Oracle® Fusion Middleware

Disaster Recovery Guide

12*c* (12.2.1.4)
F45425-07
June 2023

Oracle Fusion Middleware Disaster Recovery Guide, 12*c* (12.2.1.4)

F45425-07

Copyright © 2015, 2023, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Preface

This document provides disaster recovery solution for Oracle Fusion Middleware components.

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 `https://www.oracle.com/corporate/accessibility/`.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `https://support.oracle.com/portal/` or visit `Oracle Accessibility Learning and Support` if you are hearing impaired.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

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 WebCenter Content*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Business Intelligence*
- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity and Access Management*

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

Oracle Fusion Middleware Disaster Recovery is a disaster recovery solution that provides protection to Oracle Fusion Middleware components in different Oracle product suites.

This chapter includes the following sections:

- Overview of Oracle Fusion Middleware Disaster Recovery
  Learn about the problems Oracle Fusion Middleware Disaster Recovery solves and get familiar with the terminology.
- Setting Up Disaster Recovery for Oracle Fusion Middleware Components
  Learn how to set up your Disaster Recovery for an Oracle Fusion Middleware enterprise deployment.

Overview of Oracle Fusion Middleware Disaster Recovery

Learn about the problems Oracle Fusion Middleware Disaster Recovery solves and get familiar with the terminology.

This overview includes the following sections:

- Problem Description and Common Solutions
  Learn how to deploy an Oracle Fusion Middleware Disaster Recovery solution for enterprise deployments on Linux and UNIX operating systems, that make use of a replication technologies and Oracle Data Guard technologies.
- Terminology
  Learn about Disaster Recovery terminology.

Problem Description and Common Solutions

Learn how to deploy an Oracle Fusion Middleware Disaster Recovery solution for enterprise deployments on Linux and UNIX operating systems, that make use of a replication technologies and Oracle Data Guard technologies.

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, service migration, cluster 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.

In addition, 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, although Oracle recommends configuring symmetrical topology and capacity at both production and standby sites to prevent inconsistencies at the functional and performance levels. 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 inter-operates with other technologies and business applications. For instance, Oracle Fusion Middleware products such as ADF, Oracle BPEL Process Manager, Oracle Enterprise Service Bus, Oracle Web Services Manager, Adapters, Oracle Access Manager, Oracle Identity Governance, 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 different replication technologies for disaster protection of Oracle Fusion Middleware middle tier components. It uses storage level replication and it is compatible with third-party vendor recommended solutions. There are also other supported methods to replicate the FMW middle tier configuration, like DBFS or rsync that will be discussed in this document.

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

Terminology

Learn about Disaster Recovery terminology.

Disaster Recovery uses the following terms:

- **Disaster**

  A sudden, unplanned catastrophic event that causes unacceptable damage or loss in a site. 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**

  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**

  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.

- **Physical Host Name**

  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, that correspond to the IP addresses enabled on its network interfaces, which are used by clients to access it over the network. Further, each network name can be aliased with one or more alias host names.

- **Virtual Host Name**

  Virtual Host Name is a network addressable host name that can be mapped to one or more physical systems. This can be done by enabling the associated VIP in a node, through a load balancer or a hardware cluster.

  For load balancers, the term *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 systems by using the virtual host name.

  In a hardware cluster, a virtual host name 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.

  In the context of a single host, a virtual host name is an additional host name to access to the system besides the real network name. It is typically mapped to a virtual IP enabled in the node's network interfaces, or it can be mapped to an existing IP address in the system. In this last case it becomes an alias host name of the system in the name resolution system DNS or locally in the local host file.

- **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 system). 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 system 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 by 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 are 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 Setting Up Storage for more details about symbolic links.

- **Site Failover**

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**

  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**

  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**

  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 is also deployed at the standby site.

- **Standby Site Setup**

  Process of creating the standby site. To create the standby site by 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 Setting Up Storage 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.

- **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.

> **Note:**
>
> Oracle does not recommend using scaled down secondary systems. Non-symetric standbys can cause cascade falls if workloads are not handled properly and they can also produce miss configurations and data loss.

- **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 domain or Oracle Database.

- **System**

  A System is a set of targets (hosts, databases, application servers, and so on) that work together to host your applications. For example, 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 such as switchover and failover.

Setting Up Disaster Recovery for Oracle Fusion Middleware Components

Learn how to set up your Disaster Recovery for an Oracle Fusion Middleware enterprise deployment.

The following section describe the setup details:

- Oracle Fusion Middleware Disaster Recovery Architecture Overview
  Learn about the deployment architecture for Oracle Fusion Middleware components and the methods it supports to protect Oracle Fusion Middleware data and database content.
- Components Described in This Document
  Learn about the Oracle product suites that Oracle Fusion Middleware Disaster Recovery supports.

Oracle Fusion Middleware Disaster Recovery Architecture Overview

Learn about the deployment architecture for Oracle Fusion Middleware components and the methods it supports to protect Oracle Fusion Middleware data and database content.

The product binary files and configuration for Oracle Fusion Middleware components and applications get deployed in different directories on the middle tier. In addition, 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.

To protect Oracle Fusion Middleware data and database content, Oracle Fusion Middleware Disaster Recovery supports the following methods:

- Oracle Fusion Middleware product binary files, configuration files, and metadata files

  Use 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

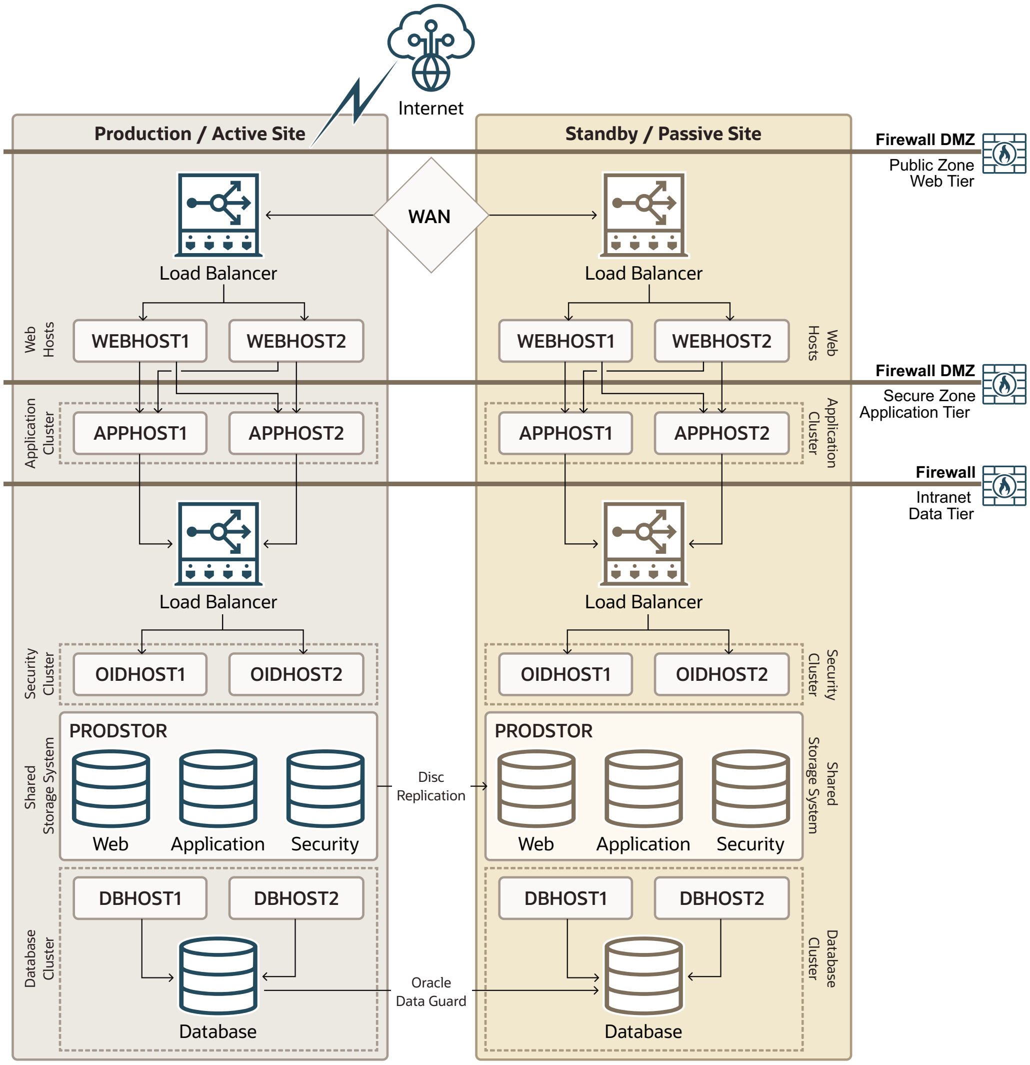

Some of the key aspects of the solution in Figure 1-1 are:

- 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 that are 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 is installed into Oracle home directories by 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 Setting Up Storage 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 Setting Up Storage 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.
- A replication technology (storage level replication, rsync, DBFS) 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. In addition, 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 Recommendations for Oracle Fusion Middleware Components for detailed Disaster Recovery recommendations for Oracle Fusion Middleware components.
- Before you force a manual synchronization, you should take a backup to capture its current state. For example, when using storage level replication, you should take a snapshot of the site. 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 needed. 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 Database Considerations.
- If your Oracle Fusion Middleware Disaster Recovery topology includes any third-party databases, use the vendor-recommended solution for those databases.
- User requests are initially routed to the production site.
- When there is a failure or planned outage of the production site, 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 or a DNS change to reroute user requests to the standby site. At this point, the standby site has assumed the production role.

Components Described in This Document

Learn about the Oracle product suites that Oracle Fusion Middleware Disaster Recovery supports.

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

- Oracle WebLogic Server

  See Recommendations for Oracle WebLogic Server for Disaster Recovery recommendations for Oracle WebLogic Server components.
- Oracle SOA Suite components:
  - Oracle SOA Service Infrastructure
  - Oracle Service Bus
  - Oracle Managed File Transfer
  - 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 Recommendations for Oracle SOA Suite for Disaster Recovery recommendations for Oracle SOA Suite components.

2
Design Considerations

Learn about the design considerations to keep in mind when you adapt an Oracle Fusion Middleware Disaster Recovery solution for your enterprise deployment.

This chapter provides instructions about how to set up an Oracle Fusion Middleware Disaster Recovery production and standby sites for the Linux and UNIX operating systems. The procedures use the Oracle SOA Suite enterprise deployment (see Figure 2-1) in the examples to illustrate how to set up the Oracle Fusion Middleware Disaster Recovery solution for that enterprise deployment. After you understand how to set up Disaster Recovery for the Oracle SOA Suite enterprise topology, use that information to set up a Disaster Recovery for your other enterprise deployments as well.

> **Note:**
>
> - You can automate disaster recovery operations such as switchover and failover, by using Oracle Site Guard. For information about the product, see Introduction to Oracle Site Guard.
> - 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.

Figure 2-1 shows the Oracle Fusion Middleware Disaster Recovery topology that uses the Oracle SOA Suite enterprise deployment at both the production site and the standby site. It 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.

While Figure 1-1 shows Oracle Fusion Middleware Disaster Recovery production and standby sites.

Figure 2-1 Deployment Used at Production and Standby Sites for Oracle Fusion Middleware Disaster Recovery

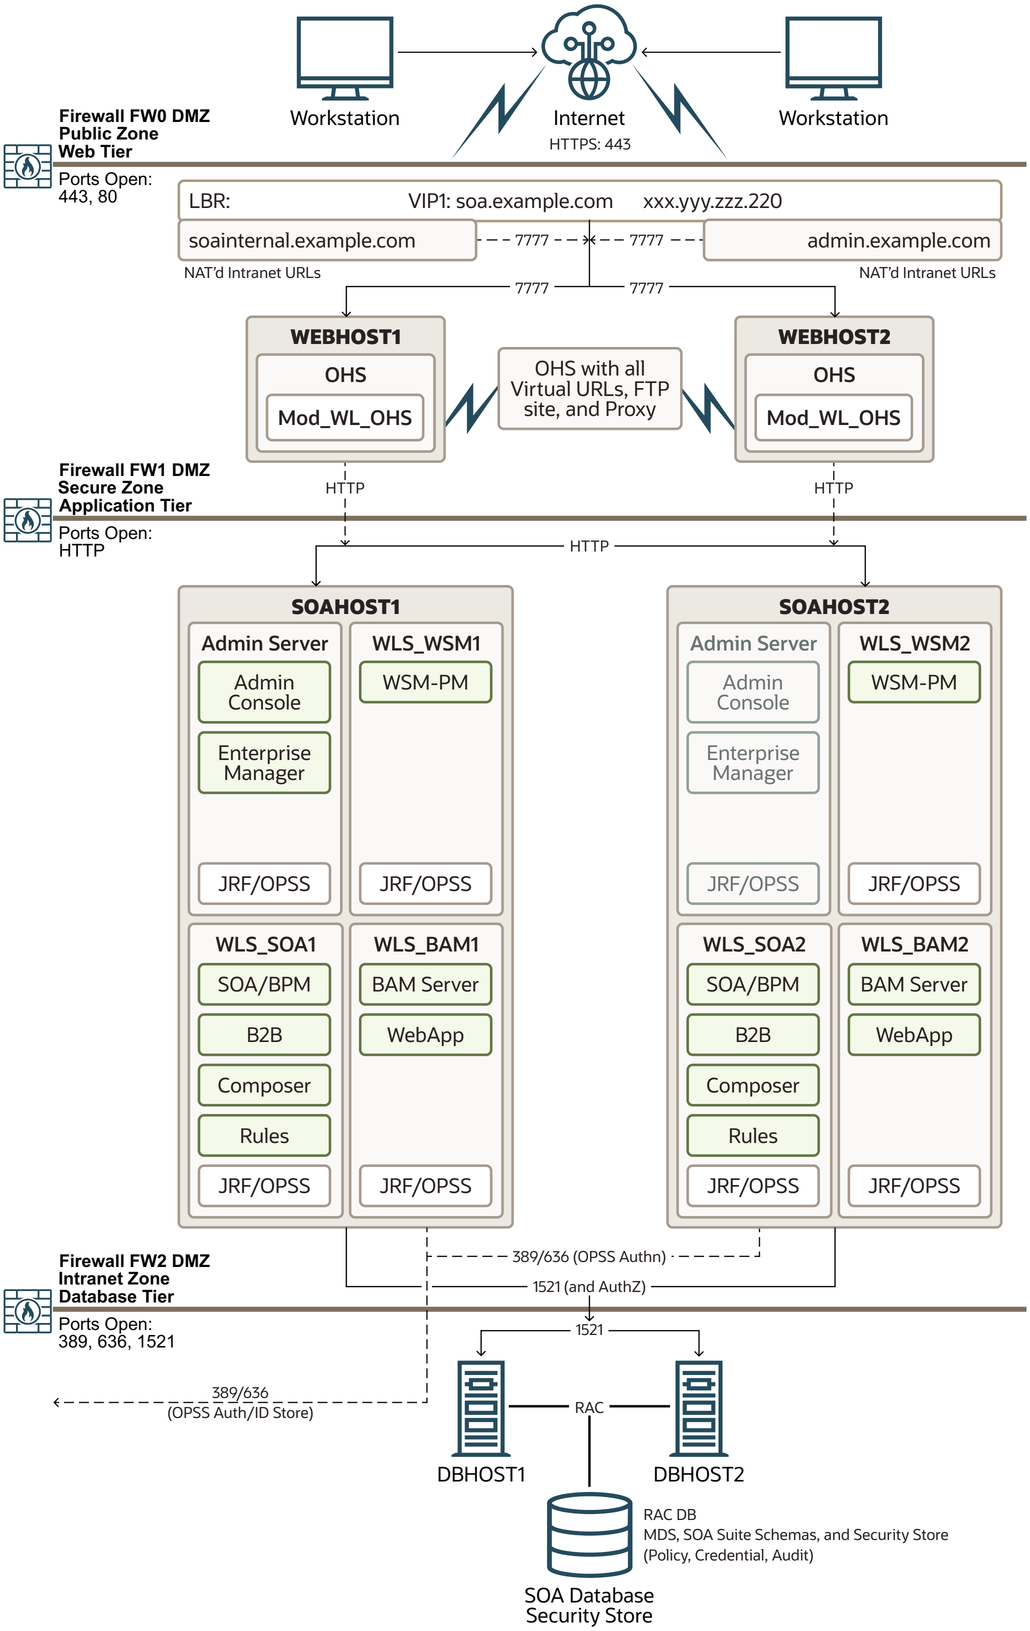

Figure 2-1 shows a diagram of the Oracle SOA, Business Process Management (BPM), and the Oracle Service Bus enterprise deployment topology. See the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for detailed information about installing and configuring an Oracle SOA Suite 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 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.
- **Security**

  The same user accounts must exist at both the production site and standby site. Also, you must configure the file system, SSL, and single sign-on 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.

This chapter includes the following topics:

- Network Considerations
  When you plan your Disaster Recovery solution, consider host names, load balance, and external clients.
- Storage Considerations
  When you design storage for your Disaster Recovery solution, consider Fusion Middleware artifacts, and storage replication.
- Database Considerations
  When you plan your Disaster Recovery solution, consider synchronizing the databases in your system with Oracle Data Guard.
- Starting Points
  When you plan your Disaster Recovery solution, you can start with an existing site or creating a new site.
- Topology Considerations
  When you plan for your Disaster Recovery solution, consider designing a symmetric or an asymmetric topology.

Network Considerations

When you plan your Disaster Recovery solution, consider host names, load balance, and external clients.

This section includes the following topics:

- Planning Host Names
  In a Disaster Recovery topology, the host name addresses used by the FMW components must be resolvable to the IP addresses of the appropriate system in each site. In production site, these addresses must be resolved to the IPs of the production hosts, and in standby site, these addresses must be resolved to the IPs of the standby site hosts.
- Virtual IP Considerations
  Starting with Oracle SOA Suite 12*c*, the SOA Suite products support Automatic Service Migration. As a result, it is no longer necessary to reserve Virtual IPs for each of the Managed Servers in the domain. Instead, a Virtual IP is required for the Administration Server only.
- Load Balancer Considerations
  In a Disaster recovery topology, both production and DR system must have a hardware load balancer, with equivalent configuration. Each load balancer will balance the traffic between the servers of its local site.
- Virtual Server Considerations
  You must configure the Virtual servers and the associated ports on the load balancer for different types of network traffic and monitoring.
- External Clients Considerations
  Systems directly accessing the servers in the topology need to be aware of the listen address that is used by the different Oracle WebLogic Server instances.
- Wide Area DNS Operations
  When a site switchover or failover is carried out, client requests must be redirected transparently to the new site that is playing the production role.

Planning Host Names

In a Disaster Recovery topology, the host name addresses used by the FMW components must be resolvable to the IP addresses of the appropriate system in each site. In production site, these addresses must be resolved to the IPs of the production hosts, and in standby site, these addresses must be resolved to the IPs of the standby site hosts.

Oracle recommends creating aliases for physical host names to isolate the real physical host names (which are different in each site) from the host names used by the FMW components (which are the same regardless the 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. If you set up an alias for the standby site, you do not need to reconfigure the host name for the host on the standby site.

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 2-1 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 you configure each component, use host-name-based configuration instead of IP-based configuration. For example, if you configure the listen address of an Oracle Fusion Middleware component to a specific IP address (such as `172.11.2.113)`, then use the host name `SOAHOST1.EXAMPLE.COM`, which resolves to `172.11.2.113`.

The following section shows how to set up host names at the Disaster Recovery production and standby sites:

> **Note:**
>
> In the examples listed, IP addresses for hosts at the initial production site have the format `172.11.x.x` and IP addresses for hosts at the initial standby site have the format `172.22.x.x`.

- Host Names for the Oracle SOA Suite Production and Standby Site Hosts
  Learn about the Oracle SOA Suite production and standby sites.

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

Learn about the Oracle SOA Suite production and standby sites.

Table 2-1 lists the IP addresses, physical host names, and aliases that are used for the Oracle SOA Suite Enterprise Deployment Guide (EDG) deployment production site hosts. Figure 2-1 shows the configuration for the Oracle SOA Suite EDG deployment at the production site.

Table 2-1 IP Addresses and Physical Host Names for SOA Suite Production Site Hosts

IP Address	Physical Host Name	Host Name Alias
172.11.2.111	prweb1.example.com	WEBHOST1
172.11.2.112	prweb2.example.com	WEBHOST2
172.11.2.113	prsoa1.example.com	SOAHOST1
172.11.2.114	prsoa2.example.com	SOAHOST2

Table 2-2 lists the IP addresses, physical host names, and aliases that are used for the Oracle SOA Suite Enterprise Deployment Guide (EDG) deployment standby site hosts.

Table 2-2 IP Addresses and Physical Host Names for SOA Suite Standby Site Hosts

IP Address	Physical Host Name	Host Name Alias
172.22.2.111	stbyweb1.example.com	WEBHOST1
172.22.2.112	stbyweb2.example.com	WEBHOST2
172.22.2.113	stbysoa1.example.com	SOAHOST1
172.22.2.114	stbysoa2.example.com	SOAHOST2

> **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. For more information about using separate DNS servers to resolve host names, see Resolving Host Names Using Separate DNS Servers.
>
> However, it is recommended to use different hostnames and same aliases to isolate the real physical host names (which are different in each site) from the host names used by the FMW components (which are the same regardless the site).

Figure 2-2 shows the physical host names that are used for the Oracle SOA Suite EDG deployment at the standby site.

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

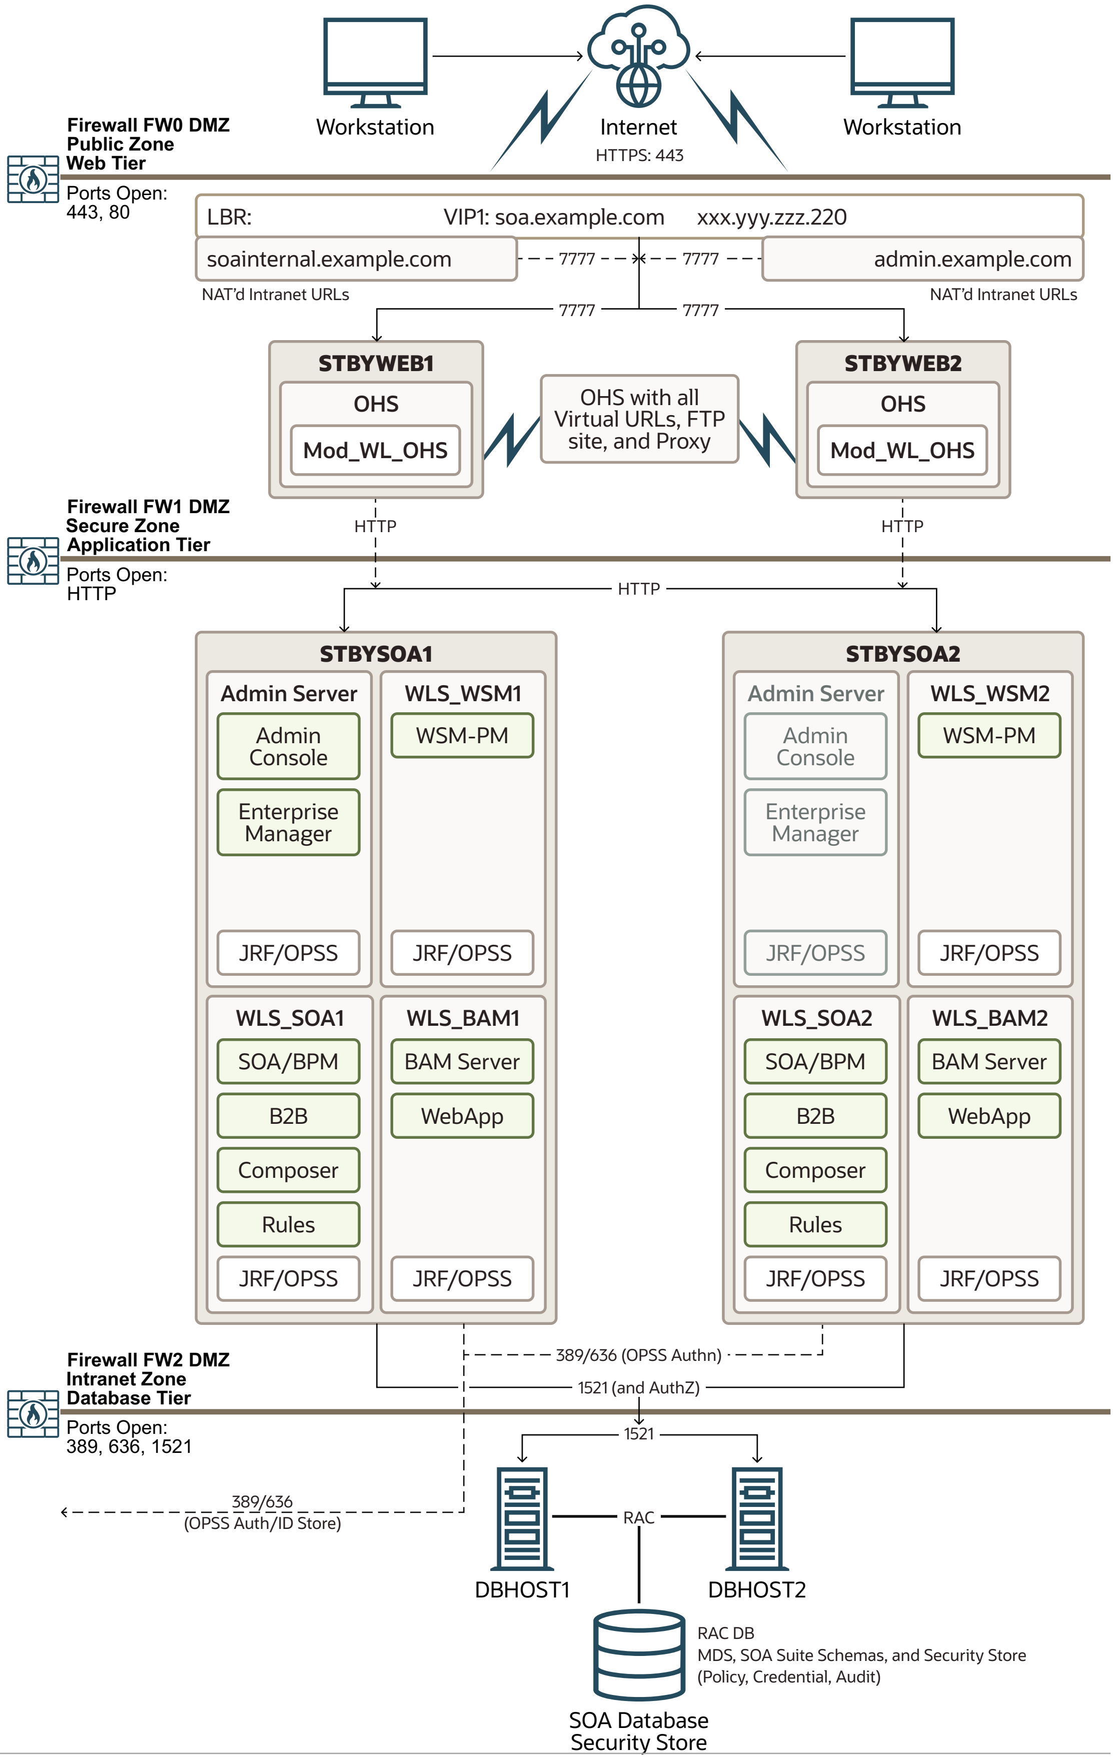

Starting Fusion Middleware version 12c, **Automatic Service Migration** is used in EDG environments as a better failover alternative to Server Migration. Contrary to Server Migration, Service Migration does not require that the managed servers use virtual IPs. Hence, only the Administration Server requires a floating IP address to be provisioned on each site. See,Table 2-3.

Ensure that you provision the floating IP addresses with the same virtual host names on the production site and the standby site.

Table 2-3 Floating IP Addresses

Virtual IP	Virtual Name	Host Name Alias

`172.11.2.134`	`prsoa-vip.example.com`
	`(in production site)`
		`ADMINVHN`

`172.22.2.134`	`stbysoa-vip.example.com`
	`(in standby site)`
		`ADMINVHN`

The following topics describe the host name resolution and testing:

- Host Name Resolution
  Host name resolution means mapping a host name to the proper IP address for communication.
- Resolving Host Names Locally
  Local host name resolution uses the host name to IP mapping that is defined in the `/etc/hosts` file of a host.
- Resolving Host Names Using Separate DNS Servers
  Use separate DNS servers to resolve host names for your Disaster Recovery topology.
- Resolving Host Names Using a Global DNS Server
  Use a global DNS server to resolve host names for your Disaster Recovery topology.
- Testing the Host Name Resolution
  Validate the host name assignment by connecting to each host at the production site and by using the `ping` command to ensure that the host can locate the other hosts at the production site.

Host Name Resolution

Host name resolution means mapping 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.

  For more information about using the `/etc/hosts` file to implement local host name file resolution, see Resolving Host Names Locally .

- Resolving host names using DNS

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

For more information about two methods of implementing DNS server host name resolution, see Resolving Host Names Using Separate DNS Servers and Resolving Host Names Using a Global DNS Server .

You must determine the method of host name resolution that you will use for your Oracle Fusion Middleware Disaster Recovery topology when you plan 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.

If you want to resolve host names locally on the host, make the `files` entry the first entry for the `hosts` parameter, as shown in Example 2-1. When `files` is the first entry for the `hosts` parameter, entries in the host `/etc/hosts` file are used first to resolve host names.

If you want to resolve host names by using DNS on the host, make the `dns` entry the first entry for the `hosts` parameter, as shown in Example 2-2. When `dns` is the first entry for the `hosts` parameter, DNS server entries are used first to resolve host names.

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.

Example 2-1 Specifying the Use of Local Host Name Resolution

```
hosts:   files   dns   nis
```

Example 2-2 Specifying the Use of DNS Host Name Resolution

```
hosts:   dns    files   nis
```

Resolving Host Names Locally

Local host name resolution uses the host name to IP mapping that is defined in the `/etc/hosts` file of a host.

When you resolve host names for your Disaster Recovery topology in this way, consider the following procedure:

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, along with the host name aliases used by the FMW components. For simplicity and ease of maintenance, Oracle recommends you to provide the same entries on all the hosts of the production site. Example 2-3 shows the `/etc/hosts` file for the production site of a SOA enterprise deployment topology.

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 host names aliases 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 2-4 shows the /etc/hosts file for the standby site of a SOA enterprise deployment topology.
4. After you set up host name resolution by 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 2-3, a ping webhost1 command on the production site returns the correct IP address (172.11.2.111) 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 2-4, a ping webhost1 command on the standby site returns the correct IP address (172.22.2.111) and it shows that the name WEBHOST1 is associated with that IP address.

Example 2-3 Making /etc/hosts File Entries for a Production Site Host

```
127.0.0.1        localhost.localdomain     localhost
172.11.2.134     prsoa-vip.example.com     prsoa-vip
ADMINVHN.EXAMPLE.COM    ADMINVHN
172.11.2.111     prweb1.example.com      prweb1      WEBHOST1.EXAMPLE.COM
WEBHOST1
172.11.2.112     prweb2.example.com      prweb2      WEBHOST2.EXAMPLE.COM
WEBHOST2
172.11.2.113     prsoa1.example.com      prsoa1      SOAHOST1.EXAMPLE.COM
SOAHOST1
172.11.2.114     prsoa2.example.com      prsoa2      SOAHOST2.EXAMPLE.COM
SOAHOST2
```

Example 2-4 Making /etc/hosts File Entries for a Standby Site Host

```
127.0.0.1         localhost.localdomain       localhost
172.22.2.134      stbysoa-vip.example.com     stbysoa-vip
ADMINVHN.EXAMPLE.COM     ADMINVHN
172.22.2.111      stbyweb1.example.com     stbyweb1     WEBHOST1.EXAMPLE.COM
WEBHOST1
172.22.2.112      stbyweb2.example.com     stbyweb2     WEBHOST2.EXAMPLE.COM
WEBHOST2
172.22.2.113      stbysoa1.example.com     stbysoa1     SOAHOST1.EXAMPLE.COM
SOAHOST1
172.22.2.114      stbysoa2.example.com     stbysoa2     SOAHOST2.EXAMPLE.COM
SOAHOST2
```

> **Note:**
>
> The subnets in the production site and standby site are different.

Resolving Host Names Using Separate DNS Servers

Use separate DNS servers to resolve host names for your Disaster Recovery topology.

The term *separate DNS servers* refers to a Disaster Recovery topology, where the production site and the standby site have separate and distinct DNS servers. When you use separate DNS servers to resolve host names for your Disaster Recovery topology, consider the following procedure:

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 2-5 shows the DNS server entries for the production site of a SOA enterprise deployment topology.
4. The DNS server entries on the standby site should have the physical and alias host names mapped to their IP addresses. Example 2-6 shows the DNS server entries for the standby site of a SOA enterprise deployment topology.
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 by using the `ping` command. For a system configured with the production site DNS entries, as shown in Example 2-5, a `ping webhost1` command on the production site returns the correct IP address (`172.11.2.111`) and indicates that the host name is fully qualified.
7. Similarly, for a system configured with the standby site DNS entries shown in Example 2-6, a `ping webhost1` command on the standby site returns the correct IP address (`172.22.2.111`) and indicates that the host name is fully qualified.

Example 2-5 DNS Entries for a Production Site Host in a Separate DNS Servers Configuration

```
PRSOA-VIP.EXAMPLE.COM    IN    A    172.11.2.134
PRWEB1.EXAMPLE.COM       IN    A    172.11.2.111
PRWEB2.EXAMPLE.COM       IN    A    172.11.2.112
PRSOA1.EXAMPLE.COM       IN    A    172.11.2.113
PRSOA2.EXAMPLE.COM       IN    A    172.11.2.114

ADMINVHN.EXAMPLE.COM    IN    A    172.11.2.134
WEBHOST1.EXAMPLE.COM    IN    A    172.11.2.111
WEBHOST2.EXAMPLE.COM    IN    A    172.11.2.112
SOAHOST1.EXAMPLE.COM    IN    A    172.11.2.113
SOAHOST2.EXAMPLE.COM    IN    A    172.11.2.114
```

Example 2-6 DNS Entries for a Standby Site Host in a Separate DNS Servers Configuration

```
STBYSOA-VIP.EXAMPLE.COM   IN    A    172.22.2.134
STBYWEB1.EXAMPLE.COM      IN    A    172.22.2.111
STBYWEB2.EXAMPLE.COM      IN    A    172.22.2.112
STBYSOA1.EXAMPLE.COM      IN    A    172.22.2.113
STBYSOA2.EXAMPLE.COM      IN    A    172.22.2.114

ADMINVHN.EXAMPLE.COM       IN    A    172.22.2.134
WEBHOST1.EXAMPLE.COM       IN    A    172.22.2.111
WEBHOST2.EXAMPLE.COM       IN    A    172.22.2.112
SOAHOST1.EXAMPLE.COM       IN    A    172.22.2.113
SOAHOST2.EXAMPLE.COM       IN    A    172.22.2.114
```

> **Note:**
>
> If you use separate DNS servers to resolve host names, then you can use the same host names for the production site hosts and standby site hosts, and you do not need to define the alias host names.

Resolving Host Names Using a Global DNS Server

Use a global DNS server to resolve host names for your Disaster Recovery topology.

The term *global DNS server* refers 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, consider the following procedure:

1. When you use 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 2-7 shows the entries for a SOA enterprise deployment topology.
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 that you have the same entries on all the hosts of the standby site. Example 2-9 shows the `/etc/hosts` file for the production site of a SOA Enterprise Deployment topology.
7. Test the host name resolution by using the `ping` command. A `ping webhost1` command on the production site returns the correct IP address (`172.11.2.111`) and indicates that the host name is fully qualified.
8. Similarly, a `ping webhost1` command on the standby site returns the correct IP address (`172.22.2.111`) and indicates that the host name is fully qualified.

Example 2-7 DNS Entries for Production Site and Standby Site Hosts when using a Global DNS Server Configuration, when the Production Site is the primary

```
PRSOA-VIP.EXAMPLE.COM       IN A 172.11.2.134
PRWEB1.EXAMPLE.COM          IN A 172.11.2.111
PRWEB2.EXAMPLE.COM          IN A 172.11.2.112
PRSOA1.EXAMPLE.COM          IN A 172.11.2.113
```

```
PRSOA2.EXAMPLE.COM         IN A 172.11.2.114

STBYSOA-VIP.EXAMPLE.COM    IN A 172.22.2.134
STBYWEB1.EXAMPLE.COM       IN A 172.22.2.111
STBYWEB2.EXAMPLE.COM       IN A 172.22.2.112
STBYSOA1.EXAMPLE.COM       IN A 172.22.2.113
STBYSOA2.EXAMPLE.COM       IN A 172.22.2.114

ADMINVHN.EXAMPLE.COM       IN A 172.11.2.134
WEBHOST1.EXAMPLE.COM       IN A 172.11.2.111
WEBHOST2.EXAMPLE.COM       IN A 172.11.2.112
SOAHOST1.EXAMPLE.COM       IN A 172.11.2.113
SOAHOST2.EXAMPLE.COM       IN A 172.11.2.114
```

Example 2-8 DNS Entries for Production Site and Standby Site Hosts when using a Global DNS Server Configuration, after a switchover to Standby Site

```
PRSOA-VIP.EXAMPLE.COM      IN A 172.11.2.134
PRWEB1.EXAMPLE.COM         IN A 172.11.2.111
PRWEB2.EXAMPLE.COM         IN A 172.11.2.112
PRSOA1.EXAMPLE.COM         IN A 172.11.2.113
PRSOA2.EXAMPLE.COM         IN A 172.11.2.114

STBYSOA-VIP.EXAMPLE.COM    IN A 172.22.2.134
STBYWEB1.EXAMPLE.COM       IN A 172.22.2.111
STBYWEB2.EXAMPLE.COM       IN A 172.22.2.112
STBYSOA1.EXAMPLE.COM       IN A 172.22.2.113
STBYSOA2.EXAMPLE.COM       IN A 172.22.2.114

ADMINVHN.EXAMPLE.COM       IN A 172.22.2.134
WEBHOST1.EXAMPLE.COM       IN A 172.22.2.111
WEBHOST2.EXAMPLE.COM       IN A 172.22.2.112
SOAHOST1.EXAMPLE.COM       IN A 172.22.2.113
SOAHOST2.EXAMPLE.COM       IN A 172.22.2.114
```

Example 2-9 Production Site /etc/hosts File Entries when using a Global DNS Server Configuration

```
127.0.0.1       localhost.localdomain     localhost
172.11.2.134    prsoa-vip.example.com     prsoa-vip    ADMINVHN.EXAMPLE.COM
ADMINVHN
172.11.2.111    prweb1.example.com       prweb1       WEBHOST1.EXAMPLE.COM     WEBHOST1
172.11.2.112    prweb2.example.com       prweb2       WEBHOST2.EXAMPLE.COM     WEBHOST2
172.11.2.113    prsoa1.example.com       prsoa1       SOAHOST1.EXAMPLE.COM     SOAHOST1
172.11.2.114    prsoa2.example.com       prsoa2       SOAHOST2.EXAMPLE.COM     SOAHOST2
```

Example 2-10 Standby Site /etc/hosts File Entries when using a Global DNS Server Configuration

```
127.0.0.1      localhost.localdomain        localhost
172.22.2.134   stbysoa-vip.example.com      stbysoa-vip    ADMINVHN.EXAMPLE.COM
ADMINVHN
172.22.2.111   stbyweb1.example.com         stbyweb1       WEBHOST1.EXAMPLE.COM
WEBHOST1
172.22.2.112   stbyweb2.example.com         stbyweb2       WEBHOST2.EXAMPLE.COM
WEBHOST2
172.22.2.113   stbysoa1.example.com         stbysoa1       SOAHOST1.EXAMPLE.COM
SOAHOST1
172.22.2.114   stbysoa2.example.com         stbysoa2       SOAHOST2.EXAMPLE.COM
SOAHOST2
```

Testing the Host Name Resolution

Validate the host name assignment by connecting to each host at the production site and by using the `ping` command to ensure that the host can locate the other hosts at the production site.

In addition, 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.

Virtual IP Considerations

Starting with Oracle SOA Suite 12*c*, the SOA Suite products support Automatic Service Migration. As a result, it is no longer necessary to reserve Virtual IPs for each of the Managed Servers in the domain. Instead, a Virtual IP is required for the Administration Server only.

In a Disaster Recovery topology, as explained in the previous section, the production site virtual IP host names aliases 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 2-1 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, by using the same ports as their counterparts at the other site.

Table 2-4 shows the virtual IP addresses and virtual host names that are used for the Oracle SOA Suite EDG deployment production site hosts. Figure 2-1 shows the configuration for the Oracle SOA Suite EDG deployment at the production site.

Table 2-4 Virtual IP Addresses and Virtual Host Names for the SOA Suite Production Site Hosts

Virtual IP Address	Virtual Host Name	Alias Host Name
`172.11.2.134`	`prsoa-vip.example.com`	`ADMINVHN`

Table 2-5 shows the virtual IP addresses, virtual host names, and alias host names that are used for the Oracle SOA Suite EDG deployment standby site hosts. Figure 2-2 shows the physical host names that are used for the Oracle SOA Suite EDG deployment at the standby site. The alias host names shown in Table 2-5 should be defined for the SOA Oracle Suite standby site hosts, as shown in Figure 2-2.

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.

For more information about using separate DNS servers to resolve host names, see Resolving Host Names Using Separate DNS Servers .

Table 2-5 Virtual IP Addresses, Virtual Host Names, and Alias Host Names for SOA Suite Standby Site Hosts

Virtual IP Address	Virtual Host Name	Host Name Alias
`172.22.2.134`	`stbysoa-vip.example.com`	`ADMINVHN`

Load Balancer Considerations

In a Disaster recovery topology, both production and DR system must have a hardware load balancer, with equivalent configuration. Each load balancer will balance the traffic between the servers of its local site.

Both primary and DR load balancer must support the wanted features of the external load balancer. For more information, see Hardware Load Balancer Requirements in *Enterprise Deployment Guide for Oracle SOA Suite* guide.

The virtual front-end name used by the production and DR load balancer must be the same. That virtual front-end name must be resolved in DNS with the IP of the load balancer of the site that has primary role in each moment.

Virtual Server Considerations

You must configure the Virtual servers and the associated ports on the load balancer for different types of network traffic and monitoring.

Configure them 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 you deal 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. When you use a single global DNS server to resolve host names, Oracle highly recommends you to create aliases for these virtual servers. Creating aliases is not required when you use separate DNS servers to resolve host names.

The virtual servers required for the various Oracle Fusion Middleware products are described in Table 2-6 and Table 2-7.

Table 2-6 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 2-7 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

External Clients Considerations

Systems directly accessing the servers in the topology need to be aware of the listen address that is used by the different Oracle WebLogic Server instances.

An appropriate host name resolution needs to be provided to the clients so that the host name 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 `ADMINVHN` aliases to correct the IP addresses for deployments to succeed.

Wide Area DNS Operations

When a site switchover or failover is carried out, 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`

This section includes the following topics:

- Using a Global Load Balancer
  A global load balancer deployed in front of the production and standby sites provides fault detection services and performance-based routing redirection for the two sites.
- Manually Changing DNS Names
  The DNS switch-over involves to manually change the name-to-IP mapping of the production site's load balancer.

Using a Global Load Balancer

A global load balancer deployed in front of the production and standby sites provides fault detection services and performance-based routing redirection for the two sites.

In addition, the load balancer can provide authoritative DNS name server equivalent capabilities.

During normal operations, you can configure the global load balancer 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.

Manually Changing DNS Names

The DNS switch-over involves to manually change the name-to-IP mapping 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 remains 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 the standby site's load balancer. It gives 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 effectively modifies 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.

If the clients that point to SOA are running on Java, another TTL property can be taken into account. Configure the DNS cache in Java for caching the successful DNS resolutions. In that case, the change in the DNS server is not refreshed until Java is restarted. This can be modified by setting the property `networkaddress.cache.ttl` to a low value:

- You can do it globally, for all the applications that are running on the JVM, by modifying the property in `JAVA_HOME/jre/lib/security/java.security file: networkaddress.cache.ttl=60`
- You can define it for a specific application only, by setting that property in the application's initialization code: `java.security.Security.setProperty("networkaddress.cache.ttl" , "60")`

Storage Considerations

When you design storage for your Disaster Recovery solution, consider Fusion Middleware artifacts, and storage replication.

This section includes the following topics:

- Oracle Fusion Middleware Artifacts
  Oracle Fusion Middleware components in a given environment are usually interdependent on one another, so it is important that the components in the topology be synchronized.
- Oracle Home and Oracle Inventory
  Oracle Fusion Middleware allows you to create multiple Oracle WebLogic Server Managed Servers from one single binary file installation.
- Setting Up Storage
  Learn about the guidelines to create volumes on a shared storage.

Oracle Fusion Middleware Artifacts

Oracle Fusion Middleware components in a given environment are usually interdependent on one another, so it is important that the components in the topology be synchronized.

This synchronization is important when you design volumes and consistency groups. Some 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.

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 and the JDBC persistent stores.
- 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.

Oracle Home and Oracle Inventory

Oracle Fusion Middleware allows you to create multiple Oracle WebLogic Server Managed Servers from one single binary file installation.

You can install binary files in a single location on a shared storage and reuse this installation by servers in different nodes. Note that, for maximum availability, Oracle recommends that you use 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 that are 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 the `ORACLE_HOME/oui/bin/attachHome.sh` file.

Setting Up Storage

Learn about the guidelines to create volumes on a 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 complete the following:

- 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.
- 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 complete the following:

- 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.
- 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.

Database Considerations

When you plan your Disaster Recovery solution, consider synchronizing the databases in your system with Oracle Data Guard.

This section provides the recommendations and considerations to set up Oracle databases that are used in an Oracle Fusion Middleware Disaster Recovery topology.

- Oracle recommends that you create 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 if the precise Oracle Fusion Middleware component supports it.

  > **Note:**
  >
  > You can use Oracle GoldenGate in an active-passive configuration only.

- The Oracle Data Guard configuration that is 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 you set up the Oracle Data Guard configuration.

- Ensure that your network is configured for low latency with sufficient bandwidth, because synchronous redo transmission can affect the response time and throughput.
- Oracle Data Guard provides three protection modes: Maximum Availability, Maximum Performance (default), and Maximum Protection. Oracle recommends to use the protection mode that better meets your availability, performance, and data protection requirements. For more information, see Oracle Data Guard Protection Modes.
- 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 recommends to use ASM as the volume manager for Oracle database files. ASM is Oracle's recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw device, and supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations.

  For additional information about RAC and ASM, see Oracle Automatic Storage Management Administrator's Guide and Real Application Clusters Installation Guide for Linux and UNIX.
- 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.

The following section explains database points:

- Setting Up DataSources in the Middle Tier
  In a Disaster Recovery topology, the following approaches can be used in the database connection string of the WebLogic datasources:

Setting Up DataSources in the Middle Tier

In a Disaster Recovery topology, the following approaches can be used in the database connection string of the WebLogic datasources:

1. Use a dataguard ready (also known as "dual") jdbc string. In this case, the db connect string includes both primary's and standby's database connect addresses, but only the database that has the primary role provides the service.

> **Note:**
>
> Provided the service name is active when the role is primary.

Example:

```
    jdbc:oracle:thin:@
    (DESCRIPTION=
    (CONNECT_TIMEOUT=15)
    (RETRY_COUNT=5)
    (RETRY_DELAY=5)
    (ADDRESS_LIST=(LOAD_BALANCE=on)(ADDRESS=(PROTOCOL=TCP)
(HOST=prmy-scan)(PORT=1521)))
    (ADDRESS_LIST=(LOAD_BALANCE=on)(ADDRESS=(PROTOCOL=TCP)
(HOST=stby-scan)(PORT=1521)))
    (CONNECT_DATA=(SERVICE_NAME=soaedg.example.com))
    )
```

Advantages:

- Since the same connect string is used in the primary and standby datasources, you need not perform modifications in the standby configuration after copying the WebLogic domain configuration from primary.

Disadvantages:

- **Delays in secondary:** The first address provided in the dual string is always tried first. When secondary system becomes primary, jdbc will first try to connect to the address that is listed in the first position. Depending how the connection is rejected (either immediately because the remote scan is not resolved, or with a delay because the remote scan name is resolved but connections are rejected by a firewall, for example), there can be delays in db connection establishment in secondary.
- **Risk of cross connections:** By default, the connections will go to the database with primary because the service will be active there. But if the standby is opened in snapshot mode and service is also active there, there would be a risk of getting connections from midtier in primary role to standby snapshot database.

This connect string is useful in the stretched cluster models. For example, as described in SOA in Best Practices for Oracle Fusion Middleware SOA 12c Multi Data Center Active-Active Deployment, where it is expected that the same SOA nodes can connect to the primary or to the secondary database, but it is not the best approach for an active-passive DR model where the cross-connection to the remote database are not expected nor recommended.

2. Another approach is to use a non-dual jdbc string, different in each site, pointing to the local database only. This approach is used for example in the SOAMP DR whitepaper.

   Example:

Connect string in datasources in primary site:

```
jdbc:oracle:thin:@
        (DESCRIPTION=
        (ADDRESS_LIST=
           (ADDRESS=(PROTOCOL=TCP)(HOST=prmy-scan)(PORT=1521)))
           (CONNECT_DATA=(SERVICE_NAME=prmy-service))
        )
```

Connect string in datasources in secondary site:

```
jdbc:oracle:thin:@
        (DESCRIPTION=
        (ADDRESS_LIST=
           (ADDRESS=(PROTOCOL=TCP)(HOST=stby-scan)(PORT=1521)))
           (CONNECT_DATA=(SERVICE_NAME=stby-service))
        )
```

> **Note:**
>
> Oracle recommends to enter the string in a single line (single line not used in these examples just for readability purpose).

Advantages:

- Given that each site points to the local database only, there is no risk of cross-connections from the mid-tier to the remote database.

Disadvantages:

- The db connection string used in the datasources is different in each site, so a replacement is required everytime that the WebLogic domain configuration is copied from primary to standby.

3. Another approach, which is the recommended in this guide, is to use a TNS alias in the datasources, as explained in Using a TNS Alias instead of a DB Connect String of the WebLogic documentation. The TNS alias is the same name in primary and secondary, hence the datasources has the same db connect string. The TNS alias is resolved with a tnsnames.ora file that is stored separately from the WebLogic domain configuration, and not replicated between sites, so you can have a different tnsnames.ora content in each site. Each site will resolve the TNS alias with the appropriate connect string in each site, pointing to the local database only.

   Example:

   Connect string in datasources in primary site:

   ```
   jdbc:oracle:thin:@soaedg
   ```

   where, `tnsnames.ora` file in primary contains:

   ```
   SOAEDG =
           (DESCRIPTION=
           (ADDRESS_LIST=
   ```

```
            (LOAD_BALANCE=ON)
            (ADDRESS=(PROTOCOL=TCP)(HOST=prmy-scan)(PORT=1521)))
            (CONNECT_DATA=(SERVICE_NAME=soaedg.example.com))
        )
```

Connect string in datasources in secondary site:

```
jdbc:oracle:thin:@soaedg
```

where, `tnsnames.ora` file in secondary contains:

```
SOAEDG =
        (DESCRIPTION=
        (ADDRESS_LIST=
            (LOAD_BALANCE=ON)
            (ADDRESS=(PROTOCOL=TCP)(HOST=stby-scan)(PORT=1521)))
            (CONNECT_DATA=(SERVICE_NAME=soaedg.example.com))
        )
```

Advantages:

- Since the same db connect string is used in the WebLogic domain config, no need to alter the WebLogic configuration after replicating the config from primary to standby.
- As each site points to the local database only, there is no risk of cross-connections from the mid-tier to the remote database.

Disadvantages:

- There are no disadvantages of using this method. It only requires a one-time setup step in order to configure the `tnsnames.ora` in each site and make the WebLogic to use it.

> **Note:**
>
> Using TNS alias in GridLink datasources is supported starting Oracle Fusion Middleware version 12.2.1.3

The TNS alias approach is used in this guide, because it provides the same advantages than the others, without any relevant disadvantage. For detailed information to configure the data sources, see Configuring Data Sources for Oracle Fusion Middleware Active-Passive Deployment.

Starting Points

When you plan your Disaster Recovery solution, you can start with an existing site or creating a new site.

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.

  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.

  Starting with a New Site 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.

This section includes the following topics:

- Starting with an Existing Site
  When you start with an existing production site, the configuration data and the Oracle binary files for the production site are already on the file system. In addition, the host names, ports, and user accounts are already defined.
- Starting with a New Site
  When you start with a new production site for an Oracle Fusion Middleware Disaster Recovery topology, consider host names and ensure that storage replication is set up to copy the configuration (based on these names) to the standby site.

Starting with an Existing Site

When you start with an existing production site, the configuration data and the Oracle binary files for the production site are already on the file system. In addition, the host names, ports, and user accounts are already defined.

When you start with an existing production site, first migrate the production site to shared storage (if not already in a shared storage), and then create a symmetric standby as described in Design Considerations for a Symmetric Topology

To migrate a production site, see the following sections:

- Migrating an Existing Production Site to Shared Storage
- Preserving the Production Hosts Hostnames as Listener Address

Migrating an Existing Production Site to Shared Storage

When you use storage replication for Oracle Fusion Middleware Disaster Recovery, the Oracle Home and middle tier configuration have to reside on the shared storage. If the production site was initially created without Disaster Recovery, the directories for the Oracle Fusion Middleware instances that comprise the site might be located on the local storage. In this scenario, the 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:

- Perform offline backup of the folder that is going to be moved to the shared storage. If the backup is performed using OS commands, it must be done as the root user and the permissions must be preserved.

  See Types of Backups and Recommended Backup Strategy in *Oracle Fusion Middleware Administering Oracle Fusion Middleware*
- Although you move the content to NFS, the path will be preserved. So the current folder that is going to be moved to an NFS (for example, `/u01/oracle/products`) will become a mount point. In preparation for this, move or rename the current folder to another path so the mount point is empty.
- Ensure that the mount point where the shared storage will be mounted exists, is empty, and has the correct ownership.
- Mount the shared storage in the appropriate mount point. The directory structure on the shared storage must be set up as described in Designing Directory Structure and Volumes. The directory structure on the shared storage must be set up as described in .
- Once the shared storage is mounted, copy the content from the backup (or from the renamed folder) to the folder, that is now in shared storage.

Example: Moving a products folder, that is in `/u01/oracle/products`, from local disk to an NFS folder.

- To backup the content the content in the local folder, a copy is performed by user root, preserving the mode:

  ```
  [root@soahost1]# cp -a /u01/oracle/products /backups/products_backup
  ```
- The current folder is moved, because the folder that will become the mount point should be empty:

  ```
  [root@soahost1]# mv /u01/oracle/products /u01/oracle/products_local
  ```
- After renaming, it is checked that the mount folder point exists (created if not), verified that it is empty, and that it has the correct ownership. The mount point is `/u01/oracle/products` in this example:

  ```
  [root@soahost1]# mkdir -p /u01/oracle/products
  [root@soahost1]# ls /u01/oracle/products
  [root@soahost1]# chown oracle:oinstall /u01/oracle/products
  ```
- The NFS volume is mounted in the mount point.

  ```
  [root@soahost1]# mount -t nfs nasfiler:VOL1/oracle/products/ /u01/
  oracle/products/
  ```
- To make this persistent, the mount is added to the mount to the `/etc/fstab` as described in Mounting the Required Shared File Systems on Each Host.
- Then, the content from the backup is copied to the mount:

  ```
  [root@soahost1]# cp -a /backups/products_backup /u01/oracle/products
  ```

Preserving the Production Hosts Hostnames as Listener Address

When a primary site is created without Disaster Protection in mind, it is possible that the FMW components are not using host name aliases as listener addresses and instead use the real physical host names of the nodes where they reside. If this is the case, you have two options:

- Modify the WebLogic domain configuration in the existing site in order to use host aliases as listener addresses for the servers, as explained in the Planning Host Names. This change has additional implications as all client points accessing the servers need to be made aware of the new hostname.
- Or, if modifying the production configuration is not feasible, preserve the configuration as it is in and continue using the physical host names as listener addresses for the FMW components and adjust secondary to use also these hostnames. This is the recommended approach. In this case, the primary physical host names must be added as aliases to the `/etc/hosts` of the standby site hosts.

  Example:

  `/etc/hosts` entries in production site hosts that do not use aliases for the components

  ```
  127.0.0.1     localhost.localdomain     localhost
  172.11.2.134  prsoa-vip.example.com     prsoa-vip
  172.11.2.111  prweb1.example.com        prweb1
  172.11.2.112  prweb2.example.com        prweb2
  172.11.2.113  prsoa1.example.com        prsoa1
  172.11.2.114  prsoa2.example.com        prsoa2
  ```

  `/etc/hosts` entries in standby site hosts

  ```
  127.0.0.1    localhost.localdomain    localhost
  172.22.2.134  stbysoa-vip.example.com  stbysoa-vip prsoa-vip.example.com 
  prsoa-vip
  172.22.2.111  stbyweb1.example.com     stbyweb1    prweb1.example.com 
  prweb1
  172.22.2.112  stbyweb2.example.com     stbyweb2    prweb2.example.com 
  prweb2
  172.22.2.113  stbysoa1.example.com     stbysoa1    prsoa1.example.com 
  prsoa1
  172.22.2.114  stbysoa2.example.com     stbysoa2    prsoa2.example.com 
  prsoa2
  ```

Starting with a New Site

When you start with a new production site for an Oracle Fusion Middleware Disaster Recovery topology, consider host names and ensure that storage replication is set up to copy the configuration (based on these names) to the standby site.

When you design a new production site, plan also the standby site, and use Oracle Universal Installer to install software on the production site. Parameters such as alias host names and software paths must be carefully designed to ensure that they are the same on both sites.

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

- Design your Oracle Fusion Middleware Disaster Recovery solution so that each host at the production site and the standby site has the desired alias host name and physical host name. For more information about host name planning, see Planning Host Names.
- 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.
- Assign ports for the Oracle Fusion Middleware installations for the production site hosts. You can also use the same ports for the standby site hosts. The same ports must be used for both sites.

  This setup is easier than checking for and resolving port conflicts between an existing production and standby sites.

Topology Considerations

When you plan for your Disaster Recovery solution, consider designing a symmetric or an asymmetric topology.

This section includes the following topics:

- Design Considerations for a Symmetric Topology
  A symmetric topology is an Oracle Fusion Middleware Disaster Recovery configuration that is identical across tiers on the production and standby sites.
- Design Considerations for an Asymmetric Topology
  An asymmetric topology is an Oracle Fusion Middleware Disaster Recovery configuration that differs across some tiers on the production and standby sites.

Design Considerations for a Symmetric Topology

A symmetric topology is an Oracle Fusion Middleware Disaster Recovery configuration that is identical across tiers on the production and standby sites.

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.

Design Considerations for an Asymmetric Topology

An asymmetric topology is an Oracle Fusion Middleware Disaster Recovery configuration that differs across some tiers on the production and standby sites.

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).

For example, consider an asymmetric topology where the standby site uses fewer Oracle Fusion Middleware instances (for example, the production site could include four Oracle Fusion Middleware instances while the standby site includes just two Oracle Fusion Middleware instances).

Another asymmetric topology includes 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.

Note:

Oracle recommends configuring symmetrical topology and capacity at both production and standby sites. Having different number of nodes or capacity can cause inconsistencies at the functional and performance levels. For example, if production has three nodes, and standby has two nodes, the **soa-infra** application may not start in standby if there is an unknown node (the additional production node does not have any equivalent node in standby site).

As a summary, having asymmetric topology is risky and generically not recommended, only applicable to special cases at customer's own risk and knowledge.

3 Setting Up and Managing Disaster Recovery Sites

Learn about the Oracle SOA Suite enterprise deployment topology that illustrates how to set up production and standby sites.

> **Note:**
>
> You can automate disaster recovery operations such as switchover and failover by using Oracle Site Guard. See *Oracle Site Guard Administrator's Guide*.

This chapter includes the following sections:

Setting Up a Site

Learn how to set up an Oracle Disaster Recovery site.

Before you start creating the production site, ensure that you:

- Set up the host name aliases for the middle tier hosts, as described in Planning Host Names.
- Create the required volumes on the shared storage on the production site, as described in Designing Directory Structure and Volumes.

- 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.

This section includes the following topics:

- Designing Directory Structure and Volumes
  Learn about the recommended directory structure in your disaster recovery topology.
- Setting Up Storage Replication
  Learn how to set up storage replication for the Oracle Fusion Middleware Disaster Recovery topology using Storage level replication, Rsync and Database file system.
- Configuring Oracle Data Guard for the FMW Database
  Learn how to install and configure Oracle Data Guard for the FMW Database in an Oracle SOA Suite enterprise deployment.

Designing Directory Structure and Volumes

Learn about the recommended directory structure in your disaster recovery topology.

You can choose a directory layout different from the one recommended in this document, but the model adopted enables maximum availability, provides the best isolation of components and symmetry in the configuration, and facilitates backup and disaster recovery.

The following 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 that are necessary to host an Oracle WebLogic Server.
- `PROD_DIR`: 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. An Oracle instance directory contains updatable files, such as configuration files, log files, and temporary files.

See Recommended Directory Structure for Oracle SOA Suite.

This section includes the following topic:

- Recommended Directory Structure for Oracle SOA Suite
  Learn about the recommended directory structures for Oracle SOA Suite.
- Recommended Volume Design for Oracle SOA Suite
  Learn about the recommended volume design for Oracle SOA Suite.

Recommended Directory Structure for Oracle SOA Suite

Learn about the recommended directory structures for Oracle SOA Suite.

Oracle Fusion Middleware allows you to create 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 that you use 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 that you use 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 that is used by the Administration Server from the domain directory that is used by Managed Servers. This allows a symmetric configuration for the domain directories that is 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 that you place the Managed Server's domain directories on a shared storage, although having them on the local file system is also supported. This is especially important when you design a production site with the disaster recovery site in mind. Figure 3-1 shows the directory structure layout for Oracle SOA Suite.

Figure 3-1 Recommended Shared Storage Directory Structure for an Enterprise Deployment

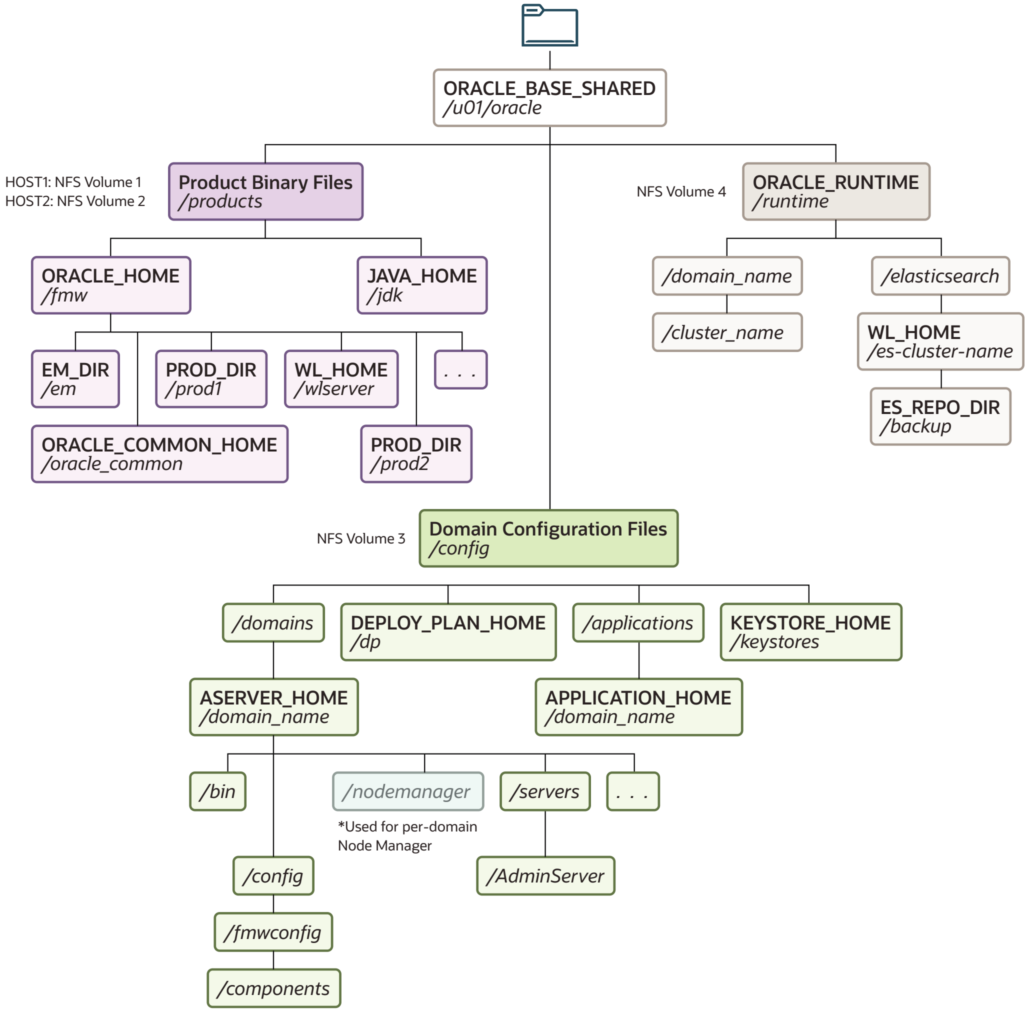

Figure 3-2 Recommended Local Storage Directory Structure for an Enterprise Deployment

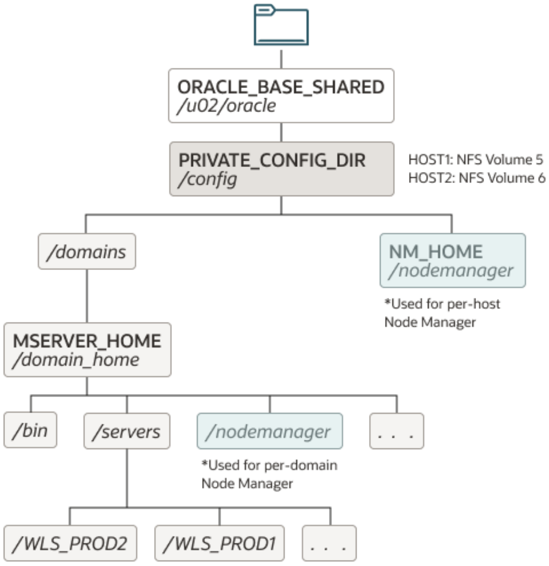

For information about setting up this directory structure, see *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

Recommended Volume Design for Oracle SOA Suite

Learn about the recommended volume design for Oracle SOA Suite.

Figure 3-3 and Figure 3-4 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 3-3 Oracle SOA Suite and Oracle Business Activity Monitoring Enterprise Topology Diagram

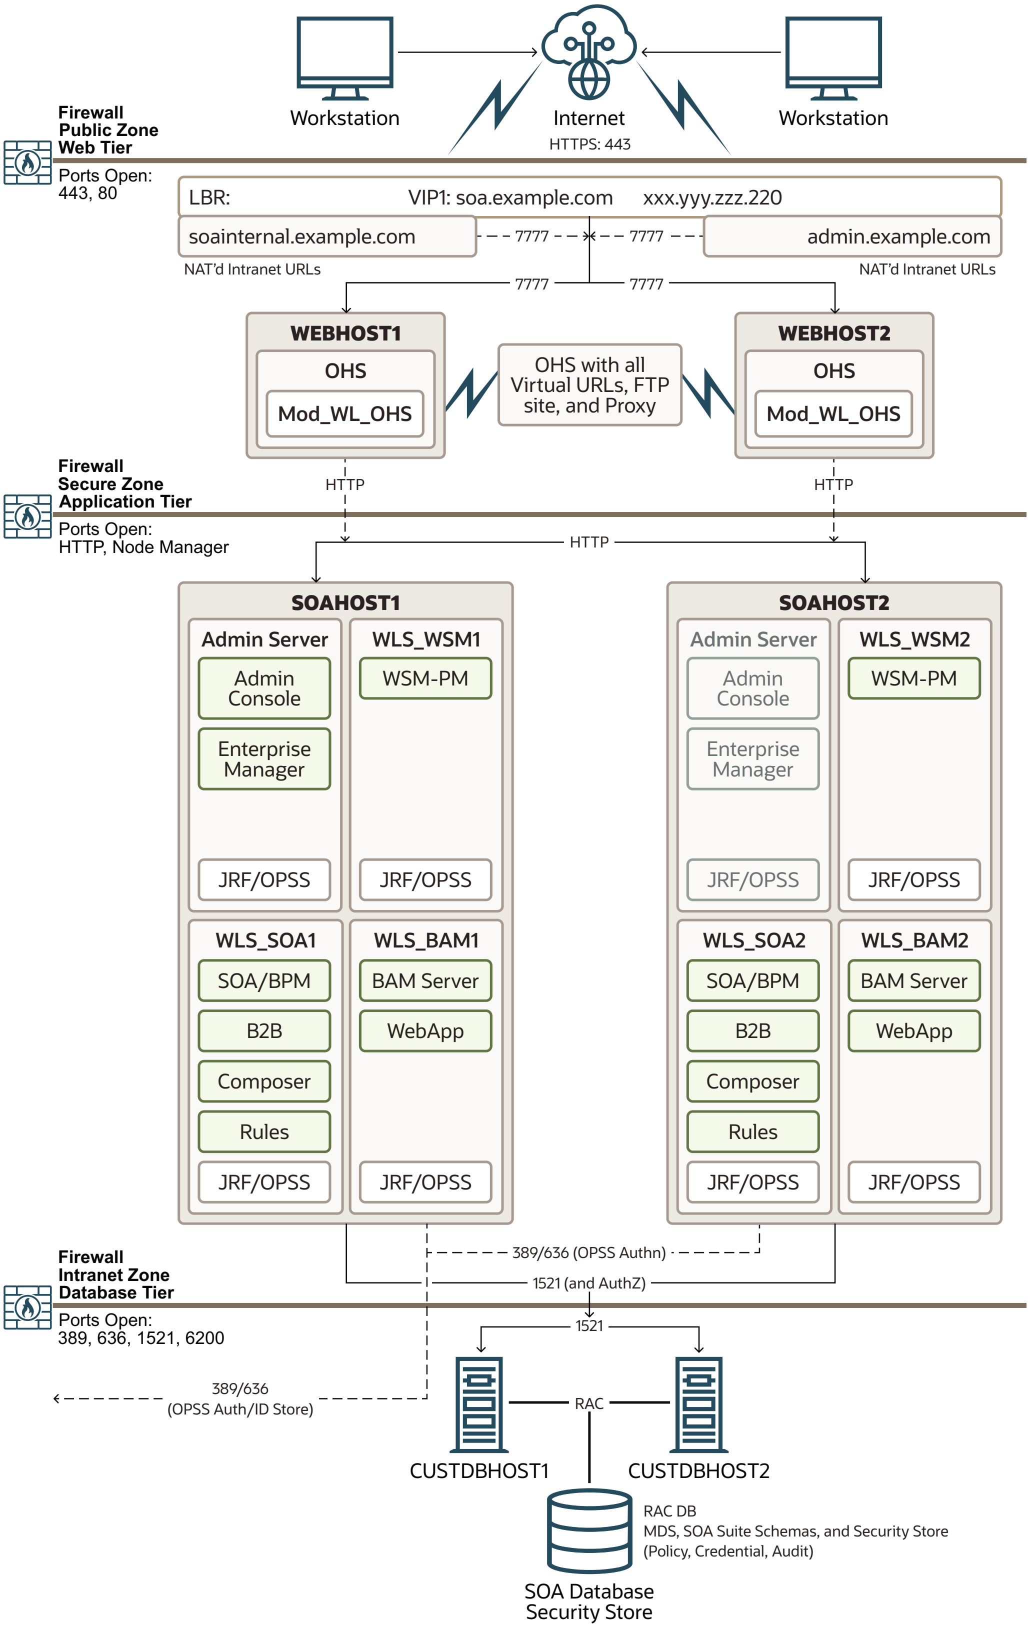

Figure 3-4 Oracle SOA Suite and Oracle Service Bus Enterprise Deployment Reference Topology Diagram

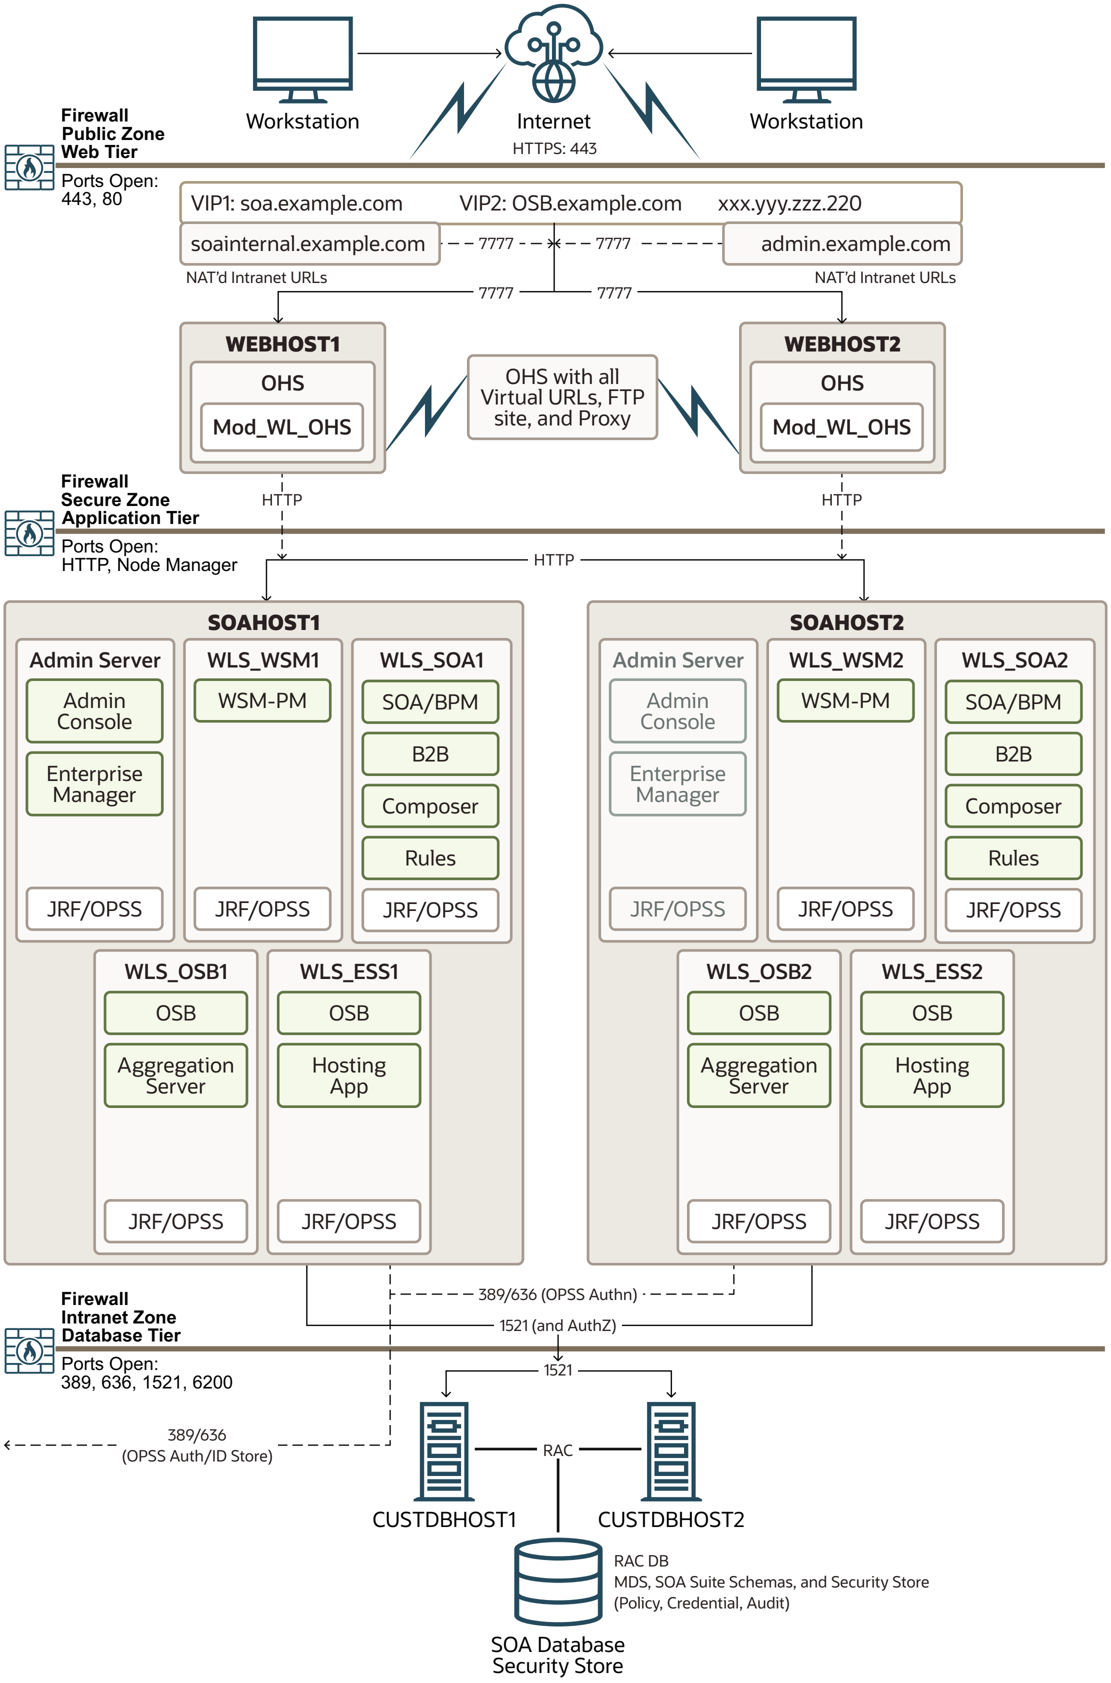

For disaster recovery of this Oracle SOA Suite topology, Oracle recommends the following volume design:

- Provision two volumes for two Oracle homes that contain the redundant product binary files (`VOLFMW1` and `VOLFMW2` in Table 3-1).
- Provision one volume for the Administration Server domain directory (`VOLADMIN` in Table 3-1).
- Provision one volume on each node for the Managed Server domain directory (`VOLSOA1` and `VOLSOA2` in Table 3-1). This directory is shared among all the Managed Servers on that node.
- Provision one volume on each node for the Oracle HTTP Server Oracle home (`VOLWEB1` and `VOLWEB2` in Table 3-1).
- Provision one volume on each node for the Oracle HTTP Server Domain Directory (`VOLOHS1` and `VOLOHS2` in Table 3-1).

> **Note:**
>
> For WebTier hosts, local storage is usually recommended. You can replicate this configuration on a regular basis to one of the other app tier volumes to sync to standby or directly from the production webhost to the standby webhost.

Table 3-1 provides a summary of Oracle recommendations for volume design for the Oracle SOA Suite topology shown in Figure 3-3 and Figure 3-4.

Table 3-1 Volume Design Recommendations for Oracle SOA Suite

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Web	VOLWEB1	WEBHOST1	/u02/oracle/products/fmw	Volume for Oracle HTTP Server installation
Web	VOLWEB2	WEBHOST2	/u02/oracle/products/fmw	Volume for Oracle HTTP Server installation
Web	VOLOHS1	WEBHOST1	/u02/oracle/config/domains/ohs_domain	Volume for Oracle HTTP Server domain directory
Web	VOLOHS2	WEBHOST2	/u02/oracle/config/domains/ohs_domain	Volume for Oracle HTTP Server domain directory
Web	VOLSTATIC1	WEBHOST1	/u02/oracle/config/static	(Optional) Volume for static HTML content
Web	VOLSTATIC2	WEBHOST2	/u02/oracle/config/static	(Optional) Volume for static HTML content
Application	VOLFMW1	SOAHOST1	/u01/oracle/products/fmw	Volume for the WebLogic Server and Oracle SOA Suite binary files
Application	VOLFMW2	SOAHOST2	/u01/oracle/products/fmw	Volume for the WebLogic Server and Oracle SOA Suite binary files

Table 3-1 (Cont.) Volume Design Recommendations for Oracle SOA Suite

Tier	Volume Name	Mounted on Host	Mount Point	Comments
Application	`VOLADMIN`	`SOAHOST1, SOAHOST2`	`/u01/oracle/config`	Volume for Administration Server domain directory and other shared configurations, such as Deployment Plans, applications, and keystores
Application	`VOLSOA1`	`SOAHOST1`	`/u02/oracle/config`	Volume for Managed Server domain directory
Application	`VOLSOA2`	`SOAHOST2`	`/u02/oracle/config`	Volume for Managed Server domain directory
Application	`VOLRUNTIME`	`SOAHOST1, SOAHOST2`	`/u01/oracle/runtime`	Volume for shared runtime content. For example, files used by file adapters, MFT transfers, and other runtime artifacts. **Note:** It is recommended to store JMS messages and TLOGS in the database, using JDBC persistent stores, instead of this folder.

For consistency group recommendations, see:

- Recommended Consistency Groups for Oracle SOA Suite
  Learn about the recommended consistency groups for Oracle SOA Suite.

Recommended Consistency Groups for Oracle SOA Suite

Learn about the recommended consistency groups for Oracle SOA Suite.

Oracle recommends the following consistency groups for the Oracle SOA Suite topology:

- Create one consistency group with the volumes that contains the domain directories for the Administration Server and Managed Servers as members (`DOMAINGROUP` in Table 3-2).
- Create one consistency group with the volume that contains the JMS file store and transaction log data as members (`RUNTIMEGROUP` in Table 3-2).
- Create one consistency group with the volume that contains the Oracle homes as members (`FMWHOMEGROUP` in Table 3-2).

Table 3-2 provides a summary of Oracle recommendations for consistency groups for the Oracle SOA Suite topology as shown in Figure 3-3.

Table 3-2 Consistency Groups for Oracle SOA Suite

Tier	Group Name	Members	Comments

Application	`DOMAINGROUP`	`VOLADMIN`
		`VOLSOA1`
		`VOLSOA2`
			Consistency group for the Administration Server, and the Managed Server domain directory

Application	`RUNTIMEGROUP`	`VOLRUNTIME`	Consistency group for the shared runtime content
Application	`FMWHOMEGROUP`	`VOLFMW1` `VOLFMW2`	Consistency group for the Oracle homes

Setting Up Storage Replication

Learn how to set up storage replication for the Oracle Fusion Middleware Disaster Recovery topology using Storage level replication, Rsync and Database file system.

- Storage Level Replication
  The Oracle Fusion Middleware Disaster Recovery solution uses storage replication technology for disaster protection of Oracle Fusion Middleware middle tier components.
- Rsync
  Alternatively, rsync can be used for replication. Rsync is a versatile copying tool, that can copy locally, or to/from another host over any remote shell. It offers a large number of options that control every aspect of its behavior, and permit very flexible specification of the set of files to be copied.
- Oracle Database File System
  Oracle Database File System (DBFS) is an additional method that can be used for replicating the configuration. Conceptually, a database file system is a file system interface placed on top of files and directories that are stored in database tables.

Storage Level Replication

The Oracle Fusion Middleware Disaster Recovery solution uses storage replication technology for disaster protection of Oracle Fusion Middleware middle tier components.

To set up storage replication for the Oracle Fusion Middleware Disaster Recovery topology:

- On the standby site, ensure that the alias host names that are created are the same as the alias host names that are used for the peer hosts at the production site.
- On the shared storage at the standby site, create the same volumes as 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:**
>
> - The symbolic links only need to be set up on the standby site if you set up symbolic links at the production site.
> - The 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 Setting Up Storage.

- It is not necessary to install the same Oracle Fusion Middleware instances at the standby site as 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 are 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 by 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 is 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 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 by using storage replication technology.

Rsync

Alternatively, rsync can be used for replication. Rsync is a versatile copying tool, that can copy locally, or to/from another host over any remote shell. It offers a large number of options that control every aspect of its behavior, and permit very flexible specification of the set of files to be copied.

Rsync implements delta-transfer algorithm, which reduces the amount of data sent over the network by sending only the differences between the source files and the existing files in the destination. Due to these advantages and easy to use ability, it is a widely used tool for backups and mirroring.

Although rsync is reliable and implements implicit retries, the network outages and other connectivity issues can still cause failures in the file synchronization. Hence, rsync can be used as an alternative to the storage level replication under these conditions:

- When the storage replication is not possible or feasible and/or does not meet the cost requirements.
- When primary and standby use a reliable and secure network connection for the copy.
- When checks are performed in the copy to make sure that it is valid.
- When the disaster recovery site is validated on a regular basis.

For more details about how to use rsync to replicate the file system artifacts, see Using Peer-To-Peer File Copy.

Oracle Database File System

Oracle Database File System (DBFS) is an additional method that can be used for replicating the configuration. Conceptually, a database file system is a file system interface placed on top of files and directories that are stored in database tables.

DBFS is similar to NFS in that it provides a shared network file system that looks like a local file system and has both a server component and a client component. The DBFS file system can be mounted from the mid-tier hosts and accessed as a regular shared file system. It could be used as an intermediate location to replicate content: any content that is copied to the DBFS mount, as it resides in the database, is automatically replicated to the standby site through the underlying Data Guard replication.

This method takes advantage of the robustness of the Data Guard replica. It has good availability through Oracle Driver’s retry logic and provides a resilient behavior. It can be used in scenarios with medium or high latencies between the data centers. However, using DBFS for configuration replication has also additional implications from the setup, database storage and lifecycle perspectives:

- It introduces some complexity due to the configuration and maintenance required by the DBFS mount. It requires the DB client to be installed in the host that is going to mount it, it requires an initial setup of some artifacts in the database (tablespace, user, and so on) and in the client (the wallet, the `tnsnames.ora`, and so on). See the script `dbfs_dr_setup_root.sh` (in Using the Application DR common scripts) as an example script that installs the database client, creates the DBFS schema in the database, configures the client artifacts, and mounts the DBFS file system in a mid-tier host.
- It requires additional capacity in the Database, because the content copied to the DBFS mount is stored in the database.
- Oracle does not recommend to store the domain configuration or the binaries directly in the DBFS mount. This would create a very strong dependency between the FMW files and the database. Instead, it is recommended to use the DBFS as an "assistance" file system: an intermediate staging file system to place the info that is going to be replicated to the standby site. Any replication to standby would have two steps: from the primary's origin folder to the intermediate DBFS mount, and then, in the standby site, from the DBFS mount to the standby's destination folder.
- The DBFS can be mounted only if the database is open. When the Data Guard is not an Active Data Guard (ADG), the standby database is in mount state. Hence, in order to access to the DBFS mount in the standby site in such cases, the database needs to be converted to snapshot standby. When ADG is used,

however, the file system can be mounted for reads and there is no need to transition to snapshot.

Due to the above stated reasons, it is not recommended to use this approach as a general purpose solution to replicate all the artifacts to the standby. For example, using DBFS to replicate the binaries is an overkill. However, this approach is suitable to replicate some dynamic artifacts during the lifecycle, like the domain shared configuration, when other methods like the storage replication or rsync are not feasible. See the Oracle WebLogic Server for Oracle Cloud Infrastructure Disaster Recovery paper to see an example on how to use DBFS to replicate the domain configuration.

Configuring Oracle Data Guard for the FMW Database

Learn how to install and configure Oracle Data Guard for the FMW Database 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 Database Considerations.

Oracle Maximum Availability Architecture (MAA) is Oracle's comprehensive architecture to reduce downtime for scheduled outages, and to prevent, detect and recover, 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.

> **Tip:**
>
> Alternatively, you can perform many of the tasks in this section by using Oracle Enterprise Manager 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 13c, see *Cloud Control Basic Installation Guide* .
>
> For more information about Setting up Oracle Data Guard using Cloud Control, see *Provisioning Oracle Standby Databases* in *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

- Prerequisites and Assumptions
- Oracle Data Guard Environment Description
- Procedure for Configuring the Data Guard
- Verifying the Data Guard Broker Configuration
- Testing Database Switchover and Switchback
  You can perform a database switchover and switchback.

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.

Oracle Data Guard Environment Description

The examples given in this section contain environment variables as described in Table 3-3.

Table 3-3 Variables Used by Primary and Standby Databases

Variable	Primary Database	Standby Database
Database names	`soa`	`soa`
SOA Database Host Names	`dbhost1.example.com`, `dbhost2.example.com`	`dbhost1stby.example.com`, `dbhost2stby.example.com`
Database unique names	`soa_pri`	`soa_stby`
Instance names	`soa1, soa2`	`soa1`, `soa2`
Service names	`soa_pri.example.com` , `soaedg.example.com`	`soa_stby.example.com` , `soaedg.example.com`

Procedure for Configuring the Data Guard

The configuration of a Data Guard involves several steps: you need to prepare the primary database with the recommended parameters, prepare the TNS aliases in primary and standby environments, create the physical standby database as a duplication of the primary database, configure the Data Guard Broker, and so on.

Oracle provides a set of sample scripts that can be used to automate most of these actions. These scripts are available here: https://github.com/oracle-samples/maa/raw/main/dg_setup_scripts/dg_setup_scripts.zip. These scripts are designed to setup a standby database for an existing primary database, using the **restore from service** feature and Data Guard Broker. They allow customization (OS user names and Oracle Homes are configurable values). They support databases with and without TDE encryption, and they work in environments with Read Only Oracle Home (ROOH) or with regular Oracle Homes. The scripts are validated in 12c (12.2), 18c, 19c and 21c RDBMS versions. Download the file and follow the instructions included in the `README.md` to configure a standby database.

If the scripts are not feasible for your specific environment, you can manually configure the Data Guard by following one of these documents:

- Creating a Physical Standby database using RMAN restore database from service (Doc ID 2283978.1)

- Creating a Physical Standby using RMAN Duplicate (RAC or Non-RAC) (Doc ID 1617946.1)

> **Note:**
>
> You can find the above documents in My Oracle Support.

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-DEC-21 17:50:45 11-DEC-21 17:50:53
         9 11-DEC-21 17:50:53 11-DEC-21 17:50:58
         10 11-DEC-21 17:50:58 11-DEC-21 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-DEC-21 17:50:45 11-DEC-21 17:50:53
        9 11-DEC-21 17:50:53 11-DEC-21 17:50:58
       10 11-DEC-21 17:50:58 11-DEC-21 17:51:03
       11 11-DEC-21 17:51:03 11-DEC-21 18:34:11

4 rows selected
```

4. Use the `show configuration` command in Data Guard Broker command line to verify that the configuration was created successfully. See Example 3-1.

Example 3-1   Verifying the Data Guard Broker Configuration

```
[oracle@dbhost1 ~]$ dgmgrl sys/'<password>'
DGMGRL for Linux: Release 19.0.0.0.0 - Production on Tue Feb 1 09:00:16 2022
Version 19.6.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.
Welcome to DGMGRL, type "help" for information.
Connected to "soa_pri"
Connected as SYSDBA.

DGMGRL> show configuration
```

```
Configuration - dg_config

Protection Mode: MaxPerformance
Databases:
soa_pri - Primary database
soa_stby - Physical standby database

Fast-Start Failover: DISABLED

Configuration Status:
SUCCESS
```

Testing Database Switchover and Switchback

You can perform a database switchover and switchback.

Performing a Switchover Operation by Using Oracle Data Guard Broker

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

1. Verify the Oracle Data Guard Broker configuration by running the following command:

   ```
   DGMGRL> show configuration;

   Configuration - dg_config

   Protection Mode: MaxPerformance
   Databases:
   soa_pri  - Primary database
   soa_stby - 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 3-1 shows how Data Guard Broker automatically shuts down and restarts the old primary database as part of the switchover operation.

   ```
   DGMGRL> switchover to 'soa_stby'
   Performing switchover NOW, please wait...
   Operation requires a connection to database "soa_stby"
   Connecting ...
   Connected to "soa_stby"
   Connected as SYSDBA.
   New primary database "soa_stby" is opening...
   Oracle Clusterware is restarting database "soa_pri" ...
   Connected to "soa_pri"
   Connected to "soa_pri"
   Switchover succeeded, new primary is "soa_stby"
   ```

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

   ```
   DGMGRL> show configuration;
   Configuration - dg_config
   Protection Mode: MaxPerformance
   Databases:
   soa_stby - Primary database
   ```

```
soa_pri  - Physical standby database
Fast-Start Failover: DISABLED
Configuration Status:
SUCCESS
```

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

Creating a Production Site on an Oracle SOA Enterprise Topology

Learn how to create a production site on an Oracle SOA enterprise deployment topology.

Before you create your production site:

- Set up the host name aliases for the middle tier hosts, as described in Planning Host Names.
- Create the required volumes on the shared storage on the production site, as described in Designing Directory Structure and Volumes.
- 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.

This section includes the following topics:

- Creating a Production Site
  This section provides information about how to create a production site for the Oracle SOA Suite topology.
- Configuring Data Sources for Oracle Fusion Middleware Active-Passive Deployment
  Configure the data sources that Oracle Fusion Middleware uses for active-passive disaster recovery.

Creating a Production Site

This section provides information about how to create a production site for the Oracle SOA Suite topology.

If you plan to create a production site for a different topology, see the appropriate Oracle Fusion Middleware Enterprise Deployment Guide listed under the *Install a Production Environment: Plan, Install & Configure an Enterprise Deployment* category.

Install and configure your production site as described in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

The following sections describe how to complete the installation and configuration of your production site:

- Creating Volumes and Consistency Groups
  Create volumes and consistency groups on the shared storage device.

- Setting Up Physical Host Names and Alias Host Names
  Set up physical host names on the production site and set up the physical host names and alias hostnames on the standby site.
- Installing and Configuring Oracle SOA Suite
  Install and configure Oracle SOA Suite.

Creating Volumes and Consistency Groups

Create volumes and consistency groups on the shared storage device.

To create volumes and consistency groups on the shared storage device, see Recommended Volume Design for Oracle SOA Suite.

Setting Up Physical Host Names and Alias Host Names

Set up physical host names on the production site and set up the physical host names and alias hostnames on the standby site.

For information about planning host names for the production and standby sites, see Planning Host Names.

Installing and Configuring Oracle SOA Suite

Install and configure Oracle SOA Suite.

To install and configure Oracle SOA Suite, see *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* and apply 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 by using the aliases to the physical and virtual host names.
3. Create SSL certificates by using the host name aliases on all of the Oracle Fusion Middleware hosts for proper Node Manager communication.

Configuring Data Sources for Oracle Fusion Middleware Active-Passive Deployment

Configure the data sources that Oracle Fusion Middleware uses for active-passive disaster recovery.

As explained in the Setting Up DataSources in the Middle Tier, Oracle recommends that you use a TNS alias in the data sources to access each site's local database. The TNS alias is the same name in production and standby. Hence, the data sources use the same db connect string. The TNS alias is resolved with a `tnsnames.ora` file stored separately from the WebLogic domain configuration and not replicated between sites. Therefore, you can have different `tnsnames.ora` content in each site. Each site will resolve the TNS alias with the appropriate connection string in each site, pointing to the local database only.

Another advantage of the active-passive deployment approach is that changes (such as retries and timeouts) to the `tnsnames.ora` file can be done without requiring a WebLogic Server restart or a WebLogic Server data source restart. At the same time, this approach reduces the repetition of addresses and settings that are typically shared by many data sources, thereby reducing the configuration creation and maintenance footprint.

If you are not already using this approach in the production system, you can use the following steps to configure it. Configure a `tns` alias in all the data sources that are used in the domain, including the data sources used by the JDBC persistence stores, leasing data sources, and custom data sources, and in the JDBC URL of the OPSS security stores.

1. Create the `tns` folder in all the middle-tier hosts.

   This folder must be readable by the Oracle user and placed in a file system that is **not replicated** between sites.

   For example:

   ```
   mkdir -p /home/oracle/tnsnames_dir
   ```

   Given that the `tns` folder is part of the configuration, you can also create it under the `config` folder that is shared by all the servers. But in that case, ensure that you exclude the `tns` folder when you copy the domain configuration from primary to standby or update `tnsnames.ora` in the standby system, after a failover or switchover, to point to the secondary database.

   > **Note:**
   >
   > Create this folder in the middle-tier hosts in both production and standby.

2. Create a `tnsnames.ora` file in the `tns` folder, with the `tns` alias that will be used in the data sources. In the production middle-tier hosts, the alias must point to the production database.

   For example:

   ```
   SOAEDG=
   (DESCRIPTION=
     (ADDRESS_LIST=
       (LOAD_BALANCE=ON)
       (ADDRESS=(PROTOCOL=TCP)(HOST=prmy-scan)(PORT=1521))
     )
     (CONNECT_DATA=(SERVICE_NAME=soaedg.example.com))
   )
   ```

   In the standby middle-tier hosts, the alias is the same name, but it is pointing to the standby database.

   For example:

   ```
   SOAEDG=
   (DESCRIPTION=
     (ADDRESS_LIST=
       (LOAD_BALANCE=ON)
       (ADDRESS=(PROTOCOL=TCP)(HOST=stby-scan)(PORT=1521))
     )
   ```

```
  (CONNECT_DATA=(SERVICE_NAME=soaedg.example.com))
)
```

If you connect to additional databases or services, ensure that you add the appropriate `tns` alias to them too.

3. Specify the `oracle.net.tns_admin` property pointing to the directory location of the `tnsnames.ora` file. Use one of the following methods:

> **Note:**
>
> Do not use a mix of these methods. It can cause unexpected behavior.

Option 1: Set the property as a data source connection property. Oracle recommends this method.

Specify the `oracle.net.tns_admin=tns_directory` property in the data source configuration. To specify this property in the WebLogic Administration Console, go to **Services**, click **Data Sources**, select a data source from the list, click **Connection Pool**, and then add it to the **Properties** text box. Repeat this step for each data source.

For example:

Add the following property to the **Properties** text box:

```
oracle.net.tns_admin=/home/oracle/tnsnames_dir
```

You must also specify this property in the OPSS security stores files `jps-config-jse.xml` and `jps-config.xm` available in the `$ASERVER_HOME/config/fmwconfig` folder. To modify these `jps` files, edit them and add the `oracle.net.tns_admin` property after the `jdbc.url` property.

For example:

```
…
<property name="jdbc.url" value="jdbc:oracle:thin:@soaedg"/>
<property name="oracle.net.tns_admin" value="tns_directory"/>
…
```

> **Note:**
>
> This property applies to the specific file (data source, `jps` file) in which it is specified.

Option 2: Set the property as a java system property.

Specify the `-Doracle.net.tns_admin=tns_directory` system property where `tns_directory` is the directory location of the `tnsnames.ora` file. To set it as a java property for the servers, edit the following files:

- `$ASERVER_HOME/bin/setUserOverrides.sh`

- `$MSERVER_HOME/bin/setUserOverrides.sh` (This file is not shared. Therefore, you should edit the file in all the SOA mid-tier hosts.)

Add the following content to these files:

```
# For using tns alias in the datasources
export EXTRA_JAVA_PROPERTIES="${EXTRA_JAVA_PROPERTIES} -
Doracle.net.tns_admin=/home/oracle/tnsnames_dir
```

> **Note:**
>
> - This property applies to all the data sources and `jps` files in WebLogic Server.
> - Before you run some of the WLST commands and the Configuration Wizard, this approach requires that you set the property in the environment. Before running WLST, you should set the property in the *WLST_PROPERTIES* environment variable. Before running the Configuration Wizard, you should add the property to the *JVM_ARGS* environment variable of the `config_internal.sh` script.

Option 3: Set the property in the `jdbc` URL.

Specify the location of the `tnsnames.ora` file as part of the connection string in the data sources and `jps` files:

```
jdbc:oracle:thin:@alias?TNS_ADMIN=tns_directory
```

> **Note:**
>
> - This property applies to the specific file (data source, `jps` file) in which it is specified.
> - You can use this method with JDBC Driver 18.3 and later. This applies to Fusion Middleware 12.2.1.4 (which uses JDBC Driver 19.3) and later.

4. Modify the URL defined in the data sources by replacing the connection string with the alias. The following is a sample JDBC URL using the `tns` alias `jdbc:oracle:thin:@soaedg`.

   You can use the WebLogic Administration Console to perform this change. Alternatively, you can perform the modification directly in the files, as described here:

   a. Sign in to `APPHOST1`.

   b. Take a backup of `$ASERVER_HOME/config/jdbc` which contains the data source config files.

   c. Run a command to replace the previous database connection string by the new one that uses the `tns` alias.

For example:

```
cp -rf $ASERVER_HOME/config/jdbc   $ASERVER_HOME/config/jdbc_bck
export
PREV_STRING='(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)
(HOST=prmy-scan)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=soaedg.example.com)))'
export NEW_STRING='soaedg'
cd $ASERVER_HOME/config/jdbc
find . -name '*.xml' | xargs sed -I 's|'${PREV_STRING}'|'$
{NEW_STRING}'|gI'
```

5. Update the JDBC URL of the OPSS security stores too by using the following steps:
   1. Sign in to `APPHOST1` and go to the `$ASERVER_HOME/config/fmwconfig` folder.
   2. Take a backup of the `jps-config-jse.xml` and `jps-config.xml` files.
   3. Edit both files to update the value of the `property name="jdbc.url"` with the appropriate JDBC URL used in the data sources.
      For example:

```
cp -rf $ASERVER_HOME/config/fmwconfig $ASERVER_HOME/config/
fmwconfig_bck
cd $ASERVER_HOME/config/fmwconfig
export
PREV_STRING='(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)
(HOST=prmy-scan)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=soaedg.example.com)))'
export NEW_STRING='soaedg'
find . -name '*.xml' | xargs sed -i 's|'${PREV_STRING}'|'$
{NEW_STRING}'|gI'
```

6. Restart all the WebLogic Servers in the domain for the changes to take effect.
   1. Stop all the WebLogic Servers in the domain (Administration Server and Managed Servers).
   2. Start the Administration Server in the domain.
   3. Start the Managed Servers.
7. Verify that the data source connections are established correctly with the database.

   Verify that the JDBC URL is updated correctly in the OPSS store. To verify this:
   1. Go to the **Enterprise Manager Console**.
   2. Navigate to **WebLogic Domain**, select **Security**, and then click **Security Provider Configuration**.
   3. Expand **Security Stores** and verify that the Database URL is updated.

Note:

Regarding the ONS host and port configuration in the data sources, these values are not required when you are using Oracle Database 12c and later because the ONS list is automatically obtained from the database by the client driver.

As per Oracle's consistent recommendation in this guide, use this feature instead of providing the scan address or the list of the RAC nodes in the ONS configuration of each data source.

Ensure that Fast Application Notification (FAN) is enabled and that the **ONS nodes** property is empty in each data source. To check this property in the WebLogic Administration Console, go to **Services**, click **Data Sources**, select **Configuration**, and then click **ONS**.

Creating a Standby Site

Learn how to create a standby site.

To create the standby site, the Oracle SOA enterprise deployment topology is used as an example.

This section includes the following topics:

- Preparing the Standby Site
  Prepare your standby site for operation.
- Validating the Standby Site Setup
  Validate your standby site.

Preparing the Standby Site

Prepare your standby site for operation.

To prepare it for operation, on your standby site:

- Set up the correct alias host names and physical host names by following the instructions in Planning Host Names.

  Ensure that each standby site host has an alias host name that is the same as the physical host name of its peer host at the production site.
- Create, on the shared storage, the same volumes that were created on the shared storage at the production 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 Setting Up Storage.
- Setting Up Middle Tier Hosts
  Middle tier hosts on a standby site do not require installation nor configuration of 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 is replicated at the standby site.

- About Self-Signed certificates and Keys on the Standby Site
  Learn about certificates and keystores.

Setting Up Middle Tier Hosts

Middle tier hosts on a standby site do not require installation nor configuration of 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 is replicated at the standby site.

To set up the middle tier hosts on the standby site:

1. Create a baseline snapshot copy of the shared storage on the production site, which sets up the replication between the storage devices. Create the initial baseline copy and subsequent snapshot copies by using asynchronous replication mode.
2. Synchronize the shared storage at the production site with the shared storage at the standby site. This transfers 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 is 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 content as the directories inside the production site volumes.

About Self-Signed certificates and Keys on the Standby Site

Learn about certificates and keystores.

As recommended in Planning Host Names the FMW components use alias host names, that are resolvable to the IP addresses of the appropriate system in each site. If the primary self-signed certificates are created with these host names, the certificates are valid both for the production and for the standby systems.

For more information about creating primary self-signed certificates, see Common Configuration and Management Tasks for an Enterprise Deployment in *Enterprise Deployment Guide for Oracle SOA Suite*.

The certificates and keystores are replicated to standby along with the configuration, so there is no need to create specific self-signed certificates or keystores for the standby site.

Validating the Standby Site Setup

Validate your standby site.

To validate a standby site:

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.

Performing Site Operations and Administration

Learn how to operate and administer your Oracle Fusion Middleware Disaster Recovery topology.

This section includes the following topics:

- Synchronizing the Production and Standby Sites
  Learn how to force a synchronization of the production and standby sites when you introduce a change in the middle tier at the production site.
- Performing a Switchover
  A switchover operation sets the standby site as the production role.
- Performing a Switchback
  A switchback operation reverts the roles of the current production and standby sites.
- Performing a Failover
  A failover operation sets the standby site as the production role when the production site becomes unavailable.
- Testing the Standby Site
  Learn how to create a clone of the read-only standby site shared storage and use it for converting database `secondary_db_unqname` to snapshot standby.
- Using Peer-To-Peer File Copy
  The `rsync` utility (which uses peer-to-peer file copy) can be used 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. The use of the `rsync` utility is explained in the context of symmetric topologies.

Synchronizing the Production and Standby Sites

Learn how to force a synchronization of the production and standby sites when you introduce a change in the middle tier at the production site.

In normal operations, 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 includes all the data up to the data contained in the last snapshot transferred from the production site before the failover or switchover.

Be sure to force a synchronization when you introduce a change to the middle tier at the production site such as, for example, when you deploy a new application at the production site. Follow the vendor-specific instructions to force a synchronization by using the storage replication technology.

The databases synchronization in an Oracle Fusion Middleware Disaster Recovery topology is managed by Oracle Data Guard.

Performing a Switchover

A switchover operation sets the standby site as the production role.

This operation is needed when you plan to take down the production site (for example, to perform maintenance) and to make the current standby site the production site.

To perform a switchover operation:

1. Shut down any processes 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 is 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.

At this point, the former standby site becomes the new production site, and you can perform maintenance at the original production site. After you have carried out the maintenance of the original production site, you can use it as either the production site or the standby site.

Note:

This note is applicable for BI-specific systems only.

After a switchover operation, the creation of an Essbase cube with CDS may fail with an error similar to the following:

```
oracle.essbase.cds.util.CDSException:
oracle.essbase.cds.util.CDSException: java.sql.SQLException: ORA-25153:
```

To work around this issue and to create an Essbase cube, on the current primary database:

- Identify the temporary tablespaces by using a select statement similar to the following (where `BIS17V1` is the Oracle Business Intelligence RCU prefix):

  ```
  select username,temporary_tablespace from dba_users where username
  like 'BIS17V1%'
  ```

  Assume that the above command returns the following list of temporary tablespaces:

  ```
  USERNAME.....TEMPORARY_TABLESPACE

  BIS17V1_IAU_VIEWER.....BIS17V1_IAS_TEMP

  BIS17V1_STB.....BIS17V1_IAS_TEMP

  BIS17V1_IAU_APPEND.....BIS17V1_IAS_TEMP

  BIS17V1_MDS.....BIS17V1_IAS_TEMP

  BIS17V1_IAU.....BIS17V1_IAS_TEMP

  BIS17V1_BIPLATFORM......BIS17V1_IAS_TEMP

  BIS17V1_OPSS.....BIS17V1_IAS_TEMP
  ```

- After the switchover, drop the tablespace `BIS17V1_IAS_TEMP` including contents and datafiles.
- Create the temporary tablespace `BIS17V1_IAS_TEMP`, as a tempfile, in the location (for example) `/work/primy/oradata/stnby/BIS17V1_IAS_TEMP.dbf`, with size 250 m.
- Issue the following alter commands (here is where you use the list temporary tablespaces):

  ```
  alter user BIS17V1_OPSS temporary tablespace  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_BIPLATFORM temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_IAU temporary tablespace  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_MDS temporary tablespace  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_IAU_APPEND temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_STB temporary tablespace  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_IAU_VIEWER temporary tablespace
  BIS17V1_IAS_TEMP ;
  ```

To use the original production site as the production site, perform a switchback as explained in Performing a Switchback.

Performing a Switchback

A switchback operation reverts the roles of the current production and standby sites.

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 is 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.

Performing a Failover

A failover operation sets the standby site as the production role when the production site becomes unavailable.

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 is 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.

Note:

After a failover operation, the creation of an Essbase cube with CDS may fail with an error similar to the following:

```
oracle.essbase.cds.util.CDSException:
oracle.essbase.cds.util.CDSException: java.sql.SQLException:
ORA-25153:
```

To work around this issue and to create an Essbase cube, on the current primary database:

- Identify the temporary tablespaces by using a select statement similar to the following, where `BIS17V1` is the Oracle Business Intelligence RCU prefix):

  ```
  select username,temporary_tablespace from dba_users where
  username like 'BIS17V1%'
  ```

  Assume that the above command returns the following list of temporary tablespaces:

  ```
  USERNAME.....TEMPORARY_TABLESPACE

  BIS17V1_IAU_VIEWER.....BIS17V1_IAS_TEMP

  BIS17V1_STB.....BIS17V1_IAS_TEMP

  BIS17V1_IAU_APPEND.....BIS17V1_IAS_TEMP

  BIS17V1_MDS.....BIS17V1_IAS_TEMP

  BIS17V1_IAU.....BIS17V1_IAS_TEMP

  BIS17V1_BIPLATFORM......BIS17V1_IAS_TEMP

  BIS17V1_OPSS.....BIS17V1_IAS_TEMP
  ```

- After the failover, drop the tablespace `BIS17V1_IAS_TEMP` including contents and datafiles.
- Create the temporary tablespace `BIS17V1_IAS_TEMP`, as a tempfile, in location. For example, `/work/primy/oradata/stnby/BIS17V1_IAS_TEMP.dbf` with size 250 m.
- Issue the following alter commands (here is where you use the list temporary tablespaces):

  ```
  alter user BIS17V1_OPSS temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_BIPLATFORM temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_IAU temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_MDS temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_IAU_APPEND temporary tablespace
  BIS17V1_IAS_TEMP ;

  alter user BIS17V1_STB temporary tablespace
  BIS17V1_IAS_TEMP ;
  ```

```
        alter user BIS17V1_IAU_VIEWER temporary tablespace
        BIS17V1_IAS_TEMP ;
```

To again use the original production site as the production site, perform a switchback as explained in Performing a Switchback.

Testing the Standby Site

Learn how to create a clone of the read-only standby site shared storage and use it for converting database `secondary_db_unqname` to snapshot standby.

A typical Oracle Fusion Middleware Disaster Recovery configuration uses:

- Storage replication 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 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.
- The standby site as the production role when the production site becomes unavailable. If the current production site becomes unavailable unexpectedly, then a failover operation (described in Performing a Failover) is performed to enable the standby site to assume the production role. Or, if the current production site is taken down intentionally (for example, for planned maintenance), then a switchover operation (described in Performing a Switchover) is performed to enable the standby site to assume the production role.

The usual method of testing a standby site is to shut down the current production site and perform a switchover operation to enable the standby site to assume the production role.

However, some enterprises may want to perform periodic testing of their Disaster Recovery standby site without shutting down the current production site and a complete switchover operation. This is possible by converting the standby database to snapshot standby. This allows the standby servers to be started in the standby site and verify the secondary system. Any change performed in the standby site database while it is in snapshot standby mode will be discarded once it is converted to physical standby again, so primary data will not be affected by secondary site validations.

To use this testing method:

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. Convert the standby database into snapshot standby by the following steps:
    a. Use Data Guard broker in primary database host to convert the secondary database to snapshot standby.

       Example:

       ```
       [oracle@primarydbhost~]$dgmgrl sys/
       your_sys_password@primary_db_unqname
       DGMGRL> convert database "secondary_db_unqname" to snapshot
       standby
       ```

    b. . Use `show configuration` command to verify that the conversion has been correctly performed.
3. 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.
4. Before you test the standby site, modify the host name resolution method for the computers that are 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.
5. Perform the standby site testing.

> **Note:**
>
> This operation must be done with caution in SOA environments: if there are pending messages or composites in the database when it is converted into snapshot, the standby site's SOA servers will process them when they start. Check that there are no pending actions in primary database when converting to snapshot standby, otherwise, remove records from runtime SOA tables in the standby database after it is converted to snapshot standby database and before starting the secondary site's SOA servers. See, Removing Records from the Runtime Tables Without Dropping the Tables.

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. Convert the standby database into snapshot standby by the following steps:

a. Use Data Guard broker in primary database host to convert the secondary database to physical standby again.

   Example:

```
[oracle@primarydbhost~]$dgmgrl sys/
your_sys_password@primary_db_unqname
DGMGRL> convert database “secondary_db_unqname” to physical standby
```

b. Use `show configuration` command to verify that the conversion has been correctly performed.

3. Before you use the original production site again, modify the host name resolution method for the computers that are 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.

Using Peer-To-Peer File Copy

The `rsync` utility (which uses peer-to-peer file copy) can be used 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. The use of the `rsync` utility is explained in the context of symmetric topologies.

For information about the conditions for using the `rsync` in Disaster Recovery environments, see the point `rsync` in the section Setting Up Storage Replication of this document.

Ensure that you are familiar with storage replication and Oracle Data Guard in an Oracle Fusion Middleware Disaster Recovery topology, because there are many similarities between using storage replication and using the `rsync` utility for disaster protection and disaster recovery of your Oracle Fusion Middleware components.

Note the following important differences between using storage replication technologies and using the `rsync` utility to replicate middle tier file systems:

- When you use storage replication, you can roll changes back to the point in time when any previous snapshot was taken at the production site.

  When you use the `rsync` utility, replicated production site data overwrites the standby site data, and you cannot roll back a replication.

- When you use 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.

  When you use the `rsync` utility, data consistency is not guaranteed.

This section includes the following topic:

- Using rsync and Oracle Data Guard in Oracle Fusion Middleware Disaster Recovery Topologies
  Learn how to use the UNIX `rsync` utility and Oracle Data Guard in your Oracle Fusion Middleware Disaster Recovery topology.

Using rsync and Oracle Data Guard in Oracle Fusion Middleware Disaster Recovery Topologies

Learn how to use the UNIX `rsync` utility and Oracle Data Guard in your Oracle Fusion Middleware Disaster Recovery topology.

> **Note:**
>
> For information about how to set up Oracle Data Guard for Oracle database, see Database Considerations.

The following sections describe how to use the `rsync` utility and Oracle Data Guard to protect and force synchronization between your production and standby sites in an Oracle Fusion Middleware Disaster Recovery topology:

- Using rsync for Oracle Fusion Middleware Middle Tier Components
  Use the UNIX `rsync` utility to protect and recover your Oracle Fusion Middleware middle tier components.
- Performing Failover and Switchover Operations
  Learn how to perform failover or switchover operations when you use the `rsync` utility.

Using rsync for Oracle Fusion Middleware Middle Tier Components

Use the UNIX `rsync` utility to protect and recover your Oracle Fusion Middleware middle tier components.

To use the `rsync` utility:

1. Install and set up `rsync` to enable replication of files from a production site host to its standby site peer host. For instructions about how to install, set up, and use the utility, see the utility man pages and also `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, 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.
   - The shared config folder, such as the WebLogic Administration Server domain directory, deployment plans, applications, keystores, and so on. As this folder is shared between the mid-tier hosts and you do not have to have a for each host.
   - The private config folders, such as the WebLogic managed server's domain, local node manager home on the host.
   - If applicable, the shared runtime folder.

- 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 files that are created by Oracle Forms on the host, and the `.rdf` deployment artifact files that are created by Oracle Reports on the host.

> **Note:**
>
> Run the `rsync` utility as root or as the OS user that is the owner of the files. If you want it 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.
>
> For an example of using `rsync` utility to copy folders to a remote node, see `rsync_copy_and_validate.sh` generic script (available here https://github.com/oracle-samples/maa/raw/main/hybrid_dr/hybrid_dr_rsync_scripts/hybrid_dr_rsync_scripts.zip). The `rsync_copy_and_validate.sh` copies a folder to a remote node, and performs a checksum validation of the copy. The zip file also includes a few examples to use this generic script for copying the typical FMW folders.

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), perform a manual synchronization of that host with its standby site peer host by using `rsync`.
5. Whenever you use `rsync` to manually synchronize an Oracle Fusion Middleware middle tier instance on a production site host with the peer standby site host, also use Oracle Data Guard to synchronize database repositories in the production site with the databases in the standby site.

Performing Failover and Switchover Operations

Learn how to perform failover or switchover operations when you use the `rsync` utility.

To perform a failover or switchover from the production site to the standby site when you use `rsync`:

1. Shut down any processes running on the production site (if applicable).
2. Stop `rsync` jobs between the production site hosts and standby site peer hosts.
3. Use Oracle Data Guard to failover or swithcover 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 `rsync` between the two sites, but configure it so that replications go now 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).

Using Oracle Site Guard for Disaster Recovery

Oracle Site Guard is a disaster-recovery solution that enables administrators to automate complete site switchover or failover. It orchestrates the coordinated failover of Oracle Fusion Middleware, Oracle Fusion Applications, and Oracle Databases.

Oracle Site Guards offers the following benefits:

- Fully automate disaster recovery operations and launch them with a single click
- Minimizes disaster-recovery time
- Reduces human errors
- Flexible and customizable
- Eliminates the need for special skills
- Use a single pane of glass to manage disaster recovery
- Assure disaster recovery readiness using on-demand or scheduled disaster recovery drills

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

Patching an Oracle Fusion Middleware Disaster Recovery Site

Apply an Oracle Fusion Middleware patch set to upgrade the Oracle homes that participate in an Oracle Fusion Middleware Disaster Recovery site.

It is assumed 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.

To apply an Oracle Fusion Middleware patch:

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 you apply 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 you apply 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 is 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.

When patching the database, check the specific patch’s documentation on how to apply the patch in a Data Guard topology.

Accessing and Managing T3/T3s External Clients in DR Scenarios

This section discusses different approaches for accessing and managing external T3/T3s clients in the context of a Disaster Recovery scenario.

> **Note:**
>
> This section does not apply to the internal T3/T3s clients, which run in the same domain as the T3/T3s servers. When connecting to a cluster in the same domain, internal clients can utilize the T3/T3s cluster syntax, such as `cluster:t3://cluster_name`. Only when the clusters are in the same domain can you utilize this cluster syntax.

WebLogic uses the T3/T3s protocol for Remote Method Invocation (RMI). It is used by several WebLogic services, such as JMS, EJB, JTA, JNDI, JMX, WLST, and so on. The external T3/T3s clients, which are those that do not run in the same domain as the T3/T3s servers, can use different ways to connect to a WebLogic cluster. For more information, see Using WebLogic RMI with T3 Protocol.

External T3/T3s clients **can connect directly to WebLogic servers' channels (default or custom)** that listen to the T3/T3s requests. The connection URL configured in the client's provider property contains the list of all the servers and ports of the cluster.

For example:

```
t3://host1.example.com:9073,host2.example.com:9073
```

The external T3/T3s clients can access to T3/T3s services through TCP **Load Balancer** (LBR) too. In this scenario, the client’s provider URL points to the load balancer service and port, and the requests are load balanced to the WebLogic Server's T3/T3s ports. In the initial contact for the JNDI context retrieval, the external clients connect to a WebLogic Server through the load balancer and download the stubs. These stubs contain the connect information that is used for the subsequent requests.

In general, the WebLogic T3/T3s is TCP/IP-based, so it can support TCP load balancing when services are homogeneous, such as JMS and EJB. For example, a JMS front-end can be configured in a WebLogic cluster in which remote JMS clients can connect to any cluster member. By contrast, a JTA subsystem cannot safely use TCP load balancing in transactions that span across multiple WebLogic domains. The JTA transaction coordinator must establish a direct RMI connection to the server instance that acts as the sub coordinator of the transaction when that transaction is either committed or rolled back. Due to this, the load balancer is normally used **only for the `JNDI initialContext`** creation. The WebLogic Server load balancing system controls the future T3/T3s requests, which connect to the WebLogic managed servers' addresses and ports (default or custom channels) indicated in the stubs retrieved during the initial context retrieval.

This section also explains **how to use the load balancer for all communications** (both initial and subsequent requests), which is applicable if JTA is not used.

> **Note:**
>
> External clients can access to T3 services using T3 tunneling over HTTP. But, this approach is not discussed in this document. This approach creates an HTTP session for each RMI session and uses standard HTTP protocols to transfer the session ID back and forth between the client and the server. This introduces some overhead and is less frequently used.

- Different Approaches to Access T3/T3s Services
  A Disaster Recovery scenario presents specific aspects which affect the configuration of external T3/T3s clients. This topic explains different approaches that you can use when accessing T3/T3s services from external clients, as well as learn how to manage them in a disaster recovery scenario.
- About T3/T3s Client’s DNS Cache
  All the approaches to access T3/T3s services mostly requires an DNS update. Since DNS update is required, you must set the limit for DNS Cache TTL (Time To Live) in DNS server and client's specific cache.

Different Approaches to Access T3/T3s Services

A Disaster Recovery scenario presents specific aspects which affect the configuration of external T3/T3s clients. This topic explains different approaches that you can use when accessing T3/T3s services from external clients, as well as learn how to manage them in a disaster recovery scenario.

> **Note:**
>
> These approaches apply to a DR scenario that complies with the MAA guidelines for Fusion Middleware and PaaS DR, so the secondary domain configuration is a mirror replica of the primary system.

- Direct T3/T3s Using Default Channels
  In this approach, the external T3/T3s client connects directly to the WebLogic managed servers' default channels. These default channels listen in the `Listen Address` and `Listen Port` specified in the general configuration of each WebLogic managed server.
- Direct T3/T3s Using Custom Channels
  In this approach, the external T3/T3s client connects directly to custom channels defined in the WebLogic servers of the cluster. The `Listen Address`, the `External Listen Address`, and the `External Port` are customizable values in the custom channels. These values can differ from the WebLogic server's default listen values.
- Using Load Balancer for Initial Lookup
  In this scenario, the client's provider URL point to the Load Balancer service.
- Using Load Balancer for All Traffic
  In this approach, not only the initial context lookup goes through the load balancer, but also the subsequent connections. There are no direct connections from the external T3/T3s client to the servers.

Direct T3/T3s Using Default Channels

In this approach, the external T3/T3s client connects directly to the WebLogic managed servers' default channels. These default channels listen in the `Listen Address` and `Listen Port` specified in the general configuration of each WebLogic managed server.

Figure 3-5 Direct T3/T3s Using Default Channels

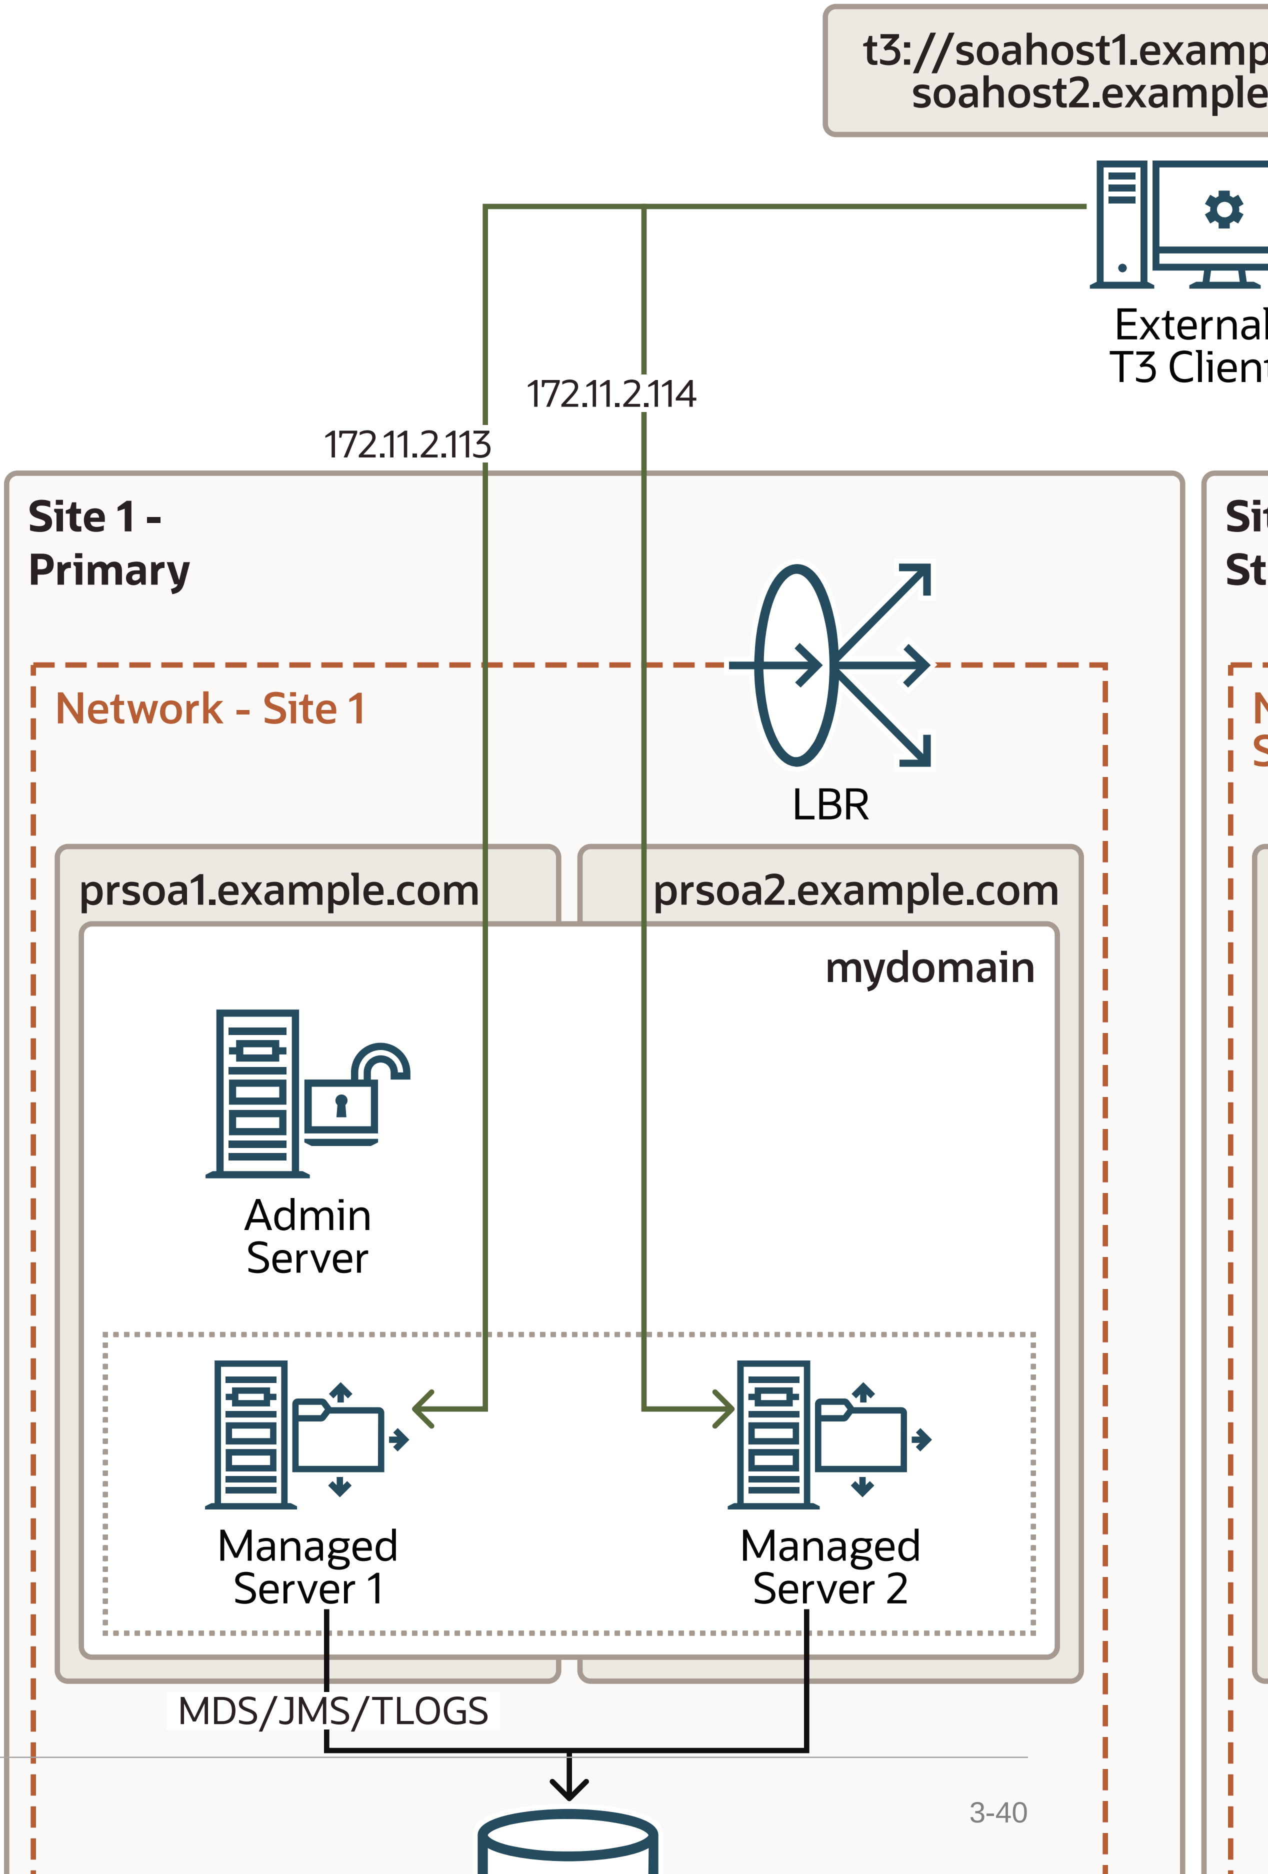

Configuration

- The provider URL in the T3 client uses the list of the WebLogic servers' default listen addresses and ports.

  Example: In the following DR example, the listener addresses of the WebLogic servers are the primary hostnames:

  ```
  t3://soahost1.example.com:8001,soahost2.example.com:8001
  ```

  In the case of using the T3s protocol, the port must be set to the server's `SSL Listen Port`.

  Example:

  ```
  t3s://soahost1.example.com:8002,soahost2.example.com:8002
  ```

- The external client must resolve these hostnames with the primary host's IPs. These IPs must be reachable from the client. It is possible to use Network Address Translating (NAT), as long as there is a NAT IP address for each server. In that case, the client will resolve each server name with the appropriate NAT IP for the server.

  Example: Naming resolution at the external client side

  This is achievable though `local /etc/hosts` or formal DNS server resolution.

  ```
  172.11.2.113    soahost1.example.com
  172.11.2.114    soahost2.example.com
  ```

Switchover

In a switchover scenario, there is no need to update the client's provider URL. Instead, you need to perform an update of the entries in the `DNS (or client's /etc/hosts)` of the client. After the switchover, the names used to connect to the servers must resolve with the IPs of the secondary servers.

Example: Naming resolution at the external client side after a switchover

```
172.22.2.113    soahost1.example.com
172.22.2.114    soahost2.example.com
```

Advantages

- You don't require additional configuration either at the server's side or at the front-end tier. All you require is only opening the ports to the client.

Disadvantages

- When a switchover happens, you need to update the `DNS (or client's /etc/hosts)` to alter the resolution of all the hostnames that the external client uses to connect to the managed servers. For more information about implications of the client's cache, see About T3/T3s Client's DNS Cache.
- Clients must be able to resolve and reach the hostnames set in WebLogic's `Listen Address`. It is not possible to use alternate names, because the default channels use the WebLogic server's default `Listen Address`.

- The WebLogic default channels not only listen for T3(s) requests, but also for HTTP(s). You cannot disable this setting. If you open the default port for the clients, direct HTTP(s) to the server is also allowed, which can result in security concerns.
- You need to modify the client's provider URL if you have to either add or remove the WebLogic server nodes from the WebLogic cluster. First contact only uses the list in the client's provider URL. So, even without updating the list, the subsequent requests can connect to any server of the cluster as the recovered stubs list all of the members. But it is a good practice that the client's provider URL matches the real list of the servers, for failover purposes in the first contact.

Direct T3/T3s Using Custom Channels

In this approach, the external T3/T3s client connects directly to custom channels defined in the WebLogic servers of the cluster. The `Listen Address`, the `External Listen Address`, and the `External Port` are customizable values in the custom channels. These values can differ from the WebLogic server's default listen values.

Figure 3-6 Direct T3/T3s Using Custom Channels

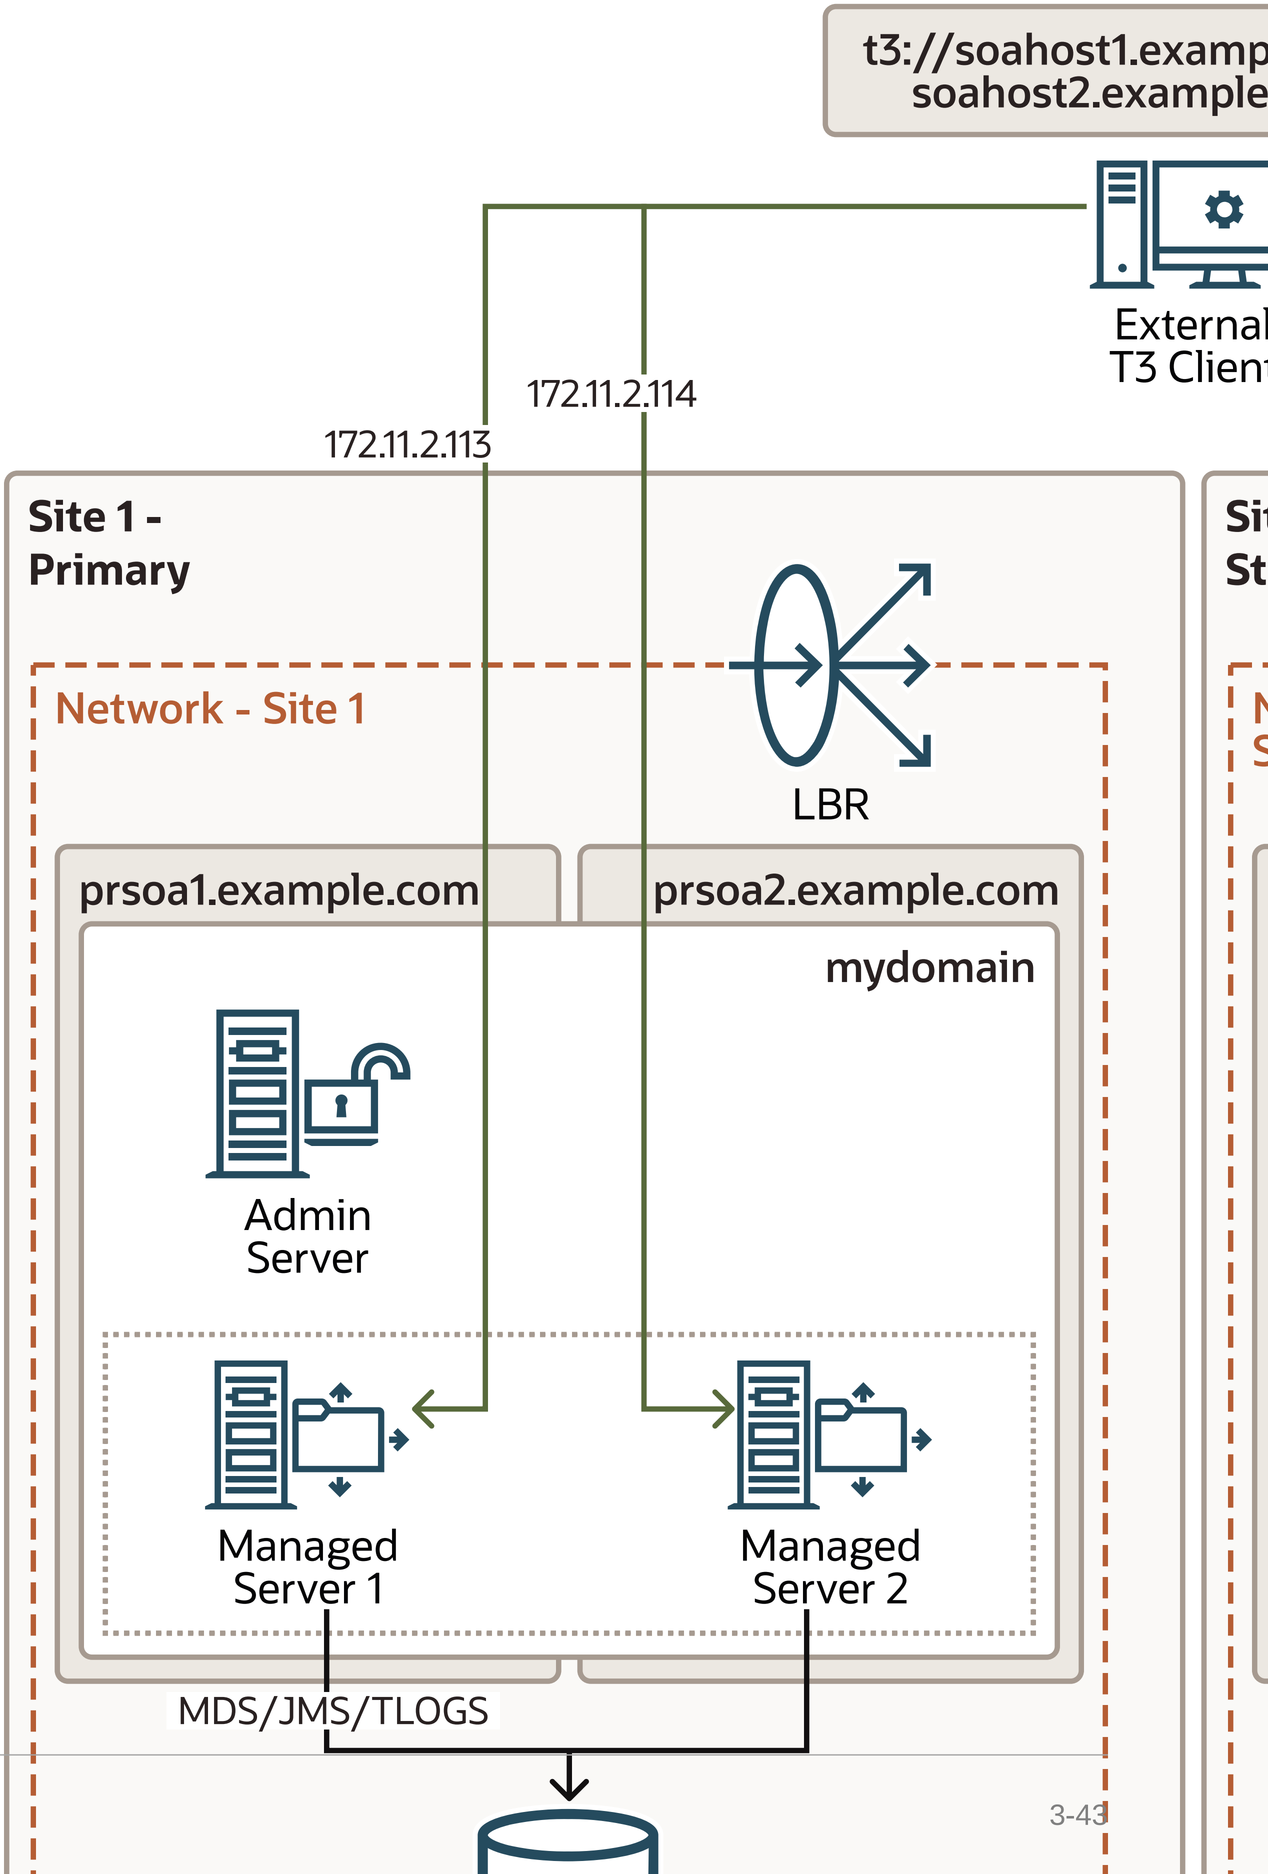

Once the T3/T3s external client has connected to the server during the initial context retrieval, the subsequent T3 calls connects directly to one of the listen addresses and ports configured in the custom channels as `External Listen address` and `External Port`. These requests will be load balanced per the mechanism specified in the connection factory defined in the WebLogic and used by the client.

This approach is similar to Direct T3/T3s Using Default Channels, with the exception that you can customize the addresses and ports used in T3/T3s calls when using WebLogic custom channels.

Configuration

- Each WebLogic server has the appropriate custom channels configured and these custom channel uses a unique `External Listen address` that points to that server node. Table 3-4 and Table 3-5 provides example of Custom Channel in Server 1 and Server 2.

  Table 3-4 Custom Channel in Server 1

Name	Protocol	Enabled	Listen Address	Listen Port	Public Address	Public Port
`t3_external_channel`	t3	true	`soahost1.example.com`	8111	`soahost1-t3ext.example.com`	8111

  Table 3-5 Custom Channel in Server 2

Name	Protocol	Enabled	Listen Address	Listen Port	Public Address	Public Port
`t3_external_channel`	t3	true	`soahost2.example.com`	8111	`soahost2-t3ext.example.com`	8111

- The external T3 client's provider URL contains the list of the external address and port of the custom channels.

  Example:

  ```
   t3://soahost1-t3ext.example.com:8111,soahost2-
  t3ext.example.com:8111
  ```

  If you are using the T3s protocol, you must create the custom channels with T3s protocol. Then the clients will connect using T3s protocol and appropriate port.

  Example:

  ```
   t3s://soahost1-t3ext.example.com:8112,soahost2-
  t3ext.example.com:8112
  ```

  In a T3s channel, you can add a specific SSL certificate for the name used as an `External Listen address`.

- The external T3/T3s client must resolve the custom channels' external hostnames with the primary host's IPs. These hostnames must be reachable from the client. It is possible to use NAT, as long as there is a NAT address for each server. In that case, the client will resolve each server name with the appropriate NAT IP for the server.

Example: Naming resolution at the external client side

```
172.11.2.113    soahost1-t3ext.example.com
172.11.2.114    soahost2-t3ext.example.com
```

With this approach, all the requests from the external client connect to `soahost1-t3ext.example.com` and `soahost2-t3ext.example.com`, both for the initial context retrieval and for the subsequent calls. You can control these requests by using the WebLogic Server load balancing mechanism.

Switchover

You don't have to update the client's provider URL if you have performed a switchover. Instead, you must update the entries in the `DNS (or in the /etc/hosts)` of the client. After the switchover, the names used to connect to the servers must resolve with the IPs of the secondary servers.

Example: Naming resolution at the external client side after a switchover

```
172.22.2.113    soahost1-t3ext.example.com
172.22.2.114    soahost2-t3ext.example.com
```

Advantages

- This method allows using specific hostnames for the external T3/T3s communication, different from the server's default listen address. This is useful if you do not want to expose the default server's `Listen Address` to the external clients for security reason.
- This method is also useful if you are using NAT IPs and you do not want to resolve the servers' default listen addresses with different IPs internally and externally.
- This method is also useful in case you want to use different names for external T3/T3s accesses just for organizational purposes. You can also use this method to isolate protocols in different interfaces or network routes.
- The protocol in the custom channel can be limited to T3/T3s only. The HTTP(s) protocol can be disabled in the custom channels.

Disadvantages

- When a switchover happens, you need to update the `DNS (or client's /etc/hosts)` at the external client side for all the hostnames used to connect to the managed servers. For more information about the implications of the client's cache, see About T3/T3s Client's DNS Cache.
- If you are using T3s, then you must create and configure specific SSL certificates for the external names in the custom channels to avoid SSL hostname verification errors in the client.
- You have to modify the client's provider URL if you have added or removed the WebLogic server nodes from the WebLogic cluster. First contact only uses the list in the client's provider URL. So, even without updating the list, the subsequent requests can connect to any server of the cluster as the recovered stubs list all of the members. But in any case, it is a good practice that the client's provider URL matches the real list of the servers, for failover purposes in the first contact.

Using Load Balancer for Initial Lookup

In this scenario, the client’s provider URL point to the Load Balancer service.

Direct T3/T3s Using Default Channels and Direct T3/T3s Using Custom Channels approaches can also use a load balancer for the initial context lookup.

However, the subsequent T3/T3s calls connects to WebLogic servers directly. If you use the default channels then the requests goes to the default channel’s listen address and port. Similarly, if you use the custom channels then the subsequent request goes to the external listen address and port defined in the customs channels.

Figure 3-7 Using Load Balancer for Initial Lookup

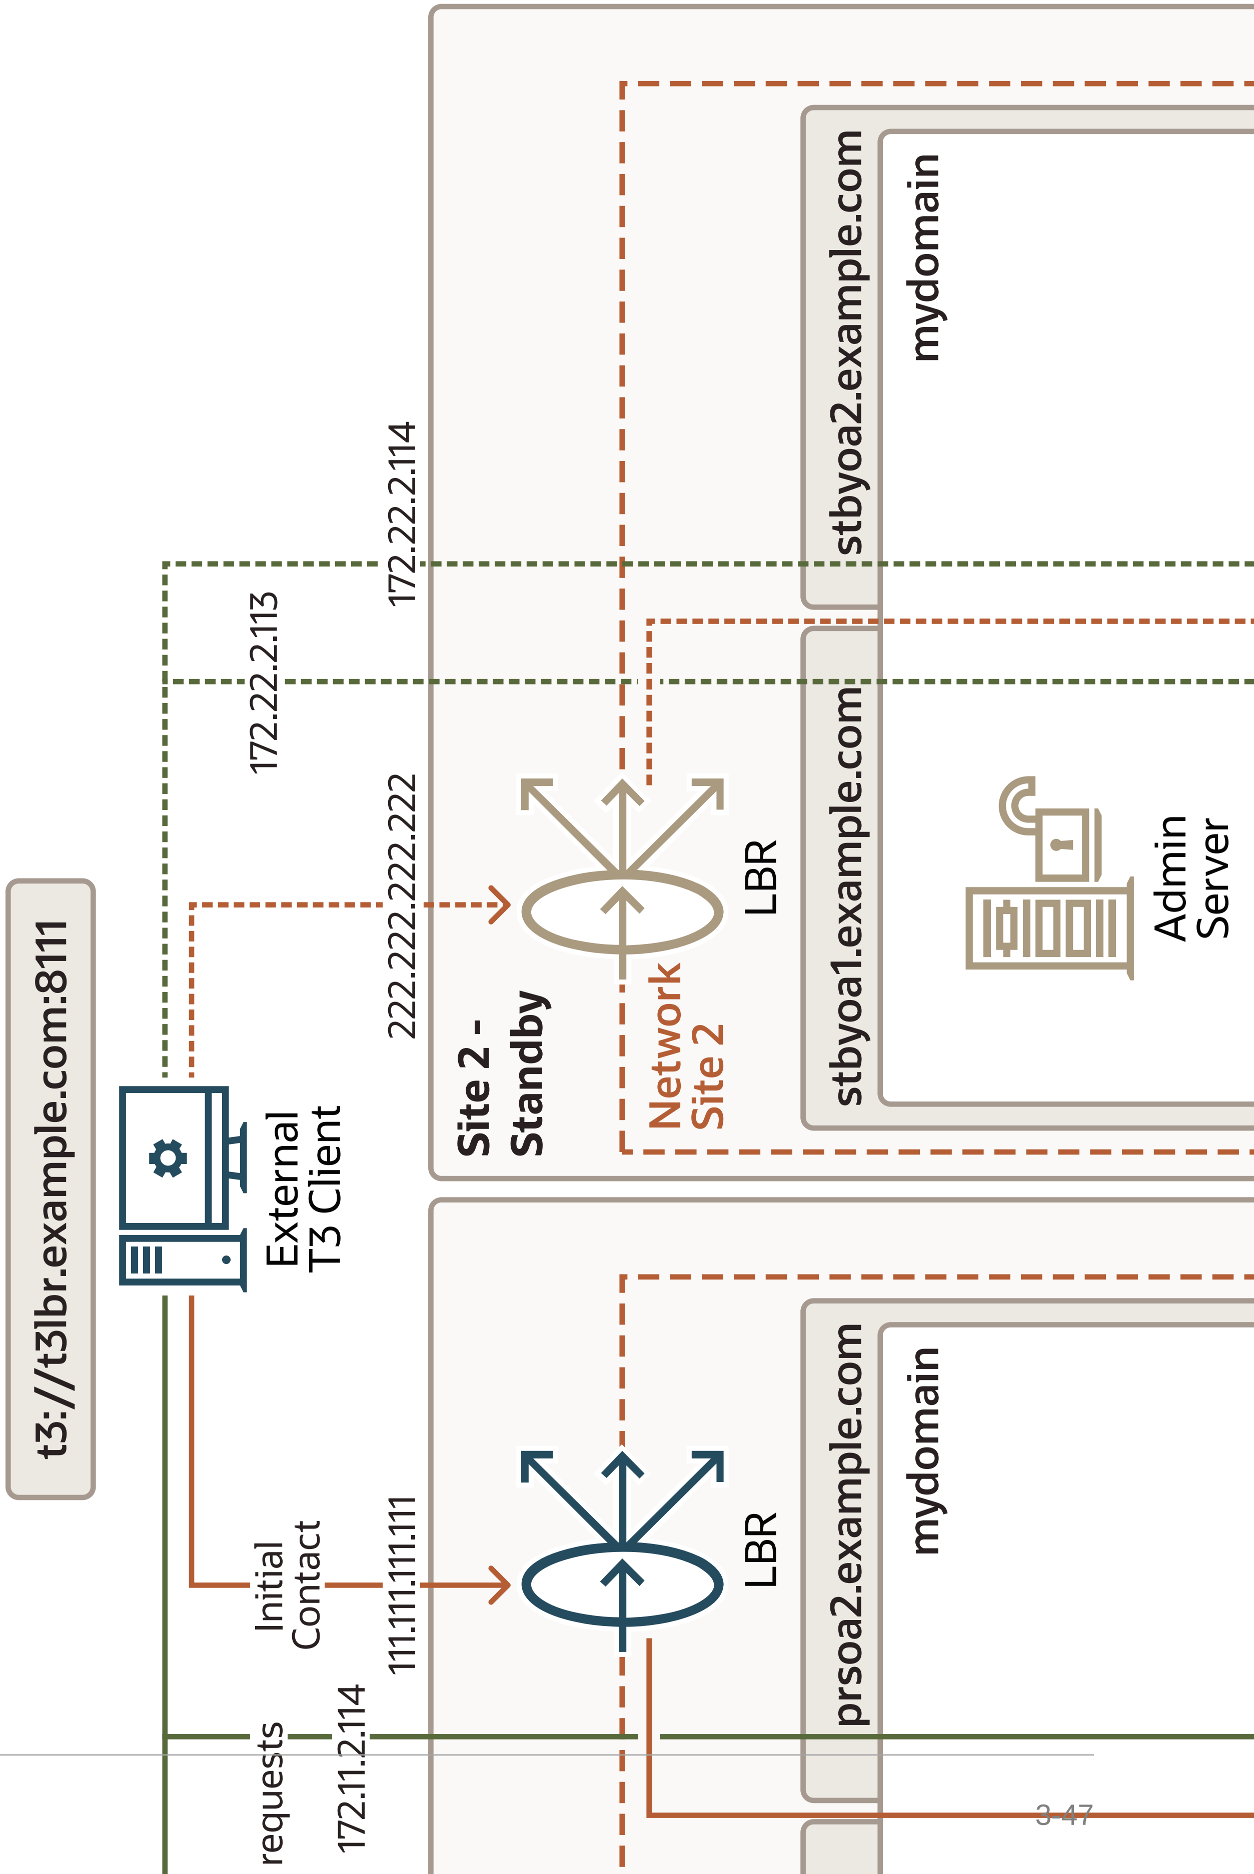

Configuration

- You will need a TCP service in the load balancer to load balances the requests to the WebLogic servers' T3/T3s ports (either to the default channels or to the custom channels when used).
- The external T3/T3s client's provider uses the front-end name and port of the load balancer as the point of contact.

  Example:

  ```
  t3://t3lbr.example.com:8111
  ```

- You can use default channels or custom channels as explained in Direct T3/T3s Using Default Channels and Direct T3/T3s Using Custom Channels scenarios. The external T3/T3s client must resolve the custom channels' external hostnames (or the default server's listeners hostnames if you are using default channels) with the primary host's IPs. The client must be able to access them. It is possible to use NAT, as long as there is a NAT address for each server. In that case, the client will resolve each server name with the appropriate NAT IP for the server.

  Example: Naming resolution at the external client side

  ```
  111.111.111.111    t3lbr.example.com
  172.11.2.113    soahost1-t3ext.example.com
  172.11.2.114    soahost2-t3ext.example.com
  ```

Switchover

In case of a complete switchover, you don't have to update the client's provider URL. Instead, you must update the entries in the `DNS (or in the /etc/hosts)` used by the client so that they resolve with the IPs of the secondary site. The name used to connect to the load balancer must resolve with the IP of the secondary load balancer, and the server names used in subsequent requests must point to the IPs of the secondary servers.

Example: Naming resolution at the external client side after a switchover

```
222.222.222.222    t3lbr.example.com
172.22.2.113    soahost1-t3ext.example.com
172.22.2.113    soahost2-t3ext.example.com
```

Advantages

You don't have to modify the client's provider URL if you add or remove the WebLogic server nodes from the WebLogic cluster.

Disadvantages

- Despite using a load balancer in front, the client still needs to reach the servers directly.
- When a switchover happens, you need to update the `DNS (or /etc/hosts)` of servers' addresses at the external client side. For more information about the implications of the client's cache, see About T3/T3s Client's DNS Cache.

- The complexity of this method is higher for the T3s cases. The client connects both through the front-end LBR and directly to the server using a secure protocol. In this case, you will need to either skip the hostname verification at the client side, or use an SSL certificate that is valid for **front** and **back** addresses (e.g. wildcard or SAN certificates).

Using Load Balancer for All Traffic

In this approach, not only the initial context lookup goes through the load balancer, but also the subsequent connections. There are no direct connections from the external T3/T3s client to the servers.

Oracle does not recommend using LBR for load balancing all types of T3/T3s communications. It is only recommended for initial context lookup. For more information, see, WebLogic RMI Integration with Load Balancers. However, there are T3/T3s use cases, where you can use Load Balancer for complete T3/T3s communication flow. WebLogic T3/T3s is TCP/IP-based protocols, so it can support TCP load balancing when services are homogeneous, such as JMS and EJB. For example, you can configure an LBR front-ended JMS subsystem in a WebLogic cluster in which remote JMS clients can connect to any cluster member.

This approach, however, will not work with external clients that use JTA connections. A JTA subsystem cannot safely use TCP load balancing in transactions that span across multiple WebLogic domains. When the transaction is either committed or rolled back, the JTA transaction coordinator must establish a direct RMI connection to the server instance that has been chosen as the transaction's sub coordinator.

This method is not suitable also for cases where you require direct connection to an specific server, like JMX or WLST when you want to connect to a particular server only.

Figure 3-8   Using Load Balancer for All Traffic

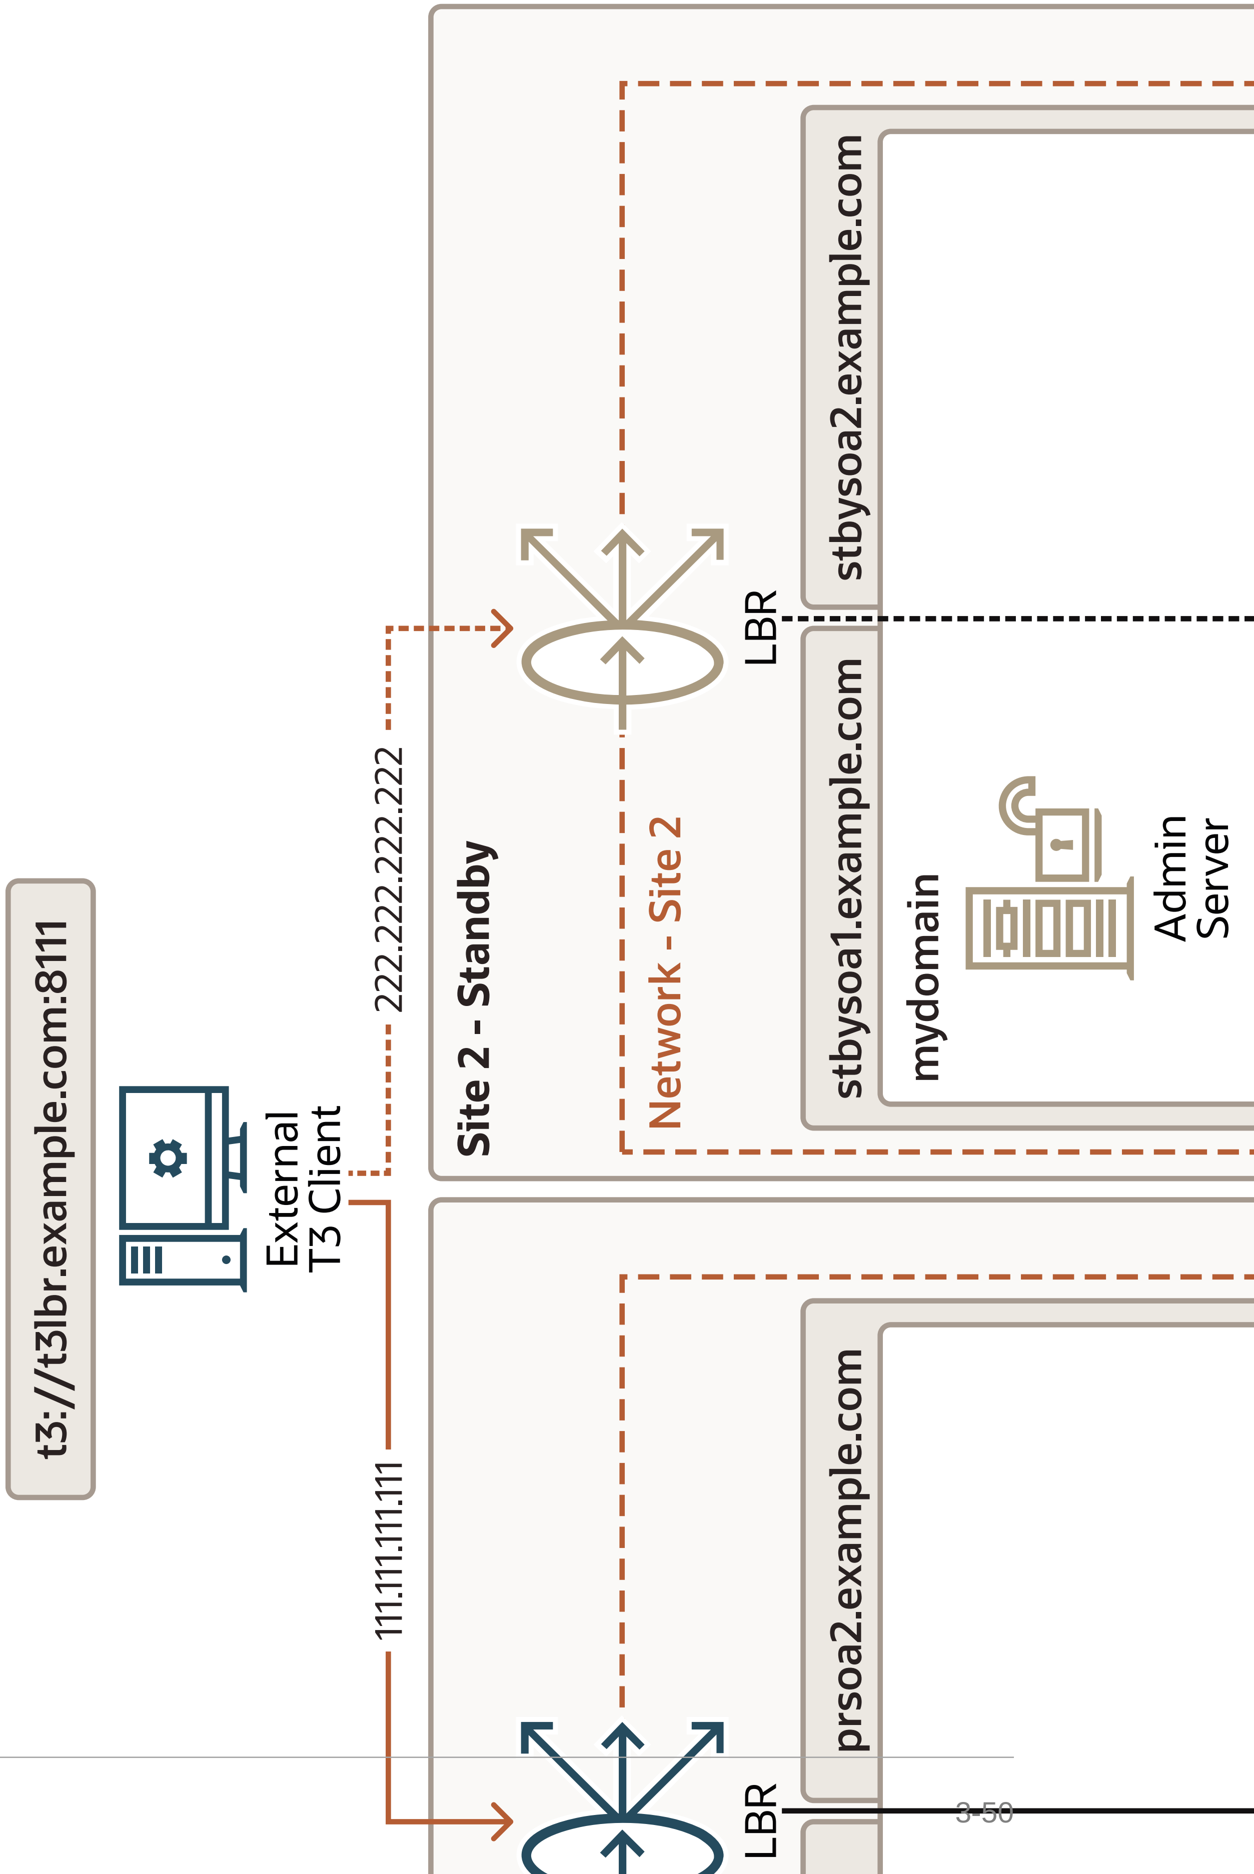

Configuration

- Load balancer requires a TCP service to load balance the requests to the WebLogic servers' T3/T3s ports defined in the custom channels.
- The external T3/T3s client's provider URL uses the front-end name and port of the load balancer as the point of contact.

  Example:

  ```
  t3://t3lbr.example.com:8111
  ```

- The external client must resolve the load balancer address with the IP of the primary site's load balancer.

  Example:

  ```
  111.111.111.111  t3lbr.example.com
  ```

- WebLogic Server requires custom channels. Configure the external listen address and port of these custom channels with the Load Balancer's address and port. Table 3-6 and Table 3-7 provides example of Custom Channel in Server 1 and Server 2.

Table 3-6 Custom Channel in Server 1

Name	Protocol	Enabled	Listen Address	Listen Port	Public Address	Public Port
`t3_external_channel`	t3	true	`soahost1.example.com`	8111	`t3blr.example.com`	8111

Table 3-7 Custom Channel in Server 2

Name	Protocol	Enabled	Listen Address	Listen Port	Public Address	Public Port
`t3_external_channel`	t3	true	`soahost2.example.com`	8111	`t3lbr.example.com`	8111

Switchover

In case of a switchover, you don't have to update the client's provider URL. Instead, you must update the entries in the `DNS (or in the /etc/hosts)` of the client. After the switchover, the names used to connect to the servers resolves with the IP of the secondary LBR service.

Example: Naming resolution at the external client side after a switchover

```
222.222.222.222  t3lbr.example.com
```

Advantages

- All the communication goes through the load balancer. The client only needs to know and reach the load balancer's service.

- You don't have to modify the client's provider URL if you have to either add or remove the WebLogic server nodes from the WebLogic cluster.
- In case of a switchover, you only have to update the load balancer's front-end name in the `DNS (or in the /etc/hosts)`.

> **Note:**
>
> Although only one DNS name is updated, you need to refresh the client's DNS cache. For more information, see About T3/T3s Client's DNS Cache.

- In T3s cases, you can use the same SSL certificate in all the custom channels, associated to the load balancer service's front-end name.

Disadvantages

- This approach is not suitable for all the T3/T3s cases. Scenarios with JTA cannot use this approach, because JTA needs to explicitly connect to the server instance that acts as the sub coordinator of the transaction.

About T3/T3s Client's DNS Cache

All the approaches to access T3/T3s services mostly requires an DNS update. Since DNS update is required, you must set the limit for DNS Cache TTL (Time To Live) in DNS server and client's specific cache.

The TTL limit in the DNS service is a setting that tells the DNS resolver how long to cache a query before requesting a new one. Note that the TTL value of the DNS entries will affect the effective RTO of the switchover; if the TTL is high (for example, 20 min), the DNS change will take that time to be effective in the clients. Using lower TTL values will make this switchover faster, however, this can cause overhead because the clients check the DNS more frequently. A good approach is to set the TTL to a low value temporarily (for example, one min), before the change in the DNS. Then, perform the change, and once the switchover procedure is completed, set the TTL to the normal value again.

Besides the DNS server's TTL, `networkaddress.cache.ttl` Java property controls the Java clients' cache TTL. This Java property indicates the caching policy for successful name lookups from the name service. Specify the value as an integer to indicate the number of seconds to cache the successful lookup. Ensure you set a limit to the `networkaddress.cache.ttl` Java property so the client's Java cache does not cache the DNS entries forever. Else, with each switchover, you might have to restart the client.

4 Recommendations for Oracle Fusion Middleware Components

Learn about the disaster protection requirements for Oracle Fusion Middleware components in different Oracle product suites and recommendations for synchronizing those components.

Oracle Fusion Middleware Disaster Recovery uses storage replication to synchronize the middle tier content, and uses Oracle Data Guard to synchronize data in Oracle databases and custom application databases that are included in your topology.

> **Note:**
>
> Certain artifacts such as Oracle Inventory, the `oratab` and the `oraInst.loc` files 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 that you place Oracle Inventory, the `oratab`, and the `oraInst.loc` files on the local disk of your systems. 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 managing Oracle Inventory, see Managing Oracle Inventory.

For related information, see the following documents:

- *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 Business Intelligence*
- *Oracle Fusion Middleware Administering Oracle Fusion Middleware*

This chapter includes the following sections:

- Recommendations for Oracle WebLogic Server
  Oracle WebLogic Server is a scalable, enterprise-ready Java Platform, Enterprise Edition (Java EE) application server.
- Recommendations for Oracle SOA Suite
  Oracle SOA Suite is an Oracle Fusion Middleware component that provides a complete set of service infrastructure components for designing, deploying, and managing SOA composite applications.
- Recommendations for Oracle WebCenter Portal
  Oracle WebCenter Portal is an Oracle Fusion Middleware component that provides a personalized, secure, and efficient way of presenting and consuming information, collaborating with others, and interacting with applications in the context of business processes.

- Recommendations for Oracle WebCenter Content
  Oracle WebCenter Content is an integrated suite of products designed to manage enterprise content.

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 the Oracle Fusion Middleware product suite.

The following artifacts and considerations apply to all 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 that you use alias host names (instead of physical host names) as the listen address for both Oracle WebLogic Administration server and the Managed Servers, as described in Network Considerations. As long as these aliases can be resolved on both the production and standby sites, there is no need to update this value after a Disaster Recovery operation.

  Configure the load balancer virtual host, which is used for accessing the WebLogic Server applications, 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)
  Learn about the Oracle WebLogic Server JMS and T-Log artifacts and what Oracle recommends for disaster recovery.
- Recommendations for Oracle Platform Security Services
  Learn about the Oracle Platform Security Services artifacts and what Oracle recommends for disaster recovery.

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

Learn about the Oracle WebLogic Server JMS and T-Log artifacts and what Oracle recommends for disaster recovery.

- **Artifacts on the File System:**

  File-based persistent stores: The file store location for the JMS and T-Log when you use a file-based persistent store.

> **Note:**
>
> Oracle recommends that you use jdbc persistence store.

- **Artifacts in the Database:**

  The schema that contains the JMS messages and the TLOGs, when you use JDBC persistent stores for JMS and JTA services.

  When automatic whole server migration or automatic service 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.
  - 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 and is one of the reasons why Oracle recommends using JDBC persistent stores both for JMS and JTA.
  - 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 use JDBC persitent stores.
  - 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 you use database-based persistent stores.
- **Recovery Recommendations:**

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

Recommendations for Oracle Platform Security Services

Learn about the Oracle Platform Security Services artifacts and what Oracle recommends for disaster recovery.

- **Artifacts in the Database:**

  Oracle Platform Security Services have database dependencies.

- **Synchronization Recommendations:**

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

  You should configure Oracle Data Guard for the Oracle database metadata repositories.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. 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. See Applying Redo Data to Physical Standby Databases.

- **Recovery Recommendations:**

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

Recommendations for Oracle SOA Suite

Oracle SOA Suite is an Oracle Fusion Middleware component that 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.

A SOA composite application consists 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 older Oracle SOA Suite releases, such as 11*g*, the soa-infra and service-engine configuration files were stored in local or shared storage files as part of the domain configuration. Starting with Oracle SOA Suite 12*c* R1, those files reside in the metadata repository database. Thus, soa-infra and service-engine configuration changes are now immediately propagated across a cluster.
>
> The Disaster Recovery recommendations for Oracle SOA Suite assume that you are using Oracle SOA Suite 12*c* R2 (12.2.1) release.

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 that contains 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 for the SOA domain.

Network Artifacts

Oracle recommends that you use aliases 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 *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite* for instructions for updating an IP address to an alias 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 you make changes in the domain-related configuration, deploy composites, and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. 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. See Applying Redo Data to Physical Standby Databases.

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.

The following sections describe Disaster Recovery recommendations for Oracle SOA Suite components:

- 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.
- Recommendations for Oracle BPEL Process Manager
The Oracle BPEL Process engine is the service engine running in an Oracle SOA Service Infrastructure that allows the execution of BPEL processes.
- Recommendations for Oracle Mediator
Oracle Mediator is a service engine within the Oracle SOA Service Infrastructure that provides the framework to mediate between providers and consumers of services and events.
- 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 and the human interaction support such as approvals, rejects, and reassign actions.
- Recommendations for Oracle B2B
Oracle B2B connects SOA composite applications to external services, applications, and technologies, and offers a multi-protocol gateway that supports industry-recognized business-to-business (B2B) standards.
- 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. Oracle Web Services Manager consists of the Policy Manager and the Agent.
- Recommendations for Oracle User Messaging Service
Oracle User Messaging Service enables two-way communication between users and deployed applications and it supports a variety of channels, such as email, IM, SMS, and text-to-voice messages.
- Recommendations for Oracle Java EE Connector Architecture (JCA) Adapters
Oracle JCA Adapters are JCA binding components that allow the Service Infrastructure to communicate by using different protocols.
- Recommendations for Oracle Business Activity Monitoring
Oracle Business Activity Monitoring (Oracle BAM) provides the tools for monitoring business services and processes in the enterprise and allows quick correlation of market indicators to the actual and changing business processes.
- Recommendations for Oracle Business Process Management
The Oracle Business Process Management (Oracle BPM) Suite provides an integrated environment for developing, administering, and using business applications centered around business processes.
- Recommendations for Oracle Managed File Transfer
Oracle Managed File Transfer (MFT) is a high performance, standards-based, end-to-end managed file gateway. It features design, deployment, and monitoring of file transfers by using a lightweight web-based design-time console that includes transfer prioritization, file encryption, scheduling, and embedded sFTP servers.

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 runtime 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

Use JDBC persistent stores for SOA JMS resources, so that, the JMS messages are replicated to the standby through the underlying Data Guard replica.

The application tier must be manually synchronized with the standby site after you make changes in the domain-related configuration, deploy composites, and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. 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.

Recommendations for Oracle BPEL Process Manager

The Oracle BPEL Process engine is the service engine running in an 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 the 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 you make domain-related configuration changes and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. 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 you perform 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 you perform Disaster Recovery operation, especially when a process is in flight.

Recommendations for Oracle Mediator

Oracle Mediator is a service engine within the Oracle SOA Service Infrastructure that provides the framework to mediate between providers and consumers of services and events.

Oracle Mediator 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 you make changes in the domain-related configuration and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . 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 application.

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 and the human interaction support such as approvals, rejects, and reassign actions.

The Human Workflow service consists of several services that handle various aspects of human interaction with a business process.

This section describes the 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, and view definitions are stored in the database.

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

Synchronization Recommendations

The application tier must be manually synchronized with the standby site after you make changes in the domain configuration and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . 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 application. The Oracle Human Workflow engine uses Oracle User Messaging Service to send and receive notifications. See Recommendations for Oracle User Messaging Service for details about Oracle User Messaging Service.

Recommendations for Oracle B2B

Oracle B2B connects SOA composite applications to external services, applications, and technologies, and 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 the various Oracle B2B artifacts and provides recommendations for disaster recovery.

Artifacts on the File System

The JMS messages of the B2B queues are stored in the file system when you use file-based persistent stores.

> **Note:**
>
> Oracle recommends that you use JDBC persistence store instead.

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.

The schema that contains the JMS messages, when you use JDBC persistent stores.

Special Considerations

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

Synchronization Recommendations

Use JDBC persistent stores for JMS, so that, they are replicated to the standby through the underlying Data Guard replica.

For information about Oracle B2B JMS queue synchronization and recovery, see Recommendations for Oracle WebLogic Server Java Message Service (JMS) and Transaction Logs (T-Logs).

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

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. 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 application. Oracle B2B stores state information within JMS queues and the SOA runtime database, so recovering the database and the Managed Server ensures that the application runs normally.

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. Oracle Web Services Manager consists of the Policy Manager and the Agent.

Oracle WSM provides capabilities to build, enforce, execute, and monitor web services policies including security, Web Services Reliable Messaging, Message Transmission Optimization Mechanism, and addressing policies.

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 you make changes in the domain-related configuration and apply patches.

You should configure Oracle Data Guard for Oracle database metadata repositories.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . 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 application. All policies are stored in the MDS repository, so recovering the database and the Managed Server ensures that the application runs normally.

Recommendations for Oracle User Messaging Service

Oracle User Messaging Service enables two-way communication between users and deployed applications and it supports a variety of channels, such as email, IM, SMS, and text-to-voice messages.

Oracle User Messaging Service is integrated with Oracle Fusion Middleware components such as Oracle BPEL 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

The JMS messages of the Oracle User Messaging Service queues are stored in the file system when you use file-based persistent stores.

> **Note:**
>
> Oracle recommends that you use JDBC persistence store.

Artifacts in the Database

The schema that contains the JMS messages of the Oracle User Messaging Service queue when using JDBC persistent stores.

Oracle User Messaging Service also uses the database repository to maintain the message and configuration state.

Special Considerations

Oracle User Messaging Service uses JMS to deliver messages among messaging applications. Oracle recommends to store these JMS artifacts in the database, using JDBC persistent stores. Therefore, they are replicated to the standby system through the Data Guard.

Synchronization Recommendations

Use JDBC persistent stores for JMS, so that, they are replicated to the standby through the underlying Data Guard replica.

The application tier must be manually synchronized with the standby site after you make changes in the configuration, deploy additional Oracle User Messaging Service drivers, and apply patches.

You should configure Oracle Data Guard for Oracle database metadata repositories.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. 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, see Synchronization Recommendations in Recommendations for Oracle WebLogic Server Java Message Service (JMS) and Transaction Logs (T-Logs).

Recommendations for Oracle Java EE Connector Architecture (JCA) Adapters

Oracle JCA Adapters are JCA binding components that allow the Service Infrastructure to communicate by using different protocols.

Oracle JCA Adapters are deployed as a JCA resource (RAR) and are not part of the Oracle SOA Service Infrastructure.

The Oracle JCA Adapters include:

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

See the *Oracle Fusion Middleware Understanding Technology Adapters* for additional information about the types of Oracle JCA Adapters.

This section describes the 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). When the file is created, the location of each `weblogic-ra.xml` file is determined by the administrator, 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.

The JMS messages of the JMS adapters are also stored in the database when they use JDBC persistent stores.

Synchronization Recommendations

Use JDBC persistent stores for JMS messages, so that, they are replicated to the standby through the underlying Data Guard replica.

The application tier must be manually synchronized with the standby site after you make changes in the domain-related configuration, that is adapter configuration changes, and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . 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.

Recommendations for Oracle Business Activity Monitoring

Oracle Business Activity Monitoring (Oracle BAM) provides the tools for monitoring business services and processes in the enterprise and allows quick correlation of market indicators to the actual and changing business processes.

Oracle BAM provides the necessary tools and runtime services to create 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 you make domain-related configuration changes and apply patches.

You should configure Oracle Data Guard for the Oracle SOA Suite database that contains the Oracle BAM schemas and the metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . 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.

Recommendations for Oracle Business Process Management

The Oracle Business Process Management (Oracle BPM) Suite provides an integrated environment for developing, administering, and using business applications centered around business processes.

Oracle BPM Suite 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 the various Oracle BPM artifacts and provides recommendations for disaster recovery.

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.

The JMS messages of the BPM JMS queues are also stored in the database when they use JDBC persistent stores.

Synchronization Recommendations

Use JDBC persistent stores for BPM JMS resources, so that, the JMS messages are replicated to the standby through the underlying Data Guard replica.

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

You should configure Oracle Data Guard 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 application.

Recommendations for Oracle Managed File Transfer

Oracle Managed File Transfer (MFT) is a high performance, standards-based, end-to-end managed file gateway. It features design, deployment, and monitoring of file transfers by using a lightweight web-based design-time console that includes transfer prioritization, file encryption, scheduling, and embedded sFTP servers.

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

Artifacts in the File System

Oracle MFT includes built-in FTP and sFTP servers, which handle many of the types of file transfers performed. These embedded servers have their own file system directories for sending and receiving files. The root directory location for both the FTP and sFTP servers resides in the file system. Also, Oracle MFT can transfer files from or to another sources (File System, SOAP, SOA, B2B), and these files can be also be stored in the File System.

Artifacts in the Database

The definition of the MFT design artifacts (sources, target and transfers) resides in the metadata repository, which is in the database.

Oracle MFT components also use JMS queues. The JMS queues reside in the database when JDBC persistent stores are configured, which is the best practice, specially in DR environments.

Synchronization Recommendations

Use JDBC persistent stores for the JMS resources, so that, the JMS messages are replicated to the standby through the underlying Data Guard replica.

The application tier must be manually synchronized with the standby site after you make changes in the domain-related configuration, deploy applications, and apply patches.

You must configure Oracle Data Guard for the Oracle MFT database and metadata repository.

If the runtime files that reside in the file system (such as the FTP files, files of File System trasfers, and so on) need to be available in the secondary site, you need to manually synchronize them as business needs.

Recovery Recommendations

The database must be recovered to the most recent point in time to ensure that the latest runtime and design time information are restored.

Recommendations for Oracle WebCenter Portal

Oracle WebCenter Portal is an Oracle Fusion Middleware component that provides a personalized, secure, and efficient way of presenting and consuming information, collaborating with others, and interacting with applications in the context of business processes.

Oracle WebCenter Portal optimizes the connections between people, information, and applications; provides a context for users to navigate, discover, and access content relevant to your business needs.

Oracle WebCenter Portal includes the following servers, applications, and services:

- Portal Managed Servers:
  - Portal application and services
  - Analytics Collector service
- Collaboration Managed Servers:
  - Discussions application and services
- Portlet Managed Servers:
  - Portlet Producer services
  - Pagelet Producer service

This section contains the following recommendations:

- Common Recommendations for All Oracle WebCenter Portal Components
  Learn about what Oracle recommends for all Oracle WebCenter Portal components.
- Recommendations for Oracle WebCenter Portal Server
  Oracle WebCenter Portal offers a single, integrated, web-based environment for social networking, communication, collaboration, and personal productivity through a robust set of services and applications.
- Recommendations for Oracle WebCenter Analytics
  Oracle WebCenter Portal allows for a single active analytics collector registration. The Analytics Collector receives usage metrics from the Portal application. All connections to the Analytics Collector are stateless.
- Recommendations for Oracle WebCenter Discussion Server
  Oracle WebCenter Discussion Server provides the ability to integrate discussion forums and announcements into your applications.
- Recommendations for Oracle WebCenter Portlet and Pagelet Services
  Oracle WebCenter Portal supports deployment and execution of both standards-based portlets (WSRP 2.0), and traditional Oracle (Portlet Developer Kit) PDK-Java based portlets. Oracle WebCenter Portal provides several ready-to-use producers.

Common Recommendations for All Oracle WebCenter Portal Components

Learn about what Oracle recommends for all Oracle WebCenter Portal components.

The following artifacts and considerations apply to all Oracle WebCenter Portal components:

Artifacts on the File System

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

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

There are no shared files ystem persistence store requirements for any Oracle WebCenter Portal runtime components.

Network Artifacts

Oracle recommends that you use a virtual alias (instead of physical host names) as the listen address for both Oracle WebLogic Administration server and the Oracle Managed Servers. As long as this alias can be resolved on both the production and standby sites, there is no need to update this value after a Disaster Recovery operation.

Oracle WebCenter Portal does not require whole server migration or service migration to be configured. No special considerations are needed in this regard.

You should configure the load balancer virtual host, which is used to access the Oracle WebCenter Portal applications and services, on both the production and standby sites.

Artifacts in the Database

Oracle WebCenter Portal stores all data and metadata in database schemas created with the Repository Creation Utility. Application runtime configuration metadata is store through the Metadata Services Repository (MDS).

Synchronization Recommendations

Synchronize the application tier manually with the standby site after you make changes in the domain-related configuration, deploy applications or shared libraries, or apply patches.

You should configure Oracle Data Guard for the Oracle WebCenter Portal database and metadata repository.

Oracle recommends that you synchronize the standby database, when the application tier synchronization is initiated on the storage. This synchronization can occur automatically when 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. See Applying Redo Data to Physical Standby Databases in *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

You must be recover the database to the most recent point in time to ensure that the latest composite definitions and inflight instances are restored.

Recommendations for Oracle WebCenter Portal Server

Oracle WebCenter Portal offers a single, integrated, web-based environment for social networking, communication, collaboration, and personal productivity through a robust set of services and applications.

The following are recommendations that are specific to Oracle WebCenter Portal Server:

Artifacts in the Database

Portal application data is stored in the `WEBCENTER` database schema. Portal customized metadata is stored in the `MDS` database schema

Synchronization Recommendations

Synchronize the application tier manually with the standby site after you change the configuration, deploy applications, or apply patches.

You should be configure Oracle Data Guard for the Oracle database containing the `WEBCENTER` schema, and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . This synchronization can occur automatically when 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. See Applying Redo Data to Physical Standby Databases in *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

You must be recover the database that contains the `WEBCENER` schema and `MDS` repository, to the most recent point in time, along with the Oracle WebCenter Portal domain.

Recommendations for Oracle WebCenter Analytics

Oracle WebCenter Portal allows for a single active analytics collector registration. The Analytics Collector receives usage metrics from the Portal application. All connections to the Analytics Collector are stateless.

Following an Enterprise Deployment Guide topology deployment, each portal server has one active local analytics collector service that is responsible for collecting metrics for all portal traffic on that portal server. The recommended configuration is to set the portal's common connection registration to `localhost`, and the default port of `31314`. Each portal server connects locally. Load-balancer configurations are not required and no configuration changes are needed for failover to the secondary site for analytics functionality.

The following are recommendations that are specific to Oracle WebCenter Analytics:

Artifacts in the Database

Analytics Collector data is stored in the Oracle WebCenter Portal's `ACTIVITIES` schema.

Synchronization Recommendations

Synchronize the application tier manually with the standby site after you make any changes in the configuration for the Analytics Collector MBeans, the Portal's Analytics service registration, or apply patches.

You should configure Oracle Data Guard for the Oracle database containing the Oracle WebCenter Portal schemas and the metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . This synchronization can occur automatically when 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. See Applying Redo Data to Physical Standby Databases in *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

You must recover the database that contains the `ACTIVITIES` schema to the most recent point in time, along with the Oracle WebCenter Portal domain.

Recommendations for Oracle WebCenter Discussion Server

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

The following are recommendations that are specific to Oracle WebCenter Discussion Server:

Artifacts in the Database

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

Synchronization Recommendations

Synchronize the application tier manually with the standby site after you change the configuration, deploy applications, or apply patches.

You should configure Oracle Data Guard for the Oracle database that contains the Oracle Discussion Server schema and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . This synchronization can occur automatically when 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. See Applying Redo Data to Physical Standby Databases in *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

You must recover the database that contains the Discussion Server schema to the most recent point in time along with the Oracle WebCenter Portal domain.

Recommendations for Oracle WebCenter Portlet and Pagelet Services

Oracle WebCenter Portal supports deployment and execution of both standards-based portlets (WSRP 2.0), and traditional Oracle (Portlet Developer Kit) PDK-Java based portlets. Oracle WebCenter Portal provides several ready-to-use producers.

The following are recommendations that are specific to Oracle WebCenter Portlet and Pagelet:

Artifacts in the Database

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

Synchronization Recommendations

Synchronize the application tier manually with the standby site after you change the configuration, deploy applications, or apply patches.

You should configure Oracle Data Guard for the Oracle database that contains the `PORTLET` schema and metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database . This synchronization can occur automatically when 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. See Applying Redo Data to Physical Standby Databases in *Oracle Data Guard Concepts and Administration*.

Recovery Recommendations

You must recover the database that contains the `PORTLET` schema and `MDS` repository to the most recent point in time along with the Oracle WebCenter Portal domain.

Recommendations for Oracle WebCenter Content

Oracle WebCenter Content is an integrated suite of products designed to manage enterprise content.

Oracle WebCenter Content enables you to leverage industry-leading document management, web content management, digital asset management, and records management functionality to build your business applications. The ability to manage the enterprise content helps you reduce costs, share content across the enterprise, minimize risk, automate manual processes, and consolidate multiple web sites onto a single platform.

Oracle Enterprise Content Management includes the following components:

- Oracle WebCenter Content
- Oracle WebCenter Content: Inbound Refinery
- Oracle WebCenter Content: Enterprise Capture
- Oracle WebCenter Content: Content User Interface

This section contains the following recommendations:

- Common Recommendations for All Oracle WebCenter Content Components
  Learn about what Oracle recommends for all Oracle WebCenter Content components.
- Recommendations for Oracle WebCenter Content
  Oracle WebCenter Content provides a unified application for several different kinds of content management.
- Recommendations for Oracle WebCenter Content Inbound Refinery
  Oracle WebCenter Content Inbound Refinery is a conversion server that manages file conversions for electronic assets such as documents, digital images, and motion video.
- Recommendations for Oracle WebCenter Content Enterprise Capture
  Oracle WebCenter Enterprise Capture (Enterprise Capture) provides organizations with a single system to capture both paper and electronic documents.
- Recommendations for Oracle WebCenter Content: Content User Interface
  WebCenter Content User Interface is a user-friendly and a feature-rich user interface for managing Oracle WebCenter Contented based on Oracle Application Development Framework.

Common Recommendations for All Oracle WebCenter Content Components

Learn about what Oracle recommends for all Oracle WebCenter Content components.

The following artifacts and considerations apply to all Oracle WebCenter Content components:

Artifacts on the File System

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

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

Domain Home—The domain home contains the Administration Server and Managed Server configuration data and Oracle WebCenter Content applications for the domain.

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

WebCenter Content Shared Directories files— In a clustered implementation of WebCenter Content as per Enterprise Deployment guidelines shared directories are required for domain configuration as well as WebCenter Content shared directory files such as data, audit, and vault directories.

Network Artifacts

Oracle recommends that you use a virtual alias (instead of physical host names) as the listen address for both Oracle WebLogic Administration server and the Oracle Managed Servers. As long as this alias can be resolved on both the production and standby sites, there is no need to update this value after a Disaster Recovery operation.

You must configure the load balancer virtual host, which is used for accessing Oracle WebCenter Content Management components, on both the production and standby sites.

Artifacts in the Database

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

Synchronization Recommendations

Synchronize the directory tier manually with the standby site after you make configuration changes, deploy applications, and apply patches.

You should configure Oracle Data Guard for the Oracle database metadata repository.

When the application tier synchronization is initiated on the storage, Oracle recommends that you synchronize the standby database. This synchronization can occur automatically when 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

You must recover the database that contains the Oracle WebCenter Content schemas to the most recent point in time along with the Identity Management component in question.

Recommendations for Oracle WebCenter Content

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

Through user-friendly interfaces, roles-based authentication and security models, Oracle WebCenter Content empowers users throughout the enterprise to view, collaborate on or retire content. This ensures that all accessible, distributed, or published information is secure, accurate, and up-to-date.

The following are recommendations that are specific to Oracle WebCenter Content:

Artifacts in the Database

The `OCS` (Oracle Content Schema) schema is part of the Oracle WebCenter Content database.

Artifacts in the File System

WebCenter Content shared directories along with any Oracle Secure Files or the setup of file-based persistent stores need to be replicated across to the standby site for having them disaster protected.

Special Considerations

You should configure the load balancer virtual host, which is required for the Oracle WebCenter Content on both the production and standby sites.

Synchronization Recommendations

Synchronize the directory tier manually with the standby site after you change the configuration, deploy applications, or apply patches.

You should configure Oracle Data Guard for the Oracle database repositories.

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

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

Recovery Recommendations

You must recover Oracle WebCenter Content with the `OCS` and `MDS` schemas to the most recent point in time.

Recommendations for Oracle WebCenter Content Inbound Refinery

Oracle WebCenter Content Inbound Refinery is a conversion server that manages file conversions for electronic assets such as documents, digital images, and motion video.

Oracle WebCenter Content Inbound Refinery provides thumbnail functionality for documents and images, storyboarding for video, and the ability to extract and use EXIF data from digital images and XMP data from electronic files that are generated from programs such as Adobe Photoshop and Adobe Illustrator. You can use Oracle WebCenter Content Inbound Refinery to convert content items stored in Oracle Content Server.

The following are recommendations that are specific to Oracle WebCenter Content Inbound Refinery:

Artifacts in the Database

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

Special Considerations

You should configure the load balancer virtual host, which is required for Oracle WebCenter Content Inbound Refinery on both the production and the standby sites.

Synchronization Recommendations

Synchronize the directory tier manually with the standby site after you change the configuration or apply patches.

Recovery Recommendations

Recover the Oracle WebCenter Content Inbound Refinery instance.

Recommendations for Oracle WebCenter Content Enterprise Capture

Oracle WebCenter Enterprise Capture (Enterprise Capture) provides organizations with a single system to capture both paper and electronic documents.

Enterprise Capture supports both centralized and distributed image capture from a user-friendly web interface capable of using high-volume, production-level scanners. Support for the industry-standard TWAIN scanning interface enables Enterprise Capture to use a wide variety of industry-leading document imaging. Enterprise Capture provides organizations with a single system to capture both paper and electronic documents scanners to digitize paper content. Existing electronic document files can be easily captured by users or automatically captured through an importing process that can monitor an email server or network folder. Once captured, documents are organized and indexed by applying metadata through manual or automated processes that use bar code recognition technology. After documents are completed, they are committed into a content management system. Enterprise Capture is fully integrated with Oracle WebCenter Content to provide organizations with one system to capture, store, manage, and retrieve their mission critical business content.

The following are recommendations that are specific to Oracle WebCenter Content Enterprise Capture:

Artifacts in the Database

Resides in Oracle WebCenter Enterprise Capture schema.

Artifacts in the File System

There are no file system artifacts for Enterprise Capture unless its JMS servers are stored in file-based persistent stores. The recommendation is to configure JMS servers with database-based persistent stores so that Oracle Data Guard is leveraged to replicate the database schema across sites.

Special Considerations

You should configure the load balancer virtual host, which is required for Oracle WebCenter Content Inbound Refinery, on both the production and the standby sites.

Synchronization Recommendations

Synchronize the directory tier manually with the standby site after you change the configuration or apply patches.

Recovery Recommendations

You must recover Enterprise Capture with the Enterprise Capture schema to the most recent point in time along with the Managed Server running the Oracle WebCenter Content Capture application, and the associated instances.

Recommendations for Oracle WebCenter Content: Content User Interface

WebCenter Content User Interface is a user-friendly and a feature-rich user interface for managing Oracle WebCenter Contented based on Oracle Application Development Framework.

The following are recommendations that are specific to Oracle WebCenter Content User Interface:

Artifacts in the Database

The Metadata Services schemas (`MDS`) are part of the Oracle WebCenter Content User Interface.

Special Considerations

You must configure load balancer virtual hosts for Oracle WebCenter Content User Interface on both the production and standby sites.

Synchronization Recommendations

Synchronize the directory tier manually with the standby site after you change the configuration or apply patches.

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

Recovery Recommendations

Recover the Managed Servers running the Oracle WebCenter Content User Interface application, along with the Administration Server.

A

Managing Oracle Inventory

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

This appendix includes the following sections:

- Updating Oracle Inventory
  Update Oracle Inventory on both the production site host and the standby site peer host.
- Updating the Windows Registry
  Export Windows Registry keys on the production site host and import them on the standby site peer host.

Updating Oracle Inventory

Update Oracle Inventory on both the production site host and the standby site peer host.

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, ensure that the same software updates are made on the standby site peer host. To accomplish this task, update Oracle inventory on the standby site peer host by executing the following script:

```
ORACLE_HOME/oui/bin/attachHome.sh
```

Updating the Windows Registry

Export Windows Registry keys on the production site host and import them on the standby site peer host.

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, 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
```

as illustrated in the following example:

```
regedit /I C:\oracleregisitry.reg
```

Alternatively, you can use the Registry Editor to import the key. See Registry Editor Help.