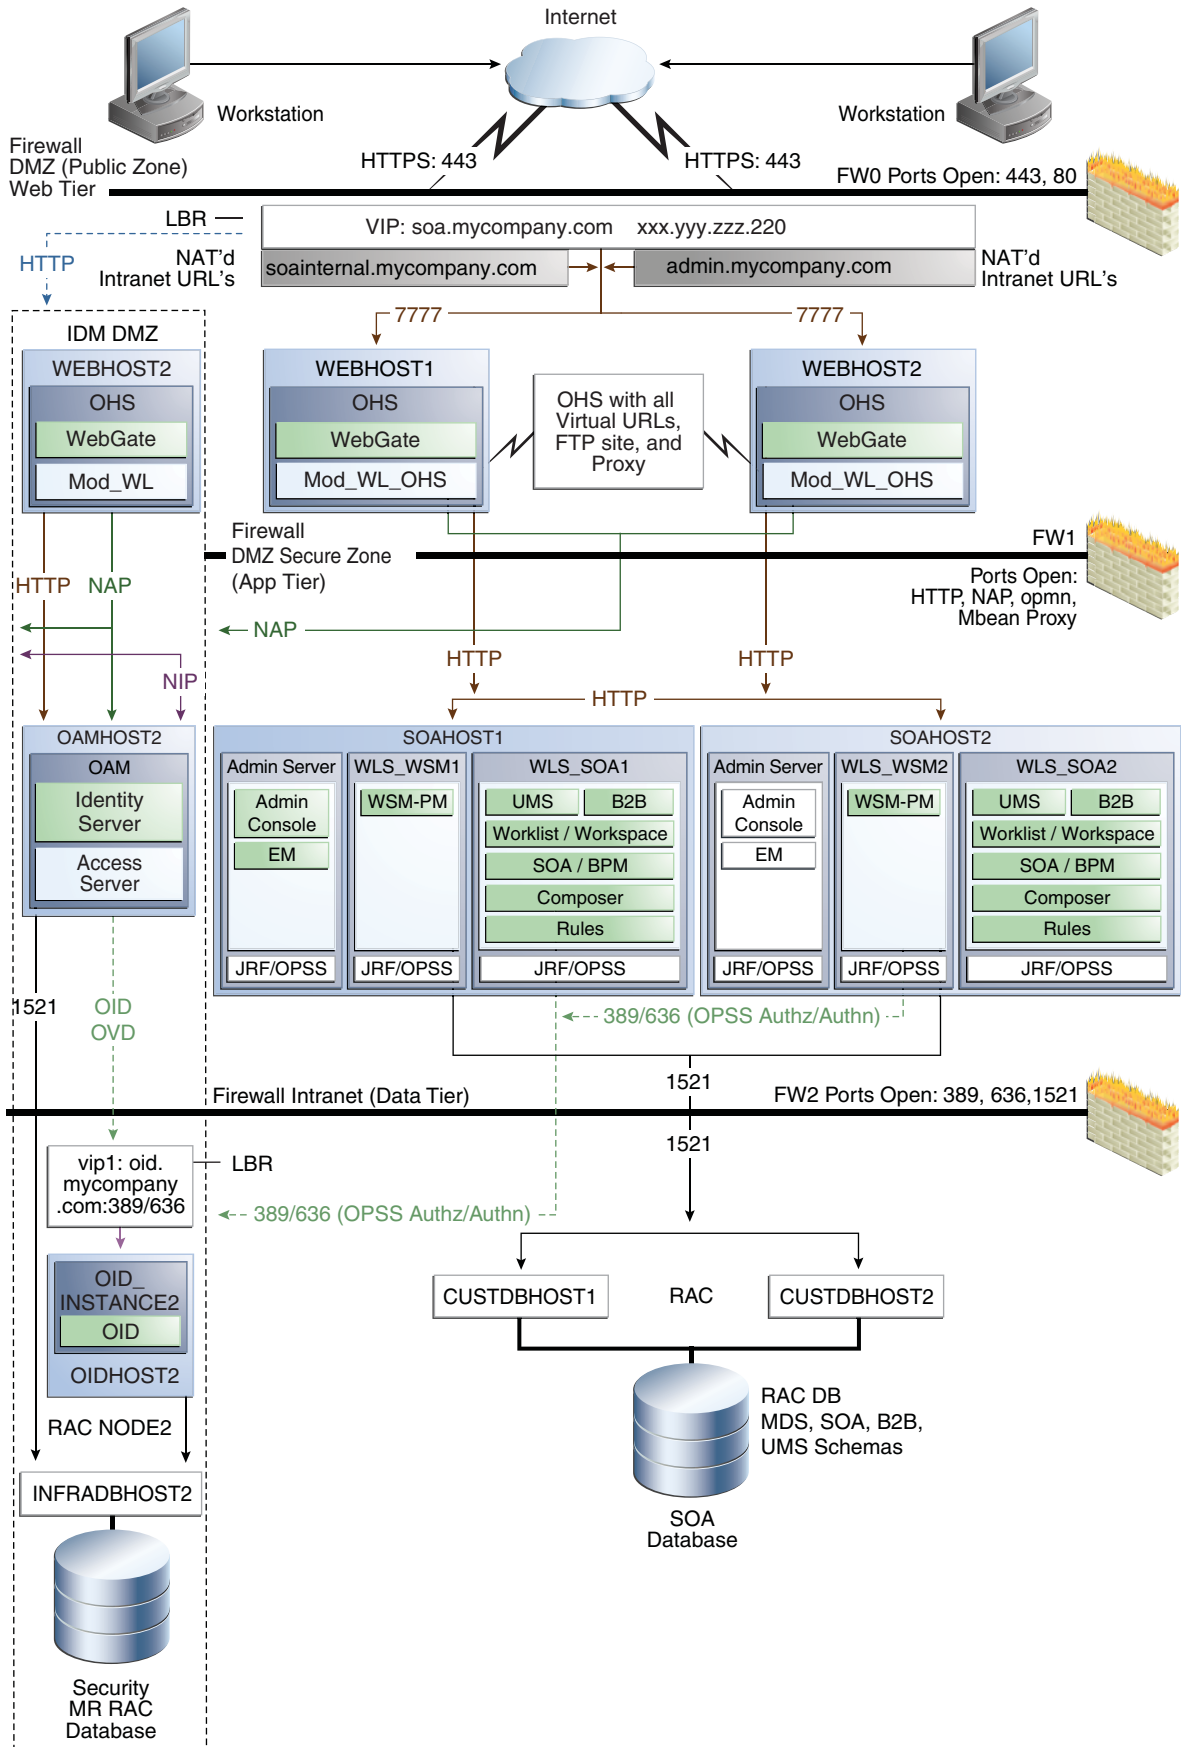


Figure 4–3 *MySOACompany Topology with Oracle Access Manager and BPM*



For disaster recovery of this Oracle SOA Suite topology, Oracle recommends the following volume design:

- Provision two volumes for two Middleware Homes that contain redundant product binaries (VOLFMW1 and VOLFMW2 in [Table 4-2](#))
- Provision one volume for the Administration Server domain directory (VOLADMIN in [Table 4-2](#))
- Provision one volume on each node for the Managed Server domain directory (VOLSOA1 and VOLSOA2 in [Table 4-2](#)). This directory is shared between all the Managed Servers on that node.
- Provision one volume for the JMS file-store and JTA transaction logs (VOLDATA in [Table 4-2](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain.
- Provision one volume on each node for the Oracle HTTP Server Oracle home (VOLWEB1 and VOLWEB2 in [Table 4-2](#)).
- Provision one volume on each node for the Oracle HTTP Server Oracle instance (VOLWEBINST1 and VOLWEBINST2 in [Table 4-2](#)).

[Table 4-2](#) provides a summary of Oracle recommendations for volume design for the Oracle SOA Suite topology shown in [Figure 4-2](#) and [Figure 4-3](#):

Table 4-2 Volume Design Recommendations for Oracle SOA Suite

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEBINST1	WEBHOST1	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLWEBINST2	WEBHOST2	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLSTATIC1 ¹	WEBHOST1	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Web	VOLSTATIC2 ²	WEBHOST2	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Application	VOLFMW1	SOAHOST1	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle SOA Suite binaries
Application	VOLFMW2	SOAHOST2	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle SOA Suite binaries.
Application	VOLADMIN	SOAHOST1	/u01/app/oracle/admin/soaDomain/admin	Volume for Administration Server domain directory

Table 4–2 (Cont.) Volume Design Recommendations for Oracle SOA Suite

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Application	VOLSOA1	SOAHOST1	/u01/app/oracle/admin/soaDomain/mng1	Volume for Managed Server domain directory
Application	VOLSOA2	SOAHOST2	/u01/app/oracle/admin/soaDomain/mng2	Volume for Managed Server domain directory
Application	VOLDATA	SOAHOST1, SOAHOST2	/u01/app/oracle/admin/soaDomain/soaCluster/jms /u01/app/oracle/admin/soaDomain/soaCluster/tlogs	Volume for transaction logs and JMS data

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.1.2 Consistency Group Recommendations for Oracle SOA Suite Oracle recommends the following consistency groups for the Oracle SOA Suite topology:

- Create one consistency group with the volumes containing the domain directories for the Administration Server and Managed Servers as members (DOMAINGROUP in [Table 4–3](#)).
- Create one consistency group with the volume containing the JMS file store and transaction log data as members (DATAGROUP in [Table 4–3](#)).
- Create one consistency group with the volume containing the Middleware Homes as members (FMWHOMEGROUP in [Table 4–3](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle homes as members (WEBHOMEGROUP in [Table 4–3](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle instances as members (WEBINSTANCEGROUP in [Table 4–3](#)).

[Table 4–3](#) provides a summary of Oracle recommendations for consistency groups for the Oracle SOA Suite topology shown in [Figure 4–2](#).

Table 4–3 Consistency Groups for Oracle SOA Suite

Tier	Group Name	Members	Comments
Application	DOMAINGROUP	VOLADMIN VOLSOA1 VOLSOA2	Consistency group for the Administration Server, Managed Server domain directory
Application	DATAGROUP	VOLDATA	Consistency group for the JMS file store and transaction log data
Application	FMWHOMEGROUP	VOLFMW1 VOLFMW2	Consistency group for the Middleware homes
Web	WEBHOMEGROUP	VOLWEB1 VOLWEB2	Consistency group for the Oracle HTTP Server Oracle homes

Table 4–3 (Cont.) Consistency Groups for Oracle SOA Suite

Tier	Group Name	Members	Comments
Web	WEBINSTANCEGROUP	VOLWEBINST1 VOLWEBINST2 VOLSTATIC1 ¹ VOLSTATIC2 ²	Consistency group for the Oracle HTTP Server Oracle instances

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

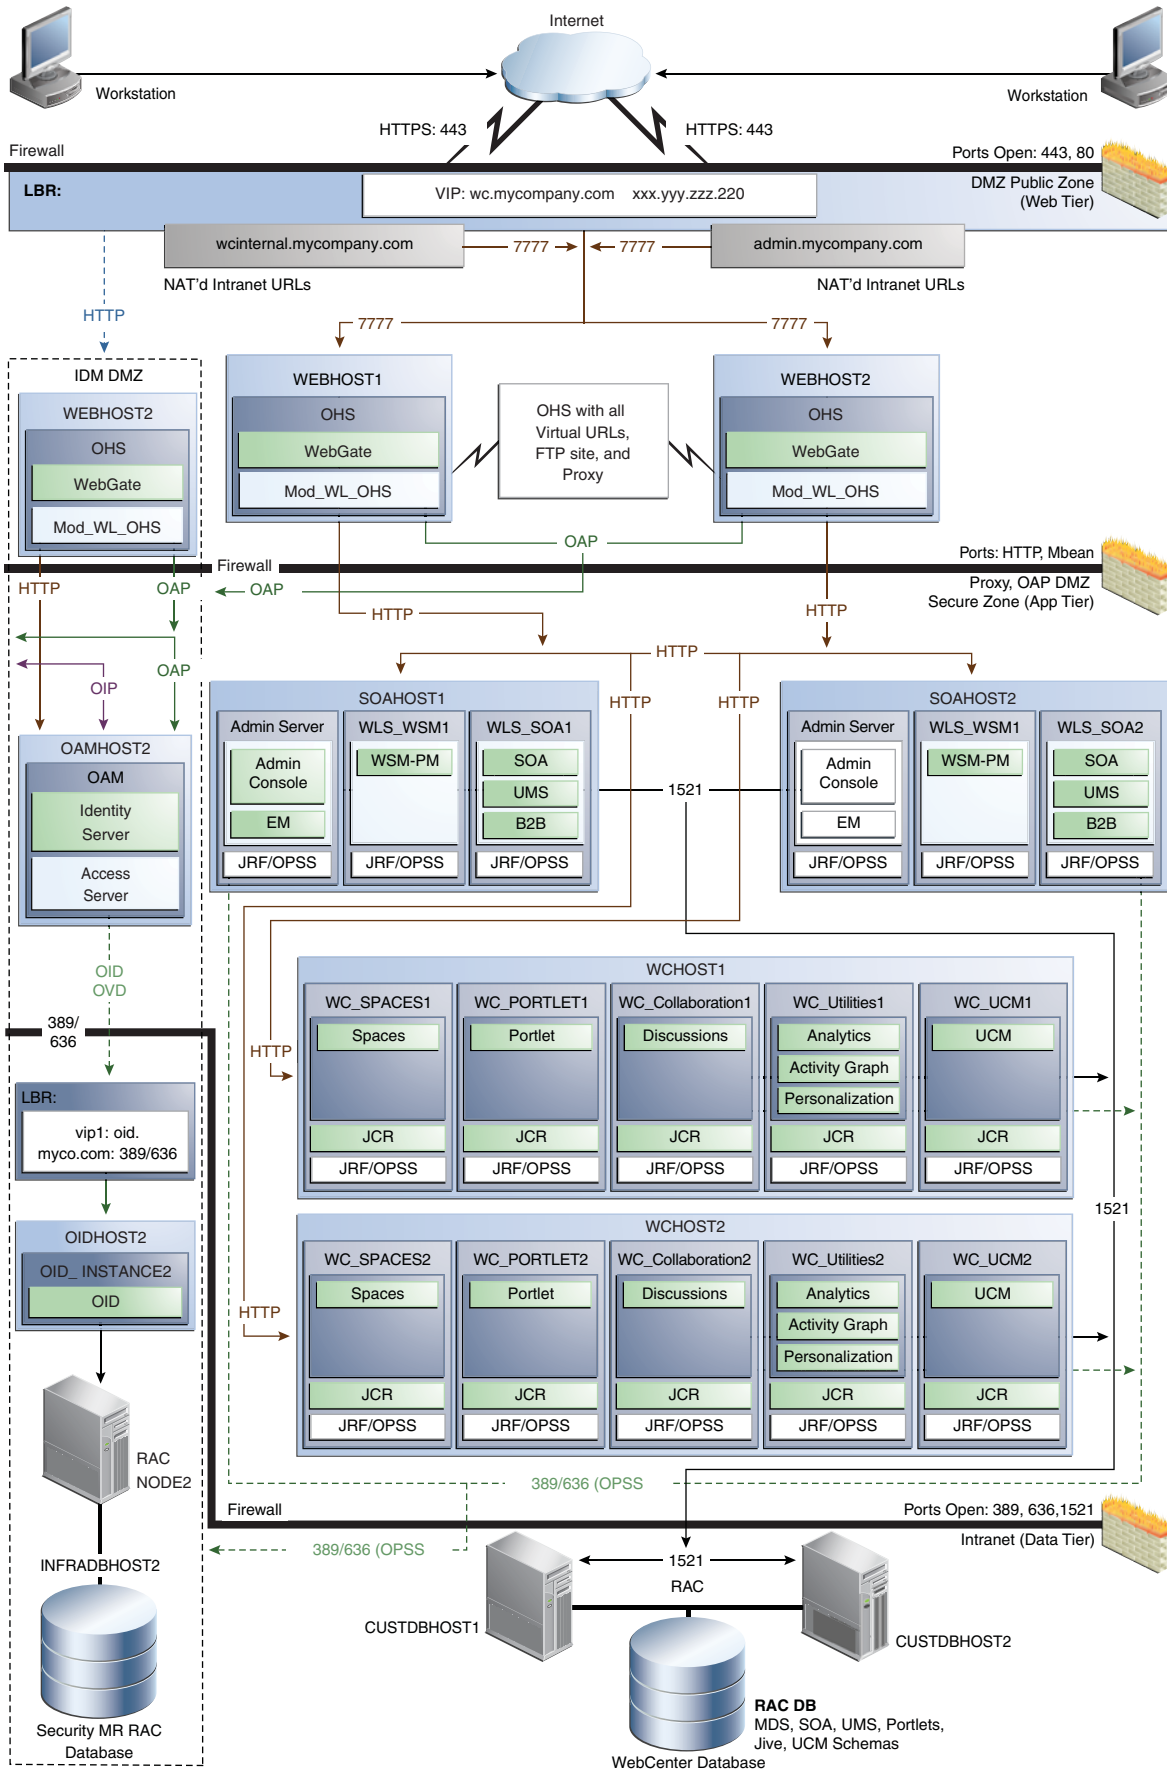
4.1.1.2 Directory Structure Recommendations for Oracle WebCenter Portal

Oracle Fusion Middleware 11g allows creating multiple Oracle WebCenter Portal Managed Servers from one single binary installation. This allows the installation of binaries in a single location on a shared storage and the reuse of this installation by the servers in different nodes. However, for maximum availability, Oracle recommends using redundant binary installations. In this model, two MW HOMEs (each of which has a WL_HOME and an ORACLE_HOME for each product suite) are installed in a shared storage. Additional servers (when scaling out or up) of the same type can use either one of these two locations without requiring more installations. Ideally, users should use two different volumes for redundant binary location, thus isolating as much as possible the failures in each volume. For additional protection, Oracle recommends using storage replication for these volumes. If multiple volumes are not available, Oracle recommends using mount points to simulate the same mount location in a different directory in the shared storage. Although this does not guarantee the protection that multiple volumes provide, it does allow protection from user deletions and individual file corruption.

Oracle also recommends separating the domain directory used by the Administration Server from the domain directory used by Managed Servers. This allows a symmetric configuration for the domain directories used by Managed Servers, and isolates the failover of the Administration Server. The domain directory for the Administration Server must reside in a shared storage to allow failover to another node with the same configuration. It is also recommended to have the Managed Servers' domain directories on a shared storage, even though having them on the local file system is also supported. This is especially true when designing a production site with the disaster recovery site in mind. [Figure 4–1](#) shows the directory structure layout for Oracle WebCenter Portal (the same directory structure layout is used for both Oracle SOA Suite and Oracle WebCenter Portal).

4.1.1.2.1 Volume Design for Oracle WebCenter Portal [Figure 4–4](#) shows an Oracle WebCenter Portal topology diagram. The volume design described in this section is for this Oracle WebCenter Portal topology. Instructions for installing and configuring this topology are provided in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal*.

Figure 4-4 *MyWCCompany Topology with Oracle Access Manager*



For disaster recovery of this Oracle WebCenter Portal topology, Oracle recommends the following volume design:

- Provision two volumes for two Middleware Homes that contain redundant product binaries (VOLFMW1 and VOLFMW2 in [Table 4-4](#))
- Provision one volume for the Administration Server domain directory (VOLADMIN in [Table 4-4](#))
- Provision one volume on each node for the Managed Server domain directory for SOA (VOLSOA1 and VOLSOA2 in [Table 4-4](#)). This directory is shared between all the Managed Servers on that node.
- Provision one volume on each node for the Managed Server domain directory for Oracle WebCenter Portal (VOLWC1 and VOLWC2 in [Table 4-4](#)). This directory is shared between all the Managed Servers on that node.
- Provision one volume for the JMS file-store and JTA transaction logs (VOLDATA in [Table 4-4](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain.
- Provision one volume on each node for the Oracle HTTP Server Oracle home (VOLWEB1 and VOLWEB2 in [Table 4-4](#)).
- Provision one volume on each node for the Oracle HTTP Server Oracle instance (VOLWEBINST1 and VOLWEBINST2 in [Table 4-4](#)).

[Table 4-4](#) provides a summary of Oracle recommendations for volume design for the Oracle WebCenter Portal topology shown in [Figure 4-4](#):

Table 4-4 Volume Design Recommendations for Oracle WebCenter Portal

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEBINST1	WEBHOST1	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLWEBINST2	WEBHOST2	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLSTATIC1 ¹	WEBHOST1	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Web	VOLSTATIC2 ²	WEBHOST2	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Application	VOLFMW1	SOAHOST1	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle SOA Suite binaries
Application	VOLFMW2	SOAHOST2	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle SOA Suite binaries.
Application	VOLADMIN	SOAHOST1	/u01/app/oracle/admin/soaDomain/admin	Volume for Administration Server domain directory

Table 4–4 (Cont.) Volume Design Recommendations for Oracle WebCenter Portal

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Application	VOLSOA1	SOAHOST1	/u01/app/oracle/admin/soaDomain/mng1	Volume for Managed Server domain directory for SOA
Application	VOLSOA2	SOAHOST2	/u01/app/oracle/admin/soaDomain/mng2	Volume for Managed Server domain directory for SOA
Application	VOLWC1	WCHOST1	/u01/app/oracle/admin/wcDomain/mng1	Volume for Managed Server domain directory for Oracle WebCenter Portal.
Application	VOLWC2	WCHOST2	/u01/app/oracle/admin/wcDomain/mng2	Volume for Managed Server domain directory for Oracle WebCenter Portal.
Application	VOLDATA	SOAHOST1, SOAHOST2, WCHOST1, WCHOST2	/u01/app/oracle/admin/soaDomain/soaCluster/jms /u01/app/oracle/admin/soaDomain/soaCluster/tlogs	Volume for transaction logs and JMS data

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.2.2 Consistency Group Recommendations for Oracle WebCenter Portal Oracle recommends the following consistency groups for the Oracle WebCenter Portal topology:

- Create one consistency group with the volumes containing the domain directories for the Administration Server and Managed Servers as members (DOMAINGROUP in [Table 4–5](#)).
- Create one consistency group with the volume containing the JMS file store and transaction log data as members (DATAGROUP in [Table 4–5](#)).
- Create one consistency group with the volume containing the Middleware Homes as members (FMWHOMEGROUP in [Table 4–5](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle homes as members (WEBHOMEGROUP in [Table 4–5](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle instances as members (WEBINSTANCEGROUP in [Table 4–5](#)).

[Table 4–5](#) provides a summary of Oracle recommendations for consistency groups for the Oracle WebCenter Portal topology shown in [Figure 4–4](#).

Table 4–5 Consistency Group for Oracle WebCenter Portal

Tier	Group Name	Members	Comments
Application	DOMAINGROUP	VOLADMIN VOLSOA1 VOLSOA2	Consistency group for the Administration Server, Managed Server domain directory
Application	DATAGROUP	VOLDATA	Consistency group for the JMS file store and transaction log data

Table 4–5 (Cont.) Consistency Group for Oracle WebCenter Portal

Tier	Group Name	Members	Comments
Application	FMWHOMEGROUP	VOLFMW1 VOLFMW2	Consistency group for the Middleware homes
Web	WEBHOMEGROUP	VOLWEB1 VOLWEB2	Consistency group for the Oracle HTTP Server Oracle homes
Web	WEBINSTANCEGROUP	VOLWEBINST1 VOLWEBINST2 VOLSTATIC1 ¹ VOLSTATIC2 ²	Consistency group for the Oracle HTTP Server Oracle instances

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.3 Directory Structure Recommendations for Oracle Identity Management

Oracle Fusion Middleware 11g allows the separation of the product binaries and the run-time artifacts for Oracle Identity Management components. The product binaries are under the `ORACLE_HOME` directory and the run-time artifacts are located under the `ORACLE_INSTANCE` directory.

In this model, for the web tier and the data tier, it is recommended to have one `ORACLE_HOME` (for product binaries) per host and one `ORACLE_INSTANCE` for an instance, installed on the shared storage. The `ORACLE_HOME` is shared among all the instances running on the host, whereas each instance has its own `ORACLE_INSTANCE` location. Additional, servers (when scaling out or up) of the same type can use either one of the same location without requiring more installations.

For the application tier, it is recommended to have one Middleware Home (`MW_HOME`) per host (each of which has a `WLS_HOME` and an `ORACLE_HOME` for each product suite) installed on the shared storage. Additional servers (when scaling out or up) of the same type can use the same location without requiring more installations.

Separation of the domain directory and the `MW_HOME` is not supported. The domain directory is under the `MW_HOME` and is shared between all the Administration Servers and Managed Servers running on the host. [Section 4–5, "Directory Structure for Oracle Identity Management"](#) shows the directory structure layout for Oracle Identity Management:

Figure 4–5 Directory Structure for Oracle Identity Management

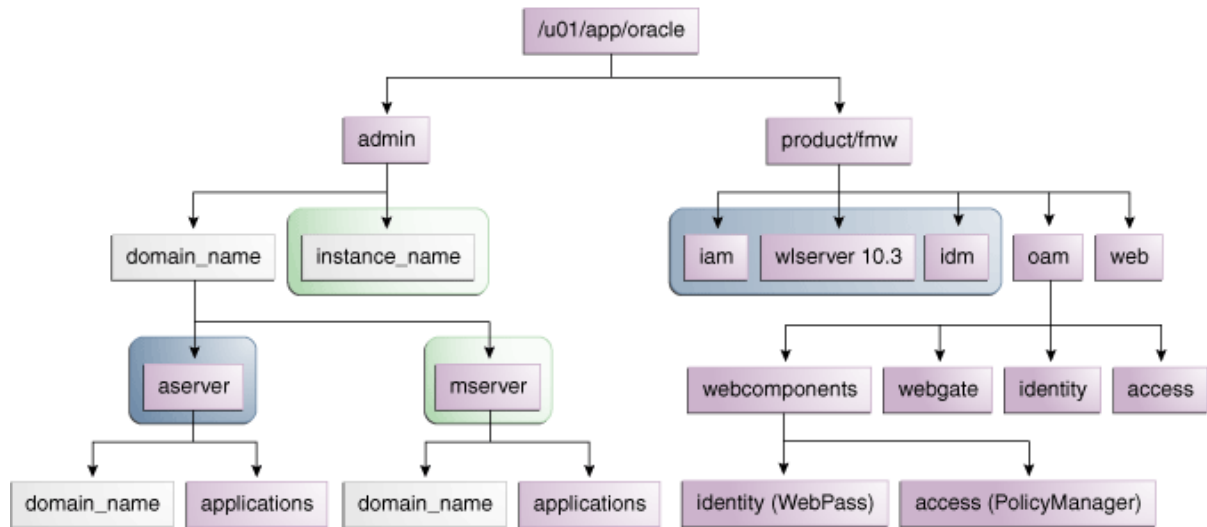






Table 4–6 Directory Structure Elements

Element	Explanation
	The Administration Server domain directories, applications, deployment plans, file adapter control directory, JMS and TX logs, and the entire MW_HOME are on a shared disk.
	The Managed Server domain directories must be on a shared disk. Further, if you want to share the Managed Server domain directories on multiple nodes, then you must mount the same shared disk location across the nodes. The <code>instance_name</code> directory for the web tier must be on a shared disk.
	Fixed name.
	Installation-dependent name.

Detailed information about setting up this directory structure is included in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management*. The directory structure in Figure 4–5 does not show other required internal directories such as `oracle_common` and `jrckit`

4.1.1.3.1 Volume Design for Oracle Identity Management Figure 4–6, Figure 4–7, and Figure 4–8 show the Oracle Identity Management topologies. The volume design described in this section is for these Oracle Identity Management topologies. Instructions for installing and configuring this topology are provided in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management*.

Figure 4–6 *MyIMCompany Topology with Oracle Access Manager*

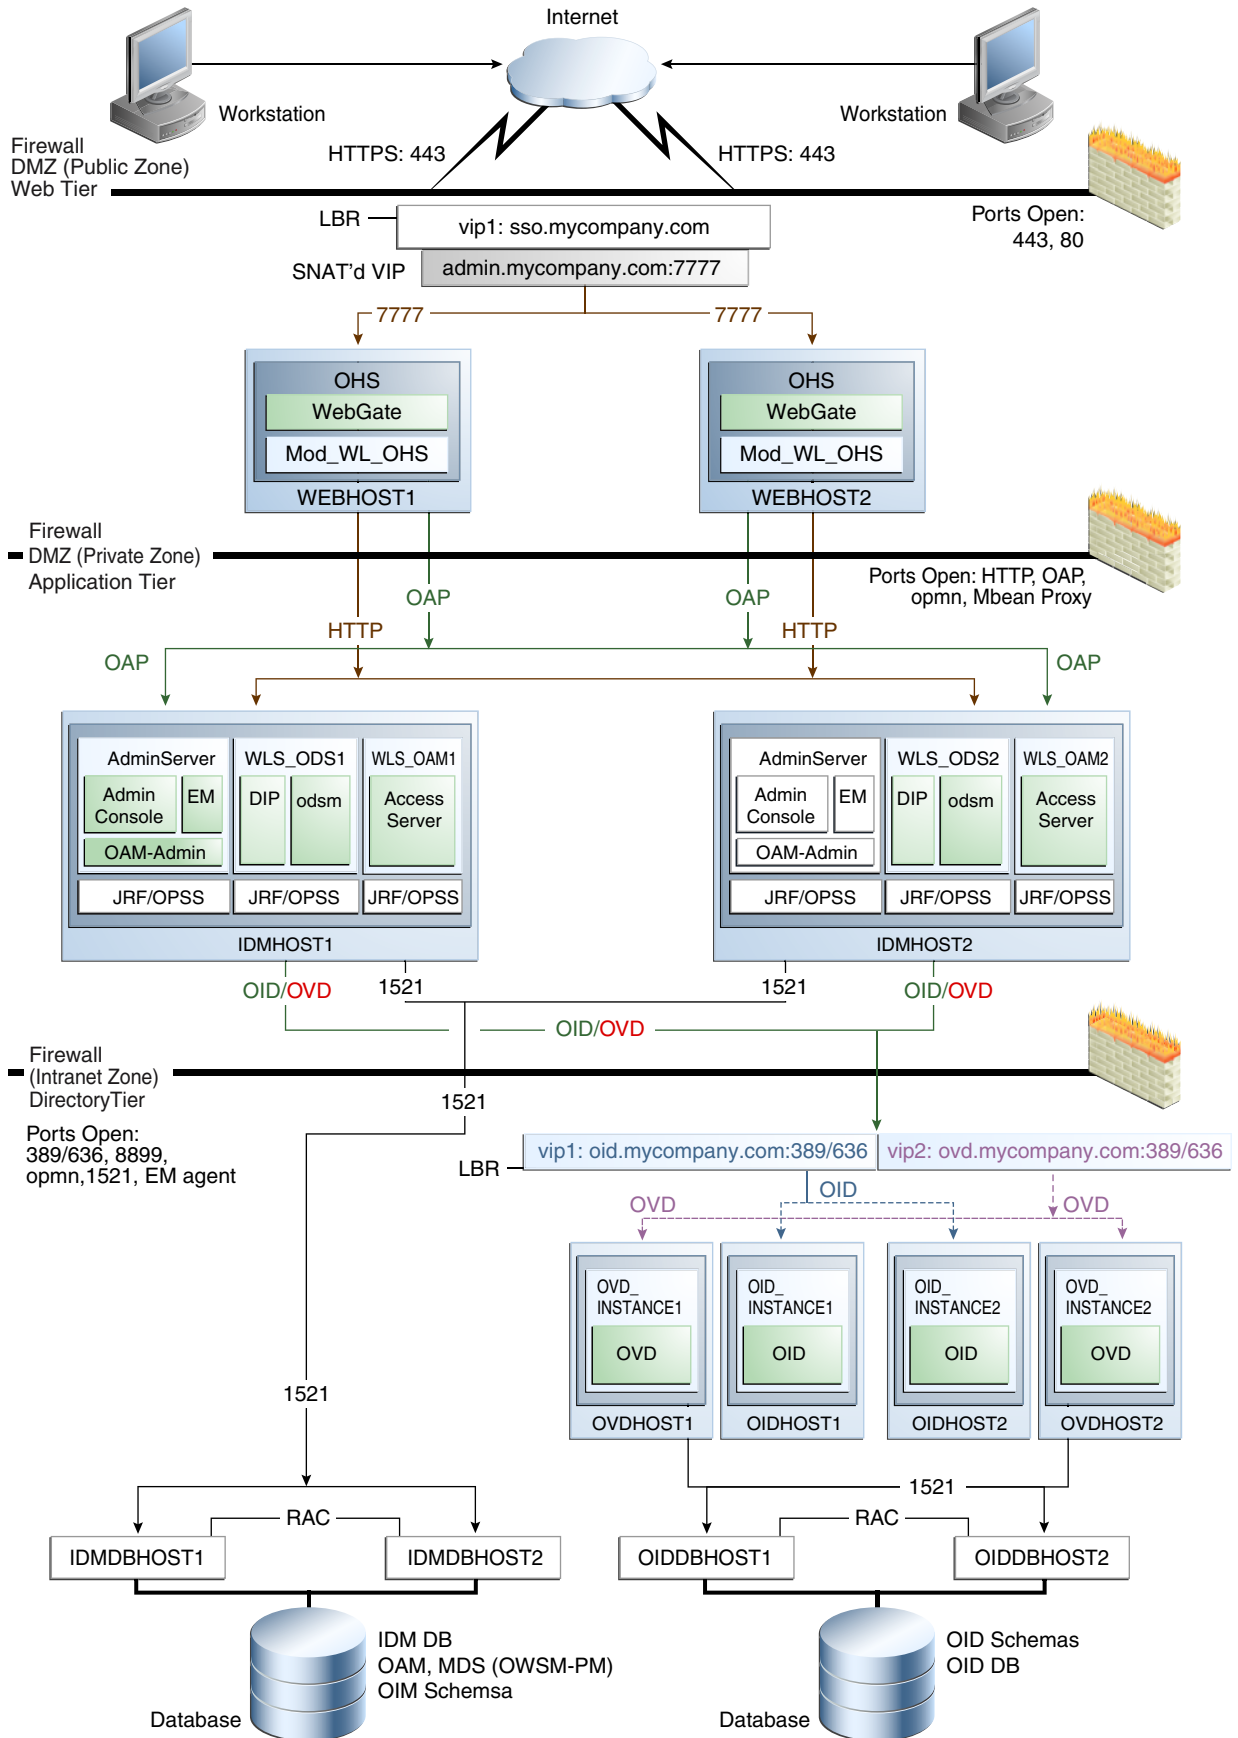


Figure 4–7 *MyIMCompany Topology with Oracle Access Manager and Oracle Identity Manager*

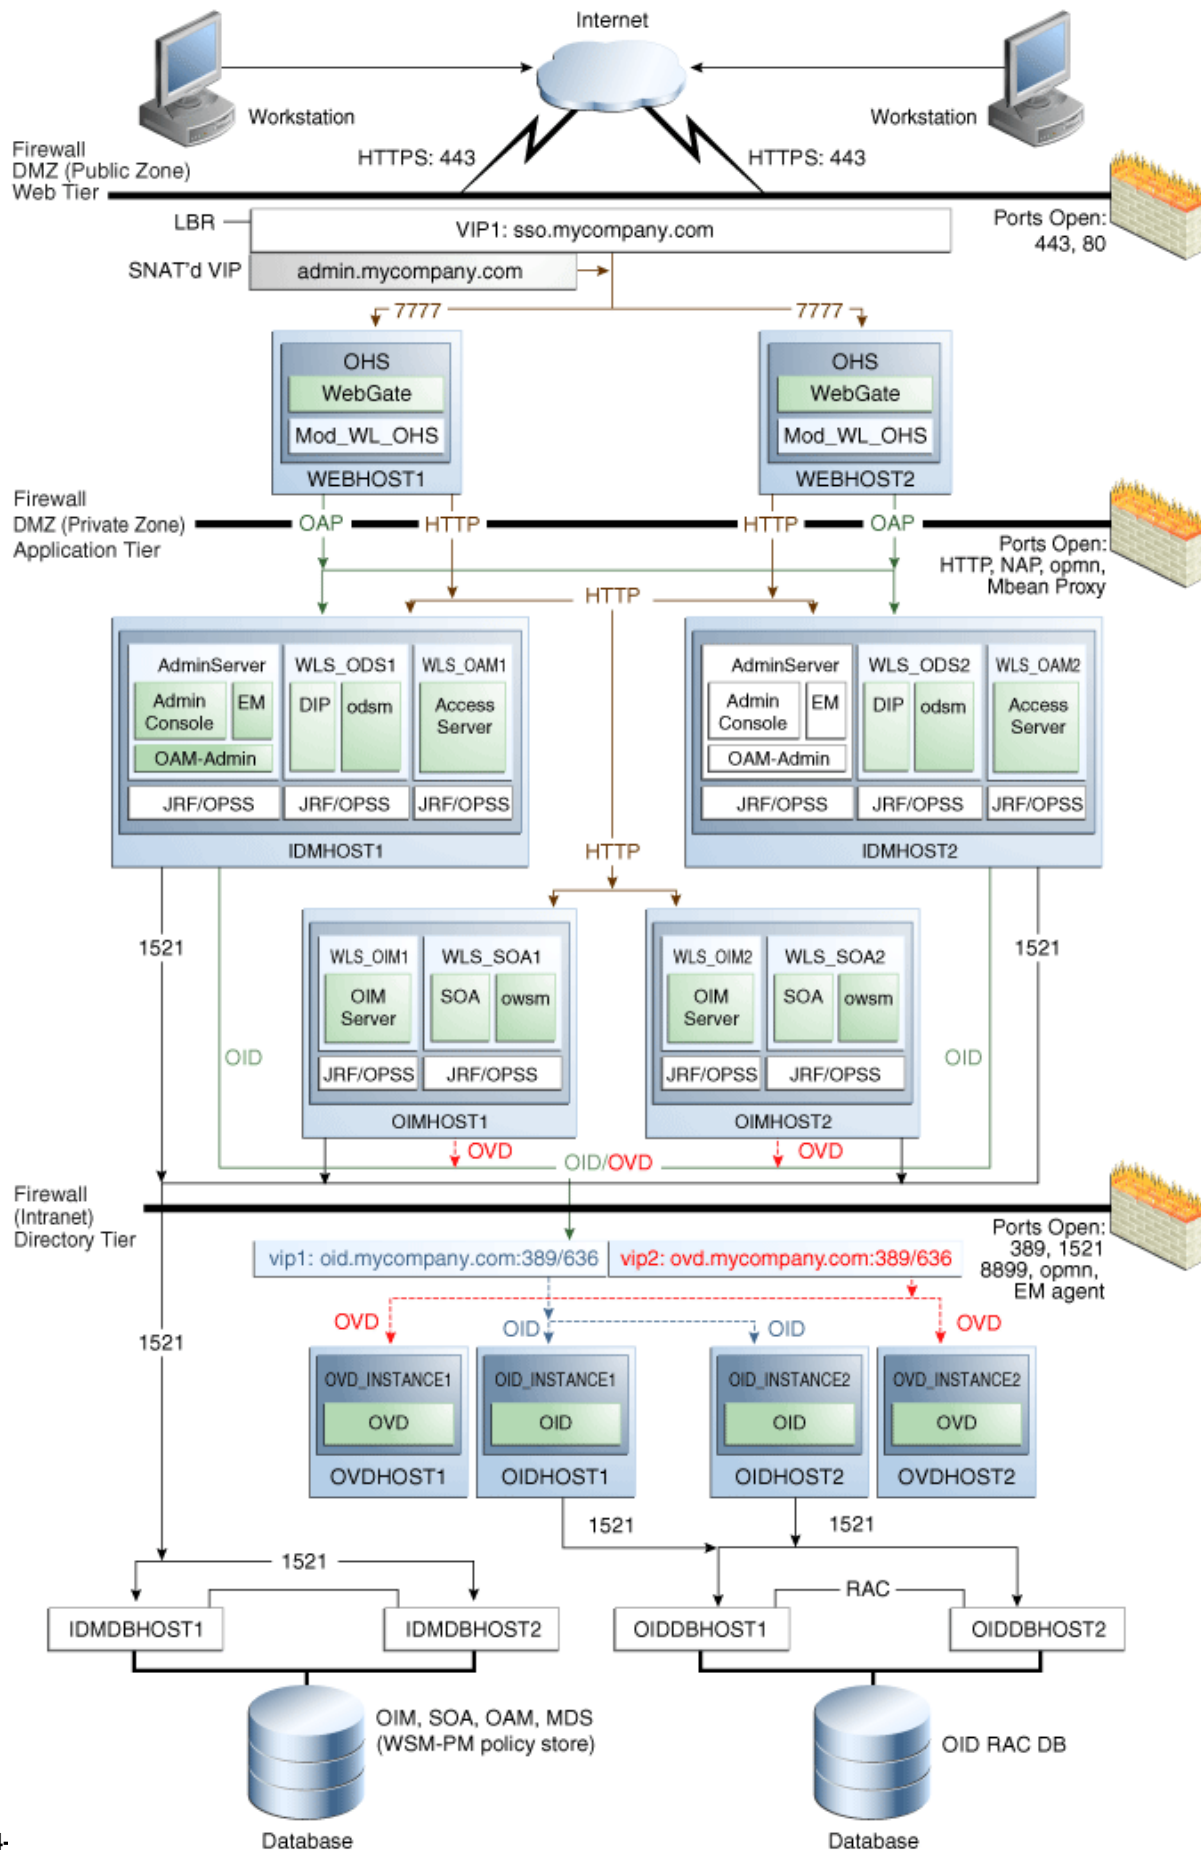
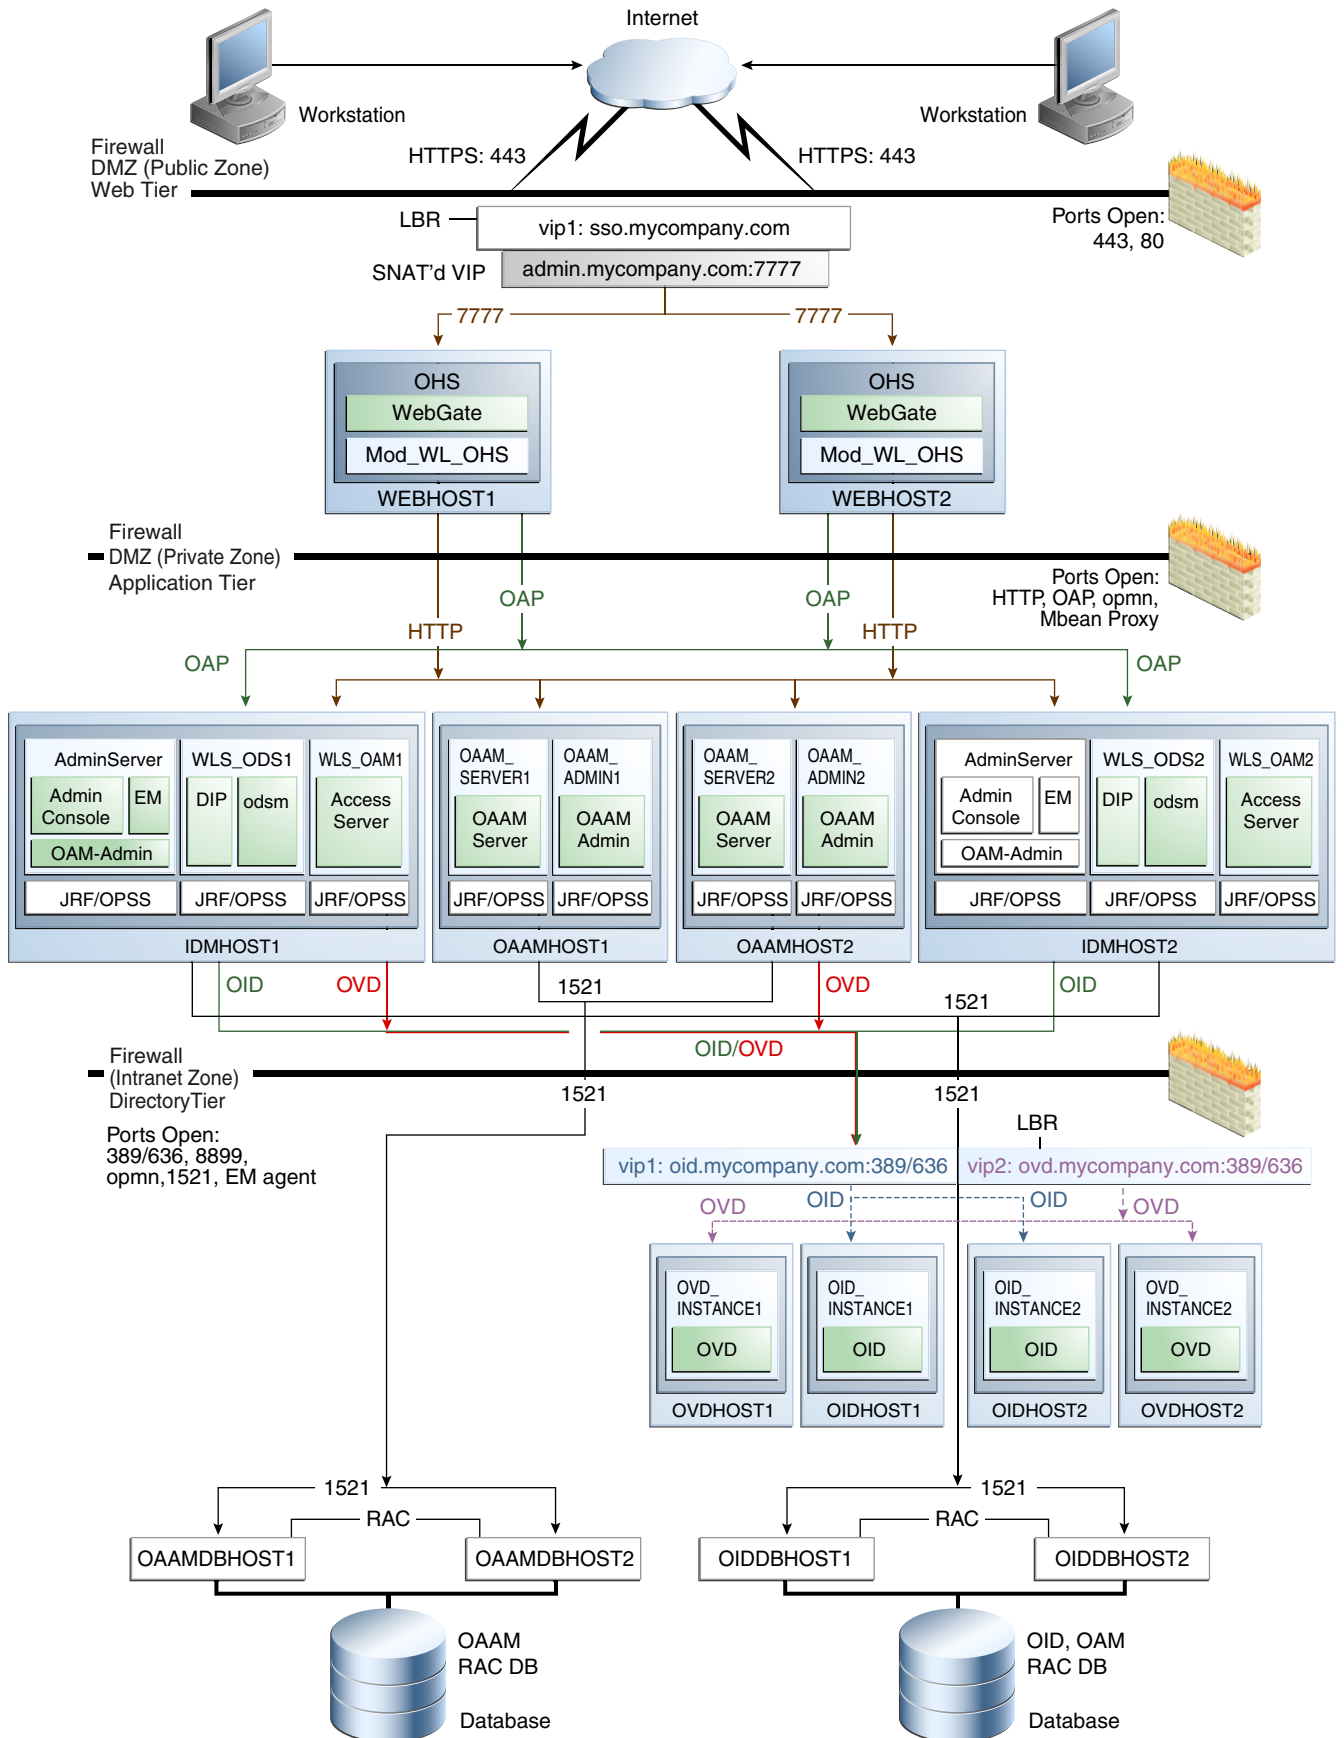


Figure 4–8 *MyIMCompany Topology with Oracle Adaptive Access Manager*



The *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management* manual describes how to set up the Oracle Identity Management enterprise deployment shown in [Figure 4-6](#), [Figure 4-7](#), and [Figure 4-8](#).

Oracle recommends the following volume design for Oracle Identity Management:

- Provision one volume on each of the Identity Management nodes for the Middleware Homes. This volume will also contain the WebLogic Server Home, Identity Management Oracle home, domain directory for the Administration Server and Managed Server running on that host. These are VOLIDM1 and VOLIDM2 in [Table 4-7](#).
- Provision one volume on each node for the Oracle homes in the directory tier and web tier. These are VOLWEB1, VOLWEB2, VOLOID1, VOLOID2, VOLOVD1, and VOLOVD2 in [Table 4-7](#).
- Provision one volume on each node for the Oracle instance home in the directory tier and web tier. These are VOLWEBINST1, VOLWEBINST2, VOLOIDINST1, VOLOIDINST2, VOLOVDINST1, and VOLOVDINST2 in [Table 4-7](#).
- Provision one volume on each node for the Identity Management Oracle instances, the Administration Server, and the Managed Server instances in the application tier. This volume is shared by the Administration Server and Managed Server instances. These are VOLIDMINST1 and VOLIDMINST2 in [Table 4-7](#).
- Provision one volume for the JMS file-store and JTA transaction logs (VOLDATA in [Table 4-7](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain.

[Table 4-7](#) provides a summary of Oracle recommendations for volume design for the Oracle Identity Management topology shown in [Figure 4-6](#):

Table 4-7 Volume Recommendations for Oracle identity Management

Tier	Volume Names	Mounted on Nodes	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installations
Web	VOLWEB2	WEBHOST2	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installations
Web	VOLWEBINST1	WEBHOST1	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instances
Web	VOLWEBINST2	WEBHOST2	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instances
Web	VOLSTATIC1 ¹	WEBHOST1	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Web	VOLSTATIC2 ²	WEBHOST2	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Application	VOLIDM1	IDMHOST1	/u01/app/oracle/product/fmw	Volume for Identity Management Middleware homes
Application	VOLIDM2	IDMHOST2	/u01/app/oracle/product/fmw	Volume for Identity Management Middleware homes

Table 4–7 (Cont.) Volume Recommendations for Oracle Identity Management

Tier	Volume Names	Mounted on Nodes	Mount Point	Comments
Application	VOLIDMINST1	IDMHOST1, OIMHOST1, and OAAMHOST1	/u01/app/oracle/admin	Volume for Oracle instances
Application	VOLIDMINST2	IDMHOST2, OIMHOST2, and OAAMHOST2	/u01/app/oracle/admin	Volume for Oracle instances
Directory	VOLOID1	OIDHOST1	/u01/app/oracle/product/fmw/idm	Volume for Oracle Internet Directory Oracle homes
Directory	VOLOID2	OIDHOST2	/u01/app/oracle/product/fmw/idm	Volume for Oracle Internet Directory Oracle homes
Directory	VOLOIDINST1	OIDHOST1	/u01/app/oracle/admin	Volume for Oracle Internet Directory Oracle instances
Directory	VOLOIDINST2	OIDHOST2	/u01/app/oracle/admin	Volume for Oracle Internet Directory Oracle instances
Directory	VOLOVD1	OVDHOST1	/u01/app/oracle/product/fmw/idm	Volume for Oracle Virtual Directory Oracle homes
Directory	VOLOVD2	OVDHOST2	/u01/app/oracle/product/fmw/idm	Volume for Oracle Virtual Directory Oracle homes
Directory	VOLOVDINST1	OVDHOST1	/u01/app/oracle/admin	Volume for Oracle Virtual Directory Oracle instances
Directory	VOLOVDINST2	OVDHOST2	/u01/app/oracle/admin	Volume for Oracle Virtual Directory Oracle instances
Application	VOLDATA	OIMHOST1, OIMHOST2	/u01/app/oracle/admin/oimDomain/soaCluster/jms /u01/app/oracle/admin/oimDomain/soaCluster/tlogs	Volume for transaction logs and JMS data

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.3.2 Consistency Group Recommendations for Oracle Identity Management

Oracle recommends the following consistency groups for the Oracle Identity Management topology:

- Create one consistency group with the volumes containing the application tier Middleware home directories as members. This is the IDMMWGROUP group in [Table 4–8](#).
- Create one consistency group with the volumes containing the application tier Oracle instances directories as members. This is the IDMINSTGROUP group in [Table 4–8](#).

- Create one consistency group with the volumes containing the Oracle Internet Directory Oracle homes as members. This is the `OIDHOMEGROUP` group in [Table 4-8](#).
- Create one consistency group with the volumes containing the Oracle Internet Directory Oracle instances as members. This is the `OIDINSTGROUP` group in [Table 4-8](#).
- Create one consistency group with the volumes containing the Oracle Virtual Directory Oracle homes as members. This is the `OVDHOMEGROUP` group in [Table 4-8](#).
- Create one consistency group with the volumes containing the Oracle Virtual Directory Oracle instances as members. This is the `OVDINSTGROUP` group in [Table 4-8](#).
- Create one consistency group with the volume containing the Oracle Access Manager Oracle homes for Oracle Access Manager Identity and Access Server components as members. This is the `OAMGROUP` group in [Table 4-8](#).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle homes as members. This is the `WEBHOMEGROUP` in [Table 4-8](#).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle instances as members. This is the `WEBINSTGROUP` group in [Table 4-8](#).
- Create one consistency group with the volume containing the JMS file store and transaction log data as members. This is the `DATAGROUP` in [Table 4-8](#).

[Table 4-8](#) provides a summary of Oracle recommendations for consistency groups for the Oracle Identity Management topology shown in [Figure 4-6](#):

Table 4-8 Consistency Groups for Oracle Identity Management

Tier	Group Name	Members	Comments
Directory	<code>OIDHOMEGROUP</code>	<code>VOLOID1</code> <code>VOLOID2</code>	Consistency group for Oracle Internet Directory Oracle homes
Directory	<code>OIDINSTGROUP</code>	<code>VOLOIDINST1</code> <code>VOLOIDINST2</code>	Consistency group for Oracle Internet Directory Oracle instances
Directory	<code>OVDHOMEGROUP</code>	<code>VOLOVD1</code> <code>VOLOVD2</code>	Consistency group for Oracle Virtual Directory Oracle homes
Directory	<code>OVDINSTGROUP</code>	<code>VOLOVDINST1</code> <code>VOLOVDINST2</code>	Consistency group for Oracle Virtual Directory Oracle instances
Application	<code>IDMMWGROUP</code>	<code>VOLIDM1</code> <code>VOLIDM2</code>	Consistency group for the Middleware homes
Application	<code>IDMINSTGROUP</code>	<code>VOLIDMINST1</code> <code>VOLIDMINST2</code>	Consistency group for the Identity Management instances
Web	<code>WEBHOMEGROUP</code>	<code>VOLWEB1</code> <code>VOLWEB2</code>	Consistency group for the Oracle HTTP Server Oracle homes

Table 4–8 (Cont.) Consistency Groups for Oracle Identity Management

Tier	Group Name	Members	Comments
Web	WEBINSTGROUP	VOLWEBINST1 VOLWEBINST2 VOLSTATIC1 ¹ VOLSTATIC2 ²	Consistency group for the Oracle HTTP Server Oracle instances

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.4 Directory Structure Recommendations for Oracle Portal, Forms, Reports, and Discoverer

Figure 4–9 shows the Oracle Portal enterprise deployment topology diagram. The volume design and consistency groups described in Section 4.1.1.4.1, "Volume Design for Oracle Portal, Forms, Reports, and Discover" and Section 4.1.1.4.2, "Consistency Group Recommendations for Oracle Portal, Forms, Reports, and Discoverer" can be used for a Disaster Recovery site that includes this Oracle Portal topology.

Detailed information about the Oracle Portal enterprise topology in Figure 4–9 is available in the 11.1.1.2 *Oracle Portal Enterprise Deployment Guide*. See Article ID 952068.1 "Oracle Fusion Middleware 11g (11.1.1.2) Enterprise Deployment Guides for Portal, Forms, Reports, and Discover" at My Oracle Support (formerly Oracle *MetaLink*) for information on obtaining the manual. The URL for My Oracle Support is:

<http://support.oracle.com>

Figure 4–9 Oracle Portal Topology Diagram

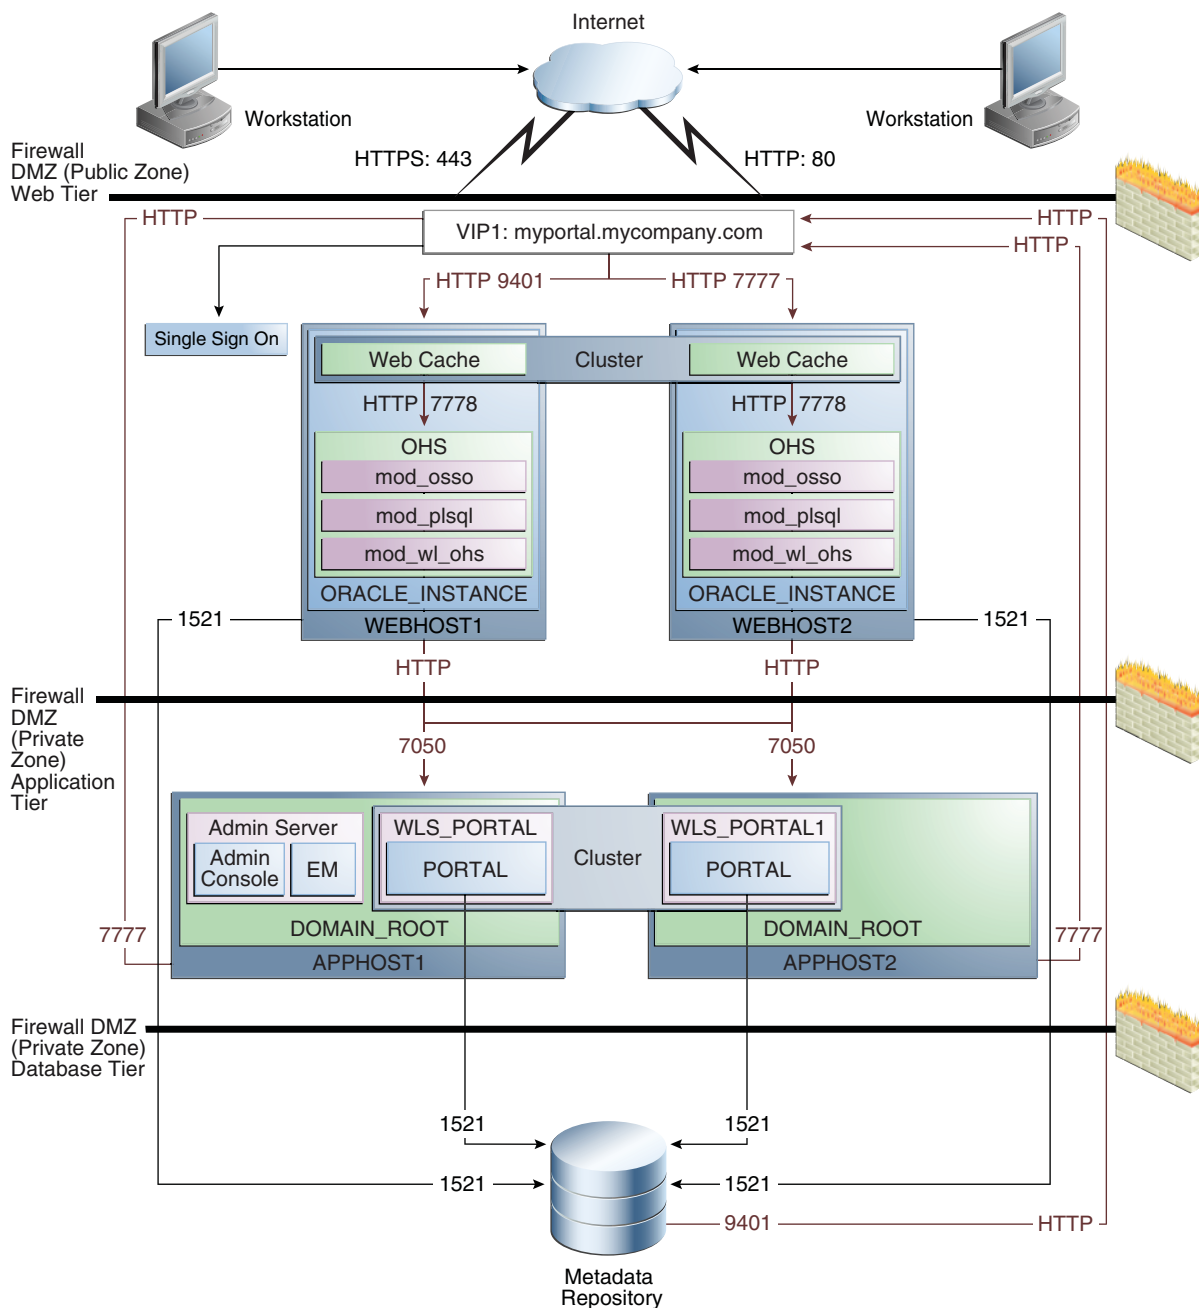
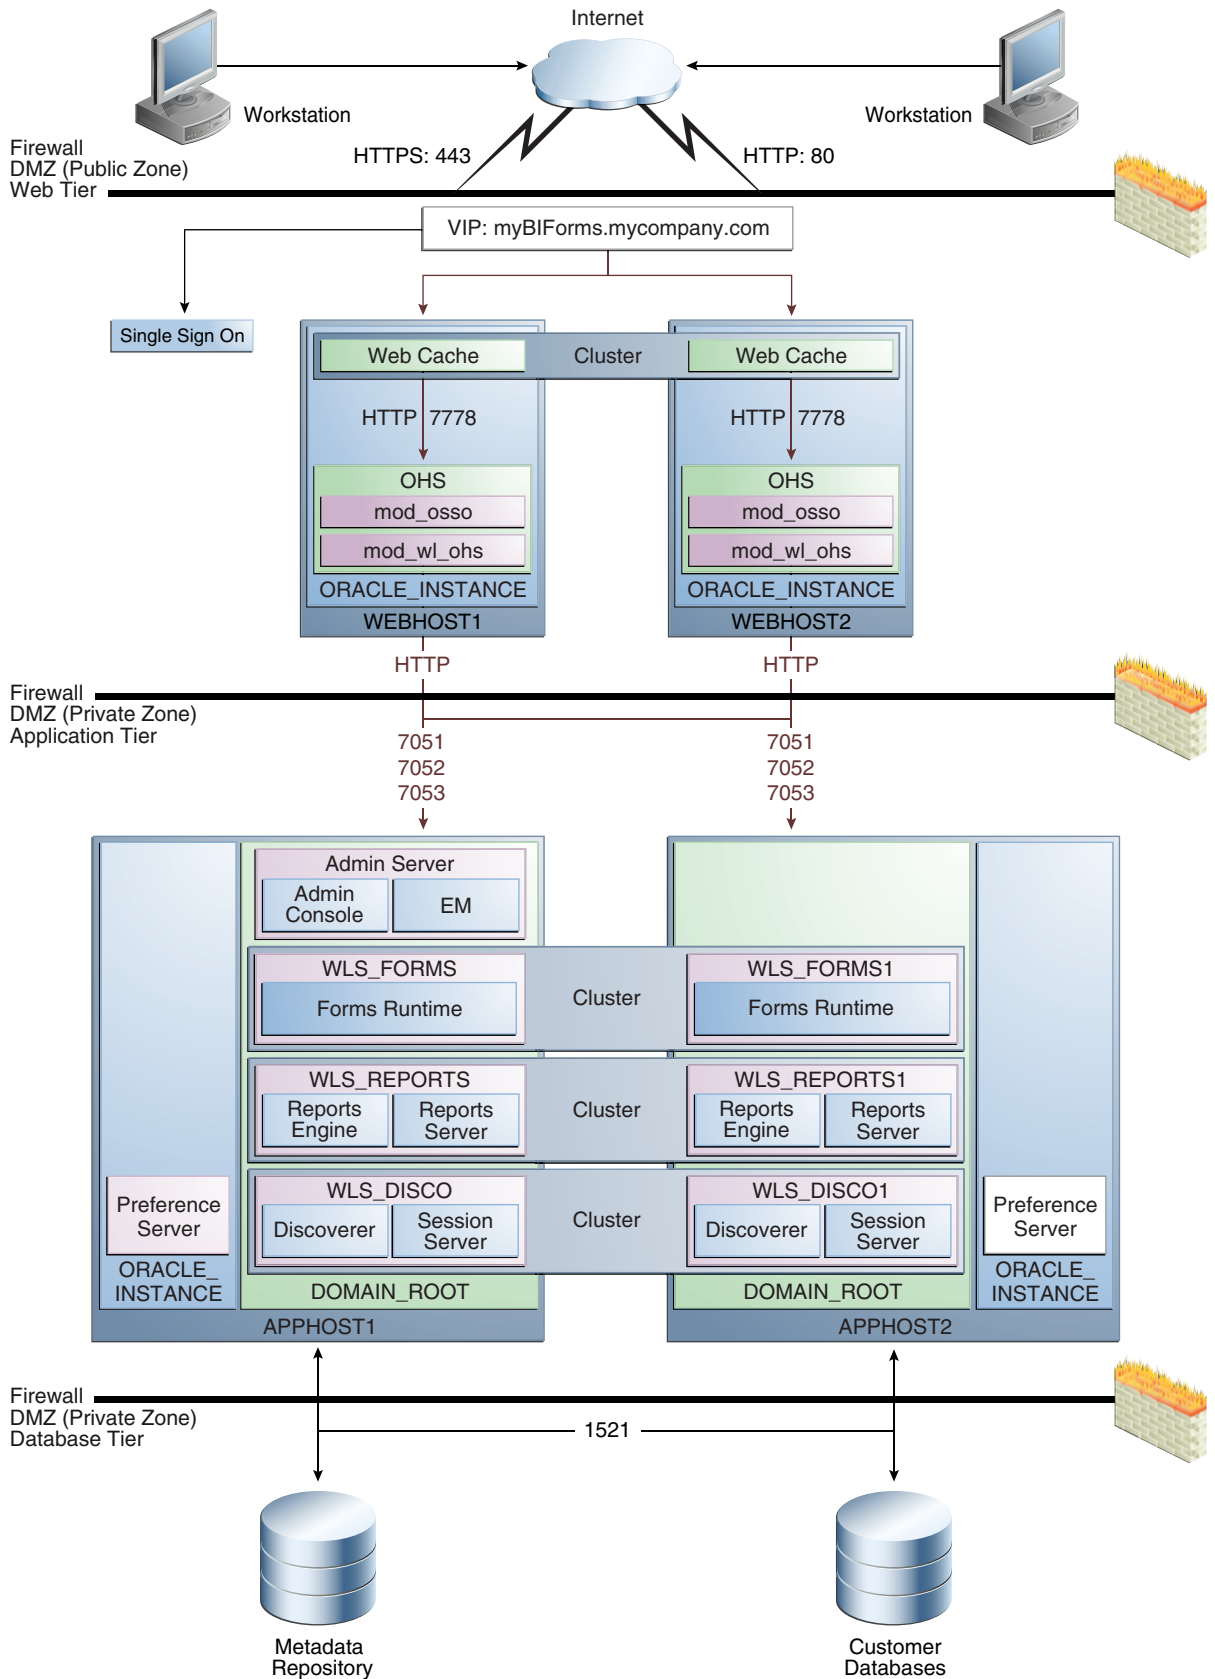


Figure 4–10 shows the Oracle Forms, Reports, and Discoverer enterprise topology diagram. The volume design and consistency groups described in Section 4.1.1.4.1, "Volume Design for Oracle Portal, Forms, Reports, and Discover" and Section 4.1.1.4.2, "Consistency Group Recommendations for Oracle Portal, Forms, Reports, and Discoverer" can be used for a Disaster Recovery site that includes this topology.

Detailed information about the Oracle Forms, Reports, and Discoverer enterprise topology in Figure 4–10 is available in the 11.1.1.2 *Oracle Forms, Reports, and Discoverer Enterprise Deployment Guide*. See Article ID 952068.1 "Oracle Fusion Middleware 11g (11.1.1.2) Enterprise Deployment Guides for Portal, Forms, Reports, and Discover" at My Oracle Support (formerly Oracle *MetaLink*) for information on obtaining the manual. The URL for My Oracle Support is:

<http://support.oracle.com>

Figure 4–10 *Oracle Forms, Reports, and Discoverer Topology*



4.1.1.4.1 Volume Design for Oracle Portal, Forms, Reports, and Discover Oracle recommends the following volume design for a Disaster Recovery site that includes both the Oracle Portal topology shown in [Figure 4-9](#) and the Oracle Forms, Reports, and Discoverer topology shown in [Figure 4-10](#):

- Provision one volume on each of the application tier hosts for the Middleware Homes. This volume will also contain the WebLogic Server Home, Oracle home for the Oracle Portal, Reports, Forms, and Discoverer components, and the domain directory for the Administration Server and Managed Server running on that host. These are VOLPFRD1 and VOLPFRD2 in [Table 4-9](#).
- Provision one volume on each node for the Oracle homes in the web tier. These are VOLWEB1 and VOLWEB2 in [Table 4-9](#).
- Provision one volume on each node for the Oracle instance homes in the directory web tier. These are VOLWEBINST1 and VOLWEBINST2 in [Table 4-9](#).
- Provision one volume on each node for the Oracle Instance homes in the application tier. This volume is shared by the Administration Server and Managed Server instances. These are VOLPFRDINST1 and VOLPFRDINST2 in [Table 4-9](#).
- Provision one volume for the Oracle Reports output directory in the application tier. This volume is mounted on all the nodes running the Oracle Reports server. This is VOLREPOUT in [Table 4-7](#).

[Table 4-9](#) provides a summary of Oracle recommendations for volume design for a Disaster Recovery site that includes both the Oracle Portal topology shown in [Figure 4-9](#) and the Oracle Forms, Reports, and Discoverer topology shown in [Figure 4-10](#):

Table 4-9 Volume Design Recommendations for Oracle Portal, Reports, Forms, and Discoverer

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEBINST1	WEBHOST1	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLWEBINST2	WEBHOST2	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLSTATIC1 ¹	WEBHOST1	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Web	VOLSTATIC2 ²	WEBHOST2	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Application	VOLPFRD1	APPHOST1	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle Portal, Forms, Reports, and Discoverer binaries.
Application	VOLPFRD2	APPHOST2	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle Portal, Forms, Reports, and Discoverer binaries.

Table 4–9 (Cont.) Volume Design Recommendations for Oracle Portal, Reports, Forms, and Discoverer

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Application	VOLPFRDINST1	APPHOST1	/u01/app/oracle/admin	Volume for Oracle instances
Application	VOLPFRDINST2	APPHOST2	/u01/app/oracle/admin	Volume for Oracle instances
Application	VOLREPOUT	APPHOST1, APPHOST2	/u01/app/oracle/admin	Volume for report output

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it

4.1.1.4.2 Consistency Group Recommendations for Oracle Portal, Forms, Reports, and Discoverer

Oracle recommends the following consistency groups for a Disaster Recovery site that includes both the Oracle Portal topology shown in [Figure 4–9](#) and the Oracle Forms, Reports, and Discoverer topology shown in [Figure 4–10](#):

- Create one consistency group with the volumes containing the web tier Oracle homes as members. This is WEBHOMEGROUP in [Table 4–10](#).
- Create one consistency group with the volumes containing the web tier Oracle Instances as members. This is WEBINSTGROUP in [Table 4–10](#).
- Create one consistency group with the volumes containing the application tier Middleware homes. This is PFRDMWGROUP in [Table 4–10](#).
- Create one consistency group with the volumes containing the application tier Oracle instance homes. This is PFRDINSTGROUP in [Table 4–10](#).
- Create one consistency group with the volume containing the Oracle Reports output directory as a member. This is REPOUTGROUP in [Table 4–10](#).

[Table 4–10](#) summarizes the consistency group recommendations for a Disaster Recovery site that includes both the Oracle Portal topology shown in [Figure 4–9](#) and the Oracle Forms, Reports, and Discoverer topology shown in [Figure 4–10](#).

Table 4–10 Consistency Groups for Oracle Portal, Forms, Reports, and Discoverer

Tier	Volume Name	Members	Comments
Application	PFRDMWGROUP	VOLPFRD2 VOLPFRD2	Consistency group for Middleware homes
Application	PFRDINSTGROUP	VOLPFRDINST1 VOLPFRDINST2	Consistency group for the instance homes
Application	REPOUTGROUP	VOLREPOUT	Consistency group for the Reports output directory
Web	WEBHOMEGROUP	VOLWEB1 VOLWEB2	Consistency group for the Oracle HTTP Server Oracle homes
Web	WEBINSTGROUP	VOLWEBINST1 VOLWEBINST2 VOLSTATIC1 ¹ VOLSTATIC2 ²	Consistency group for the Oracle HTTP Server Oracle instance

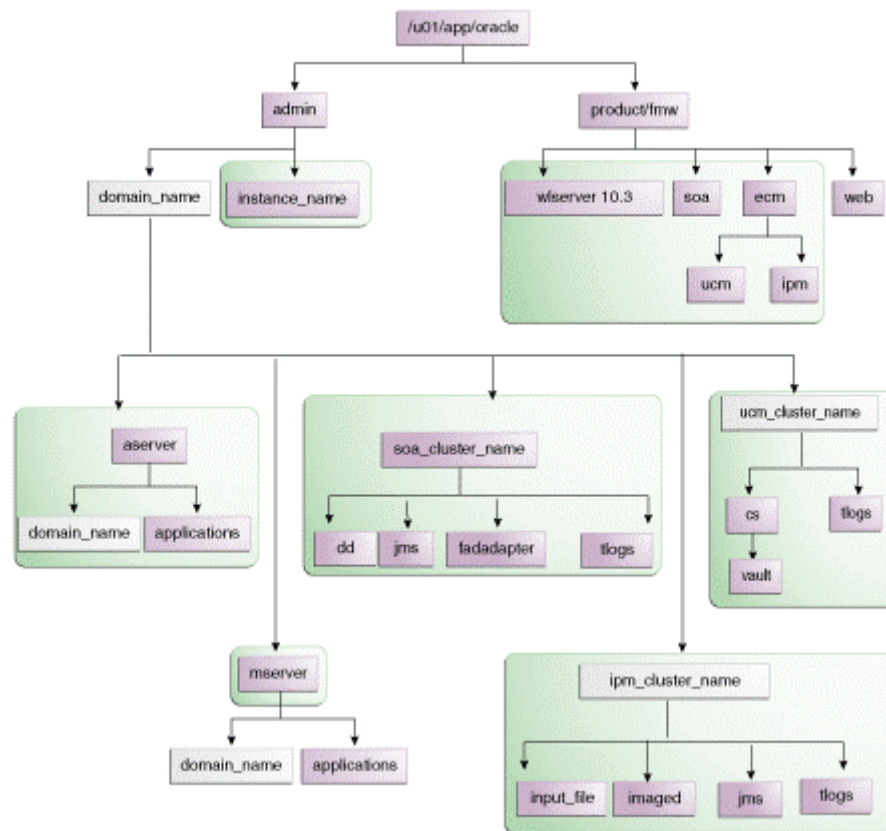
¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.5 Directory Structure Recommendations for Oracle WebCenter Content

Oracle Fusion Middleware 11g allows creating multiple Oracle Enterprise Content Management Managed Servers from one single binary installation. This allows the installation of binaries in a single location on a shared storage and the reuse of this installation by the servers in different nodes. However, for maximum availability, Oracle recommends using redundant binary installations. In this model, two MW HOMEs (each of which has a WL_HOME and an ORACLE_HOME for each product suite) are installed in a shared storage. Additional servers (when scaling out or up) of the same type can use either one of these two locations without requiring more installations. Ideally, users should use two different volumes for redundant binary location, thus isolating as much as possible the failures in each volume. For additional protection, Oracle recommends using storage replication for these volumes. If multiple volumes are not available, Oracle recommends using mount points to simulate the same mount location in a different directory in the shared storage. Although this does not guarantee the protection that multiple volumes provide, it does allow protection from user deletions and individual file corruption.




Oracle also recommends separating the domain directory used by the Administration Server from the domain directory used by Managed Servers. This allows a symmetric configuration for the domain directories used by Managed Servers, and isolates the failover of the Administration Server. The domain directory for the Administration Server must reside in a shared storage to allow failover to another node with the same configuration. It is also recommended to have the Managed Servers' domain directories on a shared storage, even though having them on the local file system is also supported. This is especially true when designing a production site with the disaster recovery site in mind. [Figure 4-11](#) represents the directory structure layout for Oracle WebCenter Content Suite.

Figure 4–11 Directory Structure for Oracle WebCenter Content

The directory structure in [Figure 4–11](#) does not show other required internal directories such as `oracle_common` and `jrockit`.

[Table 4–11](#) explains what the color-coded elements in [Figure 4–11](#) mean. The directory structure in [Figure 4–11](#) does not show other required internal directories such as `oracle_common` and `jrockit`.

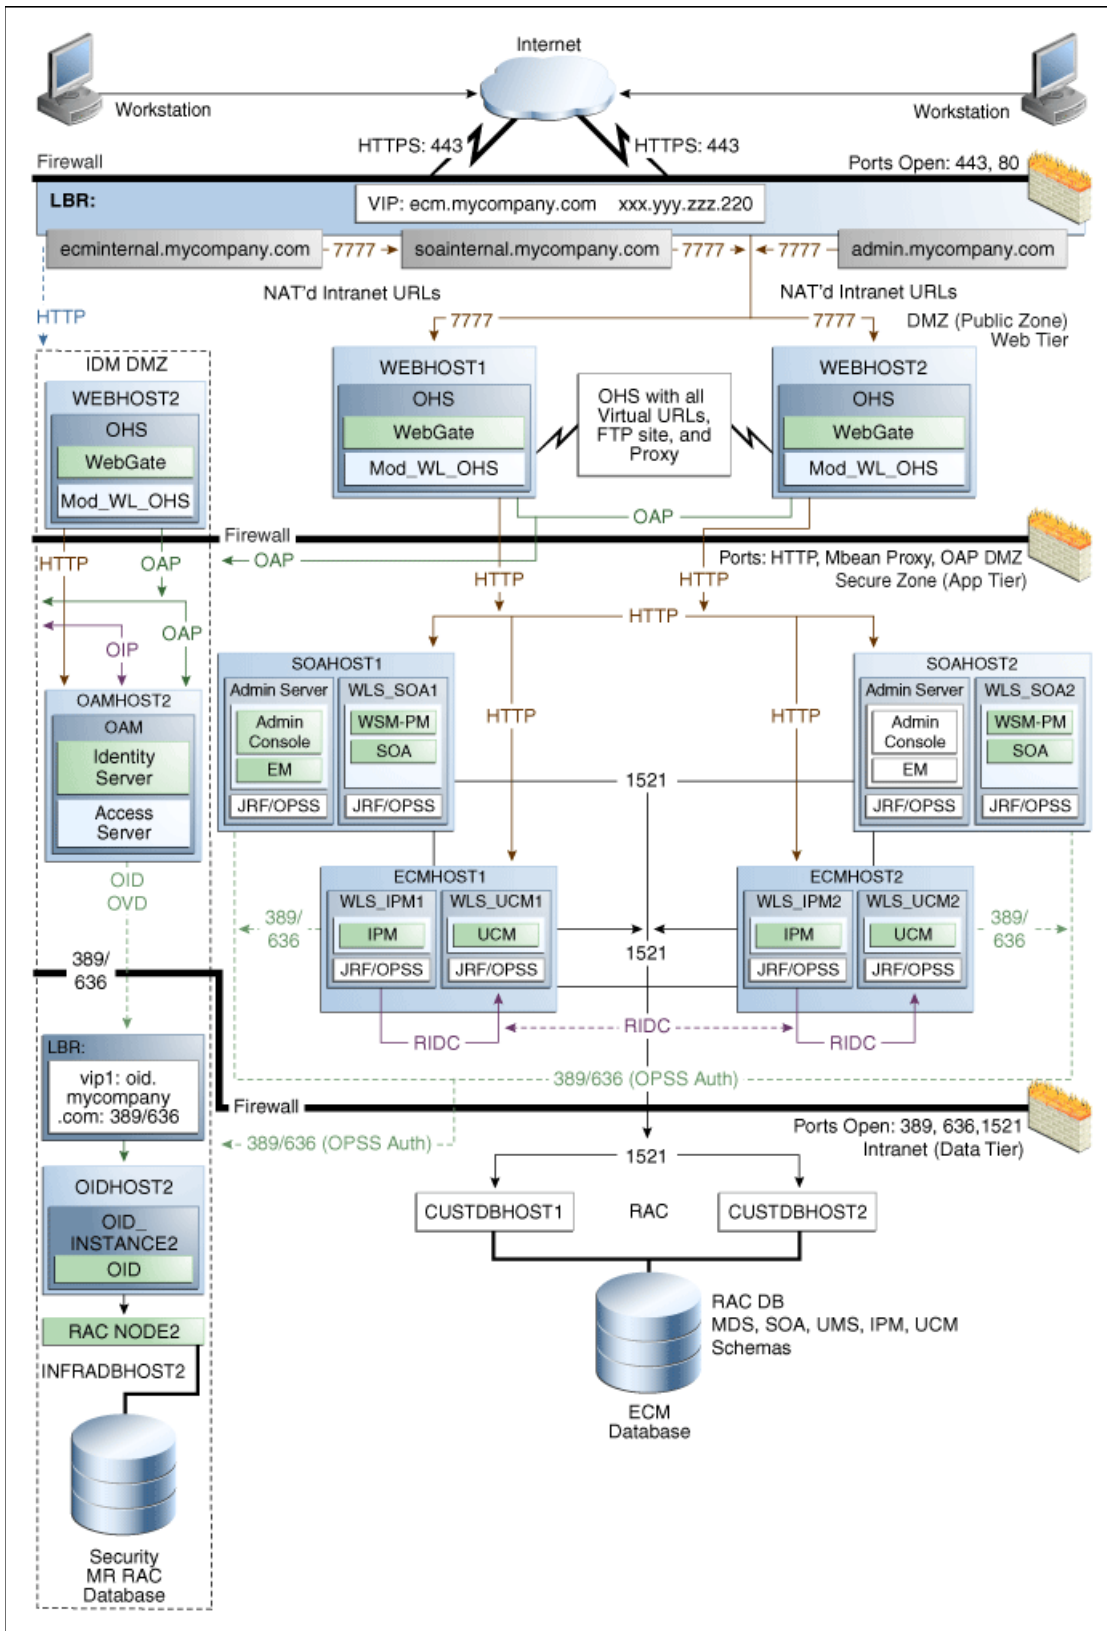
Table 4–11 Directory Structure Elements

Element	Explanation
	The Administration Server domain directories, the Managed Server domain directories, applications, deployment plans, file adapter control directory, JMS and TX logs, and the entire <code>MW_HOME</code> are on a shared disk.
	Fixed name.
	Installation-dependent name.

Detailed information about setting up this directory structure is included in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content*.

4.1.1.5.1 Volume Design for Oracle WebCenter Content [Figure 4–12](#) shows an Oracle WebCenter Content topology diagram. The volume design described in this section is for this Oracle WebCenter Content topology. Detailed instructions for installing and configuring this topology are provided in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content*.

Figure 4-12 Reference Topology for Oracle WebCenter Content Enterprise Deployment



For disaster recovery of this Oracle WebCenter Content topology, Oracle recommends the following volume design:

- Provision two volumes for two Middleware Homes that contain redundant product binaries (VOLFMW1 and VOLFMW2 in [Table 4–12](#))
- Provision one volume for the Administration Server domain directory (VOLADMIN in [Table 4–12](#))
- Provision one volume on each node for the Managed Server domain directory (VOLWCC1 and VOLWCC2 in [Table 4–12](#)). This directory is shared between all the Managed Servers on that node.
- Provision one volume for the JMS file-store and JTA transaction logs (VOLDATA in [Table 4–12](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain.
- Provision one volume on each node for the Oracle HTTP Server Oracle home (VOLWEB1 and VOLWEB2 in [Table 4–12](#)).
- Provision one volume on each node for the Oracle HTTP Server Oracle instance (VOLWEBINST1 and VOLWEBINST2 in [Table 4–12](#)).

[Table 4–12](#) provides a summary of Oracle recommendations for volume design for the Oracle Enterprise Content Management Suite topology shown in [Figure 4–12](#):

Table 4–12 Volume Design Recommendations for Oracle WebCenter Content

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEBINST1	WEBHOST1	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLWEBINST2	WEBHOST2	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLSTATIC1 ¹	WEBHOST1	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Web	VOLSTATIC2 ²	WEBHOST2	/u01/app/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Application	VOLFMW1	WCCHOST1	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle SOA Suite binaries
Application	VOLFMW2	WCCHOST2	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle SOA Suite binaries.
Application	VOLADMIN	WCCHOST1	/u01/app/oracle/admin/wccDomain/admin	Volume for Administration Server domain directory

Table 4–12 (Cont.) Volume Design Recommendations for Oracle WebCenter Content

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Application	VOLWCC1	WCCHOST1	/u01/app/oracle/admin/wccDomain/mng1	Volume for Managed Server domain directory
Application	VOLWCC2	WCCHOST2	/u01/app/oracle/admin/wccDomain/mng2	Volume for Managed Server domain directory
Application	VOLDATA	WCCHOST1, WCCHOST2	/u01/app/oracle/admin/wccDomain/soaCluster/jms /u01/app/oracle/admin/wccDomain/soaCluster/tlogs	Volume for transaction logs and JMS data Other UCM related directories is configured outside the domain Volume for UCM files (Vault and Web layout) is outside the domain

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.5.2 Consistency Group Recommendations for Oracle WebCenter Content Oracle recommends the following consistency groups for the Oracle WebCenter Content topology:

- Create one consistency group with the volumes containing the domain directories for the Administration Server and Managed Servers as members (DOMAINGROUP in [Table 4–13](#)).
- Create one consistency group with the volume containing the JMS file store and transaction log data as members (DATAGROUP in [Table 4–13](#)).
- Create one consistency group with the volume containing the Middleware Homes as members (FMWHOMEGROUP in [Table 4–13](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle homes as members (WEBHOMEGROUP in [Table 4–13](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle instances as members (WEBINSTANCEGROUP in [Table 4–13](#)).

[Table 4–13](#) provides a summary of Oracle recommendations for consistency groups for the Oracle WebCenter Content topology shown in [Figure 4–12](#).

Table 4–13 Consistency Groups for Oracle WebCenter Content

Tier	Group Name	Members	Comments
Application	DOMAINGROUP	VOLADMIN VOLWCC1 VOLWCC2	Consistency group for the Administration Server, Managed Server domain directory
Application	DATAGROUP	VOLDATA	Consistency group for the JMS file store and transaction log data

Table 4–13 (Cont.) Consistency Groups for Oracle WebCenter Content

Tier	Group Name	Members	Comments
Application	FMWHOMEGROUP	VOLFMW1 VOLFMW2	Consistency group for the Middleware homes
Web	WEBHOMEGROUP	VOLWEB1 VOLWEB2	Consistency group for the Oracle HTTP Server Oracle homes
Web	WEBINSTANCEGROUP	VOLWEBINST1 VOLWEBINST2 VOLSTATIC1 ¹ VOLSTATIC2 ²	Consistency group for the Oracle HTTP Server Oracle instances

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

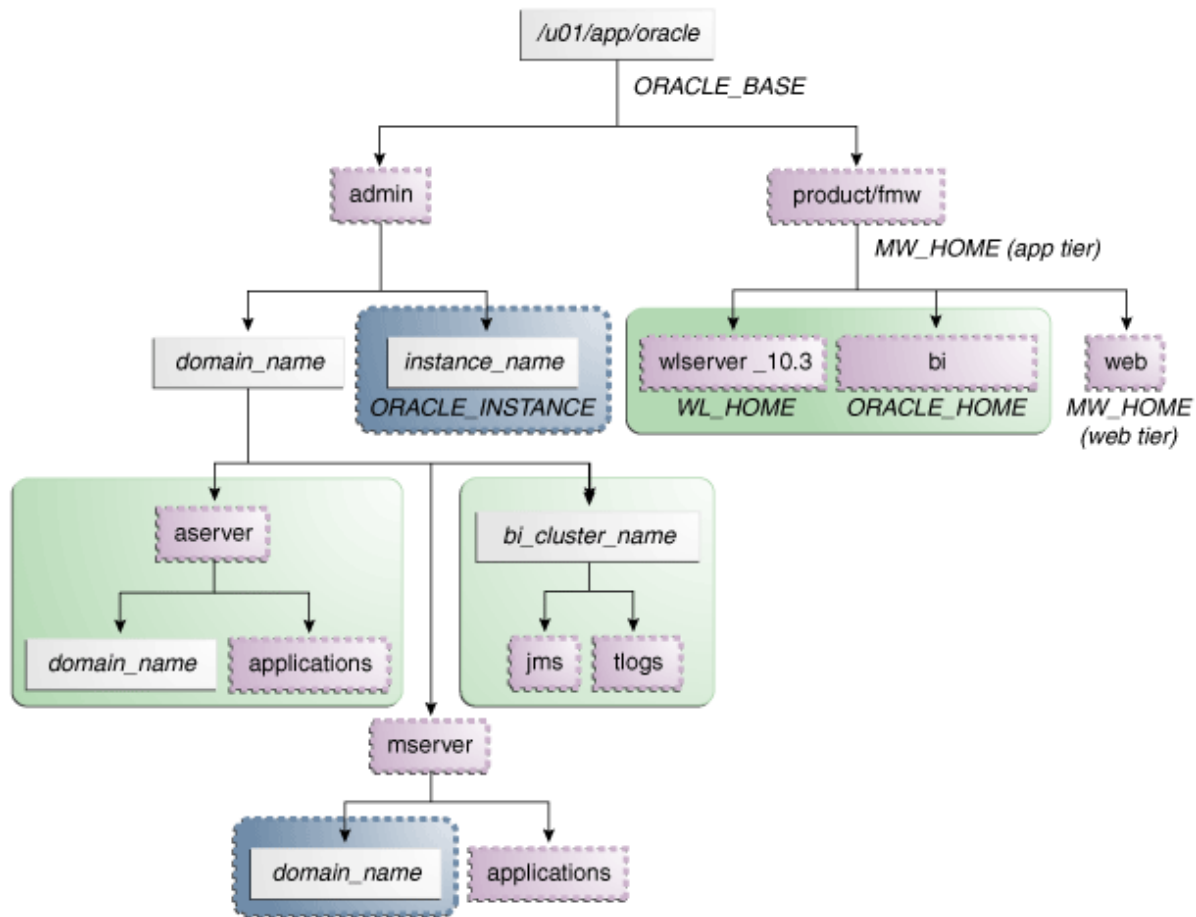
² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.6 Directory Structure Recommendations for Oracle Business Intelligence

Oracle Fusion Middleware 11g allows creating multiple Oracle Business Intelligence Managed Servers from one single binary installation. This allows the installation of binaries in a single location on a shared storage and the reuse of this installation by the servers in different nodes. However, for maximum availability, Oracle recommends using redundant binary installations. In this model, two MW HOMEs (each of which has a WL_HOME and an ORACLE_HOME for each product suite) are installed in a shared storage. Additional servers (when scaling out or up) of the same type can use either one of these two locations without requiring more installations. Ideally, users should use two different volumes for redundant binary location, thus isolating as much as possible the failures in each volume. For additional protection, Oracle recommends using storage replication for these volumes. If multiple volumes are not available, Oracle recommends using mount points to simulate the same mount location in a different directory in the shared storage. Although this does not guarantee the protection that multiple volumes provide, it does allow protection from user deletions and individual file corruption.

Oracle also recommends separating the domain directory used by the Administration Server from the domain directory used by Managed Servers. This allows a symmetric configuration for the domain directories used by Managed Servers, and isolates the failover of the Administration Server. The domain directory for the Administration Server must reside in a shared storage to allow failover to another node with the same configuration. It is also recommended to have the Managed Servers' domain directories on a shared storage, even though having them on the local file system is also supported. This is especially true when designing a production site with the disaster recovery site in mind. [Figure 4–13](#) represents the directory structure layout for Oracle Business Intelligence Suite.




Figure 4–13 Directory Structure for Oracle Business Intelligence



The directory structure in [Figure 4–13](#) does not show other required internal directories such as `oracle_common` and `jrockit`.

[Table 4–14](#) explains what the color-coded elements in [Figure 4–13](#) mean. The directory structure in [Figure 4–13](#) does not show other required internal directories such as `oracle_common` and `jrockit`.

Table 4–14 Directory Structure Elements

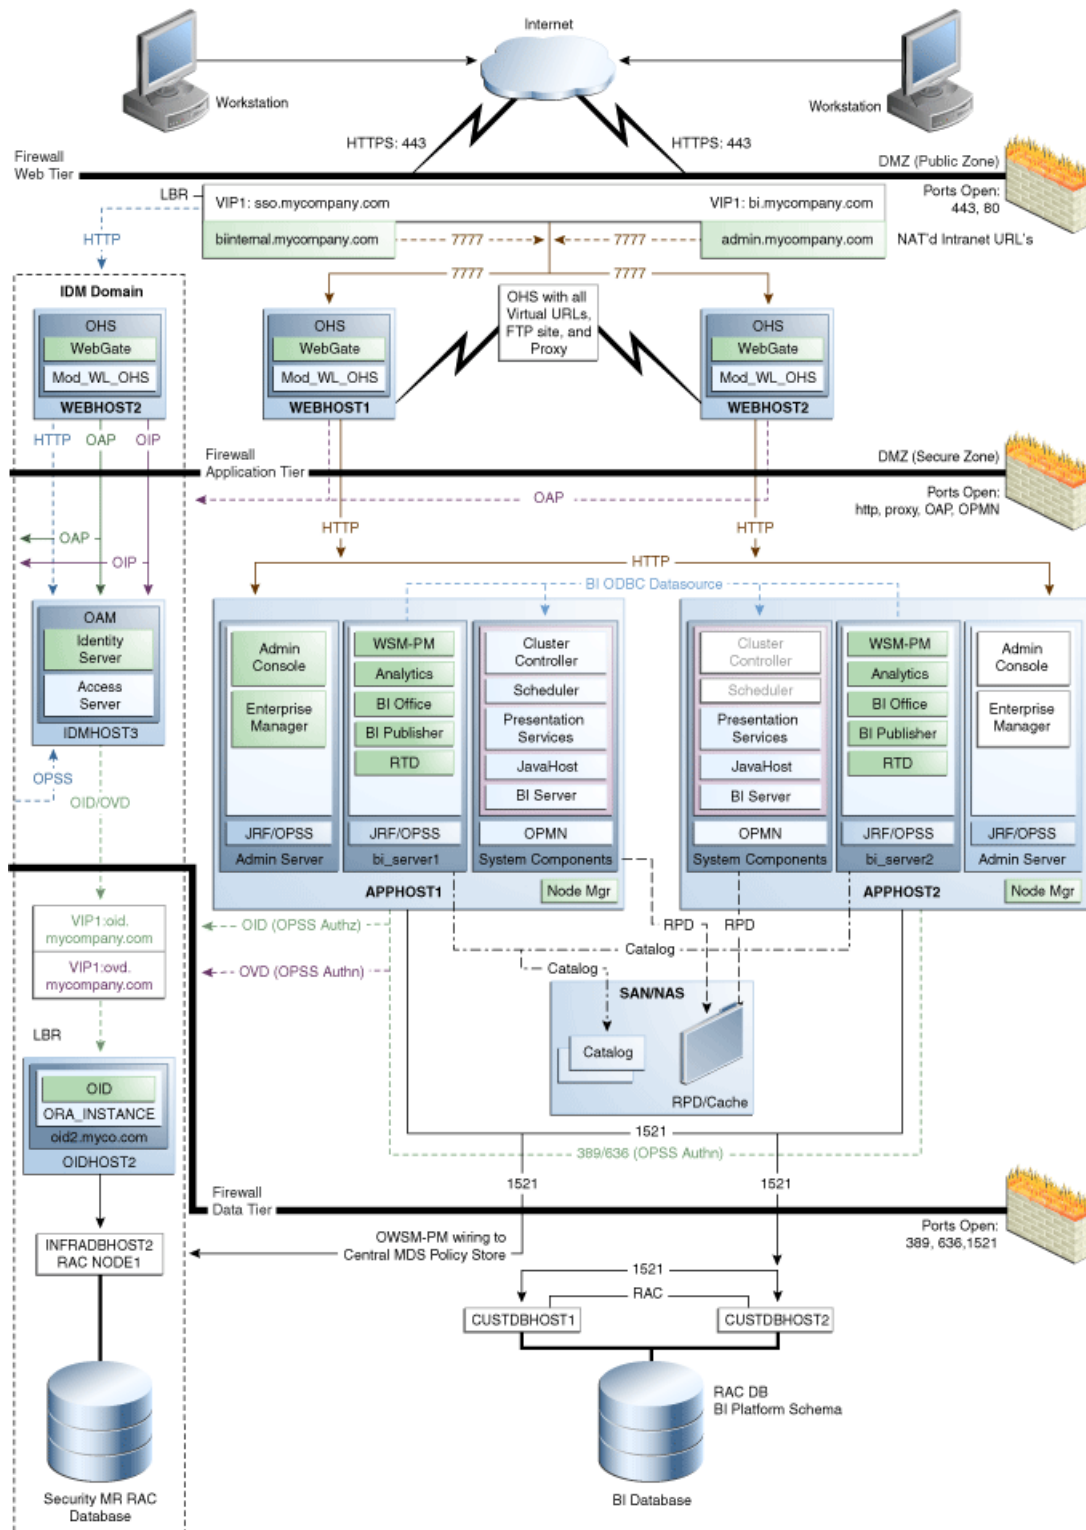
Element	Explanation
	The Administration Server domain directories, the Managed Server domain directories, applications, deployment plans, file adapter control directory, JMS and TX logs, and the entire <code>MW_HOME</code> are on a shared disk.
	Fixed name.
	Installation-dependent name.

Detailed information about setting up this directory structure is included in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence*.

4.1.1.6.1 Volume Design for Oracle Business Intelligence [Figure 4–14](#) shows an Oracle Business Intelligence topology diagram. The volume design described in this section is for this Oracle Business Intelligence topology. Detailed instructions for installing and

configuring this topology are provided in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence*.

Figure 4-14 Reference Topology for Oracle Business Intelligence Enterprise Deployment



For disaster recovery of this Oracle Business Intelligence topology, Oracle recommends the following volume design:

- Provision two volumes for two Middleware Homes that contain redundant product binaries (VOLFMW1 and VOLFMW2 in [Table 4–15](#))
- Provision one volume for the Administration Server domain directory (VOLADMIN in [Table 4–15](#))
- Provision one volume on each node for the Managed Server domain directory (VOLBI1 and VOLBI2 in [Table 4–15](#)). This directory is shared between all the Managed Servers on that node.
- Provision one volume for the JMS file-store and JTA transaction logs (VOLDATA1 in [Table 4–15](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain.
- Provision one volume for the Oracle Business Intelligence Enterprise Edition (EE) shared file and directories (VOLDATA2 in [Table 4–15](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain. This volume will contain the Oracle BI repository (RPD file), Oracle BI Presentation Catalog, Global Cache, and shared Oracle BI Scheduler scripts.
- Provision one volume for the Oracle Business Intelligence Publisher shared configuration file and directories (VOLDATA3 in [Table 4–15](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain. This volume will contain the Oracle Business Intelligence Publisher configuration files.
- Provision one volume for the Oracle Business Intelligence Publisher Catalog (VOLDATA4 in [Table 4–15](#)). There will be one volume for the entire domain that is mounted on all the nodes in the domain. This volume will contain the Oracle Business Intelligence Publisher catalog which stores the objects such as reports, data models, and style templates that you create using Oracle Business Intelligence Publisher.
- Provision one volume on each node for the Oracle HTTP Server Oracle home (VOLWEB1 and VOLWEB2 in [Table 4–15](#)).
- Provision one volume on each node for the Oracle HTTP Server Oracle instance (VOLWEBINST1 and VOLWEBINST2 in [Table 4–15](#)).
- Provision one volume on each node for the Oracle Business Intelligence Oracle instance (VOLBIINST1 and VOLBIINST2 in [Table 4–15](#)).

[Table 4–15](#) provides a summary of Oracle recommendations for volume design for the Oracle Business Intelligence Suite topology shown in [Figure 4–14](#):

Table 4–15 Volume Design Recommendations for Oracle Business Intelligence

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u01/app/oracle/product/fmw/web	Volume for Oracle HTTP Server installation
Web	VOLWEBINST1	WEBHOST1	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLWEBINST2	WEBHOST2	/u01/app/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance

Table 4–15 (Cont.) Volume Design Recommendations for Oracle Business Intelligence

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLSTATIC1 ¹	WEBHOST1	/u01/app/oracle/admin/ohs_ instance/config/static	Volume for static HTML content
Web	VOLSTATIC2 ²	WEBHOST2	/u01/app/oracle/admin/ohs_ instance/config/static	Volume for static HTML content
Application	VOLBIINST1	HOST1	/u01/app/oracle/admin/bi_ instance	Volume for Oracle Business Intelligence instance
Application	VOLBIINST2	HOST2	/u01/app/oracle/admin/bi_ instance	Volume for Oracle Business Intelligence instance
Application	VOLFMW1	HOST1	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle Business Intelligence binaries
Application	VOLFMW2	BIHOST2	/u01/app/oracle/product/fmw	Volume for the WebLogic Server and Oracle Business Intelligence binaries
Application	VOLADMIN	BIHOST1	/u01/app/oracle/admin/biDomain /admin	Volume for Administration Server domain directory
Application	VOLBI1	BIHOST1	/u01/app/oracle/admin/biDomain /mng1	Volume for Managed Server domain directory
Application	VOLBI2	BIHOST2	/u01/app/oracle/admin/biDomain /mng1	Volume for Managed Server domain directory
Application	VOLDATA1	BIHOST1, BIHOST2	/u01/app/oracle/admin/biDomain /biCluster/jms /u01/app/oracle/admin/biDomain /biCluster/tlogs	Volume for transaction logs and JMS data
Application	VOLDATA2	BIHOST1, BIHOST2	/u01/app/oracle/admin/biDomain /biCluster/repository	Volume for the Oracle BI repository (RPD file), Oracle BI Presentation Catalog, Global Cache, and shared Oracle BI Scheduler scripts
Application	VOLDATA3	BIHOST1, BIHOST2	/u01/app/oracle/admin/biDomain /biCluster/bipublisher/config	Volume for the Oracle Business Intelligence Publisher configuration files
Application	VOLDATA4	BIHOST1, BIHOST2	/u01/app/oracle/admin/biDomain /biCluster/bipublisher/reports	Volume for the objects such as reports, data models, and style templates that you create using Oracle Business Intelligence Publisher

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.1.6.2 Consistency Group Recommendations for Oracle Business Intelligence Oracle recommends the following consistency groups for the Oracle Business Intelligence topology:

- Create one consistency group with the volumes containing the domain directories for the Administration Server and Managed Servers as members (DOMAINGROUP in [Table 4-16](#)).
- Create one consistency group with the volume containing the JMS file store and transaction log data as members (DATAGROUP1 in [Table 4-16](#)).
- Create one consistency group with the volume containing the Oracle Business Intelligence shared files, directory, and Oracle Business Intelligence Publisher configuration folder as members (DATAGROUP2 in [Table 4-16](#)).
- Create one consistency group with the volume containing the Middleware Homes as members (FMWHOMEGROUP in [Table 4-16](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle homes as members (WEBHOMEGROUP in [Table 4-16](#)).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle instances as members (WEBINSTANCEGROUP in [Table 4-16](#)).
- Create one consistency group with the volumes containing the Oracle Business Intelligence Oracle instances as members (BIINSTANCEGROUP in [Table 4-16](#)).

[Table 4-16](#) provides a summary of Oracle recommendations for consistency groups for the Oracle Business Intelligence topology shown in [Figure 4-14](#).

Table 4-16 Consistency Groups for Oracle Business Intelligence

Tier	Group Name	Members	Comments
Application	DOMAINGROUP	VOLADMIN VOLBI1 VOLBI2	Consistency group for the Administration Server, Managed Server domain directory
Application	DATAGROUP1	VOLDATA1	Consistency group for the JMS file store and transaction log data
Application	DATAGROUP2	VOLDATA2 VOLDATA3 VOLDATA4	Consistency group for the Oracle Business Intelligence shared files, directory, and Oracle Business Intelligence Publisher configuration
Application	FMWHOMEGROUP	VOLFMW1 VOLFMW2	Consistency group for the Middleware homes
Application	BIINSTANCEGROUP	VOLBIINST1 VOLBIINST2 VOLSTATIC1 ¹ VOLSTATIC2 ²	Consistency group for the Oracle Business Intelligence Oracle instances
Web	WEBHOMEGROUP	VOLWEB1 VOLWEB2	Consistency group for the Oracle HTTP Server Oracle homes

Table 4–16 (Cont.) Consistency Groups for Oracle Business Intelligence

Tier	Group Name	Members	Comments
Web	WEBINSTANCEGROUP	VOLWEBINST1 VOLWEBINST2 VOLSTATIC1 ³ VOLSTATIC2 ⁴	Consistency group for the Oracle HTTP Server Oracle instances

¹ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

² This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

³ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

⁴ This volume for static HTML data is optional. Oracle Fusion Middleware will operate normally without it.

4.1.2 Storage Replication

Follow these steps to set up storage replication for the Oracle Fusion Middleware Disaster Recovery topology:

1. On the standby site, ensure that aliases host names are created that are the same as the physical host names used for the peer hosts at the production site.
2. On the shared storage at the standby site, create the same volumes as were created on the shared storage at the production site.
3. On the standby site, create the same mount points and symbolic links that you created at the production site (note that symbolic links only need to be set up on the standby site if you set up symbolic links at the production site). Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
4. It is not necessary to install the same Oracle Fusion Middleware instances at the standby site as were installed at the production site. When the production site storage is replicated to the standby site storage, the Oracle software installed on the production site volumes will be replicated at the standby site volumes.
5. Perform any other necessary configuration required by the shared storage vendor to enable storage replication between the production site shared storage and the standby site shared storage.
6. Create the baseline snapshot copy of the production site shared storage that sets up the replication between the production site and standby site shared storage. Create the initial baseline copy and subsequent snapshot copies using asynchronous replication mode. After the baseline snapshot copy is performed, validate that all the directories inside the standby site volumes have the same contents as the directories inside the production site volumes.
7. Set up the frequency of subsequent copies of the production site shared storage, which will be replicated at the standby site. When asynchronous replication mode is used, then at the requested frequency the changed data blocks at the production site shared storage (based on comparison to the previous snapshot copy) become the new snapshot copy, and the snapshot copy is transferred to the standby site shared storage.
8. Ensure that disaster protection for any database that is included in the Oracle Fusion Middleware Disaster Recovery production site is provided by Oracle Data

Guard. Do not use storage replication technology to provide disaster protection for Oracle databases.

9. The standby site shared storage receives snapshots transferred on a periodic basis from the production site shared storage. After the snapshots are applied, the standby site shared storage will include all the data up to and including the data contained in the last snapshot transferred from the production site before the failover or switchover.
10. It is strongly recommended to manually force a synchronization operation whenever a change is made to the middle tier at the production site (for example, when a new application is deployed at the production site). Follow the vendor-specific instructions for forcing a synchronization using storage replication technology.

4.1.3 Database

See [Section 3.3, "Database Considerations"](#) for recommendations and considerations for setting up Oracle databases that will be used in the Oracle Fusion Middleware Disaster Recovery topology.

4.1.3.1 Setting Up Oracle Data Guard

Oracle Data Guard should be set up between the Oracle Fusion Middleware Repository databases on the primary site and standby site. The databases on the standby site should be set up as Physical Standby Databases. This section describes the setup and configuration of the data tier on the standby site.

For more information regarding Oracle Data Guard, refer to *Oracle Data Guard Concepts and Administration* in the Oracle Database documentation set.

4.1.3.1.1 Prerequisites and Assumptions The Oracle Data Guard setup and configuration steps below assume that the following conditions are met:

- The RAC cluster and ASM instances on the standby site have been created.
- The RAC databases on the standby site and the production site are using a Flash Recovery Area.
- The database hosts on the standby site already have Oracle software installed.
- The physical path for the DB_HOME on the standby site matches that of the production site.

4.1.3.1.2 Oracle Data Guard Environment Description The Oracle Data Guard steps use the environment variables shown in [Table 4-17](#) for the SOA database at the production site.

Table 4-17 Environment Variables Used for SOA Databases at the Production Site

Variable	Value
SOA Database Host Names	soadbhost1.mycompany.com soadbhost2.mycompany.com
ORACLE_HOME	/u01/app/oracle/product/db_1
SOA_DBNAME	PSOA
SOA_DB_UNIQUE_NAME	PSOA
SOA_DB_INSTANCE_NAMES	SOA1, SOA2

Table 4–17 (Cont.) Environment Variables Used for SOA Databases at the Production

Variable	Value
SOA_SERVICE_NAMES	PSOA, SSOA
ORACLE_SID	SOA1, SOA2

The Oracle Data Guard steps use the environment variables shown in [Table 4–18](#) for the SOA database at the standby site.

Table 4–18 Environment Variables Used for SOA Databases at the Standby Site

Variable	Value
SOA Database Host Names	soadbhost1.mycompany.com soadbhost2.mycompany.com
ORACLE_HOME	/u01/app/oracle/product/db_1
SOA_DBNAME	PSOA
SOA_DB_UNIQUE_NAME	SSOA
SOA_DB_INSTANCE_NAMES	SOA1, SOA2
SOA_SERVICE_NAMES	PSOA, SSOA
ORACLE_SID	SOA1, SOA2

These high level steps for setting up Oracle Data Guard are described in detail in the following sections:

- [Gather Files and Perform Backup](#)
- [Configure Oracle Net Services on the Standby Site](#)
- [Create Instances and Database on the Standby Site](#)
- [Test Database Switchover and Switchback](#)

4.1.3.1.3 Gather Files and Perform Backup Follow these steps to gather files and perform the database backup:

1. On the SOADBHOST1 of the primary site, create a directory for staging purposes. For example:

```
$ mkdir -p /u01/app/stage/psoa
```

2. Create the exact path on SOADBHOST1 of the standby site. Follow the example shown in step 1.

3. On the SOADBHOST1 of the primary site, connect to the database instance psoa1 and create a pfile from the spfile. For example:

```
SQL > create pfile='/u01/app/stage/psoa/initpsoa.ora' from spfile;
```

4. On the SOADBHOST1 of the primary site, connect to RMAN, perform a backup of the database, and place the backup files in the stage directory. For example:

```
$ $ORACLE_HOME/bin/rman target /
```

```
RMAN> backup device type disk format '/u01/app/stage/psoa/%U' database plus  
archivelog;
```

```
RMAN> backup device type disk format '/u01/app/stage/psoa/%U' current
```

```
controlfile for standby;
```

5. Follow the steps below to validate that the backups created by RMAN are valid.
6. Connect to RMAN on SOADBHOST1 of the primary site and then list the backup summary.
7. Validate the backup sets created by RMAN in step 4:

```
RMAN> list backup summary;
using target database control file instead of recovery catalog
List of Backups
=====
Key          TY LV S Device Type Completion Time #Pieces #Copies Compressed Tag
-----
93          B A A DISK          14-MAY-07          1          1          NO
TAG20070514T122312
94          B F A DISK          14-MAY-07          1          1          NO
TAG20070514T122315
95          B F A DISK          14-MAY-07          1          1          NO
TAG20070514T122315
96          B A A DISK          14-MAY-07          1          1          NO
TAG20070514T122629
97          B F A DISK          14-MAY-07          1          1          NO
TAG20070514T123220

RMAN> validate backupset 93;
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=451 instance=psoa1 devtype=DISK
channel ORA_DISK_1: starting validation of archive log backupset
channel ORA_DISK_1: reading from backup piece /u01/app/stage/psoa/34ihmtdg_1_1
channel ORA_DISK_1: restored backup piece 1
piece handle=/u01/app/stage/psoa/34ihmtdg_1_1 tag=TAG20070514T122312
channel ORA_DISK_1: validation complete, elapsed time: 00:00:02
```

8. On SOADBHOST1 of the primary site, copy the `listener.ora`, `sqlnet.ora`, and `tnsnames.ora` files from the `$ORACLE_HOME/network/admin` directory to the staging directory.
9. Using operating system utilities, copy the contents of staging directory on SOADBHOST1 of the primary site to the staging directory on SOADBHOST1 of the standby site.

4.1.3.1.4 Configure Oracle Net Services on the Standby Site Follow these steps to configure Oracle Net Services on the standby site:

1. Copy the `listener.ora`, `sqlnet.ora`, and `tnsnames.ora` files from the staging directory on SOADBHOST1 on the primary site to the `$ORACLE_HOME/network/admin` directory on all the nodes of the standby site.
2. Modify the `listener.ora` file on each of the standby host to contain the virtual IP of that host.
3. Modify the `tnsnames.ora` file on each node, including the primary RAC nodes and standby RAC nodes, to contain all primary and standby net service names.
4. Modify the Oracle Net aliases that are used for the `local_listener` and `remote_listener` parameters to point to the listener on each standby host. The example below shows excerpts from the `tnsnames.ora` file:

```
#local_listener
PSOA =
```

```

(DESCRIPTION =
  (ADDRESS =
    (PROTOCOL = TCP)
    (HOST = soadbhost1-vip)
    (HOST = soadbhost2-vip)
    (PORT = 1521))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = psoa)
  )
)
#remote_listener
SSOA =
(DESCRIPTION =
  (ADDRESS =
    (PROTOCOL = TCP)
    (HOST = soadbhost1-vip)
    (HOST = soadbhost2-vip)
    (PORT = 1521))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = ssoa)
  )
)

```

5. Start the listeners on the standby database hosts.

4.1.3.1.5 Create Instances and Database on the Standby Site Follow these steps to create instances and the database on the standby site:

1. To enable secure transmission of redo data, make sure the databases on the primary and standby sites use a password file, and make sure the password for the SYS user is identical on every system. Create a password file on both the nodes of the standby databases. For example:

On SOADBHOST1 of the standby site

```

$ cd $ORACLE_HOME/dbs
$ orapwd file=orapwpsoa1 password=welcome1

```

On SOADBHOST2 of the standby site

```

$ cd $ORACLE_HOME/dbs
$ orapwd file=orapwpsoa2 password=welcome1

```

2. Copy and rename the pfile from the staging area to the `$ORACLE_HOME/dbs` directory on SOADBHOST1 of the standby site. For example:

```

$ cp /u01/app/stage/psoa/initpsoa.ora $ORACLE_HOME/dbs/initsoa1.ora

```

3. Modify the standby initialization parameter file copied from the primary node to include the parameters shown [Table 4-19](#):

Table 4–19 Parameters to Specify in the Standby Initialization Parameter File

Parameter	Value
RAC Parameters	<pre>*.cluster_database=true PSOA1.instance_name=SOA1 PSOA2.instance_name=SOA2 PSOA1.instance_number=1 PSOA2.instance_number=2 PSOA1.thread=1 PSOA2.thread=2 PSOA1.undo_tablespace=UNDOTBS1 PSOA2.undo_tablespace=UNDOTBS2 *.remote_listener=LISTENERS_PSOA</pre>
Data Guard Parameters	<pre>*.db_unique_name=SSOA *.log_archive_config='dg_config=(SSOA,PSOA)' *.log_archive_dest_1='LOCATION=USE_DB_RECOVERY_FILE_DEST' *.log_archive_dest_2='service=PSOA valid_for=(online_logfiles,primary_role) db_unique_name=PSOA' *.db_file_name_convert='+DATA/PSOA/', '+DATA/SSOA/', '+RECO/PSOA', '+RECO/SSOA' *.log_file_name_convert='+DATA/PSOA/', '+DATA/SSOA/', '+RECO/PSOA', '+RECO/SSOA' *.standby_file_management=auto *.fal_server='PSOA' *.fal_client='SSOA'</pre>
Miscellaneous Parameters	<pre>*.background_dump_dest=/u01/app/admin/PSOA/bdump *.core_dump_dest=/u01/app/admin/PSOA/cdump *.user_dump_dest=/u01/app/admin/PSOA/udump *.audit_file_dest=/u01/app/admin/PSOA/adump *.db_recovery_file_dest='+RECO' *.dispatchers=PSOAXDB</pre>

4. Connect to the ASM instance on SOADBHOST1 of the standby site, and create a directory within the DATA disk group that has the same name as the DB_UNIQUE_NAME of the standby database. For example:

```
SQL> alter diskgroup data add directory '+DATA/SSOA';
```

5. Connect to the standby database on SOADBHOST1 of the standby site, with the standby database in the IDLE state, and create an SPFILE in the standby DATA disk group. For example:

```
SQL> CREATE SPFILE='+DATA/SSOA/spfilepsoa.ora' FROM
PFILE='?/dbs/initsoa1.ora';
```

6. In the `$ORACLE_HOME/dbs` directory on `SOADBHOST1` and `SOADBHOST2` of the standby site, create a PFILE that contains a pointer to the SPFILE. The PFILE should follow the naming convention `init<OracleSID>.ora`. For example:

On `SOADBHOST1`:

```
$ cd $ORACLE_HOME/dbs
$ echo "SPFILE='+DATA/SSOA/spfilepsoa.ora'" > initsoa1.ora
```

On `SOADBHOST2`:

```
$ cd $ORACLE_HOME/dbs
$ echo "SPFILE='+DATA/SSOA/spfilepsoa.ora'" > initsoa2.ora
```

7. Create the dump directories on all standby hosts as referenced in the standby initialization parameter file. For example:

```
$ mkdir -p $ORACLE_BASE/admin/psoa/bdump
$ mkdir -p $ORACLE_BASE/admin/psoa/cdump
$ mkdir -p $ORACLE_BASE/admin/psoa/udump
$ mkdir -p $ORACLE_BASE/admin/psoa/adump
```

8. On `SOADBHOST1` of the standby site, set the `ORACLE_HOME`, `PATH`, `ORACLE_SID` and startup the standby database without mounting the control file. This host should have the staging directory. For example:

```
SQL > startup nomount
```

9. From `SOADBHOST1` of the primary site, duplicate the primary database as a standby into the ASM disk group by using `RMAN`. For example:

```
$ rman target / auxiliary sys/oracle@ssoa
RMAN> duplicate target database for standby;
```

10. Use `SQL*Plus` to log in to the newly created database to validate that it was created correctly. For example:

```
$ sqlplus '/as sysdba'
```

11. Connect to the standby database on `SOADBHOST1` of the standby site, and create the standby redo logs to support the standby role. For example:

```
SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 1
GROUP 5 SIZE 300M,
GROUP 6 SIZE 300M,
GROUP 7 SIZE 300M;
```

```
SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 2
GROUP 8 SIZE 300M,
GROUP 9 SIZE 300M,
GROUP 10 SIZE 300M;
```

12. On `SOADBHOST1` of the standby site, start managed recovery and real-time apply on the standby database. For example:

```
SQL> ALTER DATABASE recover managed standby database using current logfile
disconnect;
```

13. On `SOADBHOST1` and `SOADBHOST2` of the standby site, register the standby database and the database instances with the Oracle Cluster Registry (OCR) using the Server Control (SRVCTL) utility. For example:

```
$ srvctl add database -d psoa -o /u01/app/oracle/product/10.2.0/db_1
```

```
$ srvctl add instance -d psoa -i soa1 -n soadbhost1
$ srvctl add instance -d psoa -i soa2 -n soadbhost2
```

- 14.** Establish a dependency between the database and the ASM instance. For example:

```
$ srvctl modify instance -d psoa -i soa1 -s +ASM1
$ srvctl modify instance -d psoa -i soa2 -s +ASM2
$ srvctl enable asm -n stbdd03 -i +ASM1
$ srvctl enable asm -n stbdd04 -i +ASM2
```

- 15.** Configure the primary database for Oracle Data Guard by modifying/adding the Data Guard parameters in the primary initialization file from the staging directory (/u01/app/stage/psoa/initpsoa.ora) to the values shown below:

```
*.log_archive_config='dg_config=(SSOA,PSOA) '
*.log_archive_dest_1='LOCATION=USE_DB_RECOVERY_FILE_DEST'

*.log_archive_dest_2='service=SSOA valid_for=(online_logfiles,primary_role) db_
unique_name=SSOA'

*.db_file_name_convert='+DATA/SSOA/', '+DATA/PSOA/', '+RECO/SSOA', '+RECO/PSOA'

*.log_file_name_convert='+DATA/SSOA/', '+DATA/PSOA/', '+RECO/SSOA', '+RECO/PSOA'

*.standby_file_management=auto

*.fal_server='SSOA'

*.fal_client='PSOA'
```

- 16.** Connect to the primary database on SOADBHOST1 of the primary site, with the primary database in the IDLE state, and create an SPFILE in the primary DATA disk group. For example:

```
SQL> CREATE SPFILE='+DATA/PSOA/spfilepsoa.ora' FROM
PFIL=' /u01/app/stage/psoa/initpsoa.ora';
```

- 17.** In the ORACLE_HOME/dbs directory on SOADBHOST1 and SOADBHOST2 of the primary site, create a PFILE that contains a pointer to the SPFILE. The PFILE must follow the naming convention init<OracleSID>.ora. For example:

On SOADBHOST1:

```
$ cd $ORACLE_HOME/dbs
$ echo "SPFILE='+DATA/PSOA/spfilepsoa.ora'" > initsoa1.ora
```

On SOADBHOST2:

```
$ cd $ORACLE_HOME/dbs
$ echo "SPFILE='+DATA/PSOA/spfilepsoa.ora'" > initsoa2.ora
```

- 18.** Restart the primary database after modifying the parameters.

- 19.** Create the standby redo logs on the primary database to support the standby role. For example:

```
SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 1
GROUP 5 SIZE 300M,
GROUP 6 SIZE 300M,
GROUP 7 SIZE 300M;
```

```
SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 2
GROUP 8 SIZE 300M,
```



```
GROUP 9 SIZE 300M,
GROUP 10 SIZE 300M;
```

20. Verify the Oracle Data Guard configuration by querying the V\$ARCHIVED_LOG view to identify existing files in the archived redo log. For example:

```
SQL> select sequence#, first_time, next_time from v$archived_log order by
sequence#;
```

21. On the primary database, issue the following SQL statement to force a log switch and archive the current online redo log file group:

```
SQL> alter system archive log current;
```

22. On the standby database, query the V\$ARCHIVED_LOG view to verify that the redo data was received and archived on the standby database:

```
SQL> select sequence#, first_time, next_time from v$archived_log order by
sequence#;
```

4.1.3.1.6 Test Database Switchover and Switchback Follow these steps to test that the database switchover and switchback operation works correctly between the newly-created physical standby database and the primary RAC databases:

1. Shutdown all but one instance of the RAC databases (PSOA) on the primary site. For example, run the command below on SOADBHOST1 of the production site:

```
$ srvctl stop instance -d psoa -i soa2
```

2. Initiate the role transition to the physical standby on the current primary database. For example, run the command below on SOADBHOST1 of the production site:

```
SQL > ALTER DATABASE COMMIT TO SWITCHOVER TO PHYSICAL STANDBY WITH SESSION
SHUTDOWN;
```

3. Shut down the primary instance and mount the primary instance. For example, run the command below on SOADBHOST1 of the production site:

```
SQL > shutdown immediate
SQL > startup mount
```

4. At this point, both the databases are in Physical Standby mode. To verify that both the databases are in Physical Standby mode, run this SQL query on both the databases:

```
SQL> select database_role from v$database;
DATABASE_ROLE
-----
PHYSICAL_STANDBY
```

5. Switch the physical standby database role to the primary role. For example, run the command below on SOADBHOST1 of the standby site:

```
SQL> ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY WITH SESSION SHUTDOWN;
```

6. Now the physical standby database is the new primary.
7. Shut down the new primary database and start up both the RAC nodes using srvctl. For example, run the following command on the SOADBHOST1 of the standby site:

```
srvctl start database -d psoa
```

8. On the new physical standby database (the old primary) start the managed recovery of the database. For example, run the command below on SOADBHOST1 of the primary site:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE DISCONNECT FROM SESSION;
```

9. Start sending the redo data to the new physical standby database. For example, run the command below on SOADBHOST1 of the standby site:

```
SQL> ALTER SYSTEM SWITCH LOGFILE;
```

10. Check the new physical standby database to see if it is receiving the archive log files by querying the V\$ARCHIVED_LOG view.

4.1.4 Node Manager

The Node Manager communicates with the Administration Server over SSL. For this communication to work correctly on the standby site, you must create SSL certificates using the physical host names. This section includes these topics:

- [Generate Self Signed Certificates](#)
- [Create an Identity KeyStore](#)
- [Create Trust KeyStore](#)
- [Configure Node Manager for Custom KeyStores](#)

The examples in these sections show how to perform these tasks for the Oracle SOA Suite enterprise topology shown in [Figure 4-2](#).

Note: Remember that when you are setting up the Oracle SOA Suite enterprise topology shown in [Figure 4-2](#) as the production site for a Disaster Recovery topology, you must use the physical host names shown in [Table 3-1](#) for the production site hosts instead of the host names shown in [Figure 4-2](#).

The steps in this section must be performed on the application tier hosts on which WebLogic Server is installed.

4.1.4.1 Generate Self Signed Certificates

Follow these steps to generate self signed certificates:

1. Set your environment using the `setWLSenv` script located under the `$WL_HOME/server/bin` directory.
2. Create a user-defined directory for the certificates. For example, create the `certs` directory under the `$MW_HOME/user_projects/domains/SOADomain` directory.
3. Run the `utils.CertGen` tool from the user-defined directory to create the certificates for the application tier hosts on which WebLogic Server is installed. The syntax is:

```
Syntax: java utils.CertGen <key_passphrase> <cert_file_name> <key_file_name>
[export|domestic] [hostname]
```

For example, enter these commands:

```
java utils.CertGen welcome1 soahost1_cert soahost1_key domestic soahost1
```

```
java utils.CertGen welcome1 soahost2_cert soahost2_key domestic soahost2
```

4.1.4.2 Create an Identity KeyStore

Follow these steps to create an identity keystore using the `utils.ImportPrivateKey` utility:

1. Create a new identity keystore called `appIdentityKeyStore` using the `utils.ImportPrivateKey` utility.
2. Create this keystore under the same directory as the certificates, for example:
`$MW_HOME/user_projects/domains/j2eeDomain/certs`
3. The Identity Store is created (if none exists) when you import a certificate and the corresponding key into the Identity Store using the `utils.ImportPrivateKey` utility.
4. Import the certificate and private key for the application tier hosts on which WebLogic Server is installed into the Identity Store; also make sure to use a different alias for each of the certificate/key pair imported. The syntax is:

```
Syntax: java utils.ImportPrivateKey <keystore_file> <keystore_password>
<certificate_alias_to_use> <private_key_passphrase> <certificate_file>
<private_key_file> [<keystore_type>]
```

For example, enter these commands:

```
java utils.ImportPrivateKey appIdentityKeyStore.jks welcome1 appIdentity1
welcome1 $MW_HOME/user_projects/domains/SOADomain/certs/soahost1_cert.pem
$MW_HOME/user_projects/domains/SOADomain/certs/soahost1_key.pem

java utils.ImportPrivateKey appIdentityKeyStore.jks welcome1 appIdentity2
welcome1 $MW_HOME/user_projects/domains/SOADomain/certs/soahost2_cert.pem
$MW_HOME/user_projects/domains/SOADomain/certs/soahost2_key.pem
```

4.1.4.3 Create Trust KeyStore

Follow these steps to create a trust keystore:

1. Create a new trust keystore called `appTrustKeyStore` using the `keytool` utility.
2. Use the standard java keystore to create the new trust keystore since it already contains most of the root CA certificates needed. It is recommended not to modify the standard Java trust key store directly.
3. Copy the standard Java keystore `cacerts` located under the `$WL_HOME/server/lib` directory to the same directory as the certificates. For example:

```
cp $WL_HOME/server/lib/cacerts
$MW_HOME/user_projects/domains/SOADomain/certs/appTrustKeyStore.jks
```

4. The default password for the standard Java keystore is `changeit` and it is always recommended to change the default password. Use the `keytool` utility to do this. The syntax is:

```
keytool -storepasswd -new <NewPassword> -keystore <TrustKeyStore> -storepass
<Original Password>
```

For example, enter this command:

```
keytool -storepasswd -new welcome1 -keystore appTrustKeyStore.jks -storepass
changeit
```

5. The CA certificate `CertGenCA.der` is used to sign all certificates generated by `utils.CertGen` tool and is located at `$WL_HOME/server/lib` directory. This CA certificate must be imported into the `appTrustKeyStore` using the `keytool` utility. The syntax is:

```
keytool -import -v -noprompt -trustcacerts -alias <AliasName> -file
<CAFileLocation> -keystore <KeyStoreLocation> -storepass <KeyStore Password>
```

For example, enter this command:

```
keytool -import -v -noprompt -trustcacerts -alias clientCACert -file
$WL_HOME/server/lib/CertGenCA.der -keystore appTrust.jks -storepass welcome1
```

4.1.4.4 Configure Node Manager for Custom KeyStores

Configure Node Manager on each of the nodes to use the newly-created custom keystores by editing the following lines at the end of the `nodemanager.properties` file located under the `$WL_HOME/common/nodemanager` directory. These lines and their meanings are shown below:

```
KeyStores=CustomIdentityAndCustomTrust
CustomIdentityKeyStoreFileName=<Identity KeyStore>
CustomIdentityKeyStorePassPhrase=<Identity KeyStore Password>
CustomIdentityAlias=<Identity Key Store Alias>
CustomIdentityPrivateKeyPassPhrase=<Private Key used when creating Certificate>
CustomTrustKeyStoreFileName=<Trust KeyStore>
CustomTrustKeyStorePassPhrase=<Trust KeyStore Password>
```

For example, make these edits in the `nodemanager.properties` file on `SOAHOST1`:

```
KeyStores=CustomIdentityAndCustomTrust
CustomIdentityKeyStoreFileName=$MW_HOME/user_projects/domains/SOADomain/certs
/appIdentityKeyStore.jks
CustomIdentityKeyStorePassPhrase=welcome1
CustomIdentityAlias=appIdentity1
CustomIdentityPrivateKeyPassPhrase=welcome1
CustomTrustKeyStoreFileName=$MW_HOME/user_projects/domains/SOADomain/certs
/appTrust.jks
CustomTrustKeyStorePassPhrase=welcome1
```

For example, make these edits in the `nodemanager.properties` file on `SOAHOST2`:

```
KeyStores=CustomIdentityAndCustomTrust
CustomIdentityKeyStoreFileName=$MW_HOME/user_projects/domains/SOADomain/certs
/appIdentityKeyStore.jks
CustomIdentityKeyStorePassPhrase=welcome1
CustomIdentityAlias=appIdentity2
CustomIdentityPrivateKeyPassPhrase=welcome1
CustomTrustKeyStoreFileName=$MW_HOME/user_projects/domains/SOADomain/certs
/appTrust.jks
CustomTrustKeyStorePassPhrase=welcome1
```

4.2 Creating a Production Site

This section provides the steps to create the production site. The Oracle SOA enterprise deployment topology and the Oracle Identity Management Enterprise deployment topology are used as examples.

Ensure that you have performed the following prerequisites before you start creating the production site:

- Set up the host name aliases for the middle tier hosts, which was described in [Section 3.1.1, "Planning Host Names."](#)
- Create the required volumes on the shared storage on the production site, which was described in [Section 4.1.1, "Directory Structure and Volume Design."](#)
- Create the mount points and the symbolic links (if required). Refer to [Section 3.2.3, "Storage Replication"](#) to determine whether you must create symbolic links for the production site.

For more information, see the following:

- [Creating the Production Site for the Oracle SOA Suite Topology](#)
- [Creating the Production Site for the Oracle Identity Management Topology](#)
- [Creating the Production Site for the Oracle WebCenter Portal Topology](#)
- [Creating the Production Site for the Oracle WebCenter Content Topology](#)
- [Creating the Production Site for the Oracle Business Intelligence Topology](#)
- [Creating the Production Site for the Oracle Portal, Forms, Reports, and Discoverer Topology](#)
- [Validating the Production Site Setup](#)

4.2.1 Creating the Production Site for the Oracle SOA Suite Topology

The production site should be installed and configured as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* with the following variations. The steps to install and configure the production site are listed below and should be followed in the sequence listed.

1. Create volumes and consistency groups on the shared storage device, as described in [Section 4.1.1.1.1, "Volume Design for Oracle SOA Suite."](#)
2. Set up physical host names on the production site and physical host names and alias host names for the standby site. See [Section 3.1.1, "Planning Host Names"](#) for information on planning host names for the production and standby sites.
3. Install and configure Oracle SOA Suite as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* with the following modifications:
 - a. Install the Oracle SOA Suite components into the volumes created on the shared storage device.
 - b. Use the physical host names when installing and configuring WebLogic domain.
 - c. Create a separate volume on each site for the JMS stores and transaction logs.
 - d. After the installation and configuration of the production site, turn off host name verification. See the "Disabling Host Name Verification for the Oracle WebLogic Administration Server and the WLS_WSM1 Managed Server"

section in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for detailed instructions about turning off host name verification for an Administration Server and Managed Server.

- e. If you do not plan on turning host name verification off, follow the steps in [Section 4.1.4, "Node Manager"](#) to configure Node Manager communication.
- f. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.2 Creating the Production Site for the Oracle Identity Management Topology

The production site should be installed and configured as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management* with the following variations. The steps to install and configure the production site are listed below and should be followed in the sequence listed.

1. Create volumes and consistency groups on the shared storage device, as described in [Section 4.1.1.3.1, "Volume Design for Oracle Identity Management."](#)
2. Set up physical host names on the production site and physical host names and alias host names for the standby site. See [Section 3.1.1, "Planning Host Names"](#) for information on planning host names for the production and standby sites.
3. Install and configure Oracle Identity Management as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management* with the following modifications:
 - a. Install the Oracle Identity Management components into the volumes created on the shared storage device.
 - b. Use the physical host names when installing and configuring the WebLogic domain.
 - c. Create a separate volume on each site for the JMS stores and transaction logs.
 - d. After the installation and configuration of the production site, turn off host name verification. See the "Disabling Host Name Verification for the Oracle WebLogic Administration Server and the WLS_WSM1 Managed Server" section in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for detailed instructions about turning off host name verification for an Administration Server and Managed Server.
 - e. If you do not plan on turning host name verification off, follow the steps in [Section 4.1.4, "Node Manager"](#) to configure Node Manager communication.
 - f. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.3 Creating the Production Site for the Oracle WebCenter Portal Topology

The production site should be installed and configured as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal* with the following variations. The steps to install and configure the production site are listed below and should be followed in the sequence listed.

1. Create volumes and consistency groups on the shared storage device, as described in [Section 4.1.1.2.1, "Volume Design for Oracle WebCenter Portal."](#)
2. Set up physical host names on the production site and physical host names and alias host names for the standby site. See [Section 3.1.1, "Planning Host Names"](#) for information on planning host names for the production and standby sites.

3. Install and configure Oracle WebCenter Portal as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal* with the following modifications:
 - a. Install the Oracle WebCenter Portal components into the volumes created on the shared storage device.
 - b. Use the physical host names when installing and configuring WebLogic domain.
 - c. After the installation and configuration of the production site, turn off host name verification. See the "Disabling Host Name Verification for the Oracle WebLogic Administration Server and the WLS_WSM1 Managed Server" section in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal* for detailed instructions about turning off host name verification for an Administration Server and Managed Server.
 - d. If you do not plan on turning host name verification off, follow the steps in [Section 4.1.4, "Node Manager"](#) to configure Node Manager communication.
 - e. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.4 Creating the Production Site for the Oracle WebCenter Content Topology

The production site should be installed and configured as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content* with the following variations. The steps to install and configure the production site are listed below and should be followed in the sequence listed.

1. Create volumes and consistency groups on the shared storage device, as described in [Section 4.1.1.5.1, "Volume Design for Oracle WebCenter Content."](#)
2. Set up physical host names on the production site and physical host names and alias host names for the standby site. See [Section 3.1.1, "Planning Host Names"](#) for information on planning host names for the production and standby sites.
3. Install and configure Oracle Enterprise Content Management as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content* with the following modifications:
 - a. Install the Oracle Enterprise Content Management components into the volumes created on the shared storage device.
 - b. Use the physical host names when installing and configuring WebLogic domain.
 - c. After the installation and configuration of the production site, turn off host name verification. See the "Disabling Host Name Verification for the Administration Server" section in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content* for detailed instructions about turning off host name verification for an Administration Server.
 - d. If you do not plan on turning host name verification off, follow the steps in [Section 4.1.4, "Node Manager"](#) to configure Node Manager communication.
 - e. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.5 Creating the Production Site for the Oracle Business Intelligence Topology

The production site should be installed and configured as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence* with the following variations. The steps to install and configure the production site are listed below and should be followed in the sequence listed.

1. Create volumes and consistency groups on the shared storage device, as described in [Section 4.1.1.6.1, "Volume Design for Oracle Business Intelligence."](#)
2. Set up physical host names on the production site and physical host names and alias host names for the standby site. See [Section 3.1.1, "Planning Host Names"](#) for information on planning host names for the production and standby sites.
3. Install and configure Oracle Business Intelligence as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence* with the following modifications:
 - a. Install the Oracle Business Intelligence components into the volumes created on the shared storage device.
 - b. Use the physical host names when installing and configuring WebLogic domain.
 - c. After the installation and configuration of the production site, turn off host name verification. See the "Disabling Host Name Verification for the bi_server1 Managed Server" and "Disabling Host Name Verification for the bi_server2 Managed Server" in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence* for detailed instructions about turning off host name verification for an Administration Server.
 - d. If you do not plan on turning host name verification off, follow the steps in [Section 4.1.4, "Node Manager"](#) to configure Node Manager communication.
 - e. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.6 Creating the Production Site for the Oracle Portal, Forms, Reports, and Discoverer Topology

The production site should be installed and configured as described in the enterprise deployment manuals for the following products:

- Oracle Portal:

Detailed instructions for setting up and configuring a production site that uses the Oracle Portal enterprise topology shown in [Figure 4–9](#) are provided in the 11.1.1.2 *Oracle Portal Enterprise Deployment Guide*. See Article ID 952068.1 "Oracle Fusion Middleware 11g (11.1.1.2) Enterprise Deployment Guides for Portal, Forms, Reports, and Discover" at My Oracle Support (formerly Oracle *MetaLink*) for information on obtaining the manual. The URL for My Oracle Support is:

<http://support.oracle.com>
- Oracle Forms, Reports, and Discoverer

Detailed instructions for setting up and configuring a production site that uses the Oracle Forms, Reports, and Discoverer enterprise topology shown in [Figure 4–10](#) are provided in the 11.1.1.2 *Oracle Forms, Reports, and Discoverer Enterprise Deployment Guide*. See Article ID 952068.1 "Oracle Fusion Middleware 11g (11.1.1.2) Enterprise Deployment Guides for Portal, Forms, Reports, and Discover"

at My Oracle Support (formerly Oracle *MetaLink*) for information on obtaining the manual. The URL for My Oracle Support is:

<http://support.oracle.com>

Follow the installation and configuration instructions in the manuals above, except for the following variations. The following steps should be performed in the sequence listed:

1. Create volumes and consistency groups on the shared storage device, as described in [Section 4.1.1.4.1, "Volume Design for Oracle Portal, Forms, Reports, and Discover."](#)
2. Set up physical host names on the production site and physical host names and alias host names for the standby site. See [Section 3.1.1, "Planning Host Names"](#) for information on planning host names for the production and standby sites.
3. Install and configure Oracle Portal, Forms, Reports, and Discoverer as described in the white papers linked to above, with the following modifications:
 - a. Install the Oracle Portal, Forms, Reports, and Discoverer components into the volumes created on the shared storage device.
 - b. Use the physical host names when installing and configuring WebLogic domain.
 - c. After the installation and configuration of the production site, turn off host name verification. See the "Disabling Host Name Verification for the Oracle WebLogic Administration Server and the WLS_WSM1 Managed Server" section in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal* for detailed instructions about turning off host name verification for an Administration Server and Managed Server.
 - d. If you do not plan on turning host name verification off, follow the steps in [Section 4.1.4, "Node Manager"](#) to configure Node Manager communication.
 - e. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.7 Validating the Production Site Setup

To validate the production site setup for the Oracle SOA Suite enterprise topology, follow the validation steps in these sections of the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*:

- In the "Installing Oracle HTTP Server" chapter, follow the validation steps in this section:
 - "Validating Oracle HTTP Server Through the Load Balancer"
- In the "Creating a Domain" chapter, follow the validation steps in these sections:
 - "Validating the Administration Server"
 - "Starting and Validating the WLS_WSM1 Managed Server"
 - "Starting and Validating the WLS_WSM2 Managed Server"
 - "Validating Access Through Oracle HTTP Server"
 - "Validating Access to SOAHOST2 Through Oracle HTTP Server"
- In the "Extending the Domain for SOA Components" chapter, follow the validation steps in these sections:

- "Validating the WLS_SOA1 and WLS_WSM1 Managed Servers"
- "Starting and Validating the WLS_SOA2 Managed Server"
- "Validating Access Through Oracle HTTP Server"
- In the "Extending the Domain to Include BAM" chapter, follow the validation steps in this section:
 - "Validating Access Through Oracle HTTP Server"

4.3 Creating a Standby Site

This section provides the steps to create the standby site. The Oracle SOA enterprise deployment topology and the Oracle Identity Management Enterprise deployment topology are used as examples.

It includes the following topics:

- [Creating the Standby Site](#)
- [Validating the Standby Site Setup](#)

4.3.1 Creating the Standby Site

Ensure that you have performed the following prerequisites before you start creating the standby site:

- On the standby site, ensure that you set up the correct alias host names and physical host names by following the instructions in [Section 3.1.1, "Planning Host Names."](#)

Ensure that each standby site host has an alias host name that is the same as the physical host name of its peer host at the production site.
- On the shared storage on the standby site, create the same volumes that were created on the shared storage at the production site.
- On the standby site, create the same mount points and symbolic links (if required) that you created at the production site. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.

4.3.1.1 Database Setup

Oracle Data Guard should be set up between the Oracle Fusion Middleware Repository databases on the primary site and standby site. The databases on the standby site should be set up as physical standby databases. Refer to [Section 4.1.3.1, "Setting Up Oracle Data Guard"](#) for instructions on setting up Oracle Data Guard between databases running the metadata repositories on the primary and standby sites.

Also, ensure that the databases running the metadata repositories on the standby site are in the Managed Recovery mode. To enable the standby database to be in the managed recovery mode run the following SQL command (the disconnect option ends the SQL session after the command is completed successfully):

```
SQL> ALTER DATABASE RECOVERY MANAGED STANDBY DATABASE DISCONNECT FROM SESSION;
```

4.3.1.2 Middle Tier Setup

The middle tier hosts on the standby site do not require the installation or configuration of the of any Oracle Fusion Middleware or WebLogic Server software. When the production site storage is replicated to the standby site storage, the software installed on the production site volumes will be replicated at the standby site volumes.

Follow the steps below to set up the middle tier hosts on the standby site:

1. Create a baseline snapshot copy of shared storage on the production site, which sets up the replication between the storage devices. Create the initial baseline copy and subsequent snapshot copies using asynchronous replication mode.
2. Synchronize the shared storage at the production site with the shared storage at the standby site. This will transfer the initial baseline snapshot from the production site to the standby site.
3. Set up the frequency of subsequent copies of the production site shared storage, which will be replicated at the standby site. When asynchronous replication mode is used, then at the requested frequency the changed data blocks at the production site shared storage (based on comparison to the previous snapshot copy) become the new snapshot copy, and the snapshot copy is transferred to the standby site shared storage.
4. After the baseline snapshot copy is performed, validate that all the directories inside the standby site volumes have the same contents as the directories inside the production site volumes.

4.3.2 Validating the Standby Site Setup

Validate the standby site by following the steps below:

1. Shut down any processes still running on the production site. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
2. Stop the replication between the production site shared storage and the standby site shared storage.
3. Use Oracle Data Guard to fail over the databases.
4. On the standby site hosts, manually start all the processes. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
5. Use a browser client to perform post-failover testing to confirm that requests are being resolved and redirected to the standby site.

4.4 Creating an Asymmetric Standby Site

The steps in this section describe how to set up an asymmetric Oracle Fusion Middleware Disaster Recovery topology.

An asymmetric topology is a disaster recovery configuration that is different across tiers at the production site and standby site. In most asymmetric Oracle Fusion Middleware Disaster Recovery topologies, the standby site has fewer resources than the production site.

Before you read this section, be sure to read and understand the concepts and information on setting up a symmetric topology presented earlier in this manual. Many of the concepts for setting up a symmetric topology are also valid for setting up an asymmetric topology.

[Section 4.4.1, "Creating the Asymmetric Standby Site"](#) describes the basic steps for creating an asymmetric topology. It does not describe in detail applicable concepts for setting up an asymmetric topology that were previously described for symmetric topologies earlier in this chapter.

4.4.1 Creating the Asymmetric Standby Site

This section describes the high level steps for creating any type of asymmetric Oracle Fusion Middleware Disaster Recovery topology. The production site is the Oracle SOA Suite enterprise deployment shown in [Figure 4-2](#). The standby site will be different from the production site.

To create an asymmetric topology:

1. Design the production site and the standby site. Determine the resources that will be necessary at the standby site to ensure acceptable performance when the standby site assumes the production role.

Note: The ports for the standby site instances must use the same port numbers as the peer instances at the production site. Therefore, ensure that all the port numbers that will be required at the standby site are available (not in use at the standby site).

2. Create the Oracle Fusion Middleware Disaster Recovery production site by performing these operations:
 - a. Create volumes on the production site's shared storage system for the Oracle Fusion Middleware instances that will be installed for the production site. For more information, see [Section 4.1.1, "Directory Structure and Volume Design."](#)
 - b. Create mount points and symbolic links on the production site hosts to the Oracle home directories for the Oracle Fusion Middleware instances on the production site's shared storage system volumes. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links. For more information about volume design, see [Section 4.1.1.1.1, "Volume Design for Oracle SOA Suite."](#)
 - c. Create mount points and symbolic links on the production site hosts to the Oracle Central Inventory directories for the Oracle Fusion Middleware instances on the production site's shared storage system volumes. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links. For more information about the Oracle Central Inventory directories, see [Section 3.2.2, "Oracle Home and Oracle Inventory."](#)
 - d. Create mount points and symbolic links on the production site hosts to the static HTML pages directories for the Oracle HTTP Server instances on the production site's shared storage system volumes, if applicable. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
 - e. Install the Oracle Fusion Middleware instances for the production site on the volumes in the production site's shared storage system. For more information,

see [Section 4.2.1, "Creating the Production Site for the Oracle SOA Suite Topology."](#)

3. Create the same volumes with the same file and directory privileges on the standby site's shared storage system as you created for the Oracle Fusion Middleware instances on the production site's shared storage system. This step is critical because it enables you to use storage replication later to create the peer Oracle Fusion Middleware instance installations for the standby site instead of installing them using Oracle Universal Installer.

Note: When you configure storage replication, ensure that all the volumes you set up on the production site's shared storage system are replicated to the same volumes on the standby site's shared storage system.

Even though some of the instances and hosts at the production site may not exist at the standby site, you must configure storage replication for all the volumes set up for the production site's Oracle Fusion Middleware instances.

4. Perform any other necessary configuration required by the shared storage vendor to enable storage replication between the production site's shared storage system and the standby site's shared storage system. Configure storage replication to asynchronously copy the volumes in the production site's shared storage system to the standby site's shared storage system.
5. Create the initial baseline snapshot copy of the production site shared storage system to set up the replication between the production site and standby site shared storage systems. Create the initial baseline snapshot and subsequent snapshot copies using asynchronous replication mode. After the baseline snapshot copy is performed, validate that all the directories for the standby site volumes have the same contents as the directories for the production site volumes. Refer to the documentation for your shared storage vendor for information on creating the initial snapshot and enabled storage replication between the production site and standby site shared storage systems.
6. After the baseline snapshot has been taken, perform these steps for the Oracle Fusion Middleware instances for the standby site hosts:
 - a. Set up a mount point directory on the standby site host to the Oracle home directory for the Oracle Fusion Middleware instance on the standby site's shared storage system. The mount point directory you set up for the peer instance on the standby site host must be the same as the mount point directory you set up for the instance on the production site host.
 - b. Set up a symbolic link on the standby site host to the Oracle home directory for the Oracle Fusion Middleware instance on the standby site's shared storage system. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link you set up for the instance on the production site host.
 - c. Set up a mount point directory on the standby site host to the Oracle Central Inventory directory for the Oracle Fusion Middleware instance on the standby site's shared storage system. The mount point directory you set up for the peer

instance on the standby site host must be the same as the mount point directory you set up for the instance on the production site host.

- d. Set up a symbolic link on the standby site host to the Oracle Central Inventory directory for the Oracle Fusion Middleware instance on the standby site's shared storage system. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link you set up for the instance on the production site host.
- e. Set up a mount point directory on the standby site host to the Oracle HTTP Server static HTML pages directory for the Oracle HTTP Server instance on the standby site's shared storage system. The mount point directory you set up for the peer instance on the standby site host must be the same as the mount point directory you set up for the instance on the production site host.
- f. Set up a symbolic link on the standby site host to the Oracle HTTP Server static HTML pages directory for the Oracle HTTP Server instance on the standby site's shared storage system. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link you set up for the instance on the production site host.

After completing these steps, the Oracle Fusion Middleware instance installations for the production site have been replicated to the standby site. At the standby site, all of the following are true:

- The Oracle Fusion Middleware instances are installed into the same Oracle home directories on the same volumes as at the production site, and the hosts use the same mount point directories and symbolic links for the Oracle home directories as at the production site. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
- The Oracle Central Inventory directories are located in same directories on the same volumes as at the production site, and the hosts use the same mount point directories and symbolic links for the Oracle Central Inventory directories as at the production site. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
- The Oracle HTTP Server static HTML pages directories are located in same directories on the same volumes as at the production site, and the hosts use the same mount point directories and symbolic links for the Oracle HTTP Server static HTML pages directories as at the production site. Note that symbolic links are required only in cases where the storage system does not guarantee consistent replication across multiple volumes; see [Section 3.2.3, "Storage Replication"](#) for more details about symbolic links.
- The same ports are used for the standby site Oracle Fusion Middleware instances as were used for the same instances at the production site.

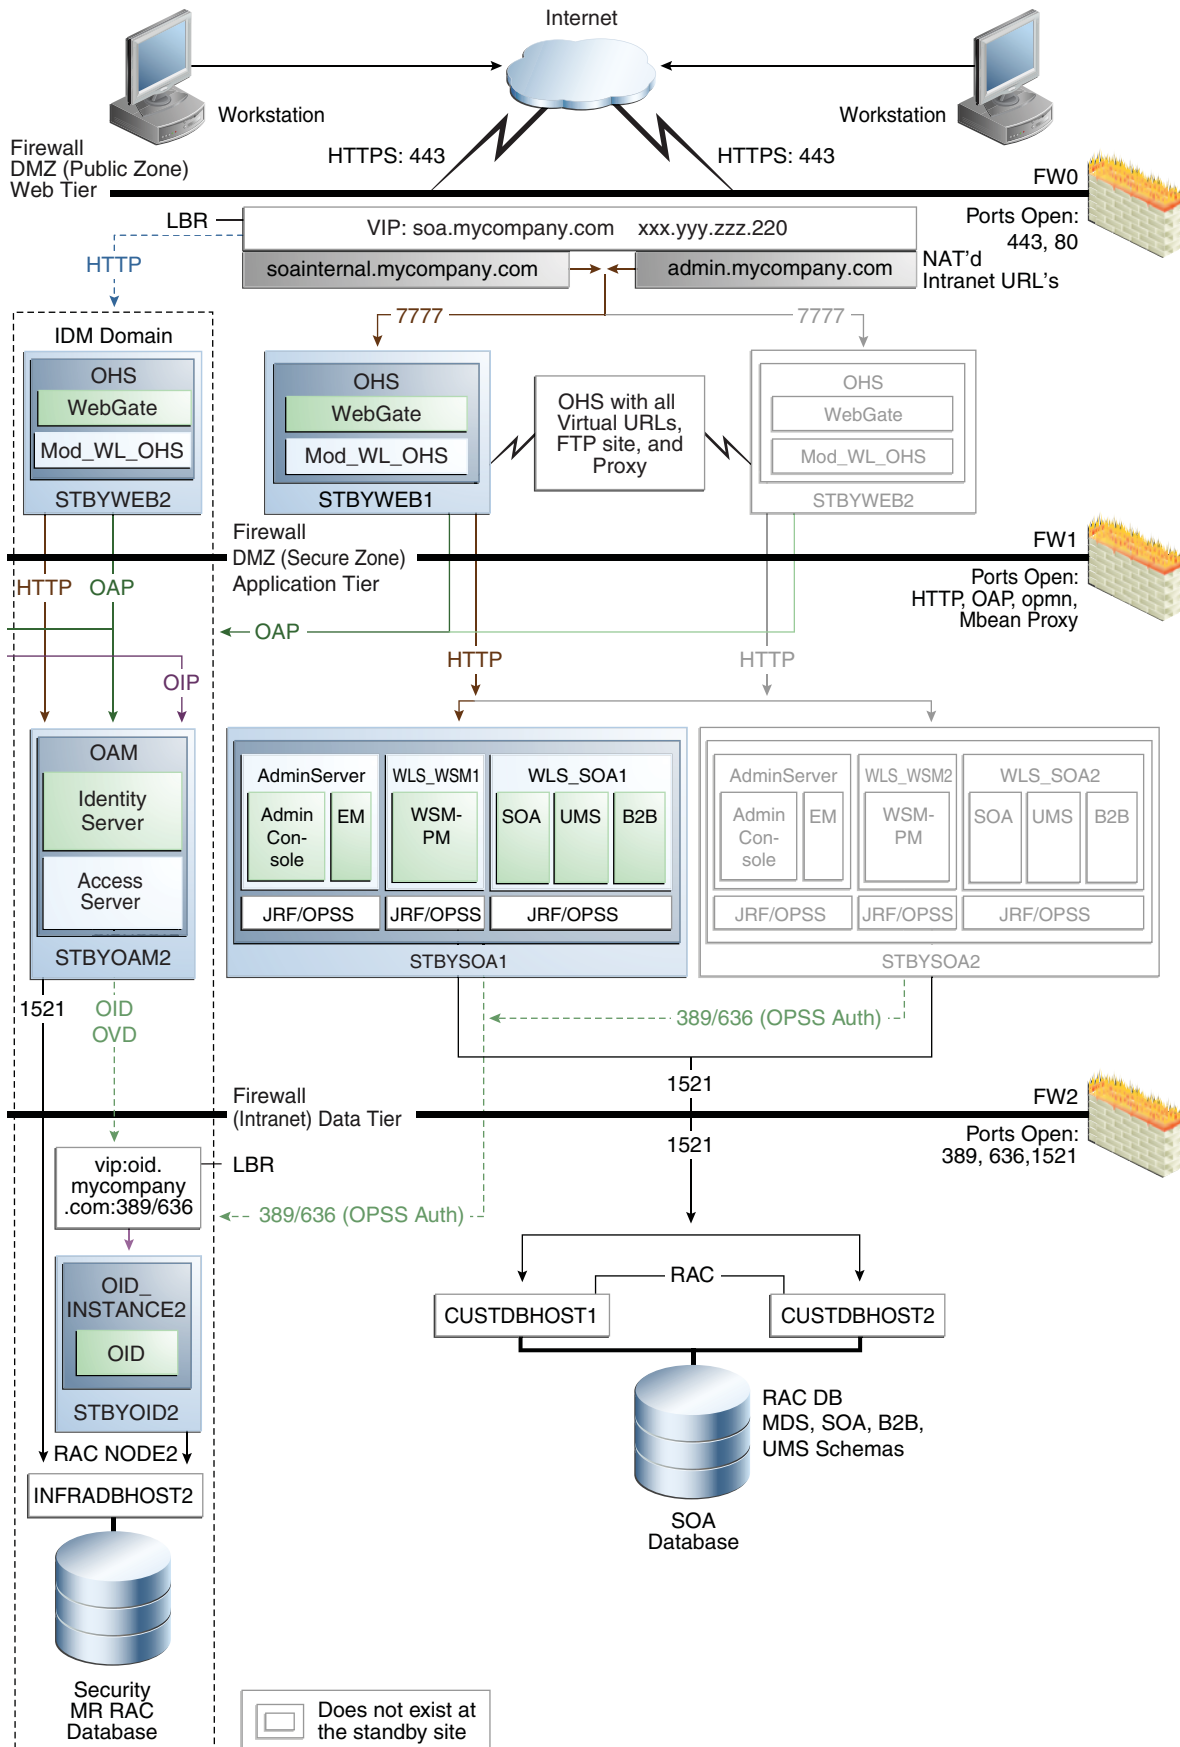
4.4.1.1 Creating an Asymmetric Standby Site with Fewer Hosts and Instances

This section describes how to create an asymmetric standby site that has fewer hosts and Oracle Fusion Middleware instances than the production site.

The production site for this Oracle Fusion Middleware Disaster Recovery topology is the Oracle SOA Suite enterprise deployment shown in [Figure 4–2](#). [Section 4.1, "Setting Up the Site"](#) through [Section 4.1.1, "Directory Structure and Volume Design"](#) describe how to set up this production site and the volumes for its shared storage system, and how to create the necessary mount points.

[Figure 4–15](#) shows the asymmetric standby site for the production site shown in [Figure 4–2](#).

Figure 4–15 *An Asymmetric Standby Site with Fewer Hosts and Instances*



The Oracle SOA Suite asymmetric standby site shown in [Figure 4-15](#) has fewer hosts and instances than the Oracle SOA Suite production site shown in [Figure 4-2](#).

The hosts WEBHOST2 and SOAHOST2 and the instances on those hosts exist at the production site in [Figure 4-2](#), but these hosts and their instances do not exist at the asymmetric standby site in [Figure 4-15](#). The standby site therefore has fewer hosts and fewer instances than the production site.

It is important to ensure that this asymmetric standby site will have sufficient resources to provide adequate performance when it assumes the production role.

When you follow the steps in [Section 4.4.1, "Creating the Asymmetric Standby Site"](#) to set up this asymmetric standby site, the standby site should be properly configured to assume the production role.

To set up the asymmetric standby site correctly, create the same volumes and consistency groups on the standby site shared storage as you did on the production site shared storage. For example, for the Oracle SOA Suite deployment, the volume design recommendations in [Table 4-4](#) and the consistency group recommendations in [Table 4-5](#) were used to set up the production site shared storage. You will use these same volume design recommendations and consistency group recommendations that you used for the production site shared storage to set up the asymmetric standby site's shared storage.

Note that at an asymmetric standby site, some hosts that exist at the production site do not exist at the standby site. For example, in the case of the asymmetric standby site for Oracle SOA Suite shown in [Figure 4-15](#), WEBHOST2 and SOAHOST2 do not exist at the standby site, therefore it is not possible or necessary for you to create mount points on these hosts to the standby site shared storage volumes.

4.4.2 Validating the Asymmetric Standby Site Setup

Validate the standby site by following the steps below:

1. Shut down any processes still running on the production site. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
2. Stop the replication between the production site shared storage and the standby site shared storage.
3. Use Oracle Data Guard to fail over the databases.
4. On the standby site hosts, manually start all the processes. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
5. Use a browser client to perform post-failover testing to confirm that requests are being resolved and redirected to the standby site.

4.5 Performing Site Operations and Administration

This section describes operations and administration to perform on your Oracle Fusion Middleware Disaster Recovery topology.

4.5.1 Synchronizing the Sites

The standby site shared storage receives snapshots transferred on a periodic basis from the production site shared storage. After the snapshots are applied, the standby site shared storage will include all the data up to and including the data contained in

the last snapshot transferred from the production site before the failover or switchover.

You should manually force a synchronization operation whenever a change is made to the middle tier at the production site (for example, when a new application is deployed at the production site). Follow the vendor-specific instructions for forcing a synchronization using storage replication technology.

The synchronization of the databases in the Oracle Fusion Middleware Disaster Recovery topology is managed by Oracle Data Guard.

4.5.2 Performing a Switchover

When you plan to take down the production site (for example, to perform maintenance) and make the current standby site the new production site, you must perform a switchover operation so that the standby site takes over the production role.

Follow these steps to perform a switchover operation:

1. Shut down any processes still running on the production site. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
2. Stop the replication between the production site shared storage and the standby site shared storage.
3. Use Oracle Data Guard to switch over the databases.
4. On the standby site hosts, manually start all the processes. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
5. Ensure that all user requests are routed to the standby site by performing a global DNS push or something similar, such as updating the global load balancer.
6. Use a browser client to perform post-switchover testing to confirm that requests are being resolved and redirected to the standby site.

At this point, the former standby site is the new production site and the former production site is the new standby site.

7. Reestablish the replication between the two sites, but configure the replication so that the snapshot copies go in the opposite direction (from the current production site to the current standby site). Refer to the documentation for your shared storage to learn how to configure the replication so that snapshot copies are transferred in the opposite direction.

After these steps have been performed, the former standby site is the new production site. At this point, you can perform maintenance at the original production site. After performing the planned tasks on the original production site, you can use it again at some point in the future as either the production site or standby site.

To use the original production site as the new production site, perform the switchback steps described in [Section 4.5.3, "Performing a Switchback."](#)

4.5.3 Performing a Switchback

After a switchover operation has been performed, a switchback operation can be performed to revert the current production site and the current standby site to the roles they had prior to the switchover operation.

Follow these steps to perform a switchback operation:

1. Shut down any processes running on the current production site. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
2. Stop the replication between the current production site shared storage and standby site shared storage.
3. Use Oracle Data Guard to switch back the databases.
4. On the new production site hosts, manually start all the processes. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
5. Ensure that all user requests are routed to the new production site by performing a global DNS push or something similar, such as updating the global load balancer.
6. Use a browser client to perform post-switchback testing to confirm that requests are being resolved and redirected to the new production site.

At this point, the former standby site is the new production site and the former production site is the new standby site.
7. Reestablish the replication between the two sites, but configure the replication so that the snapshot copies go in the opposite direction (from the new production site to the new standby site). Refer to the documentation for your shared storage to learn how to configure the replication so that snapshot copies are transferred in the opposite direction.

4.5.4 Performing a Failover

When the production site becomes unavailable unexpectedly, you must perform a failover operation so that the standby site takes over the production role.

Follow these steps to perform a failover operation:

1. Stop the replication between the production site shared storage and the standby site shared storage.
2. From the standby site, use Oracle Data Guard to fail over the databases.
3. On the standby site hosts, manually start all the processes. This includes the database instances in the data tier, Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
4. Ensure that all user requests are routed to the standby site by performing a global DNS push or something similar, such as updating the global load balancer.
5. Use a browser client to perform post-failover testing to confirm that requests are being resolved and redirected to the production site.

At this point, the standby site is the new production site. You can examine the issues that caused the former production site to become unavailable.

6. To use the original production site as the current standby site, you must reestablish the replication between the two sites, but configure the replication so that the snapshot copies go in the opposite direction (from the current production site to the current standby site). Refer to the documentation for your shared storage system to learn how to configure the replication so that snapshot copies are transferred in the opposite direction.

To use the original production site as the new production site, perform the switchback steps in [Section 4.5.3, "Performing a Switchback."](#)

4.5.5 Performing Periodic Testing of the Standby Site.

This manual describes how to set up Disaster Recovery for an Oracle Fusion Middleware production site and standby site. In a normal Oracle Fusion Middleware Disaster Recovery configuration, the following are true:

- Storage replication is used to copy Oracle Fusion Middleware middle tier file systems and data from the production site shared storage to the standby site shared storage. During normal operation, the production site is active and the standby site is passive. When the production site is active, the standby site is passive and the standby site shared storage is in read-only mode; the only write operations made to the standby site shared storage are the storage replication operations from the production site shared storage to the standby site shared storage.
- Oracle Data Guard is used to copy database data for the production site Oracle databases to the standby databases at standby site. By default, the production site databases are active and the standby databases at the standby site are passive. The standby databases at the standby site are in managed recovery mode while the standby site is in the standby role (is passive). When the production site is active, the only write operations made to the standby databases are the database synchronization operations performed by Oracle Data Guard.
- When the production site becomes unavailable, the standby site is enabled to take over the production role. If the current production site becomes unavailable unexpectedly, then a failover operation (described in [Section 4.5.4, "Performing a Failover"](#)) is performed to enable the standby site to assume the production role. Or, if the current production site is taken down intentionally (for example, for planned maintenance), then a switchover operation (described in [Section 4.5.2, "Performing a Switchover"](#)) is performed to enable the standby site to assume the production role.

The usual method of testing a standby site is to shut down the current production site and perform a switchover operation to enable the standby site to assume the production role. However, some enterprises may want to perform periodic testing of their Disaster Recovery standby site without shutting down the current production site and performing a switchover operation.

An alternate method of testing the standby site without shutting down the current production site is to create a clone of the read-only standby site shared storage and then use the cloned standby site shared storage in testing. To use this alternate testing method, perform these steps:

1. Use the cloning technology provided by the shared storage vendor to create a clone of the standby site's read-only volumes on the shared storage at the standby site. Ensure that the cloned standby site volumes are writable. If you want to test the standby site just once, then this can be a one-time clone operation, but if you want to test the standby site regularly, you can set up periodic cloning of the standby site read-only volumes to the standby site's cloned read/write volumes.
2. Perform a backup of the standby site databases, then modify the Oracle Data Guard replication between the production site and standby site databases.
 - For 10.1 databases, break the replication by following the instructions in the 10.1 Oracle Data Guard documentation.
 - For 10.2 and later databases, follow these steps to establish a snapshot standby database:
 - a. If you do not have a Flash Recovery Area, set one up.

- b. Mount the read-only shared storage volumes.

At this point, the mount commands are reset to what they were before the standby site testing was performed.

2. Configure Oracle Data Guard to perform replication between the production site databases and standby databases at the standby site. Performing this configuration puts the standby database into managed recovery mode again:
 - For 10.1 databases, reinitiate the databases by following the instructions in the 10.1 Oracle Data Guard documentation.
 - For 10.2 and later databases, follow these steps:
 - a. Revert the activated database back to a physical standby database:


```
SQL> STARTUP MOUNT FORCE;
SQL> FLASHBACK DATABASE TO POINT standbytest;
SQL> ALTER DATABASE CONVERT TO PHYSICAL STANDBY;
SQL> STARTUP MOUNT FORCE;
```
 - b. Restart managed recovery:


```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY
          DATABASE USING CURRENT LOGFILE DISCONNECT;
```
 - c. Reenable the standby destination and switch logs:


```
SQL> ALTER SYSTEM SET
          LOG ARCHIVE DEST STATE 2=ENABLE;
```
 - For 11g databases, set up the replication again by following the steps in the "Managing a Snapshot Standby Database" section in *Oracle Data Guard Concepts and Administration*.
3. Before using the original production site again, modify the host name resolution method for the computers that will be used to access the production site to ensure that the host names point to the production site computers and not the standby site computers. For example, on a Linux computer, change the `/etc/hosts` file to point to the virtual IP of the load balancer for the production site.

4.5.6 Using Peer to Peer File Copy for Testing

As an alternative to using storage replication technology for disaster protection and disaster recovery of Oracle Fusion Middleware middle tier components, you can use peer to peer file copy mechanisms in test environments to replicate middle tier file system data from a production site host to a standby site peer host in an Oracle Fusion Middleware Disaster Recovery topology. An example of a peer to peer file copy mechanism is `rsync` (an open source utility for UNIX systems).

This section describes how to use `rsync` instead of storage replication in your Oracle Fusion Middleware Disaster Recovery topology. This section discusses `rsync` in the context of symmetric topologies. For more information about symmetric topologies, refer to [Section 4.4, "Creating an Asymmetric Standby Site."](#) The information provided for `rsync` in this section also applies to other peer to peer file copy mechanisms.

Before you read this section, read the rest of this manual to ensure that you are familiar with how to use storage replication and Oracle Data Guard in an Oracle Fusion Middleware Disaster Recovery topology. There are many similarities between using storage replication and `rsync` for disaster protection and disaster recovery of your Oracle Fusion Middleware components.

Note: You can use rsync instead of storage replication technology to replicate middle tier file system data from the production site to the standby site. However, be aware that the following beneficial storage replication features are not available when you use rsync:

- With storage replication, you can roll changes back to the point in time when any previous snapshot was taken at the production site.

With rsync, replicated production site data overwrites the standby site data, and you cannot roll back a replication.

- With storage replication, the volume you set up for each host cluster in the shared storage systems ensures data consistency for that host cluster across the production site's shared storage system and the standby site's shared storage system.

With rsync, data consistency is not guaranteed.

Because of these deficiencies in comparison to storage replication, rsync is not supported for disaster recovery use in actual production environments.

4.5.6.1 Using rsync and Oracle Data Guard for Oracle Fusion Middleware Disaster Recovery Topologies

These two basic principles apply when you use rsync and Oracle Data Guard to provide disaster protection and disaster recovery for your Oracle Fusion Middleware Disaster Recovery topology:

1. Use rsync for disaster protection of your Oracle Fusion Middleware middle tier components.
2. Use Oracle Data Guard for disaster protection of Oracle databases that are used in your Oracle Fusion Middleware topology. [Section 3.3, "Database Considerations"](#) describes how to set up Oracle Data Guard to provide disaster recovery for Oracle database.

4.5.6.1.1 Using rsync for Oracle Fusion Middleware Middle Tier Components Follow these steps to use rsync to provide disaster protection and disaster recovery for your Oracle Fusion Middleware middle tier components:

1. Set up rsync to enable replication of files from a production site host to its standby site peer host. See the rsync man page for instructions on installing and setting up rsync, and for syntax and usage information. Information about rsync is also available at <http://rsync.samba.org>.
2. For each production site host on which one or more Oracle Fusion Middleware components has been installed, set up rsync to copy the following directories and files to the same directories and files on the standby site peer host:
 - The Oracle Fusion Middleware home directory and subdirectories, and all the files in them.
 - The Oracle Central Inventory directory and files for the host, which includes the Oracle Universal Installer entries for the Oracle Fusion Middleware installations.
 - If applicable, the Oracle Fusion Middleware static HTML pages directory for the Oracle HTTP Server installations on the host.

- If applicable, the .fmb and .fmx deployment artifact files created by Oracle Forms on the host, and the .rdf deployment artifact files created by Oracle Reports on the host.

Note: Run rsync as root. If you want rsync to work without prompting users for a password, set up SSH keys between the production site host and standby site host, so that SSH does not prompt for a password.

3. Set up scheduled jobs, for example, cron jobs, for the production site hosts for which you set up rsync in the previous step. These scheduled jobs enable rsync to automatically perform replication of these files from the production site hosts to the standby site hosts on a regular interval. An interval of once a day is recommended for a production site where the Oracle Fusion Middleware configuration does not change very often.
4. Whenever a change is made to the configuration of an Oracle Fusion Middleware middle tier configuration on a production site host (for example, when a new application is deployed), you should perform a manual synchronization of that host with its standby site peer host using rsync.
5. Whenever you perform a manual rsync synchronization of an Oracle Fusion Middleware middle tier instance on a production site host to the peer standby site host, you should also manually force a synchronization of any associated database repository for the production site's Oracle Fusion Middleware instance to the standby site using Oracle Data Guard. See [Section 3.3.2, "Manually Forcing Database Synchronization with Oracle Data Guard"](#) for more information on manually forcing a synchronization of an Oracle database using Oracle Data Guard.

4.5.6.1.2 Performing Failover and Switchover Operations Follow these steps to perform a failover or switchover from the production site to the standby site when you are using rsync:

1. Shut down any processes still running on the production site (if applicable).
2. Stop the rsync jobs between the production site hosts and their standby site peer hosts.
3. Use Oracle Data Guard to fail over the production site databases to the standby site.
4. On the standby site, manually start the processes for the Oracle Fusion Middleware Server instances.
5. Route all user requests to the standby site by performing a global DNS push or something similar, such as updating the global load balancer.
6. Use a browser client to perform post-failover or post-switchover testing to confirm that requests are being resolved at the standby site (current production site).
At this point, the standby site is the new production site and the production site is the new standby site.
7. Reestablish the rsync replications between the two sites, but configure the replications so that they go in the opposite direction (from the current production site to the current standby site).

To use the original production site as the new production site, you perform the steps above again, but configure the rsync replications to go in the original direction (from the original production site to the original standby site).

4.6 Patching an Oracle Fusion Middleware Disaster Recovery Site

This section describes how to apply an 11g Oracle Fusion Middleware patch set to upgrade the Oracle homes that participate in an Oracle Fusion Middleware Disaster Recovery site.

The list in this section describes the steps for applying a patch set to upgrade the 11g Oracle Fusion Middleware homes in an Oracle Fusion Middleware Disaster Recovery production site.

The following steps assume that the Oracle Central Inventory for any Oracle Fusion Middleware instance that you are patching is located on the production site shared storage, so that the Oracle Central Inventory for the patched instance can be replicated to the standby site.

Use the following procedure to upgrade 11g Oracle Fusion Middleware patch versions:

1. Perform a backup of the production site to ensure that the starting state is secured.
2. Apply the patch set to upgrade the production site instances.
3. After applying the patch set, manually force a synchronization of the production site shared storage and standby site shared storage. This replicates the production site's patched instance and Oracle Central Inventory in the standby site's shared storage.
4. After applying the patch set, use Oracle Data Guard to manually force a synchronization of the Oracle databases at the production site and standby sites. Some Oracle Fusion Middleware patch sets may make updates to repositories, so this step ensures that any changes made to production site databases are synchronized to the standby site databases.
5. The upgrade is now complete. Your Disaster Recovery topology is ready to resume processing.

Note: Patches must be applied only at the production site for an 11g Oracle Fusion Middleware Disaster Recovery topology. If a patch is for an Oracle Fusion Middleware instance or for the Oracle Central Inventory, the patch will be copied when the production site shared storage is replicated to the standby site shared storage. A synchronization operation should be performed when a patch is installed at the production site.

Similarly, if a patch is installed for a production site database, Oracle Data Guard will copy the patch to the standby database at the standby site when a synchronization is performed.

Troubleshooting Disaster Recovery

This chapter describes common situations that you might encounter when deploying and managing Oracle Fusion Middleware in Disaster Recovery topologies, and explains the steps for addressing them. It contains the following topics:

- [Troubleshooting Oracle Fusion Middleware Disaster Recovery Topologies](#)
- [Need More Help?](#)

5.1 Troubleshooting Oracle Fusion Middleware Disaster Recovery Topologies

This section describes common situations and steps to perform in Oracle Fusion Middleware configurations. It contains the following topics:

- [Verify Host Name Resolution at the Production and Standby Sites](#)
- [Resolving Issues with Components in a Disaster Recovery Topology](#)
- [Resolving Issues with Components Deployed on Shared Storage](#)

5.1.1 Verify Host Name Resolution at the Production and Standby Sites

Many issues that may arise with your Disaster Recovery topology are caused by host name resolution not having been set up properly for the production site and standby site.

Make sure that host name resolution is set up properly by performing the host name validation steps in [Section 3.1.1.5, "Testing the Host Name Resolution."](#)

5.1.2 Resolving Issues with Components in a Disaster Recovery Topology

Some issues that may arise with a component in a Disaster Recovery topology are not Disaster Recovery issues, but rather are issues with the component itself.

If you encounter problems with an Oracle Fusion Middleware component used in a Disaster Recovery topology, check the Troubleshooting section in the *Oracle Fusion Middleware High Availability Guide* for that component to see if the problem is described there.

Similarly, if your Disaster Recovery topology is based on one or more of the enterprise deployments described in the following manuals and you encounter a problem, check the Troubleshooting section of that manual to see if the problem is described there:

- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Portal*

- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management*
- Detailed information about the Oracle Portal, Forms, Reports, and Discover enterprise topology is available in the 11.1.1.2 *Oracle Portal Enterprise Deployment Guide*. See Article ID 952068.1 "Oracle Fusion Middleware 11g (11.1.1.2) Enterprise Deployment Guides for Portal, Forms, Reports, and Discover" at My Oracle Support (formerly Oracle *MetaLink*) for information on obtaining the manual. The URL for My Oracle Support is:
<http://support.oracle.com>
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle WebCenter Content*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence*

5.1.3 Resolving Issues with Components Deployed on Shared Storage

Some problems that may arise with a component in your Disaster Recovery topology are not Disaster Recovery issues, but are issues associated with deploying the component on shared storage.

If you do not find any shared storage problems described in the manuals mentioned in [Section 5.1.2, "Resolving Issues with Components in a Disaster Recovery Topology,"](#) then look for notes that describe shared storage problems in the *Oracle Fusion Middleware Release Notes*, available on the Oracle Technology Network at:

<http://www.oracle.com/technology/documentation/middleware.html>

5.2 Need More Help?

In case the information in the previous section is not sufficient, you can find more solutions on My Oracle Support (formerly Oracle *MetaLink*) at:

<http://support.oracle.com>

If you do not find a solution for your problem, log a service request.

You can also read the *Oracle Fusion Middleware Release Notes*, available on the Oracle Technology Network at:

<http://www.oracle.com/technology/documentation/middleware.html>

Managing Oracle Inventory

This appendix describes how to manage your Oracle Inventory on the production and standby sites for an Oracle Fusion Middleware Disaster Recovery topology.

It includes this topic:

- [Updating Oracle Inventory](#)
- [Updating the Windows Registry](#)

A.1 Updating Oracle Inventory

When you update the Oracle inventory (for example, by installing new Oracle software, or by applying an Oracle patch set or patch to existing Oracle software) on a production site host, you must make sure that the same software updates are made on the standby site peer host.

To do this, you must update the Oracle inventory on the standby site peer host by executing the following script:

```
ORACLE_HOME/oui/bin/attachHome.sh
```

In addition, you must update the `beahomelist` file to edit the location of a Middleware home. Edit the following file to update the Middleware home information:

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

A.2 Updating the Windows Registry

When you update the Oracle inventory (for example, by installing new Oracle software, or by applying an Oracle patch set or patch to existing Oracle software) on a production site Windows host, you need to make sure that the same software updates are made on the standby site peer host by exporting the following Windows Registry key on the production site host and importing it on the standby site peer host:

```
HKEY_LOCAL_MACHINE\Software\oracle
```

In addition, when you modify system components, such as Oracle Web Cache, you must export the following Windows Registry key on the production site host and import it on the standby site peer host:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services
```

To import a key that you have previously exported, use the following command:

```
regedit /I <FileName>
```

For example:

```
regedit /I C:\oracleregistry.reg
```

You can also use the Registry Editor to import the key. See the Registry Editor Help for more information.

Index

A

- alias host name
 - definition, 1-2
 - setting up for database hosts, 3-36
- architecture
 - for Disaster Recovery solution, 1-4
- artifacts
 - common to all Oracle Identity Management components, 2-21, 2-36, 2-41
 - common to all Oracle product suites, 2-1
 - common to all Oracle WebCenter components, 2-6
 - common to all SOA Suite components, 2-11
 - common to Oracle Portal, Forms, Reports, and Discoverer, 2-29
 - common to Oracle WebLogic Server, 2-2
 - for Oracle Fusion Middleware, 3-32
- asymmetric standby site
 - creating, 4-67
 - with fewer hosts and instances than production site, 4-71
- asymmetric standby site setup
 - validating, 4-74
- asymmetric topology
 - creating, 4-68
 - definition, 1-2
 - design considerations, 3-39
- avoiding duplicate messages
 - Oracle WebLogic Server JMS and T-Logs, 2-3

C

- certificates
 - generating self-signed certificates, 4-58
- common artifacts for all Oracle Identity Management components, 2-21, 2-36, 2-41
- common artifacts for all Oracle product suites, 2-1
- common artifacts for all Oracle WebCenter components, 2-6
- common artifacts for all SOA Suite components, 2-11
- common artifacts for Oracle Portal, Forms, Reports, and Discoverer, 2-29
- common artifacts for Oracle WebLogic Server, 2-2
- consistency group recommendations

- for Oracle Identity Management, 4-29
 - for Oracle Portal, Forms, Reports, and Discoverer, 4-31, 4-37
 - for Oracle SOA Suite, 4-11, 4-42, 4-48
 - for Oracle WebCenter, 4-17
- custom keystores
 - configuring Node Manager for, 4-60

D

- database artifacts
 - Oracle Access Manager, 2-25, 2-26, 2-27, 2-28
 - Oracle ADF, 2-5
 - Oracle B2B, 2-16
 - Oracle BPEL Process Manager, 2-13
 - Oracle Business Activity Monitoring, 2-20
 - Oracle Business Intelligence Discoverer, 2-33
 - Oracle Directory Integration Platform, 2-24, 2-39
 - Oracle Directory Services Manager, 2-25, 2-40, 2-42, 2-43
 - Oracle Forms, 2-31
 - Oracle HTTP Server, 2-34
 - Oracle Human Workflow, 2-15
 - Oracle Identity Federation, 2-24, 2-39
 - Oracle Identity Management, 2-21, 2-36, 2-41
 - Oracle Internet Directory, 2-22, 2-37
 - Oracle JCA Adapters, 2-19
 - Oracle Mediator, 2-14
 - Oracle Platform Security Services, 2-4
 - Oracle Portal, 2-30
 - Oracle Portal, Forms, Reports, and Discoverer, 2-29
 - Oracle Reports, 2-32
 - Oracle SOA Service Infrastructure, 2-13
 - Oracle SOA Suite, 2-12
 - Oracle User Messaging Service, 2-18
 - Oracle Virtual Directory, 2-23, 2-38
 - Oracle Web Cache, 2-35
 - Oracle Web Services Manager, 2-17
 - Oracle WebCenter, 2-6
 - Oracle WebCenter Discussions Server, 2-8
 - Oracle WebCenter Portlets, 2-8
 - Oracle WebCenter Spaces, 2-7
 - Oracle WebCenter Wiki and Blog Server, 2-9
 - Oracle WebLogic Server JMS and T-Logs, 2-2
- database considerations, 3-34

- database recommendations, 4-50
- databases
 - forcing manual synchronization with Oracle Data Guard, 3-35
 - making TNSNAMES.ORA entries, 3-35
 - setting up alias host names for database hosts, 3-36
- definitions of Disaster Recovery terminology, 1-2
- design considerations
 - for a symmetric topology, 3-39
 - for an asymmetric topology, 3-39
- designing
 - a Disaster Recovery topology by creating a new production site, 3-38
 - a Disaster Recovery topology from a partially existing site, 3-37
 - a Disaster Recovery topology from an existing production site, 3-37
 - a symmetric Disaster Recovery topology, 3-4
- directory structure
 - for Disaster Recovery sites, 4-2
- directory structure recommendations
 - for Oracle Identity Management, 4-18
 - for Oracle SOA Suite, 4-2, 4-38, 4-43
 - for Oracle WebCenter, 4-12
- disaster
 - definition, 1-2
- Disaster Recovery
 - active production site, 1-6
 - active/passive model, 1-1, 1-6
 - creating mount points to shared storage, 1-6
 - definition, 1-2
 - deploying components on shared storage, 1-6
 - design considerations for a symmetric topology, 3-4
 - designing a topology by creating a new production site, 3-38
 - designing a topology from a partially existing site, 3-37
 - designing a topology from an existing production site, 3-37
 - directory structure and volume design for sites, 4-2
 - DNS server resolution, 3-14
 - for third party databases, 1-8
 - forcing manual synchronization of databases, 3-35
 - forcing manual synchronization of middle tier, 4-75
 - key aspects, 1-6
 - local host name resolution, 3-14, 3-15
 - overview, 1-1
 - overview of architecture, 1-4
 - passive standby site, 1-6
 - performing site administration, 4-74
 - performing site operations, 4-74
 - problems solved by, 1-1
 - protecting Oracle databases, 1-5
 - protecting Oracle Fusion Middleware product binaries, configuration, and metadata
 - files, 1-5
 - protecting third party databases, 1-5
 - setting up a site, 4-1
 - setting up and managing sites, 4-1
 - starting points for setting up sites, 3-37
 - synchronizing the sites, 4-74
 - terminology, 1-2
 - testing host name resolution, 3-18
 - use of Oracle Data Guard, 1-2
 - use of storage replication technology, 1-2
 - using a global load balancer, 3-31
 - using an /etc/hosts file for host name resolution, 3-15
 - using peer to peer file copying in test environments, 4-79
- Disaster Recovery recommendations
 - Oracle Access Manager, 2-25, 2-26, 2-27, 2-28
 - Oracle ADF, 2-5
 - Oracle B2B, 2-15
 - Oracle BPEL Process Manager, 2-13
 - Oracle Business Activity Monitoring, 2-19, 2-20
 - Oracle Business Intelligence Discoverer, 2-32
 - Oracle Directory Integration Platform, 2-23, 2-38
 - Oracle Directory Services Manager, 2-25, 2-40
 - Oracle Forms, 2-31
 - Oracle HTTP Server, 2-33
 - Oracle Human Workflow, 2-15
 - Oracle Identity Federation, 2-24, 2-39
 - Oracle Identity Management, 2-21, 2-36
 - Oracle Internet Directory, 2-22, 2-37
 - Oracle JCA Adapters, 2-18
 - Oracle Mediator, 2-14
 - Oracle Platform Security Services, 2-4
 - Oracle Portal, 2-30
 - Oracle Portal, Forms, Reports, and Discoverer, 2-29
 - Oracle Reports, 2-32
 - Oracle SOA Service Infrastructure, 2-12
 - Oracle SOA Suite, 2-10
 - Oracle User Messaging Service, 2-17
 - Oracle Virtual Directory, 2-23, 2-38
 - Oracle Web Cache, 2-35
 - Oracle Web Services Manager, 2-16
 - Oracle WebCenter, 2-6
 - Oracle WebCenter Discussions Server, 2-8
 - Oracle WebCenter Portlets, 2-8
 - Oracle WebCenter Spaces, 2-7
 - Oracle WebCenter Wiki and Blog Server, 2-9
 - Oracle WebLogic Server JMS and T-Logs, 2-2
- DNS
 - resolving host names using global DNS servers, 3-17
 - resolving host names using separate DNS servers, 3-16
- DNS switchover
 - performing by manually changing host name to IP mapping, 3-31
 - performing using a global load balancer, 3-31
- dynamic or run-time artifacts
 - for Oracle Fusion Middleware, 3-32

E

- enterprise deployment guides
 - obtaining for Oracle Portal, Forms, Reports, and Discoverer, 4-31, 4-32, 4-64, 4-65, 5-2
- /etc/hosts file
 - using for local host name resolution, 3-14

F

- failover
 - definition, 1-3
 - steps for performing, 4-76
- file system artifacts
 - Oracle ADF, 2-5
 - Oracle B2B, 2-15
 - Oracle HTTP Server, 2-34
 - Oracle Identity Management, 2-21, 2-36, 2-37, 2-39, 2-41
 - Oracle JCA Adapters, 2-19
 - Oracle Portal, Forms, Reports, and Discoverer, 2-29
 - Oracle Reports, 2-32
 - Oracle SOA Suite, 2-11
 - Oracle User Messaging Service, 2-17
 - Oracle Web Cache, 2-35
 - Oracle WebCenter, 2-6
 - Oracle WebLogic Server, 2-2
 - Oracle WebLogic Server JMS and T-Logs, 2-2
- file-based persistent store, 3-34

G

- global DNS server
 - using to resolve host names, 3-17
- global load balancer
 - using in a Disaster Recovery topology, 3-31

H

- host name
 - planning, 3-5
- host name resolution
 - determining preference, 3-15
 - precedence defined in nsswitch.conf file, 3-15
 - testing, 3-18
 - using an /etc/hosts file for local host name resolution, 3-15
 - using DNS server resolution, 3-14
 - using global DNS server resolution, 3-17
 - using local host name resolution, 3-14, 3-15
 - using ping command to test, 3-18
 - using separate DNS servers resolution, 3-16
- host names
 - planning for production and standby sites, 3-5

I

- identity keystore
 - creating, 4-59

K

- keystore
 - configuring Node Manager for a custom keystore, 4-60
 - creating a trust keystore, 4-59
 - creating an identity keystore, 4-59

L

- LDAP store
 - Oracle Access Manager, 2-25, 2-26, 2-27, 2-28
- load balancer considerations, 3-26
- LockFile directive
 - modifying to fix Oracle HTTP Server performance problems, 2-34

M

- mount points
 - to shared storage locations, 1-6

N

- network artifacts
 - for Oracle ADF, 2-5
 - for Oracle SOA Suite components, 2-11
 - for Oracle WebCenter components, 2-6
 - Oracle Identity Management, 2-21, 2-36, 2-41
 - Oracle Portal, Forms, Reports, and Discoverer, 2-30
 - Oracle WebLogic Server, 2-2
- Node Manager
 - configuring for custom keystores, 4-60
 - setting up on the standby site, 4-58
- nsswitch.conf file
 - specifying host name resolution precedence, 3-15

O

- Oracle Access Manager
 - database artifacts, 2-25, 2-26, 2-27, 2-28
 - Disaster Recovery recommendations, 2-25, 2-26, 2-27, 2-28
 - LDAP store, 2-25, 2-26, 2-27, 2-28
 - recovery recommendations, 2-26, 2-27, 2-28, 2-29
 - special considerations, 2-26, 2-27
 - synchronization recommendations, 2-26, 2-28, 2-29
- Oracle ADF
 - database artifacts, 2-5
 - Disaster Recovery recommendations, 2-5
 - file system artifacts, 2-5
 - network artifacts, 2-5
 - recovery recommendations, 2-6
 - synchronization recommendations, 2-5
- Oracle B2B
 - database artifacts, 2-16
 - Disaster Recovery recommendations, 2-15
 - file system artifacts, 2-15
 - recovery recommendations, 2-16

- special considerations, 2-16
- synchronization recommendations, 2-16
- Oracle BPEL Process Manager
 - database artifacts, 2-13
 - Disaster Recovery recommendations, 2-13
 - recovery recommendations, 2-14
 - synchronization recommendations, 2-13
- Oracle Business Activity Monitoring
 - database artifacts, 2-20
 - Disaster Recovery recommendations, 2-19, 2-20
 - recovery recommendations, 2-20, 2-21
 - synchronization recommendations, 2-20
- Oracle Business Intelligence Discoverer
 - database artifacts, 2-33
 - Disaster Recovery recommendations, 2-32
 - recovery recommendations, 2-33
 - special considerations, 2-33
 - synchronization recommendations, 2-33
- Oracle Data Guard
 - configure Oracle Net Services on the standby site, 4-52
 - create instances and database on the standby site, 4-53
 - environment variables, 4-50
 - gather files and perform a database backup, 4-51
 - setting up, 4-50
 - setup prerequisites and assumptions, 4-50
 - test database switchover and switchback, 4-57
 - using to force manual synchronization of databases, 3-35
 - using to protect Oracle databases, 1-2
- Oracle Directory Integration Platform
 - database artifacts, 2-24, 2-39
 - Disaster Recovery recommendations, 2-23, 2-38
 - recovery recommendations, 2-24, 2-39
 - synchronization recommendations, 2-24, 2-39
- Oracle Directory Services Manager
 - database artifacts, 2-25, 2-40, 2-42, 2-43
 - Disaster Recovery recommendations, 2-25, 2-40
 - recovery recommendations, 2-25, 2-41, 2-42, 2-43, 2-44
 - special considerations, 2-25, 2-40, 2-42, 2-43
 - synchronization recommendations, 2-25, 2-40, 2-42, 2-43, 2-44
- Oracle Forms
 - database artifacts, 2-31
 - Disaster Recovery recommendations, 2-31
 - recovery recommendations, 2-31
 - special considerations, 2-31
 - synchronization recommendations, 2-31
- Oracle Fusion Middleware
 - artifacts, 3-32
 - dynamic or run-time artifacts, 3-32
 - Oracle Home and Oracle Inventory, 3-33
 - protecting product binaries, configuration, and metadata files, 1-5
 - static artifacts, 3-32
- Oracle Home and Oracle Inventory
 - Oracle Fusion Middleware, 3-33
- Oracle HTTP Server
 - database artifacts, 2-34
 - Disaster Recovery recommendations, 2-33
 - file system artifacts, 2-34
 - recovery recommendations, 2-35
 - special considerations, 2-34
 - synchronization recommendations, 2-35
- Oracle HTTP Server performance problems
 - fixing by changing the LockFile directive, 2-34
- Oracle Human Workflow
 - database artifacts, 2-15
 - Disaster Recovery recommendations, 2-15
 - recovery recommendations, 2-15
 - synchronization considerations, 2-15
- Oracle Identity Federation
 - database artifacts, 2-24, 2-39
 - Disaster Recovery recommendations, 2-24, 2-39
 - recovery recommendations, 2-24, 2-40
 - special considerations, 2-24, 2-40
 - synchronization recommendations, 2-24, 2-40
- Oracle Identity Management
 - consistency group recommendations, 4-29
 - database artifacts, 2-21, 2-36, 2-41
 - directory structure recommendations, 4-18
 - Disaster Recovery recommendations, 2-21, 2-36
 - file system artifacts, 2-21, 2-36, 2-37, 2-39, 2-41
 - network artifacts, 2-21, 2-36, 2-41
 - recovery recommendations, 2-22, 2-37, 2-42
 - synchronization recommendations, 2-22, 2-36, 2-41
 - virtual servers, 3-28
 - volume design recommendations, 4-19
- Oracle Internet Directory
 - database artifacts, 2-22, 2-37
 - Disaster Recovery recommendations, 2-22, 2-37
 - recovery recommendations, 2-23, 2-38
 - special considerations, 2-22, 2-37
 - synchronization recommendations, 2-23, 2-37
- Oracle JCA Adapters
 - database artifacts, 2-19
 - Disaster Recovery recommendations, 2-18
 - file system artifacts, 2-19
 - recovery recommendations, 2-19
 - synchronization recommendations, 2-19
- Oracle Mediator
 - database artifacts, 2-14
 - Disaster Recovery recommendations, 2-14
 - recovery recommendations, 2-14
 - synchronization recommendations, 2-14
- Oracle Platform Security Services
 - database artifacts, 2-4
 - Disaster Recovery recommendations, 2-4
 - recovery recommendations, 2-5
 - synchronization recommendations, 2-4
- Oracle Portal
 - database artifacts, 2-30
 - Disaster Recovery recommendations, 2-30
 - recovery recommendations, 2-31
 - special considerations, 2-31
 - synchronization recommendations, 2-31
- Oracle Portal, Forms, Reports, and Discoverer

- consistency group recommendations, 4-37
- database artifacts, 2-29
- directory structure recommendations, 4-31
- Disaster Recovery recommendations, 2-29
- file system artifacts, 2-29
- network artifacts, 2-30
- obtaining enterprise deployment guides for, 4-31, 4-32, 4-64, 4-65, 5-2
- recovery recommendations, 2-30
- synchronization recommendations, 2-30
- virtual servers, 3-29
- volume design recommendations, 4-36
- Oracle Reports
 - database artifacts, 2-32
 - Disaster Recovery recommendations, 2-32
 - file system artifacts, 2-32
 - recovery recommendations, 2-32
 - special considerations, 2-32
 - synchronization recommendations, 2-32
- Oracle SOA Service Infrastructure
 - database artifacts, 2-13
 - Disaster Recovery recommendations, 2-12
 - recovery recommendations, 2-13
 - synchronization recommendations, 2-13
- Oracle SOA Suite
 - consistency group recommendations, 4-11, 4-42, 4-48
 - database artifacts, 2-12
 - directory structure recommendations, 4-2, 4-38, 4-43
 - Disaster Recovery recommendations, 2-10
 - file system artifacts, 2-11
 - network artifacts, 2-11
 - recovery recommendations, 2-12
 - synchronization recommendations, 2-12
 - virtual servers, 3-27, 3-28, 3-29, 3-30, 3-31
 - volume design recommendations, 4-5, 4-39, 4-44
- Oracle User Messaging Service
 - database artifacts, 2-18
 - Disaster Recovery recommendations, 2-17
 - file system artifacts, 2-17
 - recovery recommendations, 2-18
 - special considerations, 2-18
 - synchronization recommendations, 2-18
- Oracle Virtual Directory
 - database artifacts, 2-23, 2-38
 - Disaster Recovery recommendations, 2-23, 2-38
 - recovery recommendations, 2-23, 2-38
 - special considerations, 2-23, 2-38
 - synchronization recommendations, 2-23, 2-38
- Oracle Web Cache
 - database artifacts, 2-35
 - Disaster Recovery recommendations, 2-35
 - file system artifacts, 2-35
 - recovery recommendations, 2-35
 - special considerations, 2-35
 - synchronization recommendations, 2-35
- Oracle Web Services Manager
 - database artifacts, 2-17
 - Disaster Recovery recommendations, 2-16
 - recovery recommendations, 2-17
 - synchronization recommendations, 2-17
- Oracle WebCenter
 - common artifacts and considerations for all components, 2-6
 - consistency group recommendations, 4-17
 - database artifacts, 2-6
 - directory structure recommendations, 4-12
 - Disaster Recovery recommendations, 2-6
 - file system artifacts, 2-6
 - network artifacts, 2-6
 - recovery recommendations, 2-7
 - synchronization recommendations, 2-6
 - virtual servers, 3-27, 3-30
 - volume design recommendations, 4-12
- Oracle WebCenter Analytics
 - backup and recovery recommendations, 2-9
- Oracle WebCenter Discussions Server
 - database artifacts, 2-8
 - Disaster Recovery recommendations, 2-8
 - recovery recommendations, 2-9
 - synchronization recommendations, 2-8
- Oracle WebCenter Portlets
 - database artifacts, 2-8
 - Disaster Recovery recommendations, 2-8
 - recovery recommendations, 2-8
 - synchronization recommendations, 2-8
- Oracle WebCenter Spaces
 - database artifacts, 2-7
 - Disaster Recovery recommendations, 2-7
 - recovery recommendations, 2-8
 - synchronization recommendations, 2-7
- Oracle WebCenter Wiki and Blog Server
 - backup and recovery recommendations, 2-10
 - database artifacts, 2-9
 - Disaster Recovery recommendations, 2-9
 - recovery recommendations, 2-9
 - synchronization recommendations, 2-9
- Oracle WebLogic Server
 - file system artifacts, 2-2
 - network artifacts, 2-2
- Oracle WebLogic Server JMS and T-Logs
 - avoiding duplicate messages, 2-3
 - database artifacts, 2-2
 - Disaster Recovery recommendations, 2-2
 - file system artifacts, 2-2
 - recovery recommendations, 2-3
 - special considerations, 2-2
 - synchronization recommendations, 2-3

P

- patch set
 - how to apply to an Oracle Fusion Middleware home in a Disaster Recovery topology, 4-82
- physical host name
 - definition, 1-2
- ping command
 - using to test host name resolution, 3-18
- production site

- creating, 4-61
- creating for the Oracle Identity Management topology, 4-62
- creating for the Oracle Portal, Forms, Reports, and Discoverer topology, 4-64
- creating for the Oracle SOA Suite topology, 4-61
- creating for the Oracle WebCenter topology, 4-62, 4-63, 4-64
- production site setup
 - definition, 1-3

R

- recovery recommendations
 - Oracle Access Manager, 2-26, 2-27, 2-28, 2-29
 - Oracle ADF, 2-6
 - Oracle B2B, 2-16
 - Oracle BPEL Process Manager, 2-14
 - Oracle Business Activity Monitoring, 2-20, 2-21
 - Oracle Business Intelligence Discoverer, 2-33
 - Oracle Directory Integration Platform, 2-24, 2-39
 - Oracle Directory Services Manager, 2-25, 2-41, 2-42, 2-43, 2-44
 - Oracle Forms, 2-31
 - Oracle HTTP Server, 2-35
 - Oracle Human Workflow, 2-15
 - Oracle Identity Federation, 2-24, 2-40
 - Oracle Identity Management, 2-22, 2-37, 2-42
 - Oracle Internet Directory, 2-23, 2-38
 - Oracle JCA Adapters, 2-19
 - Oracle Mediator, 2-14
 - Oracle Platform Security Services, 2-5
 - Oracle Portal, 2-31
 - Oracle Portal, Forms, Reports, and Discoverer, 2-30
 - Oracle Reports, 2-32
 - Oracle SOA Service Infrastructure, 2-13
 - Oracle SOA Suite, 2-12
 - Oracle User Messaging Service, 2-18
 - Oracle Virtual Directory, 2-23, 2-38
 - Oracle Web Cache, 2-35
 - Oracle Web Services Manager, 2-17
 - Oracle WebCenter, 2-7
 - Oracle WebCenter Discussions Server, 2-9
 - Oracle WebCenter Portlets, 2-8
 - Oracle WebCenter Spaces, 2-8
 - Oracle WebCenter Wiki and Blog Server, 2-9
 - Oracle WebLogic Server JMS and T-Logs, 2-3

S

- self-signed certificates
 - generating, 4-58
- separate DNS servers
 - using to resolve host names, 3-16
- shared storage
 - creating mount points to, 1-6
 - creating Oracle home directories on, 1-6
- site failover
 - definition, 1-3

- steps for performing, 4-76
- site switchback
 - definition, 1-3
 - steps for performing, 4-75
- site switchover
 - definition, 1-4
 - steps for performing, 4-75
- site synchronization
 - definition, 1-4
- SOA Suite
 - common artifacts and considerations for all components, 2-11
 - special considerations
 - Oracle Access Manager, 2-26, 2-27
 - Oracle B2B, 2-16
 - Oracle Business Intelligence Discoverer, 2-33
 - Oracle Directory Services Manager, 2-25, 2-40, 2-42, 2-43
 - Oracle Forms, 2-31
 - Oracle HTTP Server, 2-34
 - Oracle Identity Federation, 2-24, 2-40
 - Oracle Internet Directory, 2-22, 2-37
 - Oracle Portal, 2-31
 - Oracle Reports, 2-32
 - Oracle User Messaging Service, 2-18
 - Oracle Virtual Directory, 2-23, 2-38
 - Oracle Web Cache, 2-35
 - Oracle WebLogic Server JMS and T-Logs, 2-2
- standby site
 - creating, 4-66
 - database setup, 4-66
 - middle tier setup, 4-67
 - performing periodic testing, 4-77
 - prerequisites, 4-66
 - setting up Node Manager, 4-58
 - validating setup of, 4-67
- standby site setup
 - definition, 1-4
- static artifacts
 - for Oracle Fusion Middleware, 3-32
- storage considerations, 3-32
- storage replication, 3-33, 4-49
- storage replication technology
 - using to protect Oracle Fusion Middleware middle tier components, 1-2
- switchback
 - definition, 1-3
 - steps for performing, 4-75
- switchover
 - definition, 1-4
 - steps for performing, 4-75
- symbolic links
 - to Oracle home directories on shared storage, 1-7
- symmetric topology
 - definition, 1-4
 - design considerations, 3-39
 - design considerations for, 3-4
 - directory names and paths requirement, 3-4
 - installed software requirement, 3-4
 - load balancers and virtual server names

- requirement, 3-4
 - port numbers requirement, 3-4
 - security requirement, 3-4
- synchronization
 - manually forcing after middle tier configuration changes, 4-75
 - manually forcing for databases after middle tier configuration changes, 3-35
 - of production site and standby site, 4-74
- synchronization considerations
 - Oracle Human Workflow, 2-15
- synchronization recommendations
 - Oracle Access Manager, 2-26, 2-28, 2-29
 - Oracle ADF, 2-5
 - Oracle B2B, 2-16
 - Oracle BPEL Process Manager, 2-13
 - Oracle Business Activity Monitoring, 2-20
 - Oracle Business Intelligence Discoverer, 2-33
 - Oracle Directory Integration Platform, 2-24, 2-39
 - Oracle Directory Services Manager, 2-25, 2-40, 2-42, 2-43, 2-44
 - Oracle Forms, 2-31
 - Oracle HTTP Server, 2-35
 - Oracle Identity Federation, 2-24, 2-40
 - Oracle Identity Management, 2-22, 2-36, 2-41
 - Oracle Internet Directory, 2-23, 2-37
 - Oracle JCA Adapters, 2-19
 - Oracle Mediator, 2-14
 - Oracle Platform Security Services, 2-4
 - Oracle Portal, 2-31
 - Oracle Portal, Forms, Reports, and Discoverer, 2-30
 - Oracle Reports, 2-32
 - Oracle SOA Service Infrastructure, 2-13
 - Oracle SOA Suite, 2-12
 - Oracle User Messaging Service, 2-18
 - Oracle Virtual Directory, 2-23, 2-38
 - Oracle Web Cache, 2-35
 - Oracle Web Services Manager, 2-17
 - Oracle WebCenter, 2-6
 - Oracle WebCenter Discussions Server, 2-8
 - Oracle WebCenter Portlets, 2-8
 - Oracle WebCenter Spaces, 2-7
 - Oracle WebCenter Wiki and Blog Server, 2-9
 - Oracle WebLogic Server JMS and T-Logs, 2-3

T

- testing
 - the standby site, 4-77
- third party database
 - Disaster Recovery for, 1-8
- TNSNAMES.ORA entries
 - making for databases, 3-35
- topology
 - definition, 1-4
- trust keystore
 - creating, 4-59
- TTL (Time to Live) value, 3-32

V

- validating
 - asymmetric standby site setup, 4-74
 - production site setup
 - production site
 - validating setup of, 4-65
 - standby site setup, 4-67
- virtual IP considerations, 3-26
- virtual servers
 - for Oracle Identity Management, 3-28
 - for Oracle Portal, Forms, Reports, and Discoverer, 3-29
 - for Oracle SOA Suite, 3-27, 3-28, 3-29, 3-30, 3-31
 - for Oracle WebCenter, 3-27, 3-30
- volume design
 - for Disaster Recovery sites, 4-2
- volume design recommendations
 - for Oracle Identity Management, 4-19
 - for Oracle Portal, Forms, Reports, and Discoverer, 4-36
 - for Oracle SOA Suite, 4-5, 4-39, 4-44
 - for Oracle WebCenter, 4-12

W

- wide area DNS operations, 3-31

