Oracle® Fusion Middleware

Disaster Recovery Guide

12*c* (12.1.3)

E52348-02

July 2014

This document describes the Disaster Recovery solutions for Oracle Fusion Middleware products.

ORACLE®

Oracle Fusion Middleware Disaster Recovery Guide, 12*c* (12.1.3)

E52348-02

Primary Author: Rekha Kamath

Contributor: Susan Kornberg, Fermin Castro Alonso, Satheesh Amilineni, Ashwani Raj, Mukesh Tailor, Allwarappan Sundararaj

Contents

4 Setting Up and Managing Disaster Recovery Sites

Preface

This preface contains these sections:

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/pls/topic/lookup?ctx=acc&id=info or visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Fusion 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 Administrator's Guide*

Conventions

The following text conventions are used in this document:

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

1

Introduction to Oracle Fusion Middleware Disaster Recovery

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

It contains the following topics:

1.1 Overview of Disaster Recovery

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

It contains the following topics:

1.1.1 Problem Description and Common Solutions

Providing Oracle 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 downtime 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 periodically to the standby site. 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 usually 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. 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 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 through 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**

  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 that acts as a proxy between servers and their clients.

- **Production Site Setup**

  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 binary files 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 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 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 guide 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.

- **Target**

  Targets are core Enterprise Manager entities which represent the infrastructure and business components in an enterprise. These components need to be

monitored and managed for efficient functioning of the business. For example, Oracle Fusion Middleware farm or Oracle Database.

- **System**

  A System is the set of targets (hosts, databases, application servers, and so on) that work together to host your applications. To monitor an application in Enterprise Manager, you would first create a System, that consists of the database, listener, application server, and hosts targets on which the application runs.

- **Site**

  Site is a set of different targets in a datacenter needed to run a group of applications. For example, a site could consist of Oracle Fusion Middleware instances, databases, storage, and so on. A datacenter may have more than one site defined by Oracle Site Guard and each of them managed independently for operations like switchover and failover.

1.2 Setting Up 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:

1.2.1 Oracle Fusion Middleware Disaster Recovery Architecture Overview

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

The product binary files and configuration for Oracle Fusion Middleware components and applications get deployed in Oracle home directories on the middle tier. Additionally, most of the products also have metadata or runtime 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 the following methods for providing data protection for Oracle Fusion Middleware data and database content:

- Oracle Fusion Middleware product binary files, configuration files, 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 Sites for Oracle Fusion Middleware Disaster Recovery Topology

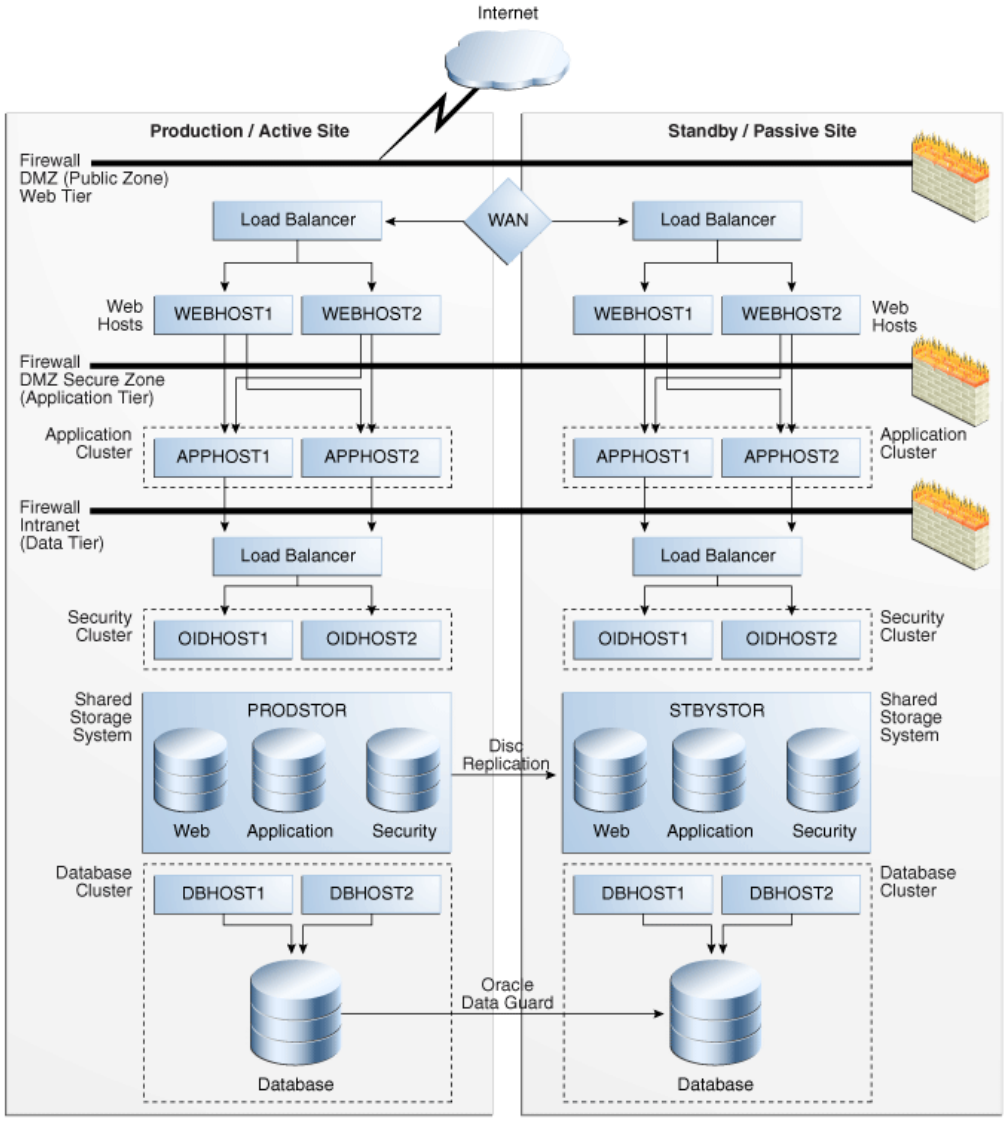

Some of the key aspects of the solution in Figure 1–1 are:

- 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 binary files 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 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 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 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 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. See Chapter 2 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.
- 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, 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 reroute 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 for Disaster Recovery recommendations for Oracle WebLogic Server 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 2.2 for Disaster Recovery recommendations for Oracle SOA Suite components.

2

Recommendations for Oracle 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 these 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 includes the following sections:

> **Note:** Certain artifacts like 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 part of the regular storage replication and synchronization activity. Oracle recommends placing 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. For information about how to manage Oracle Inventory, see Appendix A.

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

Oracle home: The Oracle home consists of a WebLogic home that has the WebLogic Server binary files.

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 sites, there is no need to update this value after a Disaster Recovery operation.

If your environment requires whole server migration to be configured, then 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, ensure 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 Java Message Service (JMS) and Transaction Logs (T-Logs)
- Recommendations for Oracle Platform Security Services

2.1.1 Recommendations for Oracle WebLogic Server Java Message Service (JMS) and Transaction Logs (T-Logs)

This section describes various Oracle WebLogic Server JMS and 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 and 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 enqueued 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 the reverse. This may delete unprocessed messages in addition to duplicate messages.

- If the store is not dedicated to JMS use, then 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 WebLogic Scripting Tools (WLST).
- When applications use both queues and topics, ensure that you manipulate both the queue and topic subscriptions.

Synchronization Recommendations

- If JMS data is critical, then synchronize transaction log data and JMS data in real time using synchronous replication. Using synchronous replication may have performance implications.
- If data consistency between tiers is important, then 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 create a consistent backup. This ensures that the persistence store is not being written to while the copy operation is being performed. In a clustered environment, shut down one server at a time, back it up, and restart it. You also can create a script to perform these operations using WLST.

Recovery Recommendations

Recover the database schemas containing the persistent stores for the Administration Server and the Managed Servers in the WLS domain to the most recent point in time.

Also, use the following recovery recommendations to avoid duplicate messages.

Avoiding Duplicate Messages

Perform the following procedure before recovery to filter messages in the JMS queue after persistent-store recovery, to avoid processing duplicate messages:

> **Note:** Do not drain and discard messages without 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 in to the Oracle WebLogic Server Administration Console.
2. Before recovery, configure JMS server to pause Production, Insertion, and Consumption operations at Startup. This ensures 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 that 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, and then select a destination.
   c. Select **Monitoring**, and then click **Show Messages**.
   d. Click **Delete All**.

Resume operations by following the instructions listed in step 2.

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

You must manually synchronize the application tier with the standby site after making configuration changes and applying patches.

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

Oracle recommends that you synchronize the standby database 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 and "Applying Redo Data to Physical Standby Databases" in *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 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 that combine multiple technology components.

SOA composite applications consist of:

- Service components: Service components are the basic building blocks of SOA composite applications. Service components implement 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 12*c* (12.1.3), 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 11gR1 (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 11g R1 (11.1.1.3), you can deploy Oracle Business Process Management (BPM) Suite on top of an Oracle SOA Suite installation.

The Disaster Recovery recommendations for Oracle SOA Suite assume that you are using Oracle SOA Suite 12c R1 (12.1.3).

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 binary files and domain-related configuration files are stored on a local or shared file system.

Common Artifacts and Considerations for All Oracle SOA Suite Components

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

Artifacts on the File System

Oracle home: The Oracle home consists of a WebLogic home that has the WebLogic Server binary files and an Oracle home containing the Oracle SOA Suite binary files.

Oracle Common Home: This is 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 sites, there is no need to update this value after a Disaster Recovery operation. See the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for instructions for updating an IP address to a virtual host name.

The load balancer virtual hosts required for accessing the Oracle 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 changes in the domain-related configuration, deploying composites, and applying patches.

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

Oracle recommends that you synchronize the standby database 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 and "Applying Redo Data to Physical Standby Databases" in *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 matching the composite definition to continue processing. For this reason, the metadata repository (where composite definitions are stored) and Oracle SOA Suite database (where the process state is maintained) must be recovered to the same point in time.

In redeployed composites, a database recovery ensures consistency between the dehydrated in-flight processes and their corresponding definition because 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 Java EE Connector Architecture (JCA) Adapters
- Recommendations for Oracle Business Activity Monitoring

- Recommendations for Oracle Business Process Management

2.2.1 Recommendations for Oracle SOA Service Infrastructure

Oracle SOA Service Infrastructure is a Java EE application that provides the foundation services for running Oracle SOA Suite. This Java EE application is a runtime engine that is automatically deployed when Oracle SOA Suite is installed. You deploy composites (the basic artifacts in a 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 lifecycle and run-time operations of the composite.

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 Oracle SOA Service Infrastructure database.

Synchronization Recommendations

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

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

Oracle recommends that you synchronize the standby database 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.

2.2.2 Recommendations for Oracle BPEL Process Manager

The Oracle BPEL Process engine is the service engine running in Oracle 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.

Oracle recommends that you synchronize the standby database 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, because no cleanup is required after performing a Disaster Recovery operation. If non-idempotent 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.

2.2.3 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. 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.

> **Note:** Sequential routing rules do not persist their messages into the database as part of the execution.

Synchronization Recommendations

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

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

Oracle recommends that you synchronize the standby database 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.

2.2.4 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 changes in the domain-related configuration and applying patches.

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

Oracle recommends that you synchronize the standby database 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. The Oracle Human Workflow engine uses Oracle User Messaging Service to send and receive notifications. See Section 2.2.7 for details about Oracle User Messaging Service.

2.2.5 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 business-to-business (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 e-mail 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, see Section 2.1.1.

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

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

Oracle recommends that you synchronize the standby database 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 runtime database, so recovering the database and the Managed Server will ensure that the application runs normally.

2.2.6 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 services policies including security, Web Services Reliable Messaging (WSRM), Message Transmission Optimization

Mechanism (MTOM) and addressing policies. Oracle Web Services Manager consists 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. The 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, gathering of runtime statistics. The agent is available on all Oracle Fusion Middleware Managed Servers and is configured on the same server as the application 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 changes in the domain-related configuration and applying patches.

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

Oracle recommends that you synchronize the standby database 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 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.

2.2.7 Recommendations for Oracle User Messaging Service

Oracle User Messaging Service (UMS) enables two-way communication between users and deployed applications. It has support for a variety of channels, such as e-mail, IM, SMS, and text-to-voice messages. Oracle User Messaging Service is integrated with Oracle Fusion Middleware components such as Oracle BPEL Process Manager, Oracle Human Workflow, Oracle Business Activity Monitoring (BAM) and Oracle WebCenter Portal. It is typically deployed with Oracle User, 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
OraSDPMAppDefRcvQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMAppDefRcvQ1
OraSDPMDriverDefSndQ1_auto	Distributed queue	OraSDPM/Queues/OraSDPMDriverDefSndQ1
OraSDPMEngineCmdQ_auto	Distributed queue	OraSDPM/Queues/OraSDPMEngineCmdQ
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 the 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 changes in the configuration, deploying additional Oracle User Messaging Service drivers, and applying patches.

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

Oracle recommends that you synchronize the standby database 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 the 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.

2.2.8 Recommendations for Oracle Java EE Connector Architecture (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 the *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 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 location of each `weblogic-ra.xml` file 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 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 changes in the domain-related configuration (that is, adapter configuration changes) and applying patches.

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

Oracle recommends that you synchronize the standby database 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 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.

2.2.9 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 correlation 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 runtime 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 Oracle BAM schemas and the metadata repository.

Oracle recommends that you synchronize the standby database 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.

2.2.10 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
- Oracle 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 Oracle Metadata Services (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 changes in the domain-related configuration 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 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.

3

Design Considerations

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

It includes the following topics:

- Section 3.1, "Network Considerations"
- Section 3.2, "Storage Considerations"
- Section 3.3, "Database Considerations"
- Section 3.4, "Starting Points"
- Section 3.5, "Topology Considerations"

Note:

- You can automate disaster recovery operations like switchover and failover, using Oracle Site Guard. For information about the product, see *Oracle Site Guard Administrator's Guide*.
- For information about installing and configuring Oracle SOA Suite components in an enterprise deployment, see *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.
- See *Oracle Fusion Middleware Release Notes for Oracle Fusion Middleware Infrastructure* for updates about errors.

This chapter provides detailed instructions for setting up an Oracle Fusion Middleware 12*c* Disaster Recovery production site and standby site for the Linux and UNIX operating systems. It primarily uses the Oracle SOA Suite enterprise deployment (see Figure 3–1) in the examples to show how to set up the Oracle Fusion Middleware 12*c* 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 12*c* enterprise deployments in this chapter to set up Disaster Recovery for those deployments as well.

Note: This chapter describes an Oracle Fusion Middleware 12*c* 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

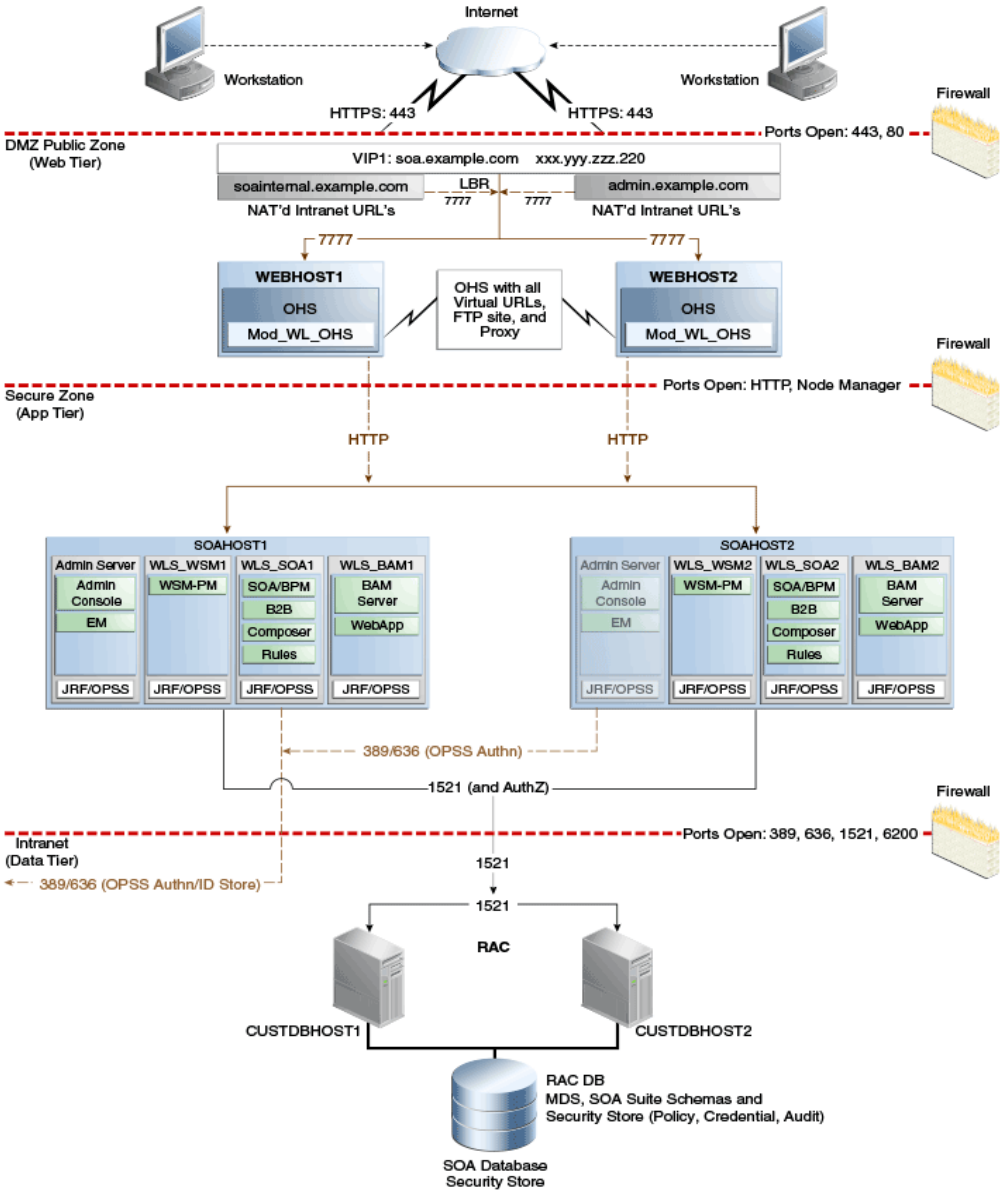

Figure 3–1 shows a diagram of the Oracle SOA, Business Process Management (BPM), and Oracle Service Bus enterprise deployment topology. See the *Oracle Fusion*

Middleware Enterprise Deployment Guide for Oracle SOA Suite for detailed information about installing and configuring an 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 the 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 must be the same at the production site hosts and their standby site peer hosts.

  Section 3.4.1 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, then 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 Host Name Considerations
- Load Balancer Considerations
- Virtual Server Considerations
- External Clients Considerations
- 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 becomes active. You do not need to reconfigure the host name for the host on the standby site if you set up an 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 host-name-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 `172.16.10.255`), then use the host name `SOAHOST1.EXAMPLE.COM`, which resolves to `172.16.10.2555`.

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

Note: In the examples listed in this book, IP addresses for hosts at the initial production site have the format 172.16.x.x and IP addresses for hosts at the initial standby site have the format 172.16.x.x.

3.1.1.1 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 Enterprise Deployment Guide (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
`172.16.2.111`	`WEBHOST1`	None
`172.16.2.112`	`WEBHOST2`	None
`172.16.2.113`	`SOAHOST1`	None
`172.16.2.114`	`SOAHOST2`	None

[1] See Section 3.1.1.1.3 and Section 3.1.1.1.4 for information about defining physical host names.

Figure 3–2 shows the physical host names used for the Oracle SOA Suite 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. See Section 3.1.1.1.3 for more information about using separate DNS servers to resolve host names.

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

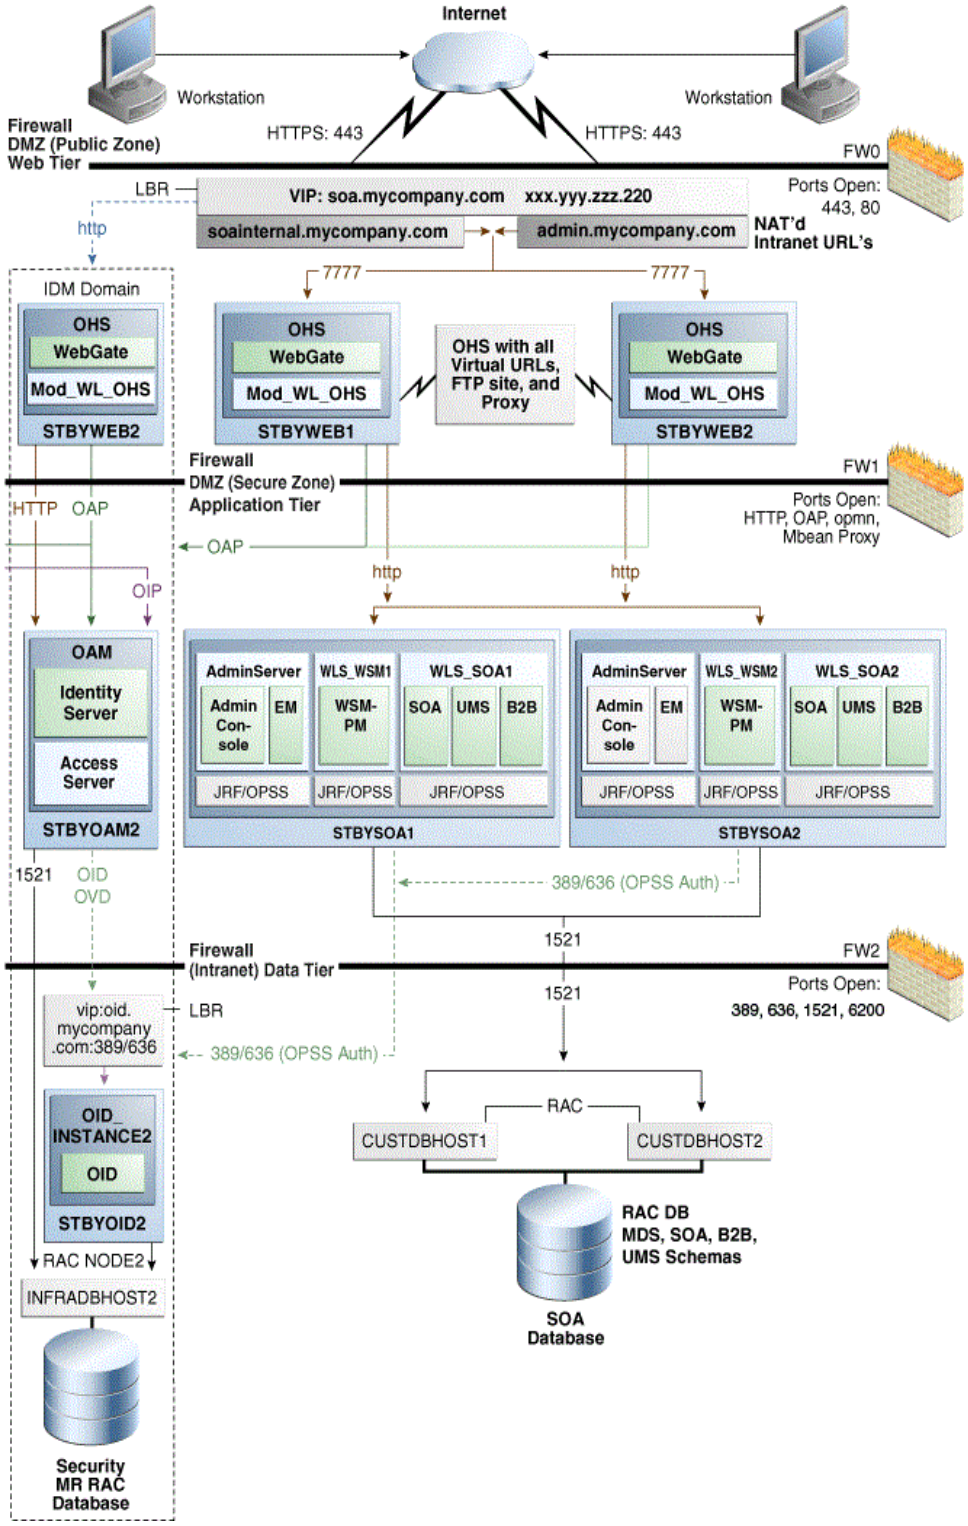

The Administration Server, the Oracle Identity Manager Managed Servers, and the SOA Managed Servers require a floating IP address to be provisioned on each site (Table 3–2). Ensure that you provision the floating IP addresses with the same virtual host names on the production site and the standby site.

Table 3–2 Floating IP Addresses

Physical Host Name	Virtual Host Name	Floating IP
AdminServer	ADMINVHN	172.16.2.134
OIMHOST1	OIMVHN1	172.16.2.135
OIMHOST2	OIMVHN2	172.16.2.136
SOAHOST1	SOAVHN1	172.16.2.137
SOAHOST2	SOAVHN2	172.16.2.138

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

3.1.1.1.1 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 3.1.1.1.2 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 3.1.1.1.3 and Section 3.1.1.1.4 for more information about two methods of implementing DNS server host name resolution.

You must determine the method of host name resolution that 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 `/etc/hosts` file are 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 are 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, Oracle recommends that all the hosts within a site (production site or standby site) should 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.

3.1.1.1.2 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 simplicity and ease of maintenance, Oracle recommends that you provide 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

```
174.0.0.1       localhost.localdomain    localhost
172.16.2.111     WEBHOST1.EXAMPLE.COM    WEBHOST1
172.16.2.112     WEBHOST2.EXAMPLE.COM    WEBHOST2
172.16.2.113     SOAHOST1.EXAMPLE.COM    SOAHOST1
172.16.2.114     SOAHOST2.EXAMPLE.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 simplicity and ease of maintenance, Oracle recommends that you have the same entries on all the hosts of the standby site. Example 3–4 shows the `/etc/hosts` file for the standby site of a SOA enterprise deployment topology.

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

```
176.0.0.1       localhost.localdomain    localhost
172.26.2.111     STBYWEB1.EXAMPLE.COM    WEBHOST1
172.26.2.112     STBYWEB2.EXAMPLE.COM    WEBHOST2
172.26.2.113     STBYSOA1.EXAMPLE.COM    SOAHOST1
172.26.2.114     STBYSOA2.EXAMPLE.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 returns the correct IP address (`172.16.2.111`) and indicates 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 returns the correct IP address (`172.26.2.111`) and it shows that the name WEBHOST1 is associated with that IP address.

3.1.1.1.3 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.EXAMPLE.COM    IN    A    172.16.2.111
WEBHOST2.EXAMPLE.COM    IN    A    172.16.2.112
SOAHOST1.EXAMPLE.COM    IN    A    172.16.2.113
SOAHOST2.EXAMPLE.COM    IN    A    172.16.2.114
```

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.EXAMPLE.COM    IN    A    172.26.2.111
WEBHOST2.EXAMPLE.COM    IN    A    172.26.2.112
SOAHOST1.EXAMPLE.COM    IN    A    172.26.2.113
SOAHOST2.EXAMPLE.COM    IN    A    172.26.2.114
```

5. Ensure that 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 returns the correct IP address (`172.16.2.111`) and indicates 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 returns the correct IP address (`172.26.2.111`) and indicates that the host name is fully qualified.

3.1.1.1.4 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, use a combination of local host name resolution and DNS host name resolution.
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.EXAMPLE.COM    IN    A    172.16.2.111
WEBHOST2.EXAMPLE.COM    IN    A    172.16.2.112
SOAHOST1.EXAMPLE.COM    IN    A    172.16.2.113
SOAHOST2.EXAMPLE.COM    IN    A    172.16.2.114
STBYWEB1.EXAMPLE.COM    IN    A    172.26.2.111
STBYWEB2.EXAMPLE.COM    IN    A    172.26.2.112
STBYSOA1.EXAMPLE.COM    IN    A    172.26.2.113
STBYSOA2.EXAMPLE.COM    IN    A    172.26.2.114
```

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 simplicity and ease of maintenance, Oracle recommends having 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

```
176.0.0.1       localhost.localdomain    localhost
172.26.2.111     STBYWEB1.EXAMPLE.COM    WEBHOST1
172.26.2.112     STBYWEB2.EXAMPLE.COM    WEBHOST2
172.26.2.113     STBYSOA1.EXAMPLE.COM    SOAHOST1
172.26.2.114     STBYSOA2.EXAMPLE.COM    SOAHOST2
```

7. Test the host name resolution using the `ping` command. A `ping webhost1` command on the production site returns the correct IP address (`172.16.2.111`) and indicates that the host name is fully qualified.
8. Similarly, a `ping webhost1` command on the standby site returns the correct IP address (`172.26.2.111`) and indicates that the host name is fully qualified.

3.1.1.1.5 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 Host Name Considerations

Virtual IP addresses and host names 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 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 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 host name for the host on the standby site if you set up 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

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

Virtual IP Address	Virtual Host Name[1]	Alias Host Name
`172.16.2.115`	`ADMINVHN`	None
`172.16.2.116`	`SOAVHN1`	None
`172.16.2.117`	`SOAVHN2`	None

[1] See Section 3.1.1.1.3 and Section 3.1.1.1.4 for information about defining physical host names.

Table 3–4 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 3–4 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.

See Section 3.1.1.1.3 for more information about using separate DNS servers to resolve host names.

Table 3–4 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
172.26.2.115	STBYADMINVHN	ADMINVHN
172.26.2.116	STBYSOAVHN1	SOAVHN1
172.26.2.117	STBYSOAVHN2	SOAVHN2

[1] See Section 3.1.1.1.3 and Section 3.1.1.1.4 for information about defining physical host names.

3.1.3 Load Balancer Considerations

Oracle Fusion Middleware components require a hardware load balancer when deployed in high availability topologies. Oracle recommends 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. 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.

- 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 URLs.
- 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 Considerations

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 ensures that incoming traffic on a given virtual host is not directed to an unavailable service in the other tiers.

Oracle recommends that you 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. Although this is supported, the deployment should consider the security implications of doing this and if 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. Oracle highly recommends that you create aliases for these virtual servers when you use 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 Table 3–5 and Table 3–6.

Table 3–5 Virtual Servers for Oracle SOA Suite Production Site

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

Table 3–6 Virtual Servers for Oracle SOA Suite Standby Site

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

3.1.5 External Clients Considerations

Systems accessing the servers used in the topology directly (that is, not using the front-end load balancers, as in the case of JMX or RMI clients) need to be aware of the listen address used by the different Oracle WebLogic Server instances. An appropriate hostname resolution needs to be provided to the clients so that the hostname alias used by the servers as listen address is correctly resolved. This is also applicable to the Oracle JDeveloper deployments. The client hosting Oracle Jdeveloper needs to map the `SOAHOSTx` and `SOAVHNx` aliases to correct the IP addresses for deployments to succeed.

3.1.6 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 serve as a front end for 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

3.1.6.1 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.

3.1.6.2 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. Note 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, 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 that the components in the topology are in sync. This is an important consideration for designing volumes and consistency groups. Some of the artifacts are static whereas others are dynamic.

Static Artifacts

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

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

Dynamic or Runtime Artifacts

Dynamic or runtime artifacts are files that change frequently. Runtime 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
- Deployment plans: Used for updating technology adapters such as file and JMS adapters. They need to be saved in a location that is accessible to all nodes in the cluster that the artifacts are being deployed to.

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 binary files 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 WebLogic home is shared by multiple servers in different nodes, Oracle recommends that you keep the Oracle Inventory and Oracle home list in those nodes updated for consistency in the installations and application of patches.

To update the inventory files in a node and attach an installation in a shared storage to it, use *ORACLE_HOME*/oui/bin/attachHome.sh.

To update the Oracle home list to add or remove a WebLogic home, edit the *user_home*/bea/beahomelist file. This is required for any nodes installed in addition 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, then:

- 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, then:

- 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. This enables the same directory structure to 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 Service (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 Storage Area Network (SAN) storage system should use either a host-based clustered or a shared file system technology such as the Oracle Clustered File System (OCFS2). OCFS2 is a symmetric shared disk cluster file system that 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.

- Oracle recommends creating Oracle Real Application Cluster (Oracle RAC) databases on both the production site and standby site as required by your topology.
- 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 use Oracle GoldenGate only in an active-passive configuration.

For more information, see 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

- 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.

  Fore more information, see *Oracle Data Guard Concepts and Administration* and related Oracle Maximum Availability Architecture collateral at the following URL:

  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

- Ensure that your network is configured for low latency with sufficient bandwidth, because synchronous redo transmission can affect the response time and throughput.
- The `LOG_ARCHIVE_DEST_n` parameter on standby site databases should have the `SYNC` or `ASYNC` attributes. `ASYNC` is the default attribute if no attributes are specified.
- The standby site database should be in Managed Recovery mode. This ensures that the standby site databases are in a constant state of media recovery. Managed Recovery mode is enabled for shorter failover times.
- The `tnsnames.ora` file on the production site and the standby site must have entries for databases on both the production and standby sites.
- Oracle strongly recommends that you force Oracle Data Guard to perform manual database synchronization whenever middle tier synchronization is performed. This is especially important for components that store configuration data in the metadata repositories.
- Oracle strongly recommends that you set up aliases for the database host names on both the production and standby sites. This enables seamless switchovers, switchbacks, and failovers.
- When one of the databases at either site is an Oracle RAC database, it is required that the single instance database at the peer site must have the same value for `instance_name`.

Note:

- The values for `ORACLE_HOME`, `home`, `ORACLE_INSTANCE`, `DOMAIN_HOME` in the middle tier must be identical.
- The values for `DB_NAME`, `INSTANCE_NAME`, `Listen Port`, and `ORACLE_SID` in the database tier must be identical.
- To avoid manipulation of the WLS data sources, the `SERVICE_NAME` specified in the Application Data Source must be identical. However, each database can have additional services defined.

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 3–7 shows the database host names and the connect strings for the databases before the alias is defined.

Table 3–7 Database Host Names and Connect Strings

Site	Database Host Name	Database Connect String
Production	`custdbhost1.example.com`	`custdbhost1.example.com:1521:orcl`
Standby	`stbycustdbhost1.example.com`	`stbycustdbhost1.example.com:1521:orcl`

In this example, all database connect strings on the production site take the form `custdbhost1.example.com:1521:orcl`. After a failover or switchover operation, this connect string must be changed to `stbycustdbhost1.example.com:1521:orcl`. However, by creating an alias of `proddb1` for the database host name as shown in Table 3–8, you can avoid manually changing the connect strings, which enables seamless failovers and switchovers.

Table 3–8 Specifying an Alias for a Database Host

Site	Database Host Name	Alias	Database Connect String
Production	custdbhost1.example.com	proddb1.example.com	proddb1.example.com:1521:orcl
Standby	stbycustdbhost1.example.com	proddb1.example.com	proddb1.example.com:1521:orcl

In this example, the production site database host name and the standby site database host name are aliased to `proddb1.example.com` and the connect strings on the production site and the standby site can take the form `proddb1.example.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, ensure that the entry for host `custdbhost1` looks like this:

```
152.68.196.213 proddb1.example.com proddb1 custdbhost1.example.com custdbhost1
```

And, in the `/etc/hosts` files at the standby site, ensure that the entry for host `stbycustdbhost1` looks like this:

```
140.87.25.40   proddb1.example.com proddb1 stbycustdbhost1.example.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 must 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 binary files 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 it is not already on shared storage, and then create either a symmetric standby site or an 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 must 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:

- 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*.
- 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*.
- Because this is a one-time operation, you should recover the entire domain.
- The directory structure on the shared storage must be set up as described in Section 4.1.1.
- For Oracle SOA Suite, see "Introducing Backup and Recovery" in *Oracle Fusion Middleware Administrator's Guide*.
- For Web Tier, see "Introducing Backup and Recovery" 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 preexisting 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.

When you create a new Oracle Fusion Middleware Disaster Recovery production site and standby site, you have the following choices:

- 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. For more information about host name planning, see Section 3.1.1.
- When you design and create your own production site, you can choose the Oracle home name and Oracle home directory for each Fusion Middleware installation.

  Designing and creating your own site is easier than modifying an existing site to meet the design requirements described in this chapter.
- You can assign ports for the Oracle 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 checking for and resolving port conflicts between an existing production site and standby site.

3.5 Topology Considerations

This section contains the following topics:

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 Oracle Fusion Middleware instances while the standby site includes two hosts with four Oracle Fusion Middleware instances). Or, in a different asymmetric topology, the standby site can use fewer Oracle Fusion Middleware instances (for example, the production site could include four Oracle Fusion Middleware instances while the standby site includes two

Oracle Fusion Middleware instances). Another asymmetric topology might include a different configuration for a database (for example, using an Oracle Real Application Clusters (Oracle RAC) database at the production site and a single instance database at the standby site).

4

Setting Up and Managing Disaster Recovery Sites

This chapter uses the Oracle SOA Suite enterprise deployment topology to illustrate the steps required to set up the production site and standby site.

This chapter includes the following sections:

- Section 4.1, "Setting Up a Site"
- Section 4.2, "Creating a Production Site"
- Section 4.3, "Creating a Standby Site"
- Section 4.4, "Creating an Asymmetric Standby Site"
- Section 4.5, "Performing Site Operations and Administration"
- Section 4.6, "Using Oracle Site Guard for Disaster Recovery"
- Section 4.7, "Patching an Oracle Fusion Middleware Disaster Recovery Site"

> **Note:** You can automate disaster recovery operations like switchover and failover using Oracle Site Guard. For more information, see *Oracle Site Guard Administrator's Guide*.

4.1 Setting Up a Site

This section provides steps to set up a site.

It includes the following topics:

- Directory Structure and Volume Design
- Storage Replication
- Database

Before you start creating the production site, ensure that you perform the following steps:

- Set up the host name aliases for the middle tier hosts, as described in Section 3.1.1.
- Create the required volumes on the shared storage on the production site as described in Section 4.1.1
- Determine the Oracle Data Guard configuration to use based on the data loss requirements of the database and network considerations such as the available bandwidth and latency when compared to the redo generation.

See *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 below which Oracle products are installed.
- Oracle_Home: This 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 an 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.
- 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 "Directory Structure Recommendations for Oracle SOA Suite".

4.1.1.1 Directory Structure Recommendations for Oracle SOA Suite

Oracle Fusion Middleware 12*c* allows the creation of multiple SOA Managed Servers from a single binary installation. This allows the installation of binary files 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 Oracle 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, this isolating the failures as much as possible 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. In addition, Oracle recommends placing the Managed Servers' domain directories on a shared storage, although having them on the local file system is also supported. This is especially important 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

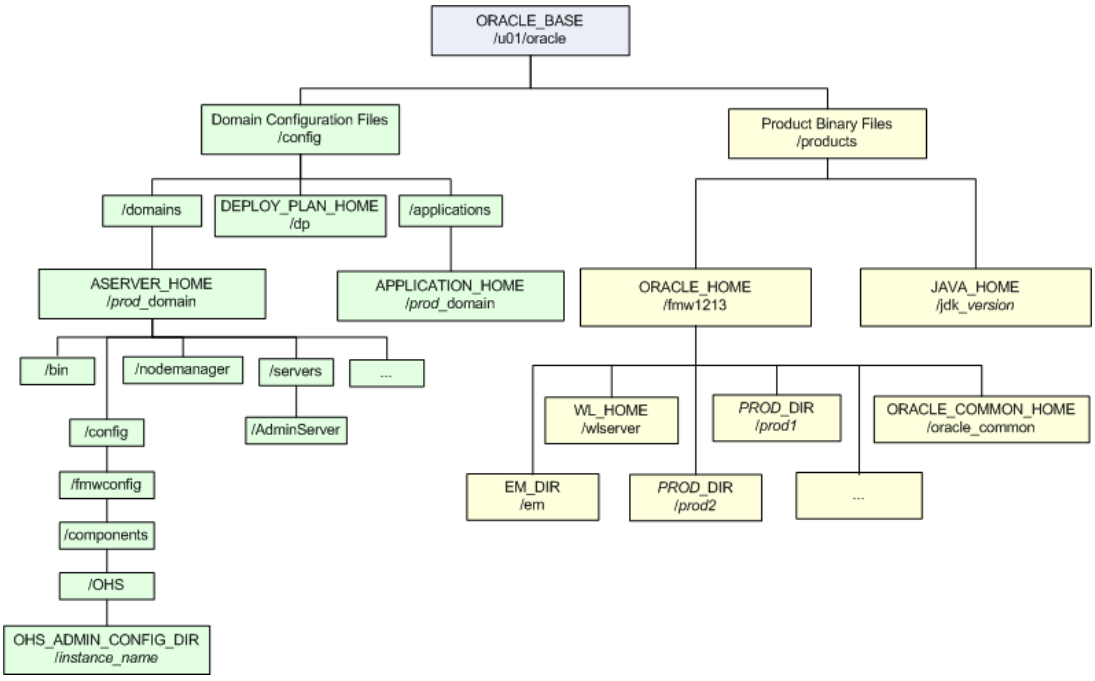

For information about setting up this directory structure, see *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

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 the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

Figure 4–2 Oracle SOA Suite and Oracle Business Activity Monitoring Enterprise Topology Diagram

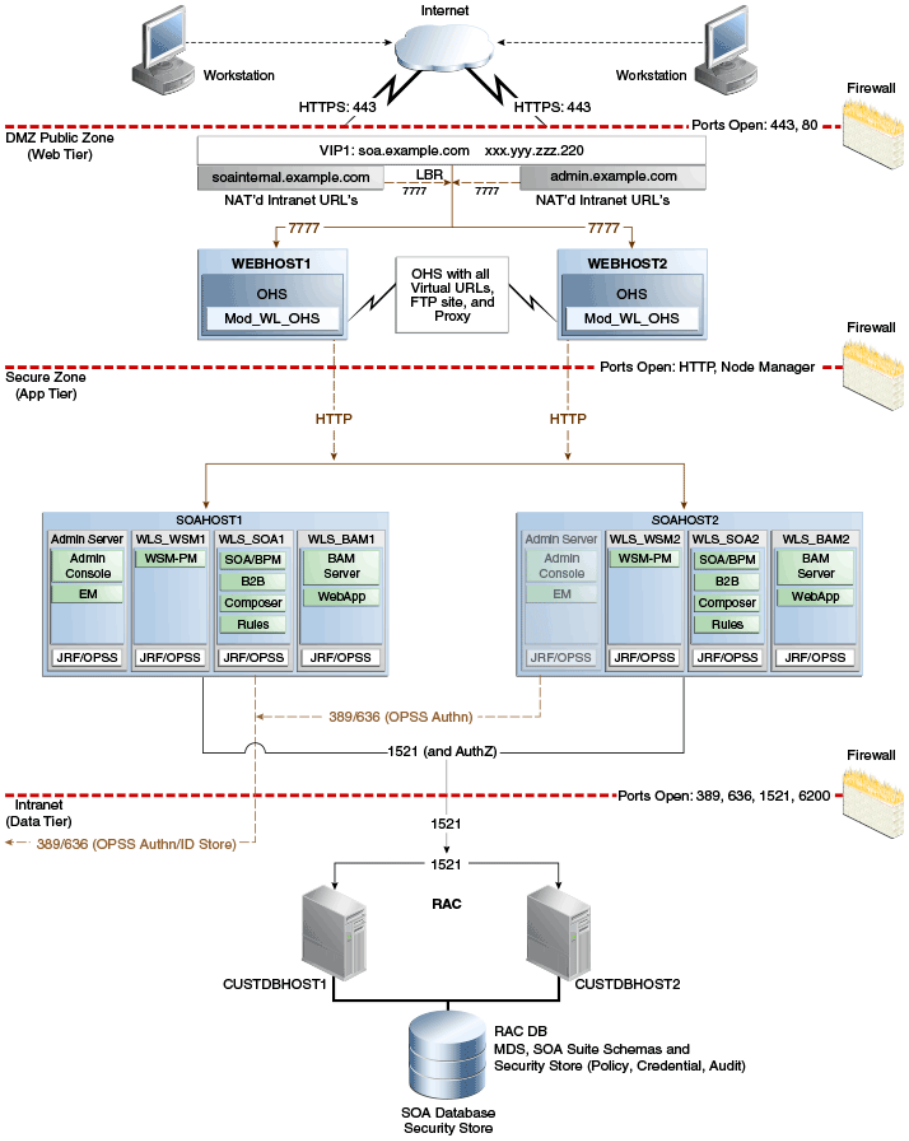

Figure 4–3 Oracle SOA Suite and Oracle Service Bus Enterprise Deployment Reference Topology Diagram

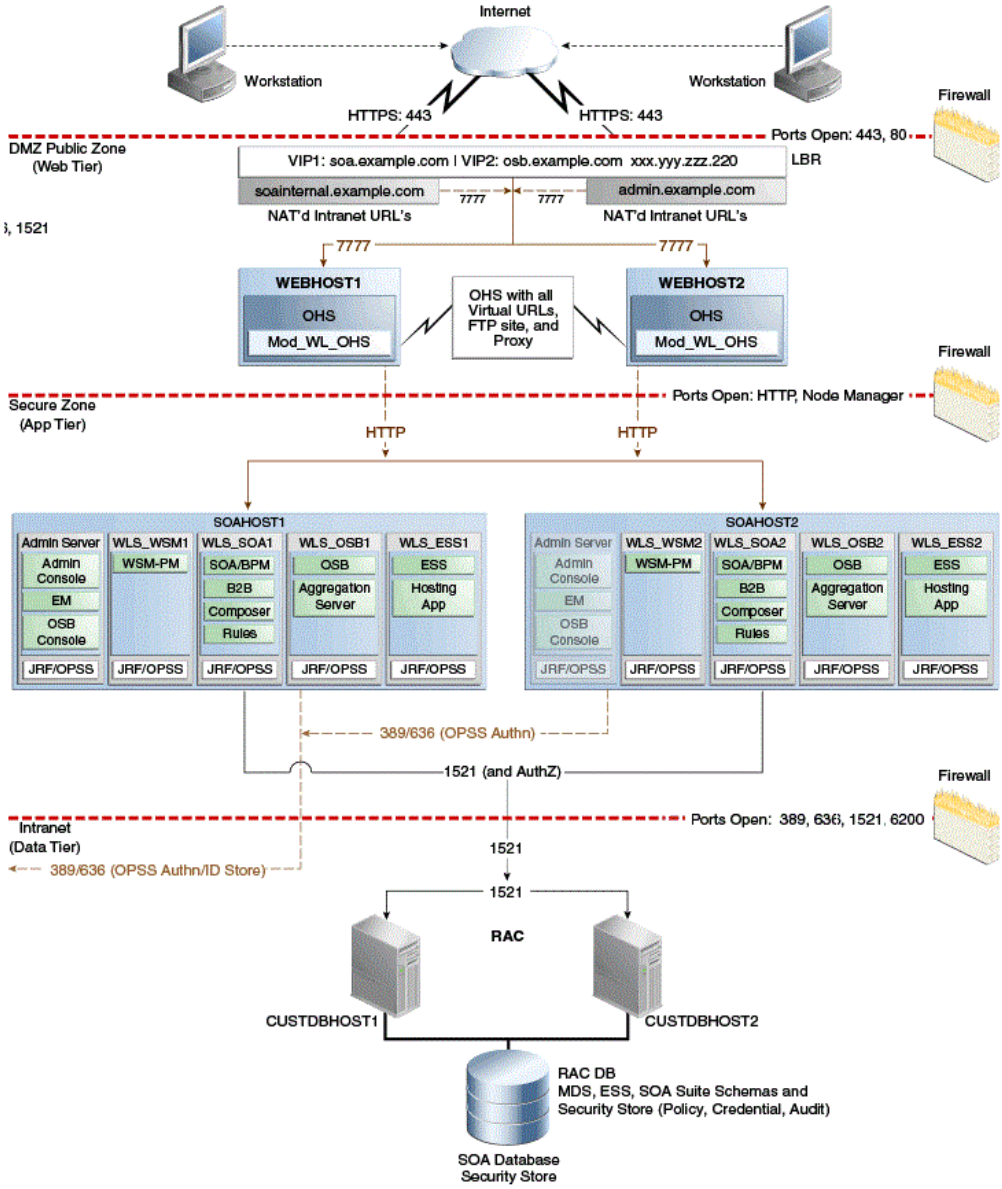

For disaster recovery of this Oracle SOA Suite topology, Oracle recommends the following volume design:

- Provision two volumes for two Oracle homes that contain redundant product binary files (`VOLFMW1` and `VOLFMW2` in Table 4–1).
- Provision one volume for the Administration Server domain directory (`VOLADMIN` in Table 4–1).
- Provision one volume on each node for the Managed Server domain directory (`VOLSOA1` and `VOLSOA2` in Table 4–1). This directory is shared among all the Managed Servers on that node.
- Provision one volume for the JMS file store and JTA transaction logs (`VOLDATA` in Table 4–1). One volume for the entire domain 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 4–1).
- Provision one volume on each node for the Oracle HTTP Server Oracle instance (VOLWEBINST1 and VOLWEBINST2 in Table 4–1).

Table 4–1 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 4–1 Volume Design Recommendations for Oracle SOA Suite

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u01/oracle/products/fmw*nnnn*/web	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u01/oracle/products/fmw*nnnn*/web	Volume for Oracle HTTP Server installation
Web	VOLWEBINST1	WEBHOST1	/u01/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLWEBINST2	WEBHOST2	/u01/oracle/admin/ohs_instance	Volume for Oracle HTTP Server instance
Web	VOLSTATIC1[1]	WEBHOST1	/u01/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Web	VOLSTATIC2[2]	WEBHOST2	/u01/oracle/admin/ohs_instance/config/static	Volume for static HTML content
Application	VOLFMW1	SOAHOST1	/u01/oracle/products/fmw*nnnn*/	Volume for the WebLogic Server and Oracle SOA Suite binary files
Application	VOLFMW2	SOAHOST2	/u01/oracle/products/fmw*nnnn*/	Volume for the WebLogic Server and Oracle SOA Suite binary files
Application	VOLADMIN	SOAHOST1	/u01/oracle/admin/soaDomain/admin	Volume for Administration Server domain directory
Application	VOLSOA1	SOAHOST1	/u01/oracle/admin/soaDomain/mng1	Volume for Managed Server domain directory
Application	VOLSOA2	SOAHOST2	/u01/oracle/admin/soaDomain/mng2	Volume for Managed Server domain directory
Application	VOLDATA	SOAHOST1, SOAHOST2	/u01/oracle/config/soaDomain/soaCluster/jms /u01/oracle/config/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 4–2).
- Create one consistency group with the volume containing the JMS file store and transaction log data as members (DATAGROUP in Table 4–2).
- Create one consistency group with the volume containing the Oracle homes as members (FMWHOMEGROUP in Table 4–2).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle homes as members (WEBHOMEGROUP in Table 4–2).
- Create one consistency group with the volumes containing the Oracle HTTP Server Oracle instances as members (WEBINSTANCEGROUP in Table 4–2).

Table 4–2 provides a summary of Oracle recommendations for consistency groups for the Oracle SOA Suite topology shown in Figure 4–2.

Table 4–2 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 Oracle 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.

4.1.2 Storage Replication

Follow these steps to set up storage replication for the Oracle Fusion Middleware Disaster Recovery topology:

- On the standby site, ensure that alias host names that are created are the same as the physical host names used for the peer hosts at the production site.
- On the shared storage at the standby site, create the same volumes as were created on the shared storage at the production site.
- 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. For more information about symbolic links, see section Section 3.2.3.

- 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.
- 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.
- 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.
- 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.
- The standby site shared storage receives snapshots transferred on periodically from the production site shared storage. After the snapshots are applied, the standby site shared storage includes all the data up to and including the data contained in the last snapshot transferred from the production site before the failover or switchover.
- Oracle strongly recommends that you 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

This section describes how to install and configure Oracle Database 11.2 or 12.1 MAA environments in an Oracle SOA Suite enterprise deployment.

For recommendations and considerations for setting up Oracle databases that are used in an Oracle Fusion Middleware Disaster Recovery topology, see Section 3.3.

4.1.3.1 Installing and Configuring Oracle Database 11.2 or 12.1 MAA Environments

Oracle Maximum Availability Architecture (MAA) is Oracle's comprehensive architecture for reducing downtime for scheduled outages, and preventing, detecting and recovering from unscheduled outages.

Real Application Clusters (RAC) and Data Guard provide the basis of the database MAA solution, where the primary site contains the RAC database, and the secondary site contains the RAC physical standby database.

This section contains the following topics:

- Prerequisites and Assumptions

- Oracle Data Guard Environment Description
- Procedure for Duplicating the Primary Database
- Procedure for Completing RAC Configuration for Standby Database
- Creating a Data Guard Broker Configuration
- Verifying the Data Guard Broker Configuration
- Testing Database Switchover and Switchback

> **Tip:** Alternatively, you can perform many of the tasks in this section using Oracle Enterprise Manager Cloud Control (Cloud Control).
>
> Setting up and managing databases using Cloud Control helps in controlling downtime and simplifies disaster recovery operations.
>
> For information about installing Enterprise Manager Cloud Control 12c, see *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.
>
> For more information about setting up Oracle Data Guard using Cloud Control, see *Set Up and Manage Oracle Data Guard using Oracle Enterprise Manager Cloud Control 12c*.

4.1.3.1.1 Prerequisites and Assumptions Ensure that the following prerequisites are met:

- The Oracle RAC cluster and Automatic Storage Management (ASM) instances on the standby site have been created.
- The Oracle RAC databases on the standby site and the production site are using a flash recovery area.
- The Oracle RAC databases are running in `archivelog` mode.
- The database hosts on the standby site already have Oracle software installed.
- In a shared `ORACLE_HOME` configuration, the `TNS_ADMIN` directory must be a local, non-shared directory.

4.1.3.1.2 Oracle Data Guard Environment Description The examples given in this section contain environment variables as described in Table 4–3.

Table 4–3 Variables Used by Primary and Standby Databases

Variable	Primary Database	Standby Database
Database names	`soa`	`soa`
SOA Database Host Names	`soadc1.dbhost1`, `soadc1.dbhost2`	`soadc2.dbhost1`, `soadc2.dbhost2`
Database unique names	`psoa`	`ssoa`
Instance names	`soa1, soa2`	`soa1, soa2`
Service names	`psoa, ssoa`	`psoa, ssoa`

4.1.3.1.3 Procedure for Duplicating the Primary Database Follow these steps to prepare the primary database for setting up Oracle Data Guard:

> **Note:** For information about prerequisites for setting up Oracle Data Guard, see "Prerequisites" in *Oracle Data Guard Broker.*

1. Enable force logging on the primary database:

```
SQL> alter database force logging;
SQL> select name, log_mode, force_logging from v$database;
NAME LOG_MODE FOR
---------- ---------------- -------
PSOA ARCHIVELOG YES
```

2. In the standby node 1, create and start a listener (for example, `LISTENER_DUPLICATE`) that offers a static `SID` entry for the standby database, and has the same value of `ORACLE_SID` as the primary (`soa1`), as shown in Example 4–1.

   Provide the following path to the `listener.ora` file:

   GRID_HOME/network/admin/listener.ora

Example 4–1 Script for Creating and Starting a Listener with Static SID

```
LISTENER_duplicate =
 (DESCRIPTION_LIST =
  (DESCRIPTION =
   (ADDRESS = (PROTOCOL = TCP)
   (HOST = soadc2.dbhost1)
   (PORT = 1521)(IP = FIRST))))

SID_LIST_LISTENER_duplicate =
 (SID_LIST =
  (SID_DESC =
   (SID_NAME = soa1)
   (ORACLE_HOME = /u01/app/oracle/product/11.2.0/db_1)))
```

3. Run the following command to start and verify the listeners:

```
lsnrctl start listener_duplicate
lsnrctl status listener_duplicate
```

4. In the standby node 2, create and start a listener (for example, `LISTENER_DUPLICATE`) that offers a static `SID` entry for the standby database, and has the same value of `ORACLE_SID` as the primary (`soa2`), as shown in Example 4–2.

   Provide the following path to the `listener.ora` file:

   GRID_HOME/network/admin/listener.ora

Example 4–2 Script for Creating and Starting a Listener with Static SID

```
LISTENER_duplicate =
 (DESCRIPTION_LIST =
  (DESCRIPTION =
   (ADDRESS = (PROTOCOL = TCP)
   (HOST = soadc2.dbhost2)
   (PORT = 1521)(IP = FIRST))))

SID_LIST_LISTENER_duplicate =
 (SID_LIST =
  (SID_DESC =
   (SID_NAME = soa2)
   (ORACLE_HOME = /u01/app/oracle/product/11.2.0/db_1)))
```

Note: Listeners are configured at the cluster level, and all nodes inherit the port and environment settings of the listener. Therefore, the TNS_ADMIN directory path have the same values on all nodes.

In a shared ORACLE_HOME configuration, the TNS_ADMIN directory must be a local, non-shared directory. These network files are included as IFILES.

Complete the following steps to set up TNS_ADMIN for a shared ORACLE_HOME in a two-node cluster, SOADC1.DBHOST1 and SOADC1.DBHOST2, with respective instances SOA1 and SOA2:

1. Create a local network directory on each node. For example, */local_network_dir*/network_admin.
2. Create a local listener.ora file in the location: */local_network_dir*/network_admin on each node.
3. In the local listener.ora file, add the values for the LISTENER_duplicate parameter.
4. In the common listener.ora file, in *GRID_HOME*/network/admin, add the IFILE parameter values, as follows:

   IFILE=/*local_network_dir*/network_admin/listener.ora

5. Run the following command to start and verify the listeners:

```
lsnrctl start listener_duplicate
lsnrctl status listener_duplicate
```

6. In the database home of the primary node, create an Oracle Net alias to connect to the listener that you created in step 2. See Example 4–3.

Example 4–3 Script for Creating an Oracle Net Alias

```
dup =
(DESCRIPTION =
 (ADDRESS =
  (PROTOCOL = TCP)
   (HOST = soadc2.dbhost1)
   (PORT = 1521))
 (CONNECT_DATA =
  (SERVER = DEDICATED)
   (SID = soa1)))
```

7. Configure Oracle password file authentication for redo transport. Make sure it meets the following requirements:
   - Set remote_login_passwordfile to EXCLUSIVE.

```
SQL> show parameter
remote_login_passwordfile
NAME TYPE VALUE
------------------------------- --------------
-------------------------------
remote_login_passwordfile string EXCLUSIVE
```

   - Ensure that primary and standby databases have the byte-identical password files.

8. In the ORACLE_HOME/dbs directory of the standby host, create a pfile, initsoa1.ora, with the following parameters:

```
db_name=soa
db_unique_name=ssoa
sga_target=5g
```

9. Create the audit directory for the `soa` database on all standby hosts:

```
mkdir -p /u01/app/oracle/admin/soa/adump
```

10. Create an Oracle Net alias on all primary hosts to reach the `ssoa` database on the standby hosts.

    Ensure that all hosts have Oracle Net alias for `psoa and ssoa`. All the aliases that you create need to reference the scan listener and not the node VIP. Also, if the `local_listener` variable is set to an alias on the primary host, then enter the details of the variable on the standby site, that point to the local listener on the primary host. See Example 4–4 and Example 4–5.

Example 4–4 Sample tnsnames.ora File on Primary Node 1 (SOADC1.DBHOST1)

```
psoa =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL= TCP)
(HOST=prmy-scan)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = psoa)
    )
  )

ssoa =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=stby-scan)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ssoa)
    )
  )

psoa_local_listener =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=prmy1-vip)(PORT = 1521))
    )
  )
```

Note:

- For primary node 2 (`SOADC1.DBHOST2`), update the `HOST` value in `psoa_local_listener` to point to the `VIP` of primary node 2, as follows:

```
psoa_local_listener =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=prmy2-vip)(PORT = 1521))
    )
  )
```

- If you are using a shared Oracle home, add the `VIP`s of the two nodes to the local_listener parameter.

  For example:

```
psoa_local_listener =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=prmy1-vip)(PORT = 1521))
    (ADDRESS=(PROTOCOL = TCP)
(HOST=prmy2-vip)(PORT = 1521))
    )
  )
```

Example 4–5 Sample tnsnames.ora File on Standby Node 1 (SOADC2.DBHOST1)

```
psoa =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL= TCP)(HOST=prmy-scan)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = psoa)
    )
  )

ssoa =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=stby-scan)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = ssoa)
    )
  )

ssoa_local_listener =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=stby1-vip)(PORT = 1521))
    )
  )
```

Note:

- For standby node 2 (SOADC2.DBHOST2), update the HOST value in ssoa_local_listener to point to the VIP of standby node 2, as follows:

```
ssoa_local_listener = (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=stby2-vip)(PORT = 1521))
    )
  )
```

- If you are using a shared Oracle home, add the VIPs of the two nodes to the local_listener parameter.

  For example:

```
ssoa_local_listener =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL = TCP)
(HOST=stby1-vip)(PORT = 1521))
    (ADDRESS=(PROTOCOL = TCP)
(HOST=stby2-vip)(PORT = 1521))
    )
  )
```

11. On the standby host, set the ORACLE_SID to the same value as that of the primary database (ORACLE_SID=soa1), and run the startup nomount command on the standby database with standby PFILE.
12. Disable the parameter cluster_interconnects on the primary host if it is set:

```
SQL> alter system reset cluster_interconnects scope=spfile sid='soa1';
SQL> alter system reset cluster_interconnects scope=spfile sid='soa2';
```

13. On the primary host, run the Recovery Manager (RMAN) script to duplicate the primary database using the command duplicate target database for standby from active database.

    Note: This command varies depending on your environment. For information about how to use the command, see "Duplicating a Database" in *Oracle Database Backup and Recovery User's Guide*.

    See Example 4–6 to understand how to duplicate between two systems with different disk-group names.

Example 4–6 Duplicating Data Between Two Systems with Different Disk-Group Names

```
rman <<EOF
connect target sys/password;
connect auxiliary sys/password@dup;
run {
allocate channel prmy1 type disk;
allocate channel prmy2 type disk;
allocate channel prmy3 type disk;
allocate channel prmy4 type disk;
allocate auxiliary channel stby type disk;
```

```
duplicate target database for standby from active database
spfile
parameter_value_convert '+DATA_prmy','+DATA_stby','+RECO_prmy','+RECO_stby'
set db_file_name_convert '+DATA_prmy','+DATA_stby'
set db_unique_name='ssoa'
set db_create_online_log_dest_1='+DATA_stby'
set db_create_file_dest='+DATA_stby'
set db_recovery_file_dest='+RECO_stby'
set log_file_name_convert '+DATA_prmy','+DATA_stby','+RECO_prmy','+RECO_stby'
set control_files='+DATA_stby/ssoa/standby.ctl'
set local_listener='ssoa_local_listener'
set remote_listener='stby-scan:1521';
} EOF
```

See Example 4–7 to understand how to duplicate between two systems that have the same disk-group name +DATA.

Example 4–7 Duplicating Data Between Two Systems with Same Disk-Group Name +DATA

```
rman <<EOF
connect target sys/password;
connect auxiliary sys/password@dup;
run {
allocate channel prmy1 type disk;
allocate channel prmy2 type disk;
allocate channel prmy3 type disk;
allocate channel prmy4 type disk;
allocate auxiliary channel stby type disk;

duplicate target database for standby from active database
spfile
set db_unique_name='ssoa'
set control_files='+DATA/ssoa/standby.ctl'
set local_listener='ssoa_local_listener'
set remote_listener='stby-scan:1521';
}
EOF
```

14. If you have disabled the parameter `cluster_interconnects` on the primary host as described in step 12, then you must set it back to the original values in the spfile.
15. (Optional) Stop and remove the listener that you created in step 2.

4.1.3.1.4 Procedure for Completing RAC Configuration for Standby Database To complete the RAC configuration on the standby database, complete the steps given in Section 4.1.3.1.3. Then, perform the following steps:

1. Create a temporary parameter file:

   ```
   SQL> create pfile='/tmp/p.ora' from spfile;
   ```

2. Create an `SPFILE` in `+DATA_stby` for the standby database:

   ```
   SQL> create spfile='+DATA_stby/ssoa/spfilessoa.ora' from pfile='/tmp/p.ora';
   ```

3. Create an `initsoan.ora` file on all the standby hosts. The file needs to point to the location of the `SPFILE` created in step 2.
4. On the standby system, restart the instances in mount state:

```
startup mount
```

5. Register the RAC database with CRS as follow:

```
srvctl add database -d ssoa -o /u01/app/oracle/product/11.2.0/db_1
srvctl add instance -d ssoa -i soa1 -n soadc2.dbhost1
srvctl add instance -d ssoa -i soa2 -n soadc2.dbhost2
srvctl modify database -d ssoa -r physical_standby
```

6. Create standby redo logs that are the same size as the online redo logs, on the primary database and standby database.

   Oracle recommends having the same number plus one additional redo logs for each thread as shown in Example 4–8.

Example 4–8 Sample Redo Log

```
SQL> alter database add standby logfile thread 1
group 9 size 500M,
group 10 size 500M,
group 11 size 500M;

SQL> alter database add standby logfile thread 2
group 12 size 500M,
group 13 size 500M,
group 14 size 500M;
```

4.1.3.1.5 Creating a Data Guard Broker Configuration This section describes the basic steps for creating a Data Guard configuration.

For complete information about Data Guard Broker, see the *Oracle Data Guard Broker* guide.

To create a Data Guard Broker configuration, complete the following steps:

1. Add the values of static SID to the local node file, `listener.ora`, that is located in the grid infrastructure home on all hosts in the configuration:

```
LISTENER_SCAN2=(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER_
SCAN2))))
LISTENER_SCAN3=(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER_
SCAN3))))
LISTENER=(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER))))

LISTENER_SCAN1=(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER_
SCAN1))))
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_SCAN1=ON
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER=ON
SID_LIST_LISTENER =
(SID_LIST =
(SID_DESC =
(GLOBAL_DBNAME = psoa_DGMGRL)
(SID_NAME = soa1)
(ORACLE_HOME = /u01/app/oracle/product/11.2.0/db_1)))
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_SCAN3=ON
```

Note:

- Ensure that you enter the appropriate values for parameters `GLOBAL_DBNAME` and `SID_NAME`, for each host.
- If `IFILE` is used for shared `ORACLE_HOME`, modify the corresponding local `listener.ora` file specified in `IFILE`.

2. Bounce the listeners on the primary site and standby site to apply the modifications made to the files:

```
srvctl stop listener
srvctl start listener
```

3. Configure the Data Guard Broker metadata files, and enable the broker on the primary site and standby site as follows:

On the primary site

```
SQL> alter system set dg_broker_config_file1='+DATA_prmy/psoa/dr1.dat'
scope=both;
SQL> alter system set dg_broker_config_file2='+DATA_prmy/psoa/dr2.dat'
scope=both;
SQL> alter system set dg_broker_start=true scope=both;
```

On the standby site

```
SQL> alter system set dg_broker_config_file1='+DATA_stby/ssoa/dr1.dat'
scope=both;
SQL> alter system set dg_broker_config_file2='+DATA_stby/ssoa/dr2.dat'
scope=both;
SQL> alter system set dg_broker_start=true scope=both;
```

4. Access `dgmgrl` on the primary host, and create the configuration. See Example 4–9.

Example 4–9 Accessing dgmgrl to Create the Data Guard Broker Configuration on Primary Host

```
dgmgrl sys/password

DGMGRL> create configuration 'dg_config' as
primary database is 'psoa'
connect identifier is psoa;

Configuration "dg_config" created with primary database "psoa"

DGMGRL> add database 'ssoa' as
connect identifier is ssoa;

Database "ssoa" added

DGMGRL> enable configuration;

Enabled.
```

4.1.3.1.6 Verifying the Data Guard Broker Configuration

Complete the following steps to verify that the Data Guard Broker configuration was created successfully.

1. 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
  2> FROM V$ARCHIVED_LOG ORDER BY SEQUENCE#;

SEQUENCE# FIRST_TIME NEXT_TIME
---------- ------------------ ------------------
         8 11-JUL-13 17:50:45 11-JUL-13 17:50:53
         9 11-JUL-1317:50:53 11-JUL-1317:50:58
        10 11-JUL-13 17:50:58 11-JUL-13 17:51:03

3 rows selected
```

2. 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;
```

3. 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
  2> FROM V$ARCHIVED_LOG ORDER BY SEQUENCE#;

SEQUENCE# FIRST_TIME NEXT_TIME
---------- ------------------ ------------------
         8 11-JUL-13 17:50:45 11-JUL-13 17:50:53
         9 11-JUL-1317:50:53 11-JUL-1317:50:58
         10 11-JUL-13 17:50:58 11-JUL-13 17:51:03
         11 11-JUL-13 17:51:03 11-JUL-1318:34:11

4 rows selected
```

4. Using the `show configuration` command, verify that the configuration was created successfully on the standby site. See Example 4–10.

Example 4–10 Verifying the Data Guard Broker Configuration

```
DGMGRL> show configuration;

Configuration - dg_config

Protection Mode: MaxPerformance
Databases:
psoa- Primary database
ssoa- Physical standby database

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS
```

4.1.3.1.7 Testing Database Switchover and Switchback This section contains the following topics:

- Performing a Switchover Operation Using Oracle Data Guard Broker
- Performing a Switchover Operation Using SQL*Plus

Performing a Switchover Operation Using Oracle Data Guard Broker

To perform a switchover operation using Oracle Data Guard Broker, complete the following tasks:

1. Verify the Oracle Data Guard Broker configuration that you created using the instructions provided in Section 4.1.3.1.5, "Creating a Data Guard Broker Configuration."

   To verify the configuration, run the following command:

```
DGMGRL> show configuration;

Configuration - dg_config

Protection Mode: MaxPerformance
Databases:
psoa- Primary database
ssoa- Physical standby database

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS
```

2. Swap the roles of the primary and standby databases by running the `SWITCHOVER` command. Example 4–11 shows how Data Guard Broker automatically shuts down and restarts the old primary database as a part of the switchover operation.

Example 4–11

```
DGMGRL> switchover to 'ssoa';
Performing switchover NOW, please wait...
New primary database "ssoa" is opening...
Operation requires shutdown of instance "psoa1" on database "psoa"
Shutting down instance "psoa1"...
ORA-01109: database not open

Database dismounted.
ORACLE instance shut down.
Operation requires startup of instance "psoa1" on database "psoa"
Starting instance "psoa1"...
ORACLE instance started.
Database mounted.
Switchover succeeded, new primary is "ssoa"
```

3. After the switchover is complete, using the `SHOW CONFIGURATION` command verify that the switchover operation was successful:

```
DGMGRL> show configuration;
Configuration - dg_config
Protection Mode: MaxPerformance
Databases:
ssoa- Primary database
psoa- Physical standby database
Fast-Start Failover: DISABLED
Configuration Status:
SUCCESS
```

Performing a Switchover Operation Using SQL*Plus

Perform the following operations to switchover or switchback databases correctly between the newly created physical standby database and the primary Oracle RAC databases:

1. Shut down all but one instance of the Oracle RAC databases (PSOA) on the primary site. For example, run the following command on `SOADC1.DBHOST1` 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 following command on `SOADC1.DBHOST1` 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 following command on `SOADC1.DBHOST1` of the production site:

   ```
   SQL > shutdown immediate
   SQL > startup mount
   ```

4. At this point, both the databases are in Physical Standby mode.

   Log in to the instances to continue.

5. 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
   ```

6. Switch the physical standby database role to the primary role. For example, run the following command on `SOADC2.DBHOST1` of the standby site:

   ```
   SQL> ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY WITH SESSION SHUTDOWN;
   ```

7. Now the physical standby database is the new primary database.

8. Shut down the new primary database and start up both the RAC nodes using `srvctl`. For example, run the following command on the new primary site `SOADC2.DBHOST1`:

   ```
   srvctl start database -d ssoa
   ```

9. On the new physical standby database (the old primary) start the managed recovery of the database. For example, run the following command on `SOADC1.DBHOST1` of the primary site:

   ```
   SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE DISCONNECT FROM SESSION;
   ```

10. Start sending the redo data to the new physical standby database. For example, run the following command on the new primary site `SOADC2.DBHOST1`:

    ```
    SQL> ALTER SYSTEM SWITCH LOGFILE;
    ```

11. Check the new physical standby database to see if it is receiving the archive log files by querying the `V$ARCHIVED_LOG` view.

12. Perform a switchback operation. A switchback operation is a subsequent switchover operation to return the roles to their original state.

> **Note:** For information about switchover and failover operation of Oracle Data Guard Broker, see "Switchover and Failover Operations" in the *Oracle Data Guard Broker*.

4.2 Creating a Production Site

This section provides the steps to create the production site on an Oracle SOA Suite enterprise deployment topology.

It contains the following topics:

- Creating the Production Site for the Oracle SOA Suite Topology
- Configuring Data Sources for Oracle Fusion Middleware SOA Active-Passive Deployment

Perform the following prerequisites before you start creating the production site:

- Set up the host name aliases for the middle tier hosts as described in Section 3.1.1.
- Create the required volumes on the shared storage on the production site, as described in Section 4.1.1.
- Determine the Oracle Data Guard configuration to use based on the data loss requirements of the database and network considerations such as the available bandwidth and latency when compared to the redo generation.

  See *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`

4.2.1 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. Complete the following tasks to install and configure the production site:

Task 1 Create volumes and consistency groups

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."

Task 2 Set up physical host names and alias host names

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 about planning host names for the production and standby sites.

Task 3 Install and configure Oracle SOA Suite

Install and configure Oracle SOA Suite as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* with the following modifications:

1. Install the Oracle SOA Suite components into the volumes created on the shared storage device.

2. Set up the production and standby sites using the aliases to the physical and virtual host names.
3. Create a separate volume on each site for the JMS stores and transaction logs.
4. After the installation and configuration of the production site, turn off host name verification. See *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.
5. If you do not plan on turning host name verification off, follow the steps in Section 4.3.2 to configure Node Manager communication.
6. Create SSL certificates using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

4.2.2 Configuring Data Sources for Oracle Fusion Middleware SOA Active-Passive Deployment

The data sources used by Oracle Fusion Middleware SOA Suite should be configured to automate failover of connections in case there is a failover or switchover of the primary database. The following data sources need to be configured properly to automate this failover:

- `EDNDataSource`
- `EDNLocalTxDataSource`
- `mds-owsm`
- `mds-soa`
- `OraSDPMDataSource`
- `SOADataSource`
- `SOALocalTxDataSource`

Additionally any persistence store using database and the leasing data source used for server migration should be configured to automate failover. The GridLink data sources must be modified:

- Update the ONS configuration to include both production and standby site ONS.

  The items in the list of ONS addresses must be separated by commas.

  For example:

  ```
  prmy-scan:6200,stby-scan:6200
  ```

  On the Test ONS Client Configuration page, review the connection parameters, and click **Test All ONS Nodes**.

  Following is an example of a successful connection notification:

  ```
  Connect test for prmy-scan:6200 succeeded.
  ```

- Update the JDBC URL to include the appropriate services in both sites.

  The following is a sample JDBC URL for the SOA Data source in an Oracle Fusion Middleware SOA Active/Passive configuration where the database uses Data Guard.

  ```
  jdbc:oracle:thin:@
  (DESCRIPTION_LIST =
          (LOAD_BALANCE = off)
  ```

```
                (FAILOVER = on)
                (DESCRIPTION =
                         (CONNECT_TIMEOUT = 10)
                         (TRANSPORT_CONNECT_TIMEOUT = 3)
                         (RETRY_COUNT = 3)
                         (ADDRESS_LIST =
                              (LOAD_BALANCE = on)
                              (ADDRESS =
                                  (PROTOCOL = TCP)(HOST = prmy-scan)(PORT = 1521)
                         )
                )
                (CONNECT_DATA =
                     (SERVICE_NAME =soaedg.example.com)
                         )
                )
(DESCRIPTION =
                (CONNECT_TIMEOUT =10)
                (TRANSPORT_CONNECT_TIMEOUT = 3)
                (RETRY_COUNT = 3)
                (ADDRESS_LIST =
                              (LOAD_BALANCE = on)
                              (ADDRESS =
                                  (PROTOCOL =TCP)(HOST = stby-scan)(PORT = 1521)
                         )
                )
                (CONNECT_DATA =
                              (SERVICE_NAME = soaedg.example.com))

                         )

)
```

In the Test GridLink Database Connection page, review the connection parameters, and click **Test All Listeners**.

Following is an example of a successful connection notification:

```
Connection test for jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=prmy-scan)(PORT=1521)))(CONNECT_
DATA=(SERVICE_NAME=soaedg.example.com))) succeeded.
```

4.3 Creating a Standby Site

This section provides the steps to create the standby site. The Oracle SOA enterprise deployment topology is used as an example.

It includes the following topics:

- Creating the Standby Site
- Updating Self-Signed Certificates and Keystore on Standby Site
- Validating the Standby Site Setup

4.3.1 Creating the Standby Site

Perform the following prerequisites before you start creating the standby site:

- On the standby site, set up the correct alias host names and physical host names by following the instructions in Section 3.1.1.

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.

  For more details about symbolic links, see Section 3.2.3.

4.3.1.1 Setting Up Middle Tier Hosts

The middle tier hosts on the standby site do not require the installation or configuration of any Oracle Fusion Middleware or Oracle WebLogic Server software. When the production site storage is replicated to the standby site storage, the software installed on the production site volumes is replicated at the standby site volumes.

Do the following 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 Updating Self-Signed Certificates and Keystore on Standby Site

If you enable the hostname verification for the Administration Server, you must also update the appropriate trust stores and key stores with the certificates of the standby site. Certificates must be specifically created for the nodes in the standby site.

For more information about creating certificates for nodes, see "Enabling SSL Communication Between the Middle Tier and the Hardware Load Balancer" in *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

This section includes the following topics:

- Generate Self-Signed Certificates
- Create an Identity KeyStore
- Create a Trust KeyStore

The examples in these sections show how to perform tasks for the Oracle SOA Suite enterprise topology shown in Figure 4–2.

Note:

- When you set 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 Oracle WebLogic Server is installed.

4.3.2.1 Generate Self-Signed Certificates

Follow these steps to generate self-signed certificates:

1. Set up your environment by running the `WL_HOME/server/bin/setWLSEnv.sh` script:

   In the Bourne shell, run the following command:

   ```
   . setWLSEnv.sh
   ```

   Verify that the `CLASSPATH` environment variable is set:

   ```
   echo $CLASSPATH
   ```

2. Create a user-defined directory for the certificates.

   ```
   mkdir home/user_projects/domains/SOADomain/certs
   ```

3. Change directory to the user-defined directory.

   ```
   cd home/user_projects/domains/SOADomain/certs
   ```

4. Run the `utils.CertGen` tool from the user-defined directory to create the certificates for both the physical hostnames and the virtual hostnames used by servers in the node:

   ```
   java utils.CertGen key_passphrase cert_file_name key_file_name
   [export|domestic] [hostname]
   ```

   For example:

   ```
   java utils.CertGen password soadc1host1_cert soahost1_key domestic soahost1
   java utils.CertGen password soadc1host2_cert soahost2_key domestic soahost2
   ```

4.3.2.2 Create an Identity KeyStore

Follow these steps to create an identity keystore using the `utils.ImportPrivateKey` utility:

1. Import the certificate and private key into the Identity Store using the following syntax. Make sure that you use a different alias for each of the certificate/key pair imported:

   ```
   java utils.ImportPrivateKey keystore_file keystore_password certificate_alias_
   to_use private_key_passphrase certificate_file private_key_file keystore_type
   ```

Note: Default value for `keystore_type` is `jks`.

For example:

```
java utils.ImportPrivateKey appIdentityKeyStore.jks password
            appIdentity1 password
            ASERVER_HOME/certs/SOAHOST1.example.com_cert.pem
            ASERVER_HOME/certs/SOAHOST1.example.com_key.pem
```

2. Repeat step 1 for all the hosts on the primary and the standby sites.

4.3.2.3 Create a Trust KeyStore

Follow these steps to create a trust keystore:

1. Copy the standard java keystore to create the new trust keystore since it already contains most of the root CA certificates needed.

   Oracle does not recommend modifying the standard Java trust key store directly.

   Copy the standard Java keystore CA certificates located under the `WL_HOME/server/lib` directory to the same directory as the certificates.

   For example:

   ```
   cp $WL_HOME/server/lib/cacerts
   $home/user_projects/domains/SOADomain/certs/appTrustKeyStore.jks
   ```

2. The default password for the standard Java keystore is `changeit`.

   Oracle recommends that you change the default password. Use the `keytool` utility to do this.

   The syntax is:

   ```
   keytool -storepasswd -new NewPassword -keystore TrustKeyStore -storepass
   OriginalPassword
   ```

   For example, enter this command:

   ```
   keytool -storepasswd -new password -keystore home//user_
   projects/domains/SOADomain/certs/appTrustKeyStore.jks -storepass
   changeit
   ```

3. The CA certificate `CertGenCA.der` is used to sign all certificates generated by the `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
   CA_file_location -keystore key_store_location -storepass key_store_password
   ```

   For example, enter this command:

   ```
   keytool -import -v -noprompt -trustcacerts -alias clientCACert -file
   $WL_HOME/server/lib/CertGenCA.der -keystore appTrust.jks -storepass password
   ```

4.3.3 Validating the Standby Site Setup

Validate the standby site by performing the following steps:

1. Shut down any processes still running on the production site. These include 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. Perform a switchback operation. A switchback operation is a subsequent switchover operation to return the roles to their original state.
4. On the standby site host, manually start all the processes. These include 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, ensure that you read and understand the concepts and information about 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.

The following sections describe the basic steps for creating an asymmetric topology:

- Creating the Asymmetric Standby Site
- Validating the Asymmetric Standby Site Setup

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.

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;

   For more details about symbolic links, see Section 3.2.3.

   For more information about volume design, see Section 4.1.1.1.1.

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;

   For more details about symbolic links, see Section 3.2.3.

   For more information about the Oracle Central Inventory directories, see Section 3.2.2.

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;

   For more details about symbolic links, see Section 3.2.3.

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.

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 about creating the initial snapshot and enabling 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 that you set up for the peer instance on the standby site host must be the same as the mount point directory that 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;

      For more details about symbolic links, see Section 3.2.3. The symbolic link that you set up for the peer instance on the standby site host must be the same as the symbolic link that 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 that you set up for the peer instance on the standby site host must be the same as the mount point directory that 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;

      For more details about symbolic links, see Section 3.2.3. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link that 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 that you set up for the peer instance on the standby site host must be the same as the mount point directory that 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 for more details about symbolic links. The symbolic link that you set up for the peer instance on the standby site host must be the same as the symbolic link that you set up for the instance on the production site host.

After you complete 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 for more details about symbolic links.
- The Oracle Central Inventory directories are located in the 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 for more details about symbolic links.
- The Oracle HTTP Server static HTML pages directories are located in the 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 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 through Section 4.1.1 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 4–4 shows the asymmetric standby site for the production site shown in Figure 4–2.

Figure 4–4 An Asymmetric Standby Site with Fewer Hosts and Instances

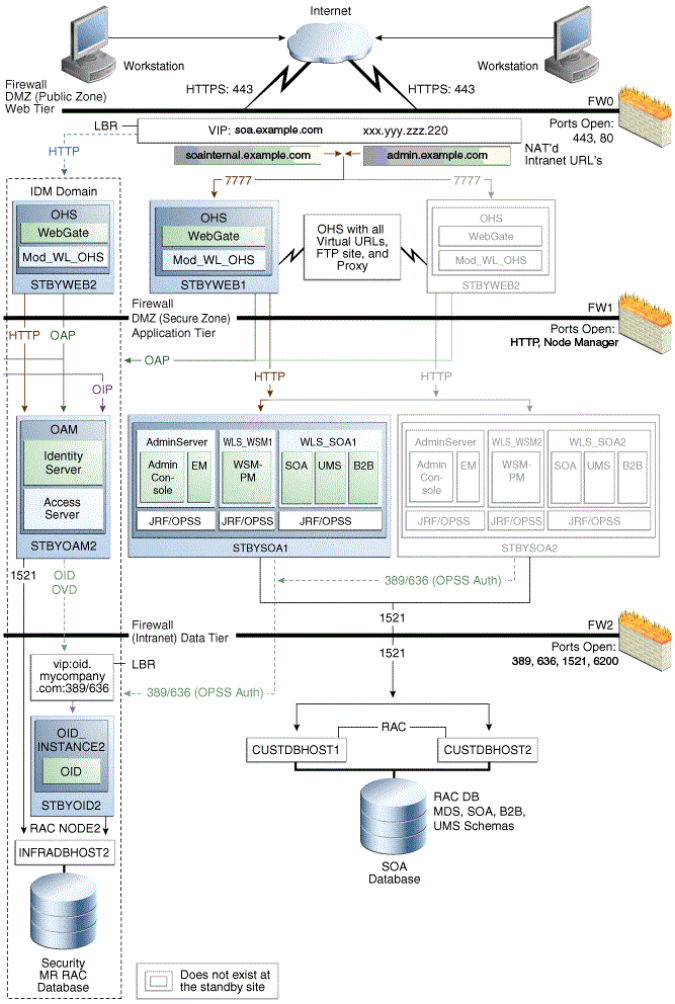

The Oracle SOA Suite asymmetric standby site shown in Figure 4–4 has fewer hosts and instances than the Oracle SOA Suite production site shown in Figure 4–2.

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 4–4. 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 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 4–1 and the consistency group recommendations in Table 4–2 were used to set up the production site shared storage. 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 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, the asymmetric standby site for Oracle SOA Suite shown in Figure 4–4, WEBHOST2 and SOAHOST2 do not exist; 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 performing the following the steps:

1. Shut down any processes still running on the production site. These include 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 switchover the databases.
4. On the standby site host, manually start all the processes. These include 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.

It contains the following sections:

- Synchronizing the Sites
- Performing a Switchover
- Performing a Switchback
- Performing a Failover
- Performing Periodic Testing of the Standby Site
- Using Peer-to-Peer File Copy for Testing

4.5.1 Synchronizing the Sites

The standby site shared storage receives snapshots transferred periodically 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. These include 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. Unmount the shared storage volume with the middle tier artifacts, on the current production site, and mount the corresponding volumes on the current standby site, which will be the new production site.
4. Use Oracle Data Guard to switch over the databases.
5. On the standby site host, manually start all the processes. These include the database instances in the data tier, Oracle Fusion Middleware instances, and any other processes in the application tier and web tier.
6. 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.
7. 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.
8. 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). See 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 in the future, you can use it as either the production site or the standby site.

To use the original production site as the new production site, perform the switchback steps described in Section 4.5.3.

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. These include 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. Unmount the shared storage volume with the middle tier artifacts, on the current production site, and mount the corresponding volumes on the current standby site, which will be the new production site.
4. Use Oracle Data Guard to switch back the databases.
5. On the new production site hosts, manually start all the processes. These include Oracle Fusion Middleware instances and any other processes in the application tier and web tier.
6. 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.
7. 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.
8. 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). See 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. Mount the shared storage volume with the middle-tier artifacts, on the current standby site, which will be the new production site.
3. From the standby site, use Oracle Data Guard to fail over the databases.
4. On the standby site hosts, manually start all the processes. These include the 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-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.
7. 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). See 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.

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) 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) 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 is to create a clone of the read-only standby site shared storage and then using 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. However, 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 10.2 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. Cancel Redo Apply:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY
     DATABASE CANCEL;
```

c. Create a guaranteed restore point:

```
SQL> CREATE RESTORE POINT standbytest
     GUARANTEE FLASHBACK DATABASE;
```

d. Archive the current logs at the primary (production) site:

```
SQL> ALTER SYSTEM ARCHIVE LOG CURRENT;
```

e. Defer the standby site destination that you will activate:

```
SQL> ALTER SYSTEM SET
     LOG_ARCHIVE_DEST_STATE_2=DEFER;
```

f. Activate the target standby database:

```
SQL> ALTER DATABASE ACTIVATE STANDBY DATABASE;
```

g. Mount the database with the Force option if the database was opened as read-only:

```
SQL> STARTUP MOUNT FORCE;
```

h. Lower the protection mode and open the database:

```
SQL> ALTER DATABASE SET STANDBY DATABASE TO
     MAXIMIZE PERFORMANCE;
SQL> ALTER DATABASE OPEN;
```

3. Use Oracle Data Guard database recovery procedures to bring the standby databases online.
4. On the standby site computers, modify the mount commands to point to the volumes on the standby site's cloned read/write shared storage by following these steps:
   a. Unmount the read-only shared storage volumes.
   b. Mount the cloned read/write volumes at the same mount point.
5. Before testing the standby site, modify the host name resolution method for the computers that will be used to perform the testing to ensure that the host names point to the standby site computers and not the production 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 standby site.
6. Perform the standby site testing.

After you complete the standby site testing, follow these steps to begin using the original production site as the production site again:

1. Modify the mount commands on the standby site computers to point to the volumes on the standby site's read-only shared storage. In other words, reset the mount commands back to what they were before the testing was performed.
   a. Unmount the cloned read/write shared storage volume.
   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 Oracle Database 10.2 and later, 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 Oracle Database 12*c*, 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. It discusses rsync in the context of symmetric topologies. The information provided for rsync in this section also applies to other peer to peer file copy mechanisms.

Before you read the information in 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 that 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

The following sections describe the basic principles that apply when you use rsync and Oracle Data Guard to provide disaster protection and disaster recovery for your Oracle Fusion Middleware Disaster Recovery topology:

- Using rsync for Oracle Fusion Middleware Middle Tier Components
- Performing Failover and Switchover Operations

Note: For information about how to set up Oracle Data Guard for Oracle database, see Section 3.3.

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 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 for more information about 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, perform the preceding steps again, but configure the rsync replications to go in the original direction (from the original production site to the original standby site).

4.6 Using Oracle Site Guard for Disaster Recovery

Oracle Site Guard primarily orchestrates switchover and failover between two disaster recovery sites. It offers the following features:

- Ensures high availability, data protection, and disaster recovery for enterprise data.
- Performs Oracle Site Guard operations like switchover and failover. If the primary site becomes unavailable due to a planned or an unplanned outage, a Switchover or Failover process needs to be initiated using Oracle Site Guard.

For more information about how to use Oracle Site Guard, see *Oracle Site Guard Administrator's Guide*.

4.7 Patching an Oracle Fusion Middleware Disaster Recovery Site

This section describes how to apply a 12*c* 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 Oracle Fusion Middleware 12*c* 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 Oracle Fusion Middleware 12*c* 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. Because some Oracle Fusion Middleware patch sets make updates to repositories, 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 Oracle Fusion Middleware 12*c* 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.

5 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 includes the following sections:

- Section 5.1, "Troubleshooting Oracle Fusion Middleware Disaster Recovery Topologies"
- Section 5.2, "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:

- Verifying 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 Verifying Host Name Resolution at the Production and Standby Sites

Many issues that arise with Disaster Recovery topology are caused by the improper setup of the host name resolution for the production site and standby site.

Ensure that host name resolution is set up properly by performing the host name validation steps in Section 3.1.1.1.5.

5.1.2 Resolving Issues with Components in a Disaster Recovery Topology

Some issues that 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, then 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, then 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 Business Intelligence*

5.1.3 Resolving Issues with Components Deployed on Shared Storage

Some problems that arise with a component in a 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, 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?

You can find more solutions on My Oracle Support at

`http://support.oracle.com`

If you do not find a solution for your problem, then 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`

A

Managing Oracle Inventory

This appendix describes how to manage Oracle Inventory on the production and standby sites for an Oracle Fusion Middleware Disaster Recovery topology.

It includes the following sections:

- Section A.1, "Updating Oracle Inventory"
- Section A.2, "Updating the Windows Registry"

A.1 Updating Oracle Inventory

When you update 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 ensure that the same software updates are made on the standby site peer host.

To do this, you must update 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 the Oracle home. Edit the following file to update the Oracle home information:

user_home/bea/beahomelist

(Windows) C:\bea\beahomelist

A.2 Updating the Windows Registry

When you update Oracle inventory (for example, by installing new Oracle software or by applying an Oracle patch set or a patch to existing Oracle software) on a production site Windows host, you must ensure 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

C

D

E

F

G

H

I

K

L

M

N

O

P

R

S

T

V

W