# Oracle® Application Server

Disaster Recovery Guide

10*g* Release 3 (10.1.3.3.0)

**E12297-02**

November 2008

ORACLE®

Oracle Application Server Disaster Recovery Guide, 10*g* Release 3 (10.1.3.3.0)

E12297-02



Primary Author: Bert Rich

Contributing Authors: Shailesh Dwivedi, Jay Feenan

Contributors: Pradeep Bhat, Fermin Castro, Gopal Kirsur, Shankar Raman, Bharath K. Reddy, Premson Rodriguez, Shilpa Shree, Prasad Vedurumudi

# Contents

# Preface

This preface contains these sections:



## Audience

This document is intended for administrators, developers, and others whose role is to deploy and manage the Oracle Application Server Disaster Recovery solution using disk mirroring technology.

## Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at `http://www.oracle.com/accessibility/`.

### Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

### Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

### TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, 7 days a week. For TTY support, call 800.446.2398. Outside the United States, call +1.407.458.2479.

## Related Documents

For more information, see the following documents in the Oracle Application Server documentation set:

- *Oracle Application Server Enterprise Deployment Guide*
- *Oracle Application Server High Availability Guide*
- *Oracle Application Server Administrator's Guide*

## Conventions

The following text conventions are used in this document:

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

# 1 Disaster Recovery Introduction

This chapter provides an introduction to the Oracle Application Server Disaster Recovery solution.

It contains the following topics:

- Disaster Recovery Overview
- Disaster Recovery for Oracle Application Server Components

## 1.1 Disaster Recovery Overview

This section provides an overview of Oracle Application Server Disaster Recovery.

It contains the following topics:

- Problem Description and Common Solutions
- Terminology

### 1.1.1 Problem Description and Common Solutions

Providing Maximum Availability Architecture is one of the key requirements for any Oracle Application Server Enterprise Deployment. Oracle Application Server includes an extensive set of High Availability features such as: Process Death Detection and Restart, Server Clustering, Load Balancing, Failover, Backup and Recovery, Rolling Upgrades, Rolling Configuration Changes, and Dynamic Discovery, which protect an Enterprise Deployment from unplanned down time and minimize planned downtime.

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

A core strategy for and a key feature of Oracle Application Server is Hot-Pluggability. Built for the heterogeneous enterprise, Oracle Application Server consists of modular component software that runs on a range of popular platforms and interoperates with middleware technologies and business applications from other software vendors such as IBM, Microsoft, and SAP. For instance, Oracle Application Server products and technologies such as ADF, Oracle BPEL Process Manager, Oracle Enterprise Service

Bus, Oracle Web Services Manager, Adapters, Oracle Access Manager, Oracle Identity Manager, Rules, Oracle TopLink, and Oracle Business Intelligence Publisher can run on non-Oracle containers such as BEA Weblogic, IBM Websphere and JBoss, in addition to running on Oracle's container (OC4J).

The Oracle Application Server Disaster Recovery solution uses disk replication technology for disaster protection of Oracle Application Server middle tier components. It supports Hot-Pluggable deployments, and it is compatible with third party vendor recommended solutions.

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

This document describes how to deploy the Oracle Application Server Disaster Recovery solution for an enterprise deployment, making use of disk replication technology and Oracle Data Guard technology.

The solution described in this document is a symmetric Oracle Application Server Disaster Recovery topology, which can be set up for Linux and UNIX operating systems.

### 1.1.2 Terminology

This section defines the following Disaster Recovery terminology:

- **Application Server host name**: This guide differentiates between the terms Application Server host name and network host name.

  The Application Server host name is the host name that Oracle Application Server uses for the host when Oracle Application Server is configured on the host. During installation, the installer automatically retrieves the Application Server host name from the current host and stores it in the Oracle Application Server configuration metadata on disk. A host can have only one Application Server host name.

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

- **asymmetric topology**: A disaster recovery configuration that is different across tiers on the production site and standby site. In an asymmetric topology, the standby site can use less hardware (for example, the production site could include four hosts with four Application Server instances while the standby site includes two hosts with four Application Server instances. Or, in a different asymmetric topology, the standby site can use fewer Application Server instances (for example, the production site could include four Application Server instances while the standby site includes two Application Server instances). Another asymmetric topology might include a different configuration for a database (for example, using a Real Application Clusters database at the production site and a single instance database at the standby site). Appendix C, "Creating an Asymmetric Topology" describes asymmetric topologies.

- **Disaster Recovery**: The ability to safeguard against natural or unplanned outages at a production site by having a recovery strategy for applications and data to a geographically separate standby site.

- **network host name**: A host name assigned to an IP address that is resolved through DNS resolution. The network host name is the host name by which a particular host is known within the host's network. A host can have the same network host name and Application Server host name. A host can have only one Application Server host name, but it can have multiple network host names.

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

- **production site setup**: The process of creating the production site. To create the production site using the procedure described in this manual, you must plan and create Application Server host names and network host names, create mount points and links on the hosts to the Oracle home directories on the shared storage where the Oracle Application Server instances will be installed, install the binaries and instances, and deploy the applications.
- **site failover**: The process of making the current standby site the new production site after the production site becomes unexpectedly unavailable (for example, due to a disaster at the production site). This book also uses the term "failover" to refer to a site failover.
- **site switchover**: The process of reversing the roles of the production site and standby site. Switchovers are planned operations done for periodic validation or to perform planned maintenance on the current production site. During a switchover, the current standby site becomes the new production site, and the current production site becomes the new standby site. This book also uses the term "switchover" to refer to a site switchover.
- **site synchronization**: The process of applying changes made to the production site at the standby site. For example, when a new application is deployed at the production site, you should perform a synchronization so that the same application will be deployed at the standby site, also.
- **standby site setup**: The process of creating the standby site. To create the standby site using the procedure described in this manual, you must plan and create Application Server host names and network host names, perform a switchover operation (which replicates the Oracle home directories and installations from the production site shared storage to the standby site shared storage), and create mount points and links to the Oracle home directories on the standby shared storage.
- **symmetric topology**: An Oracle Application Server Disaster Recovery configuration that is completely identical across tiers on the production site and standby site. In a symmetric topology, the production site and standby site have the identical number of hosts, load balancers, instances, and applications. The same ports are used for both sites. The systems are configured identically and the applications access the same data. This manual describes how to set up a symmetric Oracle Application Server Disaster Recovery topology for an enterprise configuration.
- **topology**: The production site and standby site hardware and software components that comprise an Oracle Application Server Disaster Recovery solution.

## 1.2 Disaster Recovery for Oracle Application Server Components

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

It contains the following topics:

### 1.2.1 Oracle Application Server Disaster Recovery Architecture Overview

This section describes the deployment architecture for Oracle Application Server components.

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

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

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

- Oracle Application Server product binaries, configuration, and metadata files

  Use disk 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 Application Server Disaster Recovery topology:

*Figure 1–1 Production and Standby Site for Oracle Application Server Disaster Recovery Topology*

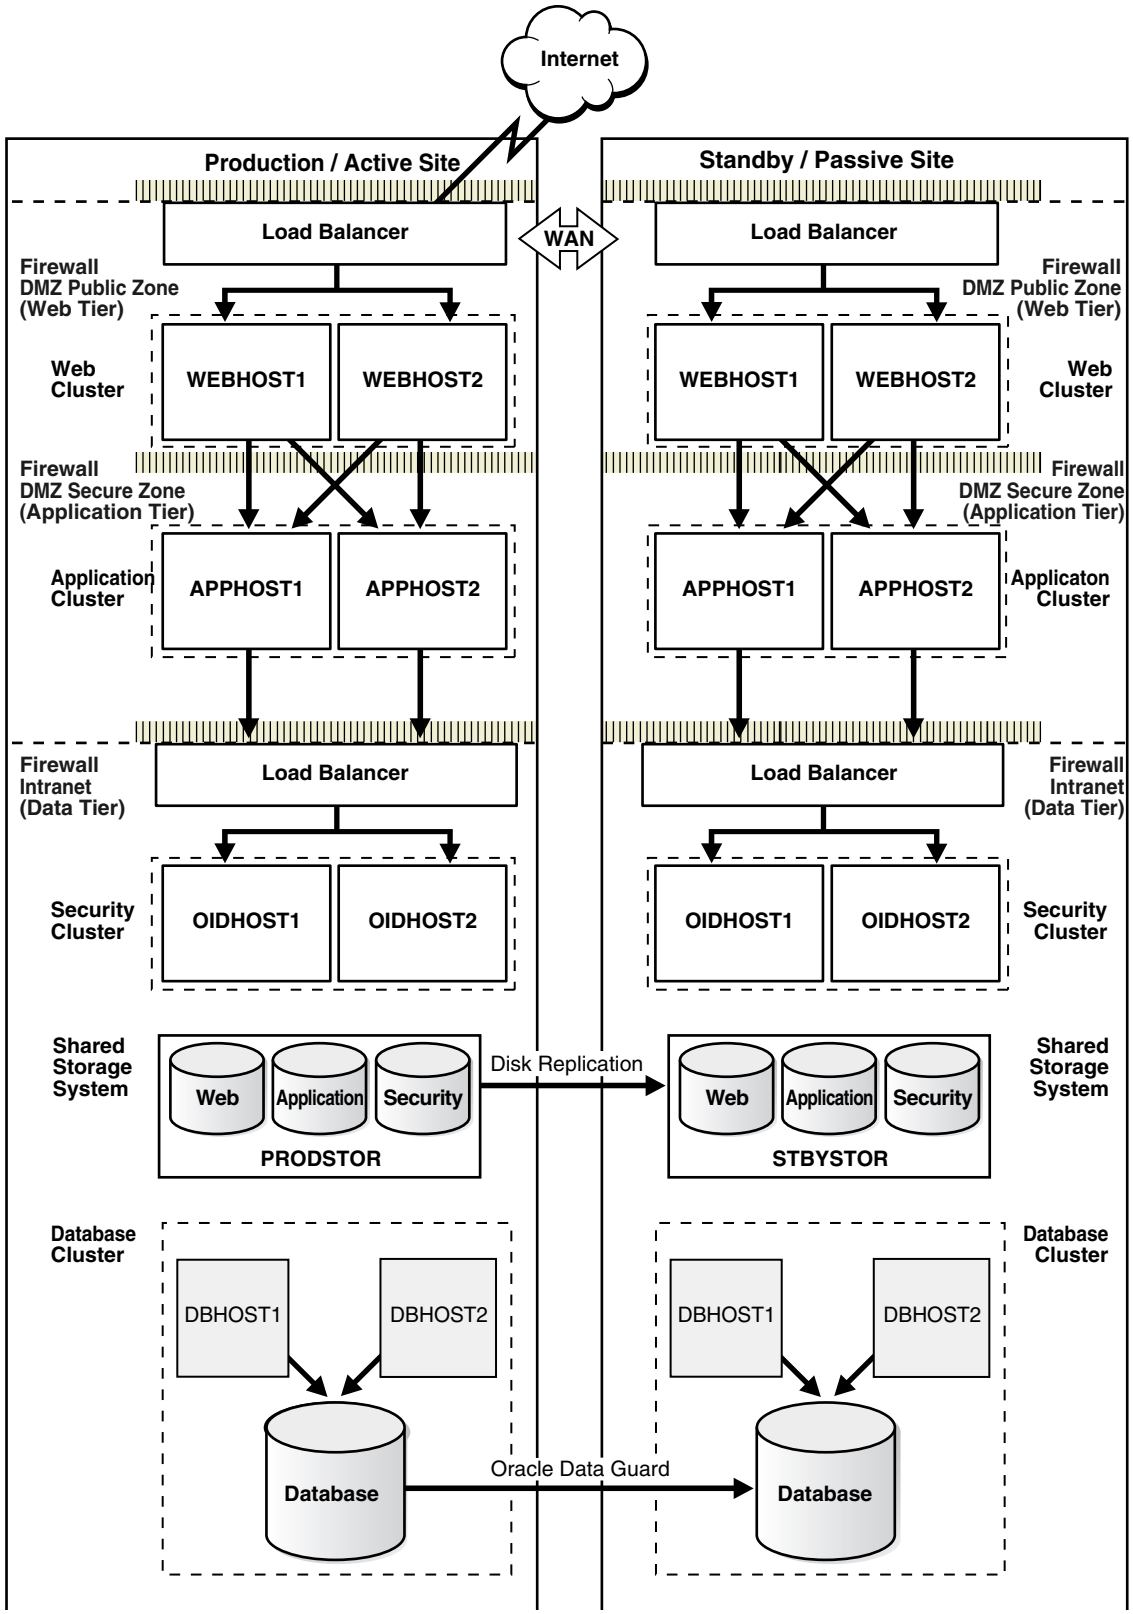


Some of the key aspects of the solution in Figure 1–1 are:

- The solution has two sites. The current production site is running and active, while the second site is serving as a standby site and is in passive mode.
- Hosts on each site have mount points defined for accessing the shared storage system for the site.
- On both sites, the Oracle Application Server components are deployed on the site's shared storage system. This involves creating all the Oracle home directories, which include product binaries and configuration data for middleware components, in volumes on the production site's shared storage and then installing the components into the Oracle home directories on the shared storage. In Figure 1–1, a separate volume is created in the shared storage for each Oracle Application Server 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 need to be created on the shared storage for the production site. The Oracle Application Server software for the production site will be installed into Oracle home directories using the mount points on the production site shared storage. Symbolic links also need to be set up on the production site hosts to the Oracle Application Server home directories on the shared storage at the production site.
- Mount points need to 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 Application Server home directories on the shared storage at the standby site. The mount points and symbolic links for the standby site hosts must be identical to those set up for the equivalent production site hosts.
- Disk replication technology is used to copy the middle tier file systems and other data from the production site's shared storage to the standby site's shared storage.
- After disk replication is enabled, application deployment, configuration, metadata, data, and product binary information is replicated from the production site to the standby site.
- It is not necessary to perform any Oracle software installations at the standby site hosts. When the production site storage is replicated at the standby site storage, the equivalent Oracle home directories and data are written to the standby site storage.
- Schedule incremental replications at a specified interval. The recommended interval is once a day for the production deployment, where the middle tier configuration does not change very often. Additionally, you should force a manual synchronization whenever you make a change to the middle tier configuration at the production site (for example, if you deploy a new application at the production site).
- Before forcing a manual synchronization, you should take a snapshot of the site to capture its current state. This ensures that the snapshot gets replicated to the standby site storage and can be used to roll back the standby site to a previous synchronization state, if desired. Recovery to the point of the previously successful replication (for which a snapshot was created) is possible when a replication fails.
- Oracle Data Guard is used to replicate all Oracle database repositories, including Oracle Application Server repositories and custom application databases. For information about using Oracle Data Guard to provide disaster protection for Oracle databases, see Section 2.2.1, "Database Considerations."
- If your Oracle Application Server Disaster Recovery topology includes any third party databases, use the vendor-recommended solution for those databases.

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

### 1.2.2 Components Described in this Document

The Oracle Application Server Disaster Recovery solution is supported for some Oracle Application Server components from these releases:

- Oracle Application Server 10.1.3 release
- Oracle Application Server 10.1.4.x Identity Management releases
- Oracle Application Server 10.1.2.x releases

Table 1–1 shows the components and applications for Oracle Application Server releases 10.1.3, 10.1.4, and 10.1.2.x Identity Management that are supported by Oracle Application Server Disaster Recovery:

***Table 1–1 Oracle Application Server Components and Applications Supported by Oracle Application Server Disaster Recovery***

| Component/Application | Oracle Application Server 10.1.3 | Oracle Application Server 10.1.4.x | Oracle Application Server 10.1.2.x |
|---|---|---|---|
| Oracle Web Cache | N/A | N/A | Y |
| Oracle HTTP Server | Y | Y | Y |
| Oracle Containers for J2EE (OC4J) | Y | Y | Y |
| Oracle BPEL Process Manager | Y | N/A | N |
| Oracle Enterprise Service Bus | Y | N/A | N/A |
| Oracle Web Services Manager | Y | N/A | N/A |
| Oracle B2B | N/A | N/A | Y |
| Oracle Business Activity Monitoring | Y | N/A | N/A |
| Oracle Business Intelligence | N | N/A | N/A |
| Oracle Identity Management (Oracle Internet Directory, Oracle Single Sign-On) | N/A | Y | N |
| Oracle Access Manager[1] | N/A | Y | N/A |
| Oracle Virtual Directory, Oracle Identity Manager | N/A | N | N |
| Oracle Discoverer, Oracle Forms, Oracle Reports, Oracle Wireless | N/A | N/A | Y |

*Table 1–1 (Cont.) Oracle Application Server Components and Applications Supported by Oracle Application Server Disaster Recovery*

| Component/Application | Oracle Application Server 10.1.3 | Oracle Application Server 10.1.4.x | Oracle Application Server 10.1.2.x |
|---|---|---|---|
| Oracle Portal | N/A | Y | Y |

[1] Oracle Access Manager ships as part of Oracle Identity Management releases instead of with Oracle Application Server releases.

## 1.2.3 Disaster Recovery Recommendations for Oracle Application Server Components

The following sections describe the disaster protection requirements for different Oracle Application Server components and provides recommendations for synchronizing these components. As mentioned previously, use disk replication to synchronize middle tier content stored in Oracle home directories, and use Oracle Data Guard to synchronize data in Oracle database repositories or custom application databases included in your Oracle Application Server Disaster Recovery topology.

### 1.2.3.1 Disaster Recovery Recommendations for Oracle Web Cache

This section describes the Oracle Web Cache data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle Web Cache server includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Web Cache is patched or upgraded.

- Product configuration

  This includes items such as `internal.xml`, `webcache.xml`, custom error pages, log files (which can be logged by Oracle Enterprise Manager for end user performance monitoring), and SSL wallets, which are updated when administrative operations (such as creating a new site definition, creating new origin servers, changing port numbers, configuring SSL, and so on) are performed.

#### Database Repository Dependencies

Oracle Web Cache does not store any configuration information in a database repository. There is no associated database repository.

#### Recommendations

For Oracle Application Server Disaster Recovery, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Web Cache middle tier file system through operations such as patching, upgrades, and configuration changes.

### 1.2.3.2 Disaster Recovery Recommendations for Oracle HTTP Server

This section describes the Oracle HTTP Server data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle HTTP Server includes the following in the middle tier file system:

- Product binaries

These typically change when Oracle HTTP Server is patched or upgraded.

- Product configuration

  This includes files such as `httpd.conf`, `ssl.conf`, and `mod_oc4j.conf`, which are updated when administrative operations (such as changing port numbers, changing the number of clients, creating virtual hosts, registering Oracle Single Sign-On, configuring SSL, and so on) are performed.

  If you are using Oracle Single Sign-On or Oracle Internet Directory with Oracle HTTP Server, read the recommendations for those components in Section 1.2.3.10, "Disaster Recovery Recommendations for Oracle Single Sign-On" and Section 1.2.3.9, "Disaster Recovery Recommendations for Oracle Internet Directory."

- Static HTML pages

  This content is deployed by administrators and is typically maintained outside the Oracle home directories. Make sure that the static HTML files are stored on the same shared storage as the Oracle HTTP Server installation so that it can be replicated from the production site to the standby site.

#### Database Repository Dependencies

Oracle HTTP Server has the following database dependencies:

- For 10.1.3 releases, all of Oracle HTTP Server's configuration is stored in the Oracle home on the middle tier files system. There is no associated database repository.
- For 10.1.2.x releases, Oracle HTTP Server is additionally maintained in the DCM repository in the Metadata Repository database. Thus, changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

For Oracle Application Server Disaster Recovery, a middle tier synchronization should be manually forced whenever a change is made to the Oracle HTTP Server middle tier configuration at the production site.

Additionally, for 10.1.2.x releases, Oracle Data Guard should be configured for Oracle database Metadata Repositories (as described later in Section 2.2.1, "Database Considerations"). Also for 10.1.2.x releases, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

### 1.2.3.3 Disaster Recovery Recommendations for Oracle Containers for J2EE

This section describes the Oracle Containers for J2EE (OC4J) data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

OC4J includes the following in the middle tier file system:

- Product binaries

  These typically change when OC4J is patched or upgraded.

- Container configuration

  This includes parameters such as JVM heap size, number of processes, replication configuration, creation of new web sites, port changes, and SSL configuration.

- Resource binding

  This includes container level configuration parameters such as data-source configuration, resource adapter configuration, and creation of JMS queues.

- Deployed applications

  This includes the binaries of the deployed applications and their related artifacts.

#### Database Repository Dependencies

OC4J has the following database dependencies:

- For 10.1.3 releases, all of the OC4J configuration is stored in the middle tier. There is no associated database repository.
- For 10.1.2.x releases, OC4J configuration is also maintained in the DCM repository in the Metadata Repository database. Thus, changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

For OC4J, a middle tier synchronization should be manually forced whenever a change is made to the OC4J middle tier file system. Changes to the OC4J middle tier file system include events such as configuration changes, resource binding changes, and application deployments.

For 10.1.2.x releases, Oracle Data Guard should be configured for Oracle database Metadata Repositories (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

### 1.2.3.4 Disaster Recovery Recommendations for Oracle BPEL Process Manager

This section describes the Oracle BPEL Process Manager data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle BPEL Process Manager includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle BPEL Process Manager is patched or upgraded.

- Process definitions

  For 10.1.2.x releases *only*, the process definitions in deployable BPEL suitcases are maintained on middle tier file systems and copied to an Oracle database repository. Thus, when a new process is deployed, the middle tier file system is updated.

  For 10.1.3 and later releases, when a suitcase is deployed to a BPEL Process Manager instance, the engine picks up the `suitcase.jar` file, writes it to the ORABPEL schema in the dehydration store database, and removes the `suitcase.jar` file from the middle tier file system. Temporary scratch files are generated during suitcase deployment, but these can be removed without consequence. All successfully loaded processes are copied to the dehydration store and removed from the middle tier file system. During engine startup, the deployed suitcases are downloaded from the dehydration store and unzipped on the middle tier file system.

- Configuration data

This includes JGroup configuration, resource adapter configuration, data source configuration, adapter configuration, and global fault policies.

#### Database Repository Dependencies

Oracle BPEL Process Manager stores the following metadata in the Oracle database dehydration store:

- Process definition

  The process definition changes when a new process is deployed.

- Process dehydrated state

  This is the run-time data generated when a process is dehydrated.

- For 10.1.2.x releases, OC4J configuration is also maintained in the DCM repository in the Oracle database Metadata Repository.

Changes in these repositories need to be replicated to the standby site's repositories.

#### Recommendations

For Oracle BPEL Process Manager, a middle tier synchronization should be manually forced whenever a change is made to the Oracle BPEL Process Manager middle tier file system. Changes to the Oracle BPEL Process Manager middle tier file system occur when you apply patches, deploy suitcases, change container configuration, change resource bindings, and deploy applications.

Oracle Data Guard should be configured for Oracle database Metadata Repositories (as described later in Section 2.2.1, "Database Considerations").

For 10.1.2.x releases, after deploying a new process, you should manually force a middle tier synchronization. You should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

For 10.1.3 releases, if new process definitions are propagated in Oracle BPEL Process Manager, you do not need to manually force a middle tier synchronization. You only need to manually force a database synchronization using Oracle Data Guard.

### 1.2.3.5 Disaster Recovery Recommendations for Oracle Enterprise Service Bus

This section describes the Oracle Enterprise Service Bus data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle Enterprise Service Bus includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Enterprise Service Bus is patched or upgraded.

- Oracle Enterprise Service Bus configuration data for OC4J

  This includes configuration data such as data source configuration, adapter configuration, and resource adapter configuration. Oracle Enterprise Service Bus cluster information is obtained from information in the repository at run time.

#### Database Repository Dependencies

Oracle Enterprise Service Bus stores the following metadata in its repository (ORAESB schema):

- Definitions such as system definitions, service definitions, and end point definitions. These get created or modified whenever a system, service, or end point is created or modified.
- JMS messages for asynchronous Oracle Enterprise Service Bus calls. These messages get created or updated with every asynchronous service invocation.
- In addition to storing JMS messages from asynchronous service invocations, database based JMS topics are used by Oracle Enterprise Service Bus to store error messages, retried messages, and instance tracking messages.

Changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

For Oracle Enterprise Service Bus, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Enterprise Service Bus middle tier configuration. Changes to the Oracle Enterprise Service Bus middle tier configuration include events such as applying patches, performing upgrades, configuration changes, and resource binding changes.

Oracle Data Guard should be configured for the Oracle database Metadata Repositories for Oracle Enterprise Service Bus (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

### 1.2.3.6 Disaster Recovery Recommendations for Oracle Web Services Manager

This section describes the Oracle Web Services Manager (Oracle WSM) data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle WSM includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle WSM is patched or upgraded.
- Configuration for gateways and agents
- Protected applications which get deployed on OC4Js

#### Database Repository Dependencies

Oracle WSM stores the following metadata in an Oracle database Metadata Repository (ORAWSM schema):

- Configuration data pertaining to deployed components, such as gateways and agents
- Policies pertaining to deployed applications (Policy Manager)
- Run-time metrics, which include metrics such as the number of requests processed, successful requests, and failed requests. These metrics are generated with every request.

Changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

For Oracle WSM, a middle tier synchronization should be manually forced whenever a change is made to the Oracle WSM middle tier file system. Changes to the Oracle

WSM middle tier file system occur when you apply patches, perform upgrades, change gateway or agent onfiguration, change container level configuration, deploy applications, and change resource bindings.

Oracle Data Guard should be configured for the Oracle database Metadata Repositories for Oracle WSM (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

#### 1.2.3.7 Disaster Recovery Recommendations for B2B

This section describes the Oracle B2B data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

##### Middle Tier Configuration

Oracle B2B includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle B2B is patched or upgraded.

- Configuration data

  This includes tip.properties configuration for Oracle B2B.

##### Database Repository Dependencies

Oracle B2B stores the following metadata in the Oracle database store:

- B2B metadata:

  This includes the document, trading partner and agreement metadata.

- B2B run-time data:

  This is the run-time data generated when a message is processed through B2B.

Changes in the database store need to be replicated to the standby site's database store.

##### Recommendations

For Oracle B2B, a middle tier synchronization should be manually forced whenever a change is made to the Oracle B2B middle tier file system. Changes to the Oracle B2B middle tier file system occur when you apply patches, deploy suitcases, change container configuration, change resource bindings, and deploy applications.

Oracle Data Guard should be configured for the Oracle database store (as described later in Section 2.2.1, "Database Considerations").

#### 1.2.3.8 Disaster Recovery Recommendations for Oracle Business Activity Monitoring

This section describes the Oracle Business Activity Monitoring data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

##### Middle Tier Configuration

Oracle Business Activity Monitoring includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Business Activity Monitoring is patched or upgraded.

- Configuration data

  This includes configuration data such as logging configuration and database connection configuration.

In Oracle Application Server 10.1.x releases, all Oracle Business Activity Monitoring data is stored in database repositories, therefore Oracle Business Activity Monitoring data protection will not be performed using disk replication.

Because disk replication is not used for Oracle Business Activity Monitoring in 10.1.x releases, when you install Oracle Business Activity Monitoring at the production site, you must perform an identical Oracle Business Activity Monitoring installation at the standby site (by specifying the same Oracle home name, the same Oracle home directory and path name, and the same configuration options).

If you patch a Oracle Business Activity Monitoring installation on the production site, you must patch the peer Oracle Business Activity Monitoring installation on the standby site.

#### Database Repository Dependencies

The following Oracle Business Activity Monitoring data is stored in the Oracle database store:

- Reports, alerts, users, roles, distribution lists, security filters, parameters, message source data, lookup data, alert history, and other data.

  These definitions change when any modifications or new artifacts are created.

#### Recommendations

See Section B, "Setting Up Oracle Business Activity Monitoring" for the steps to follow to set up Oracle Business Activity Monitoring properly in an Oracle Application Server Disaster Recovery topology.

### 1.2.3.9 Disaster Recovery Recommendations for Oracle Internet Directory

This section describes the Oracle Internet Directory data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle Internet Directory includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Internet Directory is patched or upgraded.

- Oracle Internet Directory configuration

  This includes items such as the TNSNAMES.ORA file used for database access, the SQLNET.ORA file, and wallet configuration.

#### Database Repository Dependencies

Oracle Internet Directory stores the following in an Oracle database repository:

- All of its user data
- Configuration data (such as the different LDAP server instances to run and Directory Integration Platform configuration information)

Changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

For Oracle Internet Directory, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Internet Directory middle tier file system. Changes to the Oracle Internet Directory middle tier configuration occur when you apply patches, perform upgrades, and make configuration changes for Oracle Internet Directory.

> **Note:** If your deployment uses only Oracle Internet Directory and Single Sign-On Server and no other Oracle Application Server components or applications, then you can consider using multimaster replication as a Disaster Recovery solution. However, if your deployment includes other Oracle Application Server components or applications in addition to Oracle Internet Directory and Single Sign-On Server, then you should use disk replication for the Disaster Recovery solution. See the "Deploying Identity Management with Multimaster Replication" chapter in the *Oracle Application Server High Availability Guide* for Oracle Application Server 10.1.4.0.1 for more information about multimaster replication.

Oracle Data Guard should be configured for the Oracle database repositories used by Oracle Internet Directory (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

### 1.2.3.10 Disaster Recovery Recommendations for Oracle Single Sign-On

This section describes the Oracle Single Sign-On data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle Single Sign-On includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Single Sign-On is patched or upgraded.

- Oracle HTTP Server configuration

  This includes items described in Section 1.2.3.2, "Disaster Recovery Recommendations for Oracle HTTP Server" such as `mod_osso` configuration files for partner application registration, SSL configuration (if SSL is enabled), and location of wallet files.

- OC4J_Security configuration

  This includes items described in Section 1.2.3.3, "Disaster Recovery Recommendations for Oracle Containers for J2EE" such as state replication for Delegated Administration Services, data source configuration, and changes made to `iasconfig.xml`.

#### Database Repository Dependencies

Oracle Single Sign-On stores the following in an Oracle database Metadata Repository:

- Oracle Single Sign-On metadata

  Metadata, which includes metadata for registered partner applications and external application configuration, is stored in the ORASSO schema in the Oracle Internet Directory Metadata Repository. See also the recommendations for Oracle

Internet Directory in Section 1.2.3.9, "Disaster Recovery Recommendations for Oracle Internet Directory."

- Oracle HTTP Server and OC4J configuration data is stored in the DCM repository in the Oracle database Metadata Repository.

Changes in these repositories need to be replicated to the standby site's repositories.

#### Recommendations

For Oracle Single Sign-On, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Single Sign-On middle tier file system. Changes to the Oracle Single Sign-On middle tier configuration occur when you apply patches, perform upgrades, and make configuration changes for Oracle Single Sign-On.

Oracle Data Guard should be configured for Oracle database Metadata Repositories used by Oracle Internet Directory (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

### 1.2.3.11 Disaster Recovery Recommendations for Oracle Access Manager

This section describes the Oracle Access Manager data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle Access Manager includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Access Manager is patched or upgraded.

- Oracle Access Manager configuration data

  The Oracle Access Manager configuration data store is centralized in a LDAP server directory.

  The Oracle Access Manager components read configuration data from .xml files automatically generated out of the configuration store.

  These typically change whenever configurations changes are made through the Policy Manager or when it is necessary to edit the .xml files directly.

- Oracle Access Manager policy data

  This is located in an LDAP directory server which could be the same LDAP directory server as used for Oracle Access Manager configuration.

  This changes when policy configuration changes are made through the Policy Manager.

- Audit data is generated either into files or into the database

  This is added to as users authenticate or access protected resources.

#### Recommendations

For Oracle Access Manager, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Access Manager middle tier file system. Changes to the Oracle Access Manager middle tier configuration occur as a result of applying patches, performing upgrades, and making configuration changes for Oracle Access Manager.

Oracle Data Guard should be configured for the Oracle database repositories used by Oracle Access Manager directly (for audit data only), as described later in Section 2.2.1, "Database Considerations." Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

#### 1.2.3.12 Disaster Recovery Recommendations for Oracle Discoverer

This section describes the Oracle Discoverer data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

##### Middle Tier Configuration

Oracle Discoverer includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Discoverer is patched or upgraded.

- Oracle HTTP Server configuration (if it is installed in the same Oracle home directory as Oracle Discoverer)

  This includes items described in Section 1.2.3.2, "Disaster Recovery Recommendations for Oracle HTTP Server" such as `mod_osso` configuration files for partner application registration, SSL configuration (if SSL is enabled), and location of wallet files.

- OC4J configuration

  This includes items described in Section 1.2.3.3, "Disaster Recovery Recommendations for Oracle Containers for J2EE" and Oracle Discoverer product configuration which includes items such as End User Layer database configuration and the Oracle Discoverer preference store.

##### Database Repository Dependencies

Oracle Discoverer stores the following in an Oracle database repository:

- Oracle Discoverer End User Layer
- Configuration for components generated by Oracle Discoverer

  The configuration for items such as database connections, workbooks, and worksheets is stored in the database.

Changes in this repository need to be replicated to the standby site's repository.

##### Recommendations

For Oracle Discoverer, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Discoverer middle tier configuration. Changes to the Oracle Discoverer middle tier configuration occur when you apply patches, perform upgrades, and change the configuration for Oracle Discoverer components.

Oracle Data Guard should be configured for the Oracle database Metadata Repositories for Oracle Discoverer (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

#### 1.2.3.13 Disaster Recovery Recommendations for Oracle Forms

This section describes the Oracle Forms data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

### Middle Tier Configuration

Oracle Forms includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Forms is patched or upgraded.

- Oracle Forms server configuration

  The Oracle Forms server configuration is stored in configuration files such as `formsweb.cfg` under the `forms/server` directory in the Oracle home. Oracle Forms configuration includes items such as servlet configuration, Oracle Single Sign-On configuration, and HTML templates to use.

  If you are using Oracle Single Sign-On or Oracle Internet Directory with Oracle Forms, read the recommendations for those components in Section 1.2.3.9, "Disaster Recovery Recommendations for Oracle Internet Directory" and Section 1.2.3.10, "Disaster Recovery Recommendations for Oracle Single Sign-On."

- Oracle Forms deployment artifacts

  Oracle Forms can create special files such as .fmx and .fmb files outside the Oracle home directory. The location of these Oracle Forms files is typically configured on each middle tier using the FORMS_PATH variable. When you are using Oracle Forms as part of an Oracle Application Server Disaster Recovery topology, these special files and artifacts of Oracle Forms should be created in the Oracle home.

- OC4J configuration

  This includes files such as `web.xml`, `orion-web.xml`, and `application.xml`. Read the recommendations for OC4J in Section 1.2.3.3, "Disaster Recovery Recommendations for Oracle Containers for J2EE."

- Oracle HTTP Server configuration

  This includes files such as `forms.conf`. Read the recommendations for Oracle HTTP Server in Section 1.2.3.2, "Disaster Recovery Recommendations for Oracle HTTP Server."

### Database Repository Dependencies

Oracle Forms server does not require a database repository. However, Oracle Forms server is typically used to query and modify customer data in customer databases, which should be properly protected.

### Recommendations

For Oracle Forms server, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Forms middle tier file system. Changes to the Oracle Forms middle tier file system occur when you apply patches, perform upgrades, and change the configuration for the Oracle Forms server component.

When you are using Oracle Reports with Oracle Forms, you should manually force a middle tier synchronization whenever a change is made to either the Oracle Forms or Oracle Reports configuration. Read the recommendations for Oracle Reports in Section 1.2.3.14, "Disaster Recovery Recommendations for Oracle Reports."

Oracle Data Guard should be configured for any application database used with an Oracle Forms application (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

#### 1.2.3.14 Disaster Recovery Recommendations for Oracle Reports

This section describes the Oracle Reports data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

##### Middle Tier Configuration

Oracle Reports includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Reports is patched or upgraded.

- Oracle Reports server configuration

  The Oracle Reports server configuration is stored in configuration files such as `server_name.confg`, `rwnetwork.conf`, and `rwbuilder.conf` under the `reports/conf` directory in the Oracle home. Oracle Reports configuration includes items such as initial settings for the Oracle Reports server cache, engine initialization, security, and cluster configuration.

  Oracle Reports server job details for scheduled jobs, past jobs, and current jobs are stored in the `<server-name>.dat` file in the `reports/server` directory in the Oracle home.

  If you are using Oracle Single Sign-On or Oracle Internet Directory with Oracle Reports, read the recommendations for those components in Section 1.2.3.9, "Disaster Recovery Recommendations for Oracle Internet Directory" and Section 1.2.3.10, "Disaster Recovery Recommendations for Oracle Single Sign-On."

- Oracle Reports deployment artifacts

  Oracle Reports can create special files such as .rdf files outside the Oracle home directory. The location of these Oracle Reports files is typically configured on each middle tier using the REPORTS_PATH variable. When you are using Oracle Reports as part of an Oracle Application Server Disaster Recovery topology, these special files of Oracle Reports should be created in the Oracle home.

- OC4J configuration

  The Oracle Reports servlet is deployed in the OC4J_BI_Forms container. Configuration pertaining to the Oracle Reports servlet is stored in the `rwservlet.properties` and `cgicmd.dat` files in the `$ORACLE_HOME/reports/conf` directory. These files should be protected. Also, if you have made any changes to OC4J-specific files for the OC4J_BI_Forms container (such as changes to `web.xml`, `orion-web.xml`, and `application.xml`), these files should also be protected. Read the recommendations for OC4J in Section 1.2.3.3, "Disaster Recovery Recommendations for Oracle Containers for J2EE."

- Oracle HTTP Server configuration

  This includes any configuration done specifically for the Oracle Reports application in addition to other Oracle HTTP Server-specific configuration. Read the recommendations for Oracle HTTP Server in Section 1.2.3.2, "Disaster Recovery Recommendations for Oracle HTTP Server."

##### Database Repository Dependencies

Oracle Reports server does not require a database repository. However, Oracle Reports custom applications can use a custom Oracle database repository. Data that Oracle Reports server uses for generating reports, as well as job status details, can be stored in the Oracle database repository.

Changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

For Oracle Reports server, a middle tier synchronization should be manually forced whenever a change is made to the Oracle Reports middle tier file system. Changes to the Oracle Reports middle tier file system occur when you apply patches, perform upgrades, and make configuration changes for Oracle Reports server.

When you are using Oracle Forms with Oracle Reports, you should manually force a middle tier synchronization whenever a change is made to either the Oracle Reports or Oracle Forms configuration. Read the recommendations for Oracle Forms in Section 1.2.3.13, "Disaster Recovery Recommendations for Oracle Forms."

Oracle Data Guard should be configured for any application database used with an Oracle Reports application (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

### 1.2.3.15 Disaster Recovery Recommendations for Oracle Portal

This section describes the Oracle Portal data that must be protected as part of your Oracle Application Server Disaster Recovery solution.

#### Middle Tier Configuration

Oracle Portal includes the following in the middle tier file system:

- Product binaries

  These typically change when Oracle Portal is patched or upgraded.

- Oracle Portal configuration

  The Oracle Portal configuration is stored in configuration files such as `dads.conf`, `plsql.conf`, `cache.conf`, `mod_oc4j.conf`, `cache.xml`, `portal.conf`, `portalReg.xml`, `portalconsoleConfigSNSegment.xml`, and `consoleConfigSNSegment.xml`. Other configuration items such as the Oracle Internet Directory location are stored in other configuration files such as `iasconfig.xml`.

  If you are using Oracle Single Sign-On or Oracle Internet Directory with Oracle Portal, read the recommendations for those components in Section 1.2.3.9, "Disaster Recovery Recommendations for Oracle Internet Directory" and Section 1.2.3.10, "Disaster Recovery Recommendations for Oracle Single Sign-On."

- Oracle Portal provider registration configuration

  The Oracle Portal provider registration configuration is stored in `provider.xml` files.

- OC4J configuration

  This includes files such as `web.xml`, `orion-web.xml`, and `application.xml`. Read the recommendations for OC4J in Section 1.2.3.3, "Disaster Recovery Recommendations for Oracle Containers for J2EE."

- Oracle HTTP Server configuration

  This includes security configuration for the Oracle Portal application in addition to other Oracle HTTP Server-specific configuration. Read the recommendations for Oracle HTTP Server in Section 1.2.3.2, "Disaster Recovery Recommendations for Oracle HTTP Server."

- Oracle Web Cache configuration

  This includes configuration items such as the site definition, site to server mapping, and origin servers. Read the recommendations for Oracle Web Cache in Section 1.2.3.1, "Disaster Recovery Recommendations for Oracle Web Cache."

#### Database Repository Dependencies

Oracle Portal stores all the page information, portlets, and content information in its Oracle database repository.

Changes in this repository need to be replicated to the standby site's repository.

#### Recommendations

A middle tier synchronization should be manually forced whenever a change is made to the Oracle Portal middle tier file system. Changes to the Oracle Portal middle tier file system occur when you apply patches, perform upgrades, and change configuration for the Oracle Portal component.

Oracle Data Guard should be configured for the Oracle database repository used by Oracle Portal (as described later in Section 2.2.1, "Database Considerations"). Also, you should manually force a database synchronization using Oracle Data Guard whenever a middle tier synchronization is manually forced.

# 2

# Implementing the Solution

This chapter describes how to implement Oracle Application Server Disaster Recovery solution for an enterprise deployment.

It contains the following topics:

- Design Considerations
- Environment Preparation
- Finishing the Disk Replication Setup
- Synchronization Steps and Frequency
- Failover Steps
- Switchover Steps
- Performing Periodic Testing of the Standby Site

Figure 2–1 shows the mySOACompany with Oracle Single Sign-On deployment from the *Oracle Application Server Enterprise Deployment Guide*. The *Oracle Application Server Enterprise Deployment Guide* describes how to install the Oracle Application Server instances on the five Oracle Application Server host clusters (OIDHOST1 and OIDHOST2, WEBHOST1 and WEBHOST2, APPHOST1 and APPHOST2, WEBHOST3 and WEBHOST4, and IDMHOST1 and IDMHOST2) and how to perform the necessary configuration for the deployment.

This manual describes how to set up this deployment at an Oracle Application Server Disaster Recovery production site and standby site for the Linux and UNIX operating systems.

> **Note:** This manual describes an Oracle Application Server Disaster Recovery symmetric topology that uses the deployment shown in Figure 2–1 at *both* the production site and the standby site. Figure 2–1 shows the deployment for only one site; the high level of detail shown for this deployment precludes showing the deployment for both sites in a single figure.

This manual refers to the deployment in Figure 2–1 as the EDG deployment. However, unlike the *Oracle Application Server Enterprise Deployment Guide*, this manual directs you to install the Oracle Application Server instances that will be used by the Oracle Application Server host clusters into Oracle home directories in volumes on shared storage at the production site. Shared storage will also be set up for the standby site. Then, Oracle Application Server data on the production site storage will be replicated

asynchronously to the standby site storage as part of the  Oracle Application Server Disaster Recovery solution.

***Figure 2–1 Deployment Used at Production and Standby Sites for Oracle Application Server Disaster Recovery Symmetric Topology***

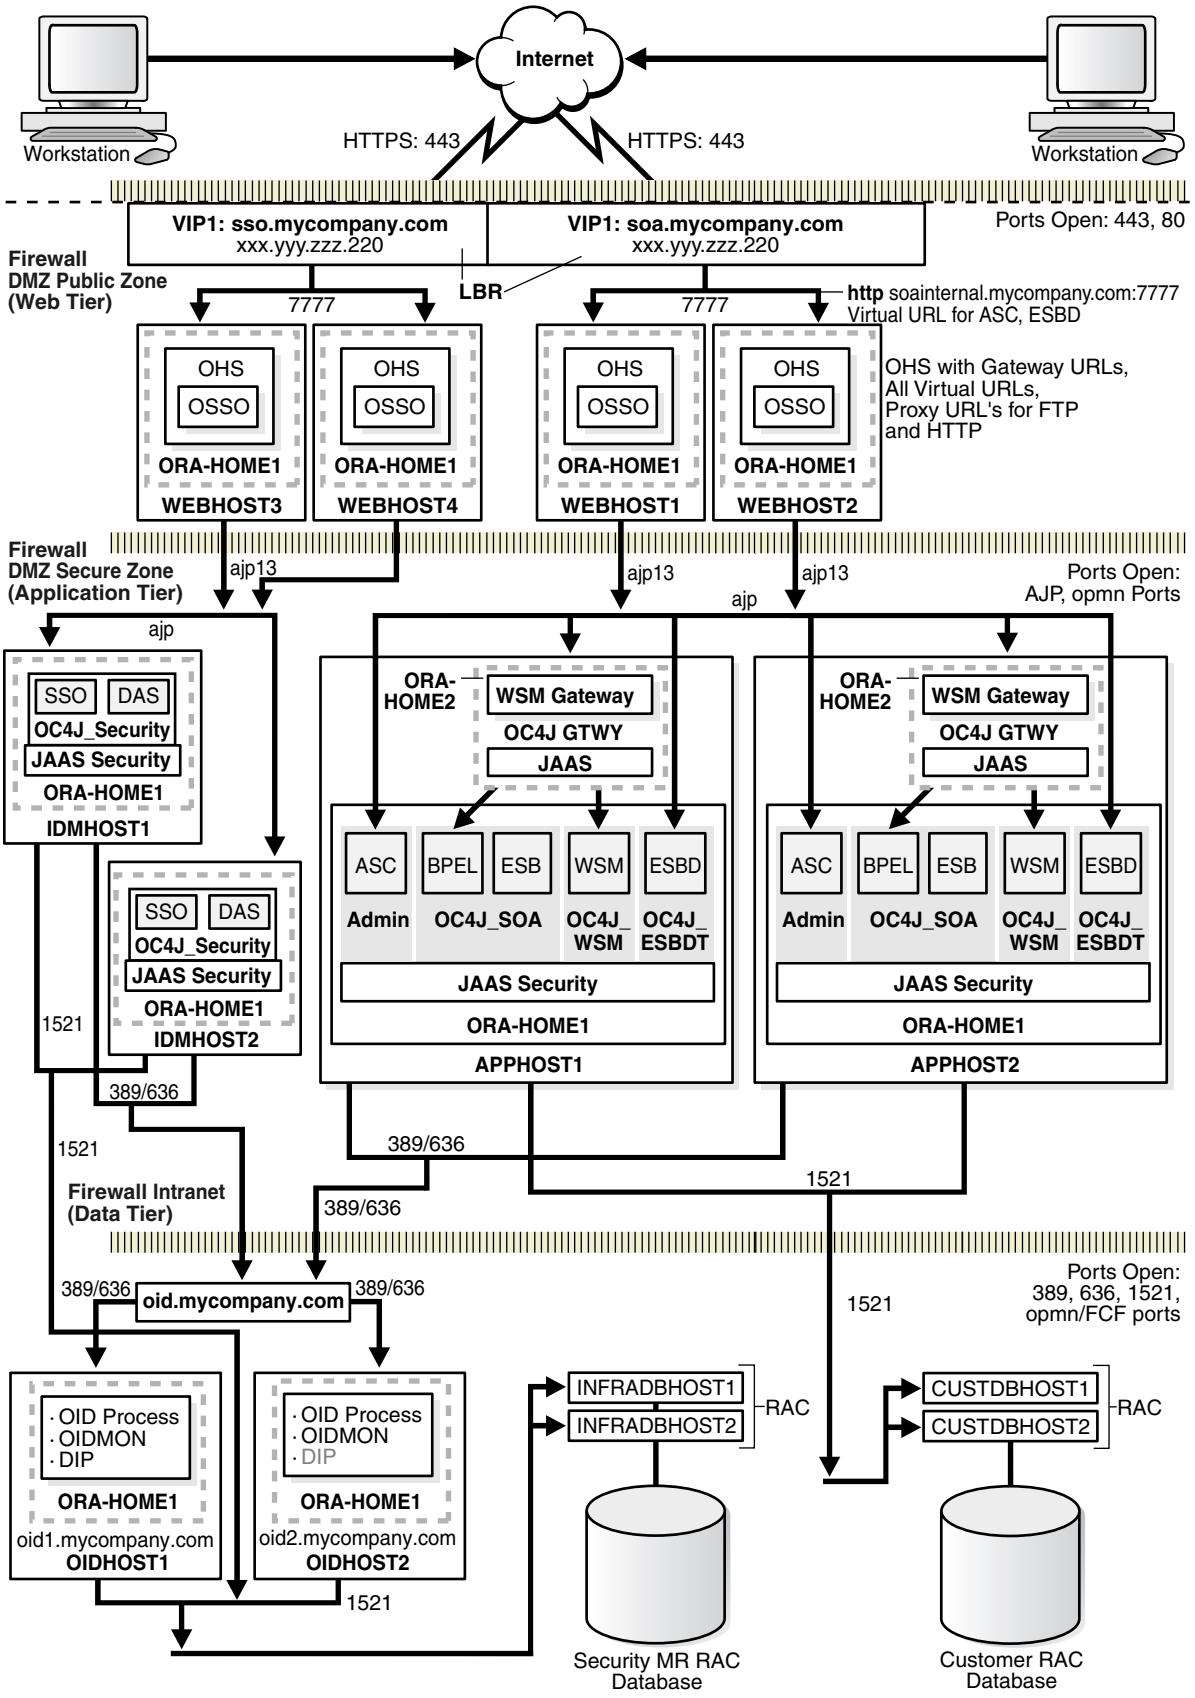
Internet
Workstation
HTTPS: 443
HTTPS: 443
Workstation
Ports Open: 443, 80
VIP1: sso.mycompany.com
xxx.yyy.zzz.220
VIP1: soa.mycompany.com
xxx.yyy.zzz.220
Firewall
DMZ Public Zone
(Web Tier)
LBR
7777
7777
http soainternal.mycompany.com:7777
Virtual URL for ASC, ESBD
OHS
OSSO
ORA-HOME1
WEBHOST3
WEBHOST4
WEBHOST1
WEBHOST2
OHS with Gateway URLs,
All Virtual URLs,
Proxy URL's for FTP
and HTTP
Firewall
DMZ Secure Zone
(Application Tier)
ajp13
ajp13
ajp13
ajp
Ports Open:
AJP, opmn Ports
ajp
SSO
DAS
OC4J_Security
JAAS Security
ORA-HOME1
IDMHOST1
IDMHOST2
ORA-
HOME2
WSM Gateway
OC4J GTWY
JAAS
ASC
BPEL
ESB
WSM
ESBD
Admin
OC4J_SOA
OC4J_
WSM
OC4J_
ESBDT
JAAS Security
ORA-HOME1
APPHOST1
APPHOST2
1521
389/636
1521
389/636
1521
Firewall Intranet
(Data Tier)
389/636
Ports Open:
389, 636, 1521,
opmn/FCF ports
389/636
oid.mycompany.com
389/636
1521
· OID Process
· OIDMON
· DIP
ORA-HOME1
oid1.mycompany.com
OIDHOST1
oid2.mycompany.com
OIDHOST2
1521
INFRADBHOST1
INFRADBHOST2
RAC
CUSTDBHOST1
CUSTDBHOST2
RAC
Security MR RAC
Database
Customer RAC
Database

The EDG deployment is described in detail in the *Oracle Application Server Enterprise Deployment Guide*.

This manual describes how to set up and deploy an Oracle Application Server Disaster Recovery topology on Linux and UNIX systems. The Oracle Application Server Disaster Recovery symmetric topology described in this manual uses the EDG deployment at both the production site and standby site.

## 2.1 Design Considerations

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

- Application Server host names

  The same Application Server host name must be used for each production site host and its standby site peer host. The Application Server host name for a host is either specified when Oracle Application Server is installed or derived from the host at installation time. A host's Application Server host name is stored in the Oracle Application Server Disaster Recovery configuration.

  It is recommended that a host have a different Application Server host name and network host name. See Section 1.1.2, "Terminology" for the definitions of Application Server host name and network host name.

  In some cases, a host will use the same host name as its Application Server host name and network host name. This situation is described in Section 2.1.2, "Designing with an Existing Production Site."

- Directory names and paths

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

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

- Port numbers

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

  Section 2.1.2, "Designing with an Existing Production Site" describes how to check for port conflicts between production site and standby site hosts.

- Security

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

- Load balancers and virtual server names

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

- Software

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

The following design topics are included in this section:

- Design Starting Point
- Designing with an Existing Production Site
- Designing a New Production Site and Standby Site
- Planning Host Names for the Production Site and Standby Site
- Host Name Resolution
- Making /etc/hosts File Entries for Local Host Name File Resolution
- Resolving Host Names Using DNS Host Name Resolution

## 2.1.1 Design Starting Point

Before setting up the standby site, the administrator must evaluate the starting point of the project. The starting point for designing a symmetrical Oracle Application Server Disaster Recovery topology is usually one of the following:

- The production site is already created and the standby site is being planned and created.

  Section 2.1.2, "Designing with an Existing Production Site" describes how to design the Oracle Application Server Disaster Recovery standby site when you have an existing production site.

- There is no existing production site or standby site. Both need to be designed and created.

  Section 2.1.3, "Designing a New Production Site and Standby Site" describes how to design the Oracle Application Server Disaster Recovery production site when you do not have an existing production site or standby site.

- Some hosts or components may exist at a current production site, but new hosts or components need to be added at that site or at a standby site to set up a functioning Oracle Application Server Disaster Recovery topology.

  Use the pertinent information in this chapter to design and implement an Oracle Application Server Disaster Recovery topology.

## 2.1.2 Designing with an Existing Production Site

When the administrator's starting point is an existing production site, the configuration data for the production site already exists. For example, for any Oracle software already installed, the Oracle home names and the paths to the Oracle home directories exist on the file system. Also, the host names, ports, and user accounts are already defined. When a production site exists, the administrator can choose to either re-create the production site from scratch (as described in Section 2.1.3, "Designing a New Production Site and Standby Site") or to create a symmetric standby site for the existing production site, as described in this section.

The following sections describe how to meet the requirements for symmetric ports, Application Server host names, storage configuration, and Oracle Central Inventories for the production site and standby site when the production site already exists.

#### 2.1.2.1 Implementing Symmetric Host Names with an Existing Production Site

When Oracle Application Server components have already been installed at an existing production site, a production site host will most likely have the same host name as its Application Server host name and its network host name. For example, if the network host name for an existing production site host is PROD001, then the Application Server host name for the host will most likely also be PROD001, because the Oracle Application Server installation will use the network name of the host as its Application Server host name.

The recommended Oracle Application Server Disaster Recovery configuration is for each production site host and each standby site host to have an Application Server host name that is different than its network host name (host name planning is described in detail in Section 2.1.4, "Planning Host Names for the Production Site and Standby Site"). However, this recommendation is difficult to implement with an existing production site where some or all of hosts were previously set up to use the same host name as the Application Server host name and network host name for the host. If you attempt to change the existing Application Server host name for a host, it can break applications that have been installed at the site and cause other issues that can be difficult to diagnose and fix. Therefore, if the same host name is being used for both the Application Server host name and network host name for a production site host, you should not change the existing Application Server host name.

When the production site already exists and the production site hosts have the same host name as both the Application Server host name and network host name, the standby site hosts must be given the same Application Server host names as their peer hosts on the production site. An example of using the Application Server host name for a production site host as the Application Server host name for the peer host at the standby site is shown in Table 2–1. Specifically, production site host PROD001 has an Application Server host name of PROD001, so its peer host STBYWEB1 on the standby site must also use PROD001 as its Application Server host name. Similarly, production site host PROD002 has an Application Server host name of PROD002, so its peer host STBYWEB2 on the standby site must also use PROD002 as its Application Server host name PROD002.

***Table 2–1 Using Production Site Application Server Host Names as the Application Server Host Names for Standby Site Hosts***

| IP Address[1] | Site | Application Server Host Name | Network Host Name[2] |
|---|---|---|---|
| 123.1.2.113 | Production | PROD001[3] | PROD001 |
| 123.1.2.114 | Production | PROD002[4] | PROD002 |
| 123.2.2.113 | Standby | PROD001 | STBYWEB1 |
| 123.2.2.114 | Standby | PROD002 | STBYWEB2 |

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

[2] See Section 2.1.7, "Resolving Host Names Using DNS Host Name Resolution" for information on defining network host names.

[3] This Application Server host name is added for illustration purposes; technically, the absence of a resolution mechanism from the Application Server host name to the network host name is a valid configuration for these hosts. (technically, an Application Server host name does not need to be assigned to a production site host, in which case the network host name is used as the Application Server host name.

[4] This Application Server host name is added for illustration purposes; technically, the absence of a resolution mechanism from the Application Server host name to the network host name is a valid configuration for these hosts. (technically, an Application Server host name does not need to be assigned to a production site host, in which case the network host name is used as the Application Server host name.

If your Oracle Application Server Disaster Recovery solution includes a pre-existing production site, then after a failover operation or switchover, the original standby site assumes the production role. In this situation, the Application Server host names for the new production site (original standby site) are the same as the Application Server and network host names of the original production site. This can cause administrative confusion, because the Application Server host names that appear in the display and in logs at the new production site are the same as the host names used at the original production site.

#### 2.1.2.2 Implementing Symmetric Ports with an Existing Production Site

Standby site hosts must use the same port numbers as their peer hosts at the production site. To prevent port conflicts between production site and standby site peer hosts, it is recommended that the administrator get a list of the port numbers for each production site host and confirm that the same ports are used for the same components at the standby site peer host.

Begin by getting a copy of `portlist.ini` file for each of the existing Oracle homes used at the production site. The contents of this file contains the port numbers that were assigned during the installation process.

Then, at the standby site peer host, you can use the netstat utility to ensure that the ports in use at the standby site are the same as those in use the production site (note that the netstat utility will only show ports in use by running software; it does not show ports used by transient run-time components if those components are not active).

For instance, the first part of Example 2–1 shows some of the ports in the `portlist.ini` file for an Oracle Application Server installation at a production site host. Note that port 7777 was assigned to Oracle HTTP Server at the production site host.

In the second part of Example 2–1, after the standby site is created, the netstat utility is used on the standby site peer host to check whether port 7777 is assigned to Oracle HTTP Server on the standby site peer host.

***Example 2–1 Checking Ports Usage for an Existing Host***

```
[Ports]
Oracle HTTP Server port =  7777
Oracle HTTP Server Listen port = 7777
Oracle HTTP Server SSL port = 4443
Oracle HTTP Server Listen (SSL) port = 4443
Oracle HTTP Server Diagnostic port = 7200
   .
   .
   .

prompt> netstat -an | grep 7777
```

Port usage for each standby site host must be identical to the port usage for the production site peer host.

#### 2.1.2.3 Implementing Symmetric Storage Configuration with an Existing Production Site

The Oracle Application Server Disaster Recovery solution relies on shared storage to implement disk replication for disaster protection of the Oracle Application Server middle tier configuration. When a production site has already been created, it is likely

that the Oracle home directories for the Oracle Application Server instances that comprise the site are not located on the shared storage. If this is the case, then these homes have to be migrated completely to the shared storage to implement the Oracle Application Server Disaster Recovery solution.

Information about setting up the shared storage for the Oracle Application Server Disaster Recovery solution is provided in Section 2.2.2, "Creating Volumes, Mount Points, and Symbolic Links."

#### 2.1.2.4 Implementing Symmetric Oracle Central Inventories with an Existing Production Site

The Oracle Central Inventory stores information about all Oracle software products installed in all the Oracle homes on a host (provided the product was installed using Oracle Universal Installer). Some Oracle components, including the Oracle Opatch utility, must have a complete inventory of the software components installed on the host.

The location of the Oracle Central Inventory for a host is recorded in the inventory pointer file, `oraInst.loc`.

By default, Oracle Universal Installer places and looks for the `oraInst.loc` file in the `/etc` directory for installations on Linux platforms and in the `/var/opt/oracle` directory for installations on Solaris platforms. You can use the `runInstaller` script with the `-invPtrLoc` option to specify another inventory pointer file.

There are two requirements regarding the location of the Oracle Central Inventory within the Oracle Application Server Disaster Recovery environment:

1. The location must be defined and accessible.
2. The location must be on shared storage that is synchronized along with the software installed on the host.

If you move the Oracle home for an Oracle Application Server instance from a production site host's local storage to the shared storage set up for the site, then you must reassociate that Oracle home with the Oracle Central Inventory on the site's shared storage. To reassociate the Oracle home with the Oracle Central Inventory, use the `runInstaller` script (located under the `$ORACLE_HOME/oui/bin` directory) with the `-attachHome` option, as shown in Example 2–2:

***Example 2–2 Attaching an Oracle Home to the Oracle Central Inventory***

```
./runInstaller -attachHome ORACLE_HOME="/u01/app/oracle/oid_oh"
ORACLE_HOME_NAME="OIDHome"
```

See *Oracle Universal Installer and OPatch User's Guide* for more information about Oracle Universal Installer, OPatch, and the Oracle Central Inventory.

### 2.1.3 Designing a New Production Site and Standby Site

This section presents the logic to implementing a new production site for an Oracle Application Server Disaster Recovery topology. It describes the planning and setup of the production site by pre-planning host names, configuring the hosts to resolve the Application Server host names and network host names, and ensuring that disk replication is set up to copy the configuration based on these names to the standby site. When you design the production site, you should also plan the standby site, which must be symmetric with the production site.

When you are designing a new production site (not using a pre-existing production site), you will use Oracle Universal Installer to install software on the production site, and parameters such as Application Server host names and software paths must be carefully designed to ensure that they are the same for both sites.

The flexibility you have when you create a new Oracle Application Server Disaster Recovery production site and standby site includes:

1. You can design your Oracle Application Server Disaster Recovery solution so that each host at the production site and at the standby site has a different Application Server host name and network host name.
2. When you design and create your production site from scratch, you can choose the Oracle home name and Oracle home directory for each Application Server installation, following the instructions recommendations in this chapter. You can also choose the location of the Oracle Central Inventory for each host, and the location of the static HTML pages directories for hosts.

   Designing and creating your site from scratch is easier than trying to modify an existing site to meet the design requirements described in this chapter.
3. You can assign ports for the Application Server installations for the production site hosts that will not conflict with the ports that will be used at the standby site hosts.

   This is easier than having to check for and resolve port conflicts between an existing production site and standby site.

### 2.1.4 Planning Host Names for the Production Site and Standby Site

This section describes how to plan Application Server host names and network host names for the middle tier hosts that use the Oracle Application Server instances at the production site and standby site.

Table 2–2 shows the IP addresses, Application Server host names, and network host names that will be used for the EDG deployment production site hosts that use the Oracle Application Server instances in the production site. Figure 2–1 shows the configuration for the EDG deployment at the production site.

***Table 2–2 IP Addresses, Application Server Host Names, and Network Host Names for Production Site Hosts and Shared Storage***

| IP Address | Application Server Host Name | Network Host Name[1] |
|---|---|---|
| 123.1.2.111 | OIDHOST1 | PRODOID1 |
| 123.1.2.112 | OIDHOST2 | PRODOID1 |
| 123.1.2.113 | WEBHOST1 | PRODWEB1 |
| 123.1.2.114 | WEBHOST2 | PRODWEB2 |
| 123.1.2.115 | APPHOST1 | PRODAPP1 |
| 123.1.2.116 | APPHOST2 | PRODAPP2 |
| 123.1.2.117 | WEBHOST3 | PRODWEB3 |
| 123.1.2.118 | WEBHOST4 | PRODWEB4 |
| 123.1.2.119 | IDMHOST1 | PRODIDM1 |
| 123.1.2.120 | IDMHOST2 | PRODIDM2 |
| 123.1.2.121 | N/A | PRODSTOR |

[1] See Section 2.1.7, "Resolving Host Names Using DNS Host Name Resolution" for information on defining network host names.

Table 2–3 shows the IP addresses, Application Server host names, and network host names that will be used for the EDG deployment standby site hosts that use the Oracle Application Server instances in the standby site.

***Table 2–3 IP Addresses, Application Server Host Names, and Network Host Names for Standby Site Hosts and Shared Storage***

| IP Address | Application Server Host Name | Network Host Name[1] |
|---|---|---|
| 123.2.2.111 | OIDHOST1 | STBYOID1 |
| 123.2.2.112 | OIDHOST2 | STBYOID2 |
| 123.2.2.113 | WEBHOST1 | STBYWEB1 |
| 123.2.2.114 | WEBHOST2 | STBYWEB2 |
| 123.2.2.115 | APPHOST1 | STBYAPP1 |
| 123.2.2.116 | APPHOST2 | STBYAPP2 |
| 123.2.2.117 | WEBHOST3 | STBYWEB3 |
| 123.2.2.118 | WEBHOST4 | STBYWEB4 |
| 123.2.2.119 | IDMHOST1 | STBYIDM1 |
| 123.2.2.120 | IDMHOST2 | STBYIDM2 |
| 123.2.2.121 | N/A | STBYSTOR |

[1] See Section 2.1.7, "Resolving Host Names Using DNS Host Name Resolution" for information on defining network host names.

The Application Server host names are resolved locally at the production site or standby site to the correct IP address. In an Oracle Application Server Disaster Recovery topology, there are two ways to configure host name resolution for the Application Server host names planned for the production site and standby site in Table 2–2 and Table 2–3. The two ways to configure Application Server host name resolution are described in Section 2.1.5, "Host Name Resolution."

## 2.1.5 Host Name Resolution

Host name resolution is the process of resolving a host name to the proper IP address for communication. Oracle Application Server Disaster Recovery supports two methods of configuring host name resolution.

The two ways to configure the Application Server host name resolution are:

- Local host name resolution

  Local host name resolution uses the Application Server host name to IP address mapping that is specified by the `/etc/hosts` file on each host.

  See Section 2.1.6, "Making /etc/hosts File Entries for Local Host Name File Resolution" for more information about using the `/etc/hosts` file to implement local host name file resolution.

- DNS server resolution

  DNS server resolution is a method of centralizing the Application Server host names to IP addresses in a database and configuring each host to use a DNS server to resolve host names.

  See Section 2.1.7, "Resolving Host Names Using DNS Host Name Resolution" for more information about implementing DNS server host name resolution.

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

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

The Application Server host names are not actually assigned to the production site hosts (included in the Oracle Application Server configuration for the host) until the Oracle Application Server middle tier components are installed at the production site host.

#### Host Name Resolution Precedence

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

As shown in Example 2–3, if local host name file resolution is used for the host, the `files` entry will be the first entry for the `hosts` parameter. When `files` is the first entry for the `hosts` parameter, entries in the host's `/etc/hosts` file will be used first to resolve host names:

***Example 2–3 Specifying the Use of Local Host Name Resolution***

```
hosts:   files   dns   nis
```

As shown in Example 2–4, if DNS server host name resolution is used for the host, the `dns` entry will be the first entry for the `hosts` parameter. When `dns` is the first entry for the `hosts` parameter, DNS server entries will be used first to resolve host names:

***Example 2–4 Specifying the Use of DNS Host Name Resolution***

```
hosts:   dns   files   nis
```

All the hosts in an Oracle Application Server Disaster Recovery production site and standby site must have the same values for the hosts parameter in the host's `/etc/nsswitch.conf` file. Also, the nis entry must be the last entry for the host's parameter for all production site and standby site hosts. Oracle Application Server Disaster Recovery does not support nis host name resolution.

### 2.1.6 Making /etc/hosts File Entries for Local Host Name File Resolution

If you decide to use local host name file resolution to resolve Application Server host names, then the `/etc/hosts` file for each host at the production site should include entries for the Application Server host names of other hosts at the production site. Make sure that the first entry in the `/etc/hosts` file for each production site host is the entry for the Application Server host name for that host.

Example 2–5 shows the entries in the `/etc/hosts` file for OIDHOST1 at the standby site. Each entry provides the IP address of a production site host and specifies the Application Server host name that will be used at the production site for that host:

***Example 2–5 Making /etc/hosts File Entries for a Production Site Host***

```
127.0.0.1       localhost.localdomain  localhost
123.1.2.111     OIDHOST1.ORACLE.COM    OIDHOST1
123.1.2.112     OIDHOST2.ORACLE.COM    OIDHOST2
123.1.2.113     WEBHOST1.ORACLE.COM    WEBHOST1
123.1.2.114     WEBHOST2.ORACLE.COM    WEBHOST2
```

```
123.1.2.115    APPHOST1.ORACLE.COM    APPHOST1
123.1.2.116    APPHOST2.ORACLE.COM    APPHOST2
123.1.2.117    WEBHOST3.ORACLE.COM    WEBHOST3
123.1.2.118    WEBHOST4.ORACLE.COM    WEBHOST4
123.1.2.119    IDMHOST1.ORACLE.COM    IDMHOST1
123.1.2.120    IDMHOST2.ORACLE.COM    IDMHOST2
```

Technically, each host in a site only needs entries for the other hosts it directly communicates with at the site. However, for consistency and ease of maintenance, it is recommended that the `/etc/hosts` file for each host at a site include entries for all the other hosts at the site.

You can copy the `/etc/hosts` file for one host to each the other hosts at the site. Then, for each host, edit the `/etc/hosts` file so that the entry for that host is the second entry (just after the 127.0.0.1 entry) in the file.

Similarly, if you use local host name file resolution to resolve Application Server host names, then the entries for each host at the standby site should include entries for the Application Server host names of other hosts at the standby site. Make sure that the second entry in the `/etc/hosts` file for each standby site host is the entry for the Application Server host name of that host.

Example 2–6 shows the entries in the `/etc/hosts` file for OIDHOST1 at the standby site. Each entry provides the IP address of a standby site host and specifies the Application Server host name that will be used at the standby site for that host:

***Example 2–6 Making /etc/hosts File Entries for a Standby Site Host***

```
127.0.0.1      localhost.localdomain  localhost
123.2.2.111    OIDHOST1.ORACLE.COM    OIDHOST1
123.2.2.112    OIDHOST2.ORACLE.COM    OIDHOST2
123.2.2.113    WEBHOST1.ORACLE.COM    WEBHOST1
123.2.2.114    WEBHOST2.ORACLE.COM    WEBHOST2
123.2.2.115    APPHOST1.ORACLE.COM    APPHOST1
123.2.2.116    APPHOST2.ORACLE.COM    APPHOST2
123.2.2.117    WEBHOST3.ORACLE.COM    WEBHOST3
123.2.2.118    WEBHOST4.ORACLE.COM    WEBHOST4
123.2.2.119    IDMHOST1.ORACLE.COM    IDMHOST1
123.2.2.120    IDMHOST2.ORACLE.COM    IDMHOST2
```

In this example, the entries in the `/etc/hosts` file for OIDHOST1 at the standby site are very similar to the entries in the `/etc/hosts` file for OIDHOST1 at the production site. The difference is that all the IP addresses for the production site hosts begin with 123.1.n.n and all the IP addresses for the standby site hosts begin with 123.2.n.n.

The `/etc/hosts` files for the other hosts at the standby site would include these same entries, but the second entry (just after the 127.0.0.1 entry) in each `/etc/hosts` file at the standby site would be the entry for the Application Server host name for that standby site host.

When you are using local host name file resolution for Application Server host name resolution, the network host names for the production site shown in Table 2–2 and the network host names for the standby site shown in Table 2–3 are not defined in the `/etc/hosts` file. Instead, the network host names for the production and standby site hosts must be defined in the corporate Domain Name System (DNS) that includes the production and standby site.

When you are using local host name file resolution, make sure that `files` is specified as the first method of host name resolution for the `hosts` parameter in the `/etc/nsswitch.conf` file for each middle tier host, as shown in Example 2–3.

See Section 2.1.7, "Resolving Host Names Using DNS Host Name Resolution" for information about defining network host names.

### Basic Rules for /etc/hosts File Entries

Follow these basic rules for creating host name entries in `/etc/hosts` files:

1. Each host name entry must include an IP address for a host, the fully qualified name (host name and domain) of the host, and the short name of the host.
2. Add the entry that specifies the Application Server host name for the local host immediately after the `127.0.0.1 localhost.localdomain` entry. During installations, the Oracle Universal Installer looks for the Application Server host name in the entry after the `127.0.0.1 localhost.localdomain` entry.
3. If you want to create multiple Application Server host names for a given IP address, enter the primary Application Server host name that you want returned for that address before the additional entries for that IP address.
4. If you would like to include multiple Application Server host names in a single entry, add them at the end of the entry.

Example 2–7 shows `/etc/hosts` file entries that could be made for production site host WEBHOST1. These entries demonstrate the basic rules described earlier:

***Example 2–7 Making Valid /etc/hosts File Entries***

```
127.0.0.1       localhost.localdomain  localhost
123.1.2.113     WEBHOST1.ORACLE.COM    WEBHOST1   MYHOST
123.1.2.114     WEBHOST2.ORACLE.COM    WEBHOST2
123.1.2.114     WEBHOSTPEER.ORACLE.COM WEBHOSTPEER
```

The rules that are demonstrated in Example 2–7 are:

- Rule 1 is followed for all the entries in WEBHOST1's `/etc/hosts` file.
- Rule 2 is followed by the second entry. This is the `/etc/hosts` file for WEBHOST1 and the entry for WEBHOST1 is just after the `127.0.0.1 localhost.localdomain` entry.
- Rule 3 is illustrated by the third and fourth entries. The third entry (for WEBHOST2) appears before the fourth entry (for WEBHOSTPEER.ORACLE.COM) because WEBHOST2.ORACLE.COM is the primary Application Server host name for IP address 123.1.2.114. Although WEBHOSTPEER will be recognized as an Application Server host name for 123.1.2.114, the first host name returned for IP address 123.1.2.114 on most operating systems will be WEBHOST2.ORACLE.COM.
- Rule 4 is illustrated by the second entry for WEBHOST1. In this entry, MYHOST has been specified at the end of the entry, so that MYHOST can be used as an additional Application Server host name in addition to the primary Application Server host name, WEBHOST1.US.ORACLE.COM.

After you set up host name resolution for your production site and standby site hosts 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 for WEBHOST1 in Example 2–7, a `ping webhost1` command would return the correct IP address (123.1.2.113) and also indicate that the host name is fully qualified. Similarly, a `ping webhostpeer` command will return the correct IP address (123.1.2.114) and it will also show that the name WEBHOST2 is also associated with that IP address.

### 2.1.7 Resolving Host Names Using DNS Host Name Resolution

The corporate DNS server includes entries that map host names to IP addresses. The corporate DNS server entries map network host names to the IP addresses of the hosts that those network host names are assigned to. The corporate DNS must include an entry for each production site host and standby site host, which will maps the network host name for the host to the IP address of the host. These network host names are visible in the network that includes the Oracle Application Server Disaster Recovery production site and standby site.

Regardless of whether you are using local host name file resolution (`/etc/hosts` file entries) or you are using DNS host name resolution exclusively (no `/etc/hosts` file entries) for your Oracle Application Server Disaster Recovery topology, you must make sure that the corporate DNS includes an entry for each production site host and standby site host that maps the network host name for the host to the IP address of the host.

When you are using DNS host name resolution exclusively for all host name resolution (in other words, if you are not using `/etc/hosts` files for local host name resolution), then in addition to making network host names entries for the production site hosts and the standby site hosts in the corporate DNS, you must also have a DNS server for the production site and a DNS server for the standby site. You will make Application Server host name entries for the production site hosts in the production site DNS and make Application Server host name entries for the standby site hosts in the standby site DNS. Logically, the entries in the production site DNS and standby site DNS serve the same purpose as the `/etc/hosts` file entries when local host name file resolution is used; the difference is that all the Application Server host names for a site are consolidated into one location (one DNS server).

---

**Note:** For the DNS configuration described in this section to work properly, these requirements must be met:

- The production site DNS server and standby site DNS server must not be aware of each other. They should make non-authoritative lookup requests to the corporate DNS servers only if they fail to resolve a host name within their specific site.
- The production site DNS server and standby site DNS server must contain entries for only Application Server host names used within their own site.
- The overall corporate DNS servers contain network host names and any other aliases or virtual names that needed for the site currently serving in the production role.
- If you are using DNS host name resolution exclusively (not using `/etc/hosts` file entries to implement local host name file resolution), then there should be no entries in the `/etc/hosts` file for any host at the production site or standby site to any other host in either site.

---

Section 2.1.7.1, "Making Host Name Entries in the Corporate DNS" describes how to make network host name entries for the production site hosts and standby site hosts in the corporate DNS server.Section 2.1.7.3, "Making Host Name Entries in the Standby Site DNS" describes how to make host name entries for the standby site.

Section 2.1.7.2, "Making Host Name Entries in the Production Site DNS" describes how to make Application Server host name entries for production site hosts in the production site DNS.

Section 2.1.7.3, "Making Host Name Entries in the Standby Site DNS" describes how to make Application Server host name entries for standby site hosts in the standby site DNS.

### 2.1.7.1 Making Host Name Entries in the Corporate DNS

You must make sure that the corporate DNS includes an entry for each production site host and standby site host that maps the network host name for the host to the IP address of the host. Example 2–8 shows the entries to make in the corporate DNS server for the Oracle Application Server Disaster Recovery topology that uses the EDG deployment in Figure 2–1 for the production site and standby site:

***Example 2–8 Network Host Name Entries for Production Site Hosts and Standby Site Hosts in the Corporate DNS***

```
PRODOID1.ORACLE.COM    IN    A    123.1.2.111
PRODOID2.ORACLE.COM    IN    A    123.1.2.112
PRODWEB1.ORACLE.COM    IN    A    123.1.2.113
PRODWEB2.ORACLE.COM    IN    A    123.1.2.114
PRODAPP1.ORACLE.COM    IN    A    123.1.2.115
PRODAPP2.ORACLE.COM    IN    A    123.1.2.116
PRODWEB3.ORACLE.COM    IN    A    123.1.2.117
PRODWEB4.ORACLE.COM    IN    A    123.1.2.118
PRODIDM1.ORACLE.COM    IN    A    123.1.2.119
PRODIDM2.ORACLE.COM    IN    A    123.1.2.120
STBYOID1.ORACLE.COM    IN    A    123.2.2.111
STBYOID2.ORACLE.COM    IN    A    123.2.2.112
STBYWEB1.ORACLE.COM    IN    A    123.2.2.113
STBYWEB2.ORACLE.COM    IN    A    123.2.2.114
STBYAPP1.ORACLE.COM    IN    A    123.2.2.115
STBYAPP2.ORACLE.COM    IN    A    123.2.2.116
STBYWEB3.ORACLE.COM    IN    A    123.2.2.117
STBYWEB4.ORACLE.COM    IN    A    123.2.2.118
STBYIDM1.ORACLE.COM    IN    A    123.2.2.119
STBYIDM2.ORACLE.COM    IN    A    123.2.2.120
```

### 2.1.7.2 Making Host Name Entries in the Production Site DNS

If your Oracle Application Server Disaster Recovery topology uses *only* DNS resolution for Application Server host names and network host names, configure the production site DNS server to use the corporate DNS server as the forwarding server for unresolved requests.

Make entries for the production site Application Server host names in the production site DNS to configure host name resolution for these host names. Example 2–9 shows the Application Server host name entries to make in the production site DNS server:

***Example 2–9 Application Server Host Name Entries for Production Site Hosts in the Production Site DNS***

```
OIDHOST1.ORACLE.COM    IN    A    123.1.2.111
OIDHOST2.ORACLE.COM    IN    A    123.1.2.112
WEBHOST1.ORACLE.COM    IN    A    123.1.2.113
WEBHOST2.ORACLE.COM    IN    A    123.1.2.114
APPHOST1.ORACLE.COM    IN    A    123.1.2.115
```

```
APPHOST2.ORACLE.COM    IN    A    123.1.2.116
WEBHOST3.ORACLE.COM    IN    A    123.1.2.117
WEBHOST4.ORACLE.COM    IN    A    123.1.2.118
IDMHOST1.ORACLE.COM    IN    A    123.1.2.119
IDMHOST2.ORACLE.COM    IN    A    123.1.2.120
```

#### 2.1.7.3 Making Host Name Entries in the Standby Site DNS

If your Oracle Application Server Disaster Recovery topology uses *only* DNS resolution for Application Server host names and network host names, configure the standby site DNS server to use the corporate DNS server as the forwarding server for unresolved requests.

Make entries for the standby site Application Server host names in the standby site DNS to configure host name resolution for these host names. Example 2–10 shows the Application Server host name entries to make in the standby site DNS server:

***Example 2–10 Application Server Host Name Entries for Standby Site Hosts in the Standby Site DNS***

```
OIDHOST1.ORACLE.COM    IN    A    123.2.2.111
OIDHOST2.ORACLE.COM    IN    A    123.2.2.112
WEBHOST1.ORACLE.COM    IN    A    123.2.2.113
WEBHOST2.ORACLE.COM    IN    A    123.2.2.114
APPHOST1.ORACLE.COM    IN    A    123.2.2.115
APPHOST2.ORACLE.COM    IN    A    123.2.2.116
WEBHOST3.ORACLE.COM    IN    A    123.2.2.117
WEBHOST4.ORACLE.COM    IN    A    123.2.2.118
IDMHOST1.ORACLE.COM    IN    A    123.2.2.119
IDMHOST2.ORACLE.COM    IN    A    123.2.2.120
```

## 2.2 Environment Preparation

Follow these steps to prepare for setting up the Oracle Application Server Disaster Recovery topology:

1. On the production site, install the Oracle databases that will be used in the Oracle Application Server Disaster Recovery topology that you are setting up. Then create standby databases on the standby site and configure Oracle Data Guard for the production and standby databases. For more information, see Section 2.2.1, "Database Considerations."
2. The Oracle home directories for each of the Application Server instances used at the production site (the Application Server instances shown in Figure 2–1) must be on volumes in the shared storage for the production site. The Oracle homes for these Application Server instances *cannot* be in local storage for the middle tier hosts that use the instances. Therefore, begin preparing the environment by planning where (on which volumes in the production site shared storage) you will create the Oracle home directories for the production site Application Server instances. See Section 2.2.2, "Creating Volumes, Mount Points, and Symbolic Links" for more information about planning and creating volumes for the Oracle Application Server instances in the production site shared storage.
3. Then plan the mount points and symbolic links to create to the Oracle home directories that will be used for the Oracle Application Server instances on the production site shared storage. No Oracle software is installed in local storage for the middle tier hosts that use these instances at either the production site or

standby site. For more on planning and creating mount points and symbolic links for the Oracle Application Server homes on the shared storage, see Section 2.2.2, "Creating Volumes, Mount Points, and Symbolic Links."

4. Validate the environment by connecting to each host at the production site and using the ping command to ensure that the host can locate the other hosts at the production site. Then connect to each host at the standby site and use the ping command to ensure that the host can locate the other hosts at the standby site. See Section 2.2.3, "Testing the Host Name Resolution" for more information.

The rest of this section describes how to perform these steps for the Oracle Application Server Disaster Recovery topology for the deployment shown in Figure 2–1.

### 2.2.1 Database Considerations

This section describes how to set up the Oracle databases that will be used in the Oracle Application Server Disaster Recovery topology that you are setting up. These are the tasks to perform for databases:

1. Install the Oracle databases that will be used in the Oracle Application Server Disaster Recovery topology that you are setting up.

   For more information about this step, refer to Section A.1, "Installing the Oracle Databases in the Production Site."

2. Create a standby database for each database included in your Oracle Application Server Disaster Recovery production site.

   For more information about this step, refer to Section A.2, "Creating Standby Databases at the Standby Site."

3. Set up the Oracle Data Guard configuration for databases at the production site and standby site in your Oracle Application Server Disaster Recovery topology.

   For more information about this step, refer to Section A.3, "Setting Up the Oracle Data Guard Configuration."

4. Set up TNSNAMES.ORA entries to enable the production site databases and standby site databases to reference each other.

   For more information about this step, refer to Section A.3, "Setting Up the Oracle Data Guard Configuration."

5. For Oracle Application Server components that store middle tier configuration data in Oracle database repositories, use Oracle Data Guard to manually force a database synchronization whenever a middle tier synchronization is performed.

   For more information about this step, refer to Section A.5, "Manually Forcing Database Synchronization with Oracle Data Guard."

6. Optionally, set up database host name aliases for the databases at your production site and standby site. The alias must be defined in DNS or in the `/etc/hosts` file on each node running a database instance.

   For more information about this step, refer to Section A.6, "Setting Up Database Host Name Aliases."

### 2.2.2 Creating Volumes, Mount Points, and Symbolic Links

This section includes the following topics:

- Creating Volumes for the Application Server Host Clusters

This section describes how to create a volume on the shared storage for each of the Application Server host clusters in the EDG deployment in Figure 2–1.

- Creating Mount Points, Symbolic Links, and Oracle Home Directories

  This section describes how to create Oracle home directories for the Application Server instances for a host cluster on the cluster's volume, and how to create mount points and symbolic links on the hosts to the Oracle home directories.

- Creating Mount Points, Symbolic Links, and Oracle Central Inventory Directories

  This section describes how to create Oracle Central Inventory directories for the Application Server instances for a host cluster on the cluster's volume, and how to create mount points and symbolic links on the hosts to the Oracle Central Inventory directories.

- Creating Mount Points, Symbolic Links, and Static HTML Pages Directories

  This section describes how to create static HTML pages directories for Oracle HTTP Server instances for a host cluster on the cluster's volume, and how to create mount points and symbolic links on the hosts to the static HTML pages directories.

#### 2.2.2.1 Creating Volumes for the Application Server Host Clusters

There are five Oracle Application Server host clusters in the EDG deployment shown in Figure 2–1. The Oracle Application Server instances for each of these host clusters must be installed on the volume created for that host cluster. The reason for using a separate volume for each host cluster is to ensure data consistency for the cluster across the production site shared storage and standby site shared storage.

---

**Note:** This section describes how to create volumes, mount points, and symbolic links for the EDG deployment production site shown in Figure 2–1. These are preliminary steps that you perform before installing the production site Oracle Application Server instances.

Do not install any production site Oracle Application Server instances until you are directed to do so later in this manual.

---

Table 2–4 shows the volume that will be created for each host cluster. In Table 2–4, the host clusters are shown in the order in which the Oracle Application Server instances for the host cluster are installed and configured in the *Oracle Application Server Enterprise Deployment Guide*.

***Table 2–4 Volumes To Be Created for Each Oracle Application Server Host Cluster in the EDG Deployment***

| Host Cluster | Volume for Host Cluster[1] |
|---|---|
| OIDHOST1 and OIDHOST2 | /vol/voloid |
| WEBHOST1 and WEBHOST2 | /vol/volweb1 |
| APPHOST1 and APPHOST2 | /vol/volapp |
| WEBHOST3 and WEBHOST4 | /vol/volweb2 |
| IDMHOST1 and IDMHOST2 | /vol/volidm |

[1] These examples use /vol in the volume name. Refer to vendor-specific storage documentation for approved naming of volumes.

Create a volume on the shared storage for each of the Application Server host clusters. The examples later in this section assume that the volumes for the host clusters in Table 2–4 have been created.

> **Note:** Some shared storage (for example, NAS, NFS, or SAN storage) may not provide the reliable file locking that Oracle HTTP Server requires. When a 10.1.2.x or 10.1.3.x Oracle HTTP Server instance is installed on shared storage that does not provide reliable file locking, the Oracle HTTP Server instance may experience performance problems.
>
> In this situation, perform the steps in Section G.1.3, "Changing the LockFile Directive for 10.1.2.x and 10.1.3.x Oracle HTTP Server Instances" and Section G.1.4, "Use the DEFAULT_DMS_DIR Environment Variable for Oracle HTTP Server 10.1.3.x Instances" to resolve the Oracle HTTP Server performance issues.

#### 2.2.2.2 Creating Mount Points, Symbolic Links, and Oracle Home Directories

In this section, mount points and symbolic links are set up on the Oracle Application Server host clusters to the Oracle home directories on the storage volume where the Application Server instances will be installed. These mount points and symbolic links are set up so that the same directory structure can be used on each Oracle Application Server host in the cluster. Using the mount points and symbolic links on the host simplifies the installation, management and maintenance of the Oracle Application Server instances for the hosts in the clusters.

Table 2–5 shows each Oracle Application Server host in the EDG deployment in Figure 2–1 and the type of Oracle Application Server instance or instances that need to be installed for each host. For each Oracle Application Server instance, the table shows the mount point directory to create on the host to the Oracle home directory on the storage volume, the symbolic link to create on the host to the Oracle home directory on the volume, and the actual Oracle home directory to create on the volume for the instance. The reason for creating the symbolic links on each host is to ensure that the same Oracle home directory path is used for the instances in each host cluster.

***Table 2–5 Mount Points, Links, and Oracle Homes for OracleAS Instances***

| Host | Application Server Instance Type | Host Mount Point Directory to Oracle Home Directory on Volume | Link on Host to Oracle Home Directory on Volume | Oracle Home Directory on Volume |
|---|---|---|---|---|
| OIDHOST1 | Oracle Internet Directory | /u02/voloidmount/oid1_oh | /u01/app/oracle/oid_oh | /vol/voloid/oid1_oh |
| OIDHOST2 | Oracle Internet Directory | /u02/voloidmount/oid2_oh | /u01/app/oracle/oid_oh | /vol/voloid/oid2_oh |
| WEBHOST1 | Oracle HTTP Server | /u02/volweb1mount/web1_oh | /u01/app/oracle/web_oh | /vol/volweb1/web1_oh |
| WEBHOST2 | Oracle HTTP Server | /u02/volweb1mount/web2_oh | /u01/app/oracle/web_oh | /vol/volweb1/web2_oh |
| APPHOST1 | First OC4J Instance | /u02/volappmount/app1_oh1 | /u01/app/oracle/app_oh | /vol/volapp/app1_oh1 |
| APPHOST1 | Second OC4J Instance | /u02/volappmount/app1_oh2 | /u01/app/oracle/app_oh2 | /vol/volapp/app1_oh2 |

*Table 2–5 (Cont.) Mount Points, Links, and Oracle Homes for OracleAS Instances*

| Host | Application Server Instance Type | Host Mount Point Directory to Oracle Home Directory on Volume | Link on Host to Oracle Home Directory on Volume | Oracle Home Directory on Volume |
|---|---|---|---|---|
| APPHOST2 | First OC4J Instance | /u02/volappmount/app2_oh1 | /u01/app/oracle/app_oh | /vol/volapp/app2_oh1 |
| APPHOST2 | Second OC4J Instance | /u02/volappmount/app2_oh2 | /u01/app/oracle/app_oh2 | /vol/volapp/app2_oh2 |
| WEBHOST3 | Oracle HTTP Server | /u02/volweb2mount/web3_oh | /u01/app/oracle/web_oh | /vol/volweb2/web3_oh |
| WEBHOST4 | Oracle HTTP Server | /u02/volweb2mount/web4_oh | /u01/app/oracle/web_oh | /vol/volweb2/web4_oh |
| IDMHOST1 | Oracle Identity Management | /u02/volidmmount/idm1_oh | /u01/app/oracle/idm_oh | /vol/volidm/idm1_oh |
| IDMHOST2 | Oracle Identity Management | /u02/volidmmount/idm2_oh | /u01/app/oracle/idm_oh | /vol/volidm/idm2_oh |

Example 2–11 uses hosts OIDHOST1 and OIDHOST2 to show the basic steps for creating the Oracle home directories for a host cluster's Application Server instances on its storage volume, and for setting up mount points and symbolic links on the host cluster to the Oracle home directories on the volume.

Figure 2–2 shows the storage after the Oracle home directories for the Oracle Application Server instances for OIDHOST1 and OIDHOST2 have been created on the `/vol/voloid` volume , and the mount points and symbolic links to the Oracle home directories have been set up on the hosts, as described in the following steps.

***Figure 2–2 Links on OIDHOST1 and OIDHOST2 to Oracle Home Directories for Application Server Instances on Storage***

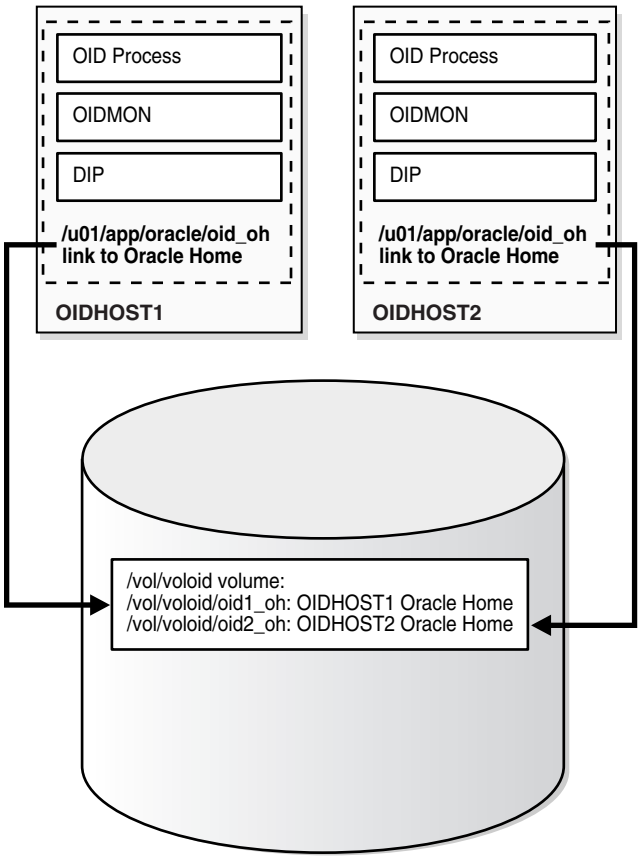


***Example 2–11 Setting Up Mount Points and Symbolic Links to Oracle Home Directories***

1. Log in as root on OIDHOST1 and create the following directories:

```
prompt> mkdir /u01/app/oracle
prompt> mkdir /u02/voloidmount
```

2. On OIDHOST1, mount the `/u02/voloidmount` directory to the `/vol/voloid` volume on the storage, and set up the mount point permissions and ownership as necessary. Refer to vendor-specific information for the shared storage to perform this step.

3. While logged in as root on OIDHOST1, mount the storage volume:

```
prompt> mount /u02/voloidmount
```

4. On OIDHOST1, create the Oracle home directory for the Oracle Application Server instance in the storage:

```
prompt> cd /u02/voloidmount
prompt> mkdir oid1_oh
```

5. On OIDHOST1, create a symbolic link named `oid_oh` in the `/u01/app/oracle` directory to the `/u02/voloidmount/oid1_oh` directory on the storage:

```
prompt> cd /u01/app/oracle
prompt> ln -s /u02/voloidmount/oid1_oh oid_oh
```

6. On OIDHOST1, the following command changes the working directory to the `/vol/voloid/oid1_oh` directory on the storage.

```
prompt> cd /u01/app/oracle/oid1_oh
```

7. You do not install the Oracle Application Server instance for OIDHOST1 now. The Application Server installations are performed later. For more information about installing Oracle Application Server instances on the storage volumes, see Section 2.3, "Installing the Oracle Application Server Instances for the Production Site."

8. Log in as root on OIDHOST2 and create the following directories:

```
prompt> mkdir /u01/app/oracle
prompt> mkdir /u02/voloidmount
```

9. On OIDHOST2, mount the `/u02/voloidmount` directory to the /vol/voloid volume on the storage, and set up the mount point permissions and ownership as necessary. Refer to vendor-specific information for the shared storage to perform this step.

10. While logged in as root on OIDHOST2, mount the storage volume:

```
prompt> mount /u02/voloidmount
```

11. On OIDHOST2, create the Oracle home directory for the Oracle Application Server instance in the storage:

```
prompt> cd /u02/voloidmount
prompt> mkdir oid2_oh
```

12. On OIDHOST2, create a symbolic link named oid_oh in the `/u01/app/oracle` directory to the `/u02/voloidmount/oid2_oh` directory on the storage:

```
prompt> cd /u01/app/oracle
prompt> ln -s /u02/voloidmount/oid2_oh oid_oh
```

13. On OIDHOST2, the following command changes the working directory to the `/vol/voloid/oid2_oh` directory on the storage.

```
prompt> cd /u01/app/oracle/oid_oh
```

14. You do not install the Oracle Application Server instance for OIDHOST2 now. The Application Server installations are performed later. For more information about installing Oracle Application Server instances on the storage volumes, see Section 2.3, "Installing the Oracle Application Server Instances for the Production Site."

Use similar steps for the other host clusters in Table 2–4 to create the storage volumes and Oracle home directories on the host cluster's volume, and then to create the mount points and symbolic links on the hosts to the Oracle Application Server homes on the volume. Use the information in Table 2–5 to perform this step.

Figure 2–3 shows the shared storage after the Oracle home directories for the Oracle Application Server instances for WEBHOST1 and WEBHOST2 have been created on the `/vol/volweb1` volume, and the mount points and symbolic links to the Oracle home directories have been set up on the hosts.

***Figure 2–3 Links on WEBHOST1 and WEBHOST2 to Oracle Home Directories for Application Server Instances on Storage***

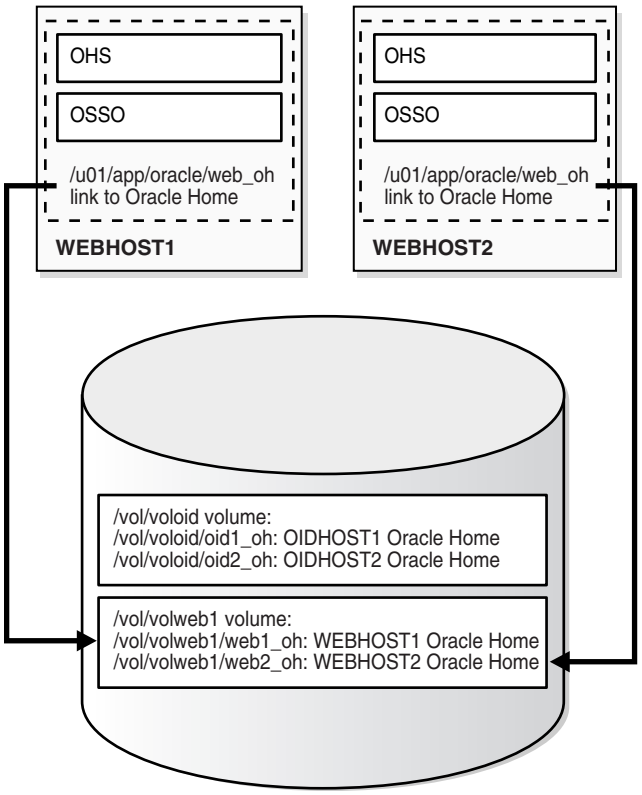


Figure 2–4 shows the shared storage after the Oracle homes for the Oracle Application Server instances for APPHOST1 and APPHOST2 have been created on the `/vol/volapp` volume, and the mount points and symbolic links to the Oracle home directories have been set up on the hosts.

***Figure 2–4 Links on APPHOST1 and APPHOST2 to Oracle Home Directories for Application Server Instances on Storage***

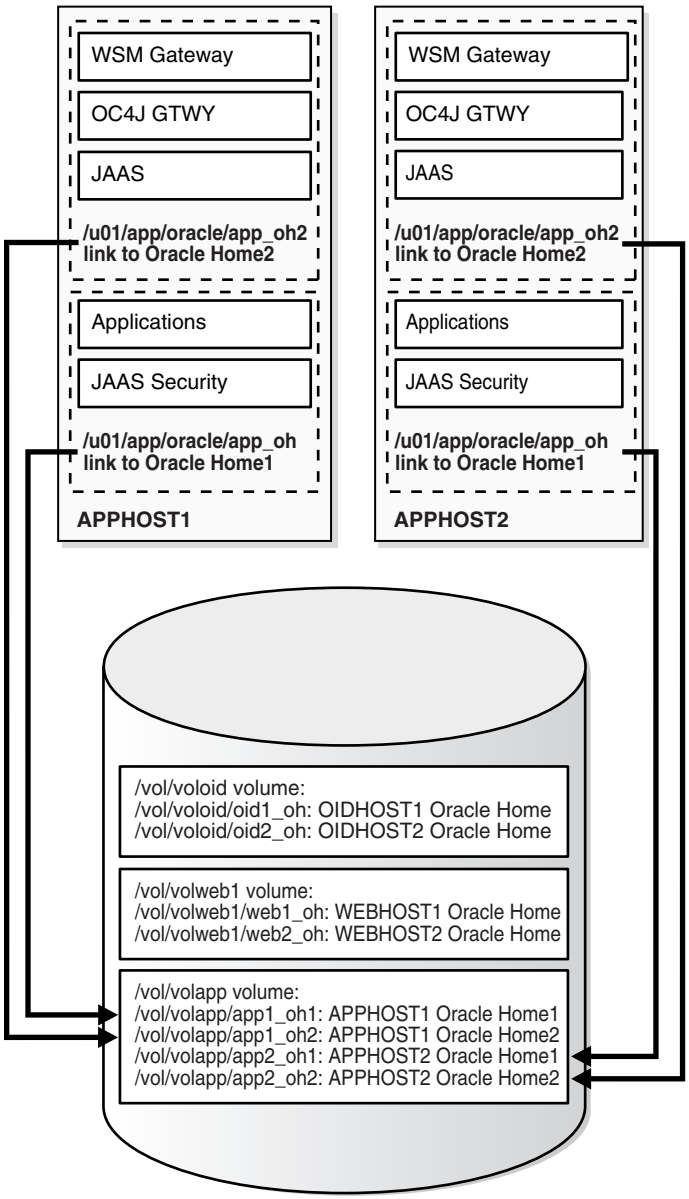


Figure 2–5 shows the shared storage after the Oracle home directories for the Oracle Application Server instances for WEBHOST3 and WEBHOST4 have been created on the `/vol/volweb2` volume, and the mount points and symbolic links to the Oracle home directories have been set up on the hosts.

***Figure 2–5 Links on WEBHOST3 and WEBHOST4 to Oracle Home Directories for Application Server Instances on Storage***

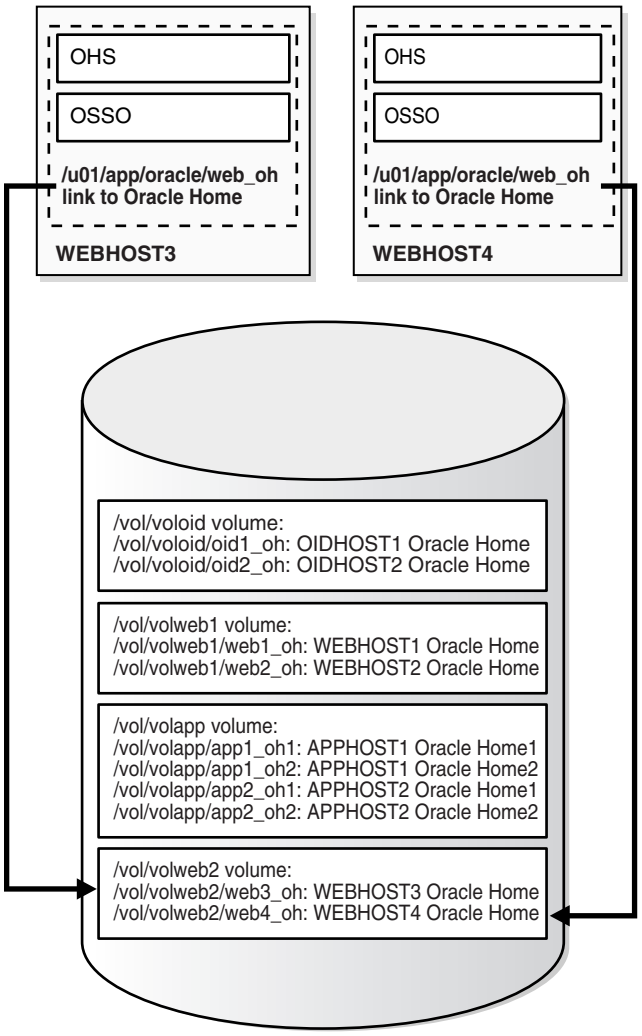


Figure 2–6 shows the shared storage after the Oracle home directories for the Oracle Application Server instances for IDMHOST1 and IDMHOST2 have been created on the `/vol/volidm` volume, and the mount points and symbolic links to the Oracle home directories have been set up on the hosts.

***Figure 2–6 Links on IDMHOST1 and IDMHOST2 to Oracle Home Directories for Application Server Instances on Storage***

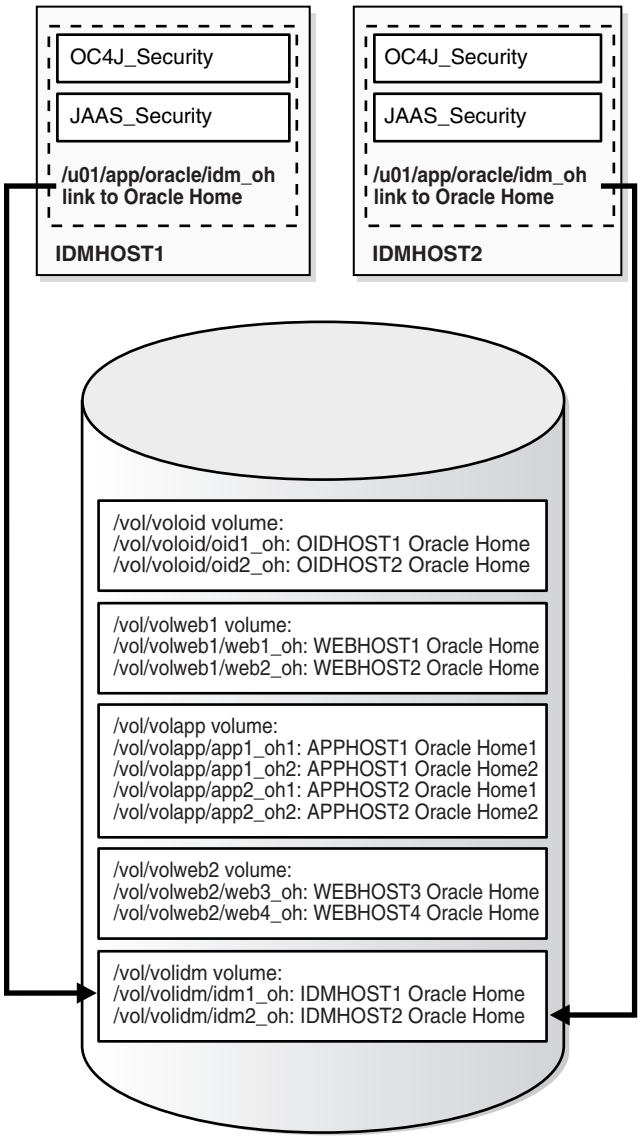


### 2.2.2.3 Creating Mount Points, Symbolic Links, and Oracle Central Inventory Directories

This section describes how to create a directory on the volume for a host cluster to store the Oracle Central Inventory for each host in the cluster, and how to set up a mount point directory and symbolic link on the host to the Oracle Central Inventory directory on the volume.

If the Oracle Central Inventory for a host cluster is on the cluster's volume on the shared storage, the Oracle Central Inventory will be updated when patches or patch sets are applied to an Application Server instance for the host cluster. Then, when the host cluster volume is replicated from the production site shared storage to the standby site shared storage, the updated Oracle Central Inventory will be replicated to the standby site shared storage. Therefore, you only need to apply patches and patch sets at the production site; any updates to the Oracle Central Inventory at the

production site will be replicated to the corresponding host cluster volume at the standby site.

Table 2–6 shows where to create directories to store the Oracle Central Inventory for each host on the host's volume in the shared storage. It also shows the mount points and symbolic links to set up on the host to access the Oracle Central Inventory directory on the volume.

***Table 2–6 Mount Points and Symbolic Links to Oracle Central Inventory Directories on Volumes***

| Host | Host Mount Point Directory to Oracle Central Inventory Directory on Volume | Link on Host to Oracle Central Inventory Directory on Volume | Oracle Central Inventory Directory on Volume |
|---|---|---|---|
| OIDHOST1 | /u02/voloidmount/oid_orainv | /u01/app/oracle/oid_orainv | /vol/voloid/oid_orainv |
| OIDHOST2 | /u02/voloidmount/oid_orainv | /u01/app/oracle/oid_orainv | /vol/voloid/oid_orainv |
| WEBHOST1 | /u02/volweb1mount/web_orainv | /u01/app/oracle/web_orainv | /vol/volweb1/web_orainv |
| WEBHOST2 | /u02/volweb1mount/web_orainv | /u01/app/oracle/web_orainv | /vol/volweb1/web_orainv |
| APPHOST1 | /u02/volappmount/app_orainv | /u01/app/oracle/app_orainv | /vol/volapp/app_orainv |
| APPHOST2 | /u02/volappmount/app_orainv | /u01/app/oracle/app_orainv | /vol/volapp/app_orainv |
| WEBHOST3 | /u02/volweb2mount/web_orainv | /u01/app/oracle/web_orainv | /vol/volweb2/web_orainv |
| WEBHOST4 | /u02/volweb2mount/web_orainv | /u01/app/oracle/web_orainv | /vol/volweb2/web_orainv |
| IDMHOST1 | /u02/volidmmount/idm_orainv | /u01/app/oracle/idm_orainv | /vol/volidm/idm_orainv |
| IDMHOST2 | /u02/volidmmount/idm_orainv | /u01/app/oracle/idm_orainv | /vol/volidm/idm_orainv |

Example 2–12 shows the steps for creating the Oracle Central Inventory directory on the host cluster volume for hosts OIDHOST1 and OIDHOST2. The example assumes that the /u02/voloidmount and /u01/app/oracle directories were already created on the hosts (see Example 2–11).

***Example 2–12 Setting Up Mount Points and Symbolic Links to Oracle Inventory Directories***

1. On OIDHOST1, log in as root and create the directory for the Oracle Central Inventory in the storage:

```
prompt> cd /u02/voloidmount
prompt> mkdir oid_orainv
```

2. On OIDHOST1, create a symbolic link named `oid_orainv` in the `/u01/app/oracle` directory to the `/u02/voloidmount/oid_orainv` Oracle Central Inventory directory on the storage:

```
prompt> cd /u01/app/oracle
prompt> ln -s /u02/voloidmount/oid_orainv oid_orainv
```

3. On OIDHOST1, the following command changes the working directory to the `/vol/voloid/oid_orainv` directory where the host's Oracle Central Inventory will be located on the storage:

```
prompt> cd /u01/app/oracle/oid_orainv
```

4. On OIDHOST2, log in as root and create the directory for the Oracle Central Inventory in the storage:

```
prompt> cd /u02/voloidmount
prompt> mkdir oid_orainv
```

5. On OIDHOST2, create a symbolic link named `oid_orainv` in the `/u01/app/oracle` directory to the `/u02/voloidmount/oid_orainv` Oracle Central Inventory directory on the storage:

```
prompt> cd /u01/app/oracle
prompt> ln -s /u02/voloidmount/oid_orainv oid_orainv
```

6. On OIDHOST2, the following command changes the working directory to the `/vol/voloid/oid_orainv` directory where the host's Oracle Central Inventory will be located on the storage:

```
prompt> cd /u01/app/oracle/oid_orainv
```

Use similar steps for the other host clusters in Table 2–4 to create the Oracle Central Inventory directories on the host cluster's volume, and then to create the mount points and symbolic links on the hosts to the Oracle Central Inventory on the volume. Use the information in Table 2–6 to perform this step.

#### 2.2.2.4 Creating Mount Points, Symbolic Links, and Static HTML Pages Directories

As described in Section 1.2.3.2, "Disaster Recovery Recommendations for Oracle HTTP Server," the static HTML pages for Oracle HTTP Server are deployed by administrators and are typically maintained outside Oracle home directories. In an Oracle Application Server Disaster Recovery topology, you must store these static HTML pages in a directory on the shared storage, and set up mount points and symbolic on the hosts that access the static HTML directories in the shared storage.

If the static HTML pages for a production site host cluster are on the cluster's volume on the shared storage, any changes to those pages will be replicated from the production site shared storage to the corresponding host cluster volume at the standby site.

Table 2–7 shows where to create directories to store the static HTML pages for the production site hosts that use Oracle HTTP Server on the host's volume in the shared storage. It also shows the mount points and symbolic links to set up on the hosts to access the static HTML pages directory on the volume.

***Table 2–7 Mount Points and Symbolic Links to Static HTML Directories on Volumes***

| Host | Host Mount Point Directory to Static HTML Directory on Volume | Link on Host to Static HTML Directory on Volume | Static HTML Directory on Volume |
|---|---|---|---|
| WEBHOST1 | /u02/volweb1mount/web1_sthtml | /u01/app/oracle/web_sthtml | /vol/volweb1/web1_sthtml |
| WEBHOST2 | /u02/volweb1mount/web2_sthtml | /u01/app/oracle/web_sthtml | /vol/volweb1/web2_sthtml |
| WEBHOST3 | /u02/volweb2mount/web3_sthtml | /u01/app/oracle/web_sthtml | /vol/volweb2/web3_sthtml |
| WEBHOST4 | /u02/volweb2mount/web4_sthtml | /u01/app/oracle/web_sthtml | /vol/volweb2/web4_sthtml |

Example 2–13 shows the steps for creating the static HTML pages directory on the host cluster volume for hosts WEBHOST1 and WEBHOST2. The example assumes that the /u02/volweb1mount and /u01/app/oracle directories were already created on the hosts.

***Example 2–13 Setting Up Mount Points and Symbolic Links to Static HTML Pages Directories***

1. On WEBHOST1, log in as root and create the directory for the static HTML pages in the storage:

```
prompt> cd /u02/volweb1mount
prompt> mkdir web1_sthtml
```

2. On WEBHOST1, create a symbolic link named `web_sthtml` in the `/u01/app/oracle` directory to the `/u02/volweb1mount/web1_sthtml` static HTML pages directory on the storage:

```
prompt> cd /u01/app/oracle
prompt> ln -s /u02/volweb1mount/web1_sthtml web_sthtml
```

3. On WEBHOST1, the following command changes the working directory to the `/vol/volweb1/web1_sthtml` directory where the host's static HTML pages will be located on the storage:

```
prompt> cd /u01/app/oracle/web_sthtml
```

4. On WEBHOST2, log in as root and create the directory for the static HTML pages in the storage:

```
prompt> cd /u02/volweb1mount
prompt> mkdir web2_sthtml
```

5. On WEBHOST2, create a symbolic link named `oid_sthtml` in the `/u01/app/oracle` directory to the `/u02/volweb1mount/web2_sthtml` static HTML pages directory on the storage:

```
prompt> cd /u01/app/oracle
prompt> ln -s /u02/volweb1mount/web2_sthtml web_sthtml
```

6. On WEBHOST2, the following command changes the working directory to the `/vol/volweb1/web2_sthtml` directory where the host's static HTML pages will be located on the storage:

```
prompt> cd /u01/app/oracle/web_sthtml
```

Use similar steps for the other host clusters in Table 2–4 that require static HTML pages directories (in the EDG deployment, this would be the WEBHOST3 and WEBHOST4 host cluster). After you create the static HTML pages directories on the host cluster's volume, create the mount points and symbolic links on the hosts to the static HTML pages directories on the volume. Use the information in Table 2–7 to perform this step.

### 2.2.3 Testing the Host Name Resolution

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

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

## 2.3 Installing the Oracle Application Server Instances for the Production Site

After you have performed all the steps in Section 2.2.2 through Section 2.2.3, you can begin to install the Oracle Application Server instances into the Oracle home directories on the production site shared storage.

The topics in this section are:

- Assigning the Application Server Host Name During Installation

### 2.3.1 Assigning the Application Server Host Name During Installation

During the planning of your Oracle Application Server Disaster Recovery topology, you planned the Application Server host names to use for your production site and standby site hosts. The Application Server host name for a host is included in the Oracle Application Server Disaster Recovery topology when you install an Oracle Application Server instance for the host. The Application Server host name for a host is assigned as follows during installation:

1. You can choose the Application Server host name to use for a host by specifying it as the second entry (after the `127.0.0.1 localhost.domain` entry) in the `/etc/hosts` file for the host for which you are installing the instance. For example, specifying the following entry as the second entry in the `/etc/hosts` file for a host with an IP address of 123.1.2.111 will result in an Application Server host name of OIDHOST1 being used for that host:

   ```
   123.1.2.111    OIDHOST1.ORACLE.COM    OIDHOST1
   ```

2. Before installing an Oracle Application Server instance for release 10.1.3 for a host, you can use the `VIRTUAL_HOST_NAME` environment variable to specify the Application Server host name you want to use for the host (even though the name of this variable is VIRTUAL_HOST_NAME, it does set the Application Server host name for the host). For example, if you set the following environment variable before installing an Oracle Application Server 10.1.3 instance for a host, an Application Server host name of OIDHOST2 will be assigned to the host:

   ```
   setenv VIRTUAL_HOST_NAME OIDHOST2
   ```

   Setting the `VIRTUAL_HOST_NAME` variable before a 10.1.3 installation causes the Application Server host name specified with the variable to take precedence over the second entry in the `/etc/hosts` file.

3. If the second entry in the `/etc/hosts` file is not the IP address of the host for which the installation is being performed and the VIRTUAL_HOST_NAME environment variable is not specified (for 10.1.3 installations only), then the Application Server host name that will be used for the host is the name returned when the `hostname` command is issued on the host.

### 2.3.2 Specifying the Oracle Home Directory During Installation

During the installation of the Oracle Application Server instance for the host, the installation procedure will prompt you for the Oracle home directory into which to install the instance. To install the software properly, specify the symbolic link that you created for the Oracle home. in Table 2–5 in Section 2.2.2, "Creating Volumes, Mount Points, and Symbolic Links."

For example, when you are installing the Oracle Application Server instance for OIDHOST1 and the installation prompts you for the Oracle home directory into which to install the instance, you would specify the `/u01/app/oracle/oid_oh` directory. This is a symbolic link to the `/vol/voloid/oid1_oh` directory on the storage, so the Oracle Internet Directory instance will be installed into the `/vol/voloid/oid1_oh` directory on the storage. See Table 2–5 for the names of the symbolic links set up for the Oracle home directories needed for the EDG deployment in Figure 2–1.

Similarly, when you are installing the Oracle Application Server instance for OIDHOST2 and the installation prompts you for the Oracle home directory into which to install the instance, you would specify the `/u01/app/oracle/oid_oh` directory.

This is a symbolic link to the `vol/voloid/oid2_oh` directory on the storage, so the Oracle Internet Directory instance will be installed into the `/vol/voloid/oid2_oh` directory on the storage. Again, see Table 2–5 for the names of the symbolic links set up for the Oracle home directories needed for the EDG deployment in Figure 2–1.

## 2.4 Finishing the Disk Replication Setup

Follow these steps to finish setting up disk replication for the Oracle Application Server Disaster Recovery topology:

1. On the standby site, make sure the same Application Server host names are used for middle tier hosts as were used for the middle tier hosts at the production site.
2. On the shared storage at the standby site, create the same volumes as were created on the shared storage at the production site.
3. On the standby site, create the same mount points and symbolic links that you created at the production site.

> **Note:** It is not necessary to install the same Oracle Application Server instances at the standby site as were installed at the production site. When the production site storage is replicated to the standby site storage, the Oracle software installed on the production site volumes will be replicated at the standby site volumes.

4. Perform any other necessary configuration required by the shared storage vendor to enable disk replication between the production site shared storage and the standby site shared storage.
5. Create the baseline snapshot copy of the production site shared storage that sets up the replication between the production site and standby site shared storage. Create the initial baseline copy and subsequent snapshot copies using asynchronous replication mode. After the baseline snapshot copy is performed, validate that all the directories inside the standby site volumes have the same contents as the directories inside the production site volumes.
6. Set up the frequency of subsequent copies of the production site shared storage, which will be replicated at the standby site. When asynchronous replication mode is used, then at the requested frequency the changed data blocks at the production site shared storage (based on comparison to the previous snapshot copy) become the new snapshot copy, and the snapshot copy is transferred to the standby site shared storage.
7. Make sure that disaster protection for any database that is included in the Oracle Application Server Disaster Recovery production site is provided by Oracle Data Guard, as described earlier in Section 2.2.1, "Database Considerations." Do *not* use disk replication technology to provide disaster protection for Oracle databases.

## 2.5 Synchronization Steps and Frequency

The standby site shared storage receives snapshots transferred on a periodic basis from the production site shared storage. After the snapshots are applied, the standby site shared storage will include all the data up to and including the data contained in the last snapshot transferred from the production site before the failover or switchover.

You should manually force a synchronization operation whenever a change is made to the middle tier at the production site (for example, when a new application is

deployed at the production site). Follow the vendor-specific instructions for forcing a synchronization using disk replication technology.

The synchronization of the databases in the OracleAS Disaster Recovery topology is managed by Oracle Data Guard.

## 2.6 Failover Steps

When the production site becomes unavailable unexpectedly, you must perform a failover operation so that the standby site takes over the production role.

Follow these steps to perform a failover operation (these steps assume that the host names and mount points have been configured properly) :

1. Shut down any processes still running on the production site (if applicable).
2. Stop the replication between the production site shared storage and the standby site shared storage.
3. Use Oracle Data Guard to fail over the databases.
4. Your production site hosts may include Oracle Application Server instances from different Oracle Application Server releases, such as release 10.1.2.x and 10.1.3.x. For any OC4J instances from Oracle Application Server 10.1.3.x releases, download and run the script provided with Oracle*MetaLink* patch 6785728 in the Oracle home directories for the 10.1.3.x OC4J instances to remove persistent OC4J lock files. The URL for Oracle*MetaLink* is:

   `https://metalink.oracle.com`

   **Note:** Run the script in the Oracle homes for 10.1.3.x OC4J instances only.

5. On the standby site hosts, manually start up the processes for the Application Server instances.
6. Ensure that all user requests are routed to the standby site (using a global DNS push or by updating the global load balancer).
7. Use a browser client to perform post-failover testing to confirm that requests are being resolved and redirected to the new failed over site.

After these steps have been performed, the former standby site has become the currrent production site. At this point, the issues that caused the original production site to become unavailable can be examined. Then work can begin on resolving those issues so that the original production site can be used again at some point in the future as either the production site or standby site, if desired.

To use the original production site as the current standby site, you need to reestablish the replication between the two sites, but configure the replication so that the snapshot copies go in the opposite direction (from the current production site to the current standby site). Refer to the documentation for your shared storage system to learn how to configure the replication so that snapshot copies are transferred in the opposite direction.

To use the original production site as the new production site, you should perform the switchover steps in Section 2.7, "Switchover Steps."

## 2.7 Switchover Steps

When you plan to take down the production site (for example, to perform maintenance) and make the current standby site the new production site, you must perform a switchover operation so that the standby site takes over the production role.

Follow these steps to perform a switchover operation:

1. Shut down any processes still running on the production site.
2. Stop the replication between the production site shared storage and the standby site shared storage.
3. Use Oracle Data Guard to switch over the databases.
4. Your production site hosts may include Oracle Application Server instances from different Oracle Application Server releases, such as release 10.1.2.x and 10.1.3.x. For any OC4J instances from Oracle Application Server 10.1.3.x releases, download and run the script provided with Oracle*MetaLink* patch 6785728 in the Oracle home directories for the 10.1.3.x OC4J instances to remove persistent OC4J lock files. The URL for Oracle*MetaLink* is:

   `https://metalink.oracle.com`

   **Note:** Run the script in the Oracle homes for 10.1.3.x OC4J instances only.

5. On the standby site hosts, manually start up the processes for the Application Server instances.
6. Ensure that all user requests are routed to the standby site (using a global DNS push or by 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 new switched over site.

After these steps have been performed, the former standby site has become the currrent production site. At this point, you can perform maintenance at the original production site. After performing the planned tasks on the original production site, you can use it again at some point in the future as either the production site or standby site.

To use the original production site as the current standby site, you need to reestablish the replication between the two sites, but configure the replication so that the snapshot copies go in the opposite direction (from the current production site to the current standby site). Refer to the documentation for your shared storage to learn how to configure the replication so that snapshot copies are transferred in the opposite direction.

To use the original production site as the new production site, perform the switchover steps described in this section again.

## 2.8 Performing Periodic Testing of the Standby Site

This manual describes how to set up Disaster Recovery for an Oracle Application Server production site and standby site. In a normal Oracle Application Server Disaster Recovery configuration, the following are true:

- Disk replication is used to copy Oracle Application Server 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 disk replication operations from the production site shared storage to the standby site shared storage.

- Oracle Data Guard is used to copy database data for the production site Oracle databases to the standby databases at standby site. By default, the production site databases are active and the standby databases at the standby site are passive. The standby databases at the standby site are in managed recovery mode while the standby site is in the standby role (is passive). When the production site is active, the only write operations made to the standby databases are the database synchronization operations performed by Oracle Data Guard.
- When the production site becomes unavailable, the standby site is enabled to take over the production role. If the current production site becomes unavailable unexpectedly, then a failover operation (described in Section 2.6, "Failover Steps") is performed to enable the standby site to assume the production role. Or, if the current production site is taken down intentionally (for example, for planned maintenance), then a switchover operation (described in Section 2.7, "Switchover Steps") is performed to enable the standby site to assume the production role.

The usual method of testing a standby site is to shut down the current production site and perform a switchover operation to enable the standby site to assume the production role. However, some enterprises may want to perform periodic testing of their Disaster Recovery standby site without shutting down the current production site and performing a switchover operation.

An alternate method of testing the standby site without shutting down the current production site is to create a clone of the read-only standby site shared storage and then use the cloned standby site shared storage in testing. To use this alternate testing method, perform these steps:

1. Use the cloning technology provided by the shared storage vendor to create a clone of the standby site's read-only volumes on the shared storage at the standby site. Make sure that the cloned standby site volumes are writable. If you want to test the standby site just once, then this can be a one-time clone operation, but if you want to test the standby site regularly, you can set up periodic cloning of the standby site read-only volumes to the standby site's cloned read/write volumes.
2. Perform a backup of the standby site databases, then modify the Oracle Data Guard replication between the production site and standby site databases.
   - For 10.1 databases, break the replication by following the instructions in the 10.1 Oracle Data Guard documentation.
   - For 10.2 and later databases, follow these steps to establish a snapshot standby database:
     a. If you do not have a Flash Recovery Area, set one up.
     b. Cancel Redo Apply:

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

     c. Create a guaranteed restore point:

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

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

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

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

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

f. Activate the target standby database:

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

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

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

h. Lower the protection mode and open the database:

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

- For 11*g* databases, use the procedure to establish a snapshot standby database in the "Managing a Snapshot Standby Database" section in the 11*g Oracle Data Guard Concepts and Administration*.

3. Use Oracle Data Guard database recovery procedures to bring the standby databases online.
4. On the standby site computers, modify the mount commands to point to the volumes on the standby site's cloned read/write shared storage by following these steps:
   a. Unmount the read-only shared storage volumes.
   b. Mount the cloned read/write volumes at the same mount point.
5. Before doing the standby site testing, modify the hostname resolution method for the computers that will be used to perform the testing to ensure that the host names point to the standby site computers and not the production site computers. For example, on a Linux computer, change the `/etc/hosts` file to point to the virtual IP of the load balancer for the standby site.
6. Perform the standby site testing.

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

1. Modify the mount commands on the standby site computers to point to the volumes on the standby site's read-only shared storage: In other words, reset the mount commands back to what they were before the testing was performed.
   a. Unmount the cloned read/write shared storage volume.
   b. Mount the read-only shared storage volumes.

   At this point, the mount commands are reset to what they were before the standby site testing was performed.
2. Configure Oracle Data Guard to perform replication between the production site databases and standby databases at the standby site. Performing this configuration puts the standby database into managed recovery mode again:
   - For 10.1 databases, reinstantiate the databases by following the instructions in the 10.1 Oracle Data Guard documentation.

- For 10.2 and later databases, follow these steps:

  **a.** Revert the activated database back to a physical standby database:

  ```
  SQL> STARTUP MOUNT FORCE;
  SQL> FLASHBACK DATABASE TO POINT standbytest;
  SQL> ALTER DATABASE CONVERT TO PHYSICAL STANDBY;
  SQL> STARTUP MOUNT FORCE;
  ```

  **b.** Restart managed recovery:

  ```
  SQL> ALTER DATABASE RECOVER MANAGED STANDBY
       DATABASE USING CURRENT LOGFILE DISCONNECT;
  ```

  **c.** Reenable the standby destination and switch logs:

  ```
  SQL> ALTER SYSTEM SET
       LOG ARCHIVE DEST STATE 2=ENABLE;
  ```

- For 11*g* databases, set up the replication again by following the steps in the "Managing a Snapshot Standby Database" section in the 11*g Oracle Data Guard Concepts and Administration*.

**3.** Before using the original production site again, modify the hostname resolution method for the computers that will be used to access the production site to ensure that the host names point to the production site computers and not the standby site computers. For example, on a Linux computer, change the `/etc/hosts` file to point to the virtual IP of the load balancer for the production site.

# A

# Using Databases in the OracleAS Disaster Recovery Solution

This appendix provides information about the steps to follow to use databases in the Oracle Application Server Disaster Recovery solution.

It includes the following topics:



## A.1 Installing the Oracle Databases in the Production Site

Install the Oracle databases that will be used at the production site for the Oracle Application Server Disaster Recovery topology that you are setting up.

These databases must exist before you install the Oracle Application Server components that store metadata in database repositories.

If you are setting up the EDG deployment shown in Figure 2–1, for example, you would install the Security Metadata Repository Real Application Clusters (RAC) database on hosts INFRADBHOST1 and INFRADBHOST2 at the production site, and you would also install the Customer RAC database on hosts CUSTDBHOST1 and CUSTDBHOST2 at the production site.

## A.2 Creating Standby Databases at the Standby Site

After you install the databases at the production site, the next step is to create a standby database at the standby site for each database included in your Oracle Application Server Disaster Recovery production site.

Create the standby database by following the instructions in the "Creating a Standby Database that Uses OMF or ASM" section of *Oracle Data Guard Concepts and Administration* in the Oracle Database documentation set.

## A.3 Setting Up the Oracle Data Guard Configuration

After creating the production site and standby site database, set up the Oracle Data Guard configuration that Oracle Application Server Disaster Recovery recommends for these databases.

Follow these recommendations to set up the Oracle Data Guard configuration for production site and standby site databases in your Oracle Application Server Disaster Recovery topology:

1. Set up the production site database with the maximum availability data protection mode by issuing the SQL statement in Example A–1:

***Example A–1 Setting Up a Production Site Database in Maximum Availability Data Protection Mode***

```
ALTER DATABASE SET STANDBY DATABASE DATABASE TO MAXIMIZE AVAILABILITY;
```

> **Note:** Do NOT set up the production site database in maximum protection mode.

2. Set up the standby database with the LGRW SYNC and AFFIRM archive attributes for the LOG_ARCHIVE_DEST_*n* parameter.
3. Place the standby database at the standby site in managed recovery mode. This puts the standby database in a constant state of media recovery. Placing the standby database in managed recovery mode is not a requirement of maximum availability, but it provides for shorter failover times.

   On the standby database, issue the SQL statement in Example A–2 to place the database in managed recovery mode. Add the optional `disconnect from session` clause if you want to end the session after the command:

***Example A–2 Placing a Standby Database in Managed Recovery Mode***

```
ALTER DATABASE RECOVERY MANAGED STANDBY DATABASE DISCONNECT FROM SESSION;
```

## A.4 Making TNSNAMES.ORA Entries for Databases

Because Oracle Data Guard is used to synchronize production and standby databases, the production database and standby database must be able to reference each other.

Oracle Data Guard uses `tnsnames.ora` file entries to direct requests to the production and standby databases, so entries for production and standby databases must be made to the `tnsnames.ora` file. See *Oracle Data Guard Concepts and Administration* in the Oracle Database documentation set for more information about using tnsnames.ora files with Oracle Data Guard.

## A.5 Manually Forcing Database Synchronization with Oracle Data Guard

For Oracle Application Server components that store middle tier configuration data in Oracle database repositories, use Oracle Data Guard to manually force a database synchronization whenever a middle tier synchronization is performed. Use the SQL `alter system archive log all` statement to switch the logs, which forces the synchronization of the production site and standby site databases.

Example A–3 shows the SQL statement to use to force the synchronization of a production site database and standby site database.

***Example A–3 Manually Forcing an Oracle Data Guard Database Synchronization***

```
ALTER SYSTEM ARCHIVE LOG ALL;
```

## A.6 Setting Up Database Host Name Aliases

Optionally, you can set up database host name aliases for the databases at your production site and standby site. The alias must be defined in DNS or in the `/etc/hosts` file on each node running a database instance.

In a Disaster Recovery environment, the site that actively accepts connections is the production site. At the completion of a successful failover or switchover operation, the standby site becomes the new production site.

This section includes an example of defining an alias for database hosts named stajo01 and stajo02. Table A–1 shows the database host names and the connect strings for the databases before the alias is defined

***Table A–1 Database Host Names and Connect Strings***

| Site | Database Host Name | Database Connect String |
|---|---|---|
| Production | stajo01.us.oracle.com | stajo01.us.oracle.com:1521:orcl |
| Standby | stajo02.us.oracle.com | stajo02.us.oracle.com:1521:orcl |

In this example, all database connect strings on the production site take the form "stajo01.us.oracle.com:1521:orcl." After a failover or switchover operation, this connect string must be changed to "stajo02.us.oracle.com:1521:orcl." However, by creating an alias of "proddb1" for the database host name as shown in Table A–2, you can avoid manually changing the connect strings, which enables seamless failovers and switchovers:

***Table A–2 Specifying an Alias for a Database Host***

| Site | Database Host Name | Alias | Database Connect String |
|---|---|---|---|
| Production | stajo01.us.oracle.com | proddb1.us.oracle.com | proddb1.us.oracle.com:1521:orcl |
| Standby | stajo02.us.oracle.com | proddb1.us.oracle.com | proddb1.us.oracle.com:1521:orcl |

In this example, the production site database host name and the standby site database host name are aliased to "proddb1.us.oracle.com" and the connect strings on the production site and the standby site can take the form "proddb1.us.oracle.com:1521:orcl". On failover and switchover operations, the connect string does not need to change, thus enabling a seamless failover and switchover.

The format for specifying aliases in /etc/hosts file entries is:

```
<IP>    <ALIAS WITH DOMAIN> <ALIAS>    <HOST NAME WITH DOMAIN> <HOST NAME>
```

In this example, you create a database host name alias of proddb1 for host stajo01 at the production site and for host stajo02 at the standby site. The hosts file entry should specify the fully qualified database host name alias with the `<ALIAS WITH DOMAIN>` parameter, the short database host name alias with the `<ALIAS>` parameter, the fully

qualified host name with the <HOST NAME WITH DOMAIN> parameter, and the short host name with the <HOST NAME> parameter.

So, in the /etc/hosts files at the production site, make sure the entry for host stajo01 looks like this:

```
152.68.196.213   proddb1.us.oracle.com proddb1   stajo01.us.oracle.com stajo01
```

And, in the /etc/hosts files at the standby site, make sure the entry for host stajo02 looks like this:

```
140.87.25.40   proddb1.us.oracle.com proddb1   stajo02.us.oracle.com stajo02
```

# B

# Setting Up Oracle Business Activity Monitoring

Disk replication is used for disaster protection for most Oracle Application Server components. However, disk replication is not used for disaster protection for Oracle Business Activity Monitoring.

This appendix describes how to set up data protection for Oracle Business Activity Monitoring in an Oracle Application Server Disaster Recovery topology.

## B.1 Setting Up Oracle Business Activity Monitoring in an Oracle Application Server Disaster Recovery Topology

Follow these steps to set up data protection for Oracle Business Activity Monitoring in an Oracle Application Server Disaster Recovery topology:

1. Use Oracle Data Guard for database recovery. The database at the production site is the production database and the database at the standby site is the standby database. Configure Oracle Data Guard to keep the production and standby databases synchronized. See Section 2.2.1, "Database Considerations" for more information about setting up Oracle Data Guard for data protection of Oracle databases.
2. Perform a complete Oracle Business Activity Monitoring installation on both the production site host and standby site host.
3. Use the same database during both installations. For example, specify the same production database for Oracle Business Activity Monitoring with the Sagent schemas (orabam and orasagent, by default).
4. When you need to install a patch for Oracle Business Activity Monitoring, install it on both the production site host and the standby site host.
5. Shut down Oracle Business Activity Monitoring on the standby site.
6. For production, use Oracle Business Activity Monitoring on the production site. At regular intervals, check that the Oracle Data Guard synchronization is working properly, keeping the standby database updated.
7. If the production site fails, follow these steps to recover Oracle Business Activity Monitoring at the standby site:
    a. Use Oracle Data Guard database recovery procedures to bring the standby database online.
    b. Locate the `tnsnames.ora` file in the Oracle Business Activity Monitoring database client location, which has a default location of:

```
C:\OracleBAM\ClientForBAM\NETWORK\ADMIN\tnsnames.ora
```

**c.** In the `tnsnames.ora` file, change the name of the database host to the name of the standby database host.

**d.** In the `tnsnames.ora` file, change the database port number to the port number for the standby database.

**e.** In the `tnsnames.ora` file, change the database name to the standby database name.

**f.** Start Oracle Business Activity Monitoring using the shortcut on the desktop.

---

**Note:** As a best practice, when using JMS messaging with Oracle Business Activity Monitoring, use the Message Integrity feature (also known as the Guaranteed Messaging feature).

For information on setting up Message Integrity for Oracle Business Activity Monitoring, see the "Message Integrity Setup" section of the *Oracle Application Server High Availability Guide* for Oracle Application Server release 10.1.3.1.0.

---

# C

# Creating an Asymmetric Topology

Chapter 2, "Implementing the Solution" of this manual describes how to set up a symmetric Oracle Application Server Disaster Recovery topology, which is a configuration that is completely identical across tiers on the production site and standby site.

Oracle Application Server Disaster Recovery also supports asymmetric topologies. An asymmetric topology is a disaster recovery configuration that is different across tiers at the production site and standby site. In most asymmetric OracleAS Disaster Recovery topologies, the standby site has fewer resources than the production site.

Before you read this appendix, read Chapter 2, "Implementing the Solution" to ensure that you understand the concepts and information presented in that chapter. Many of the concepts for setting up a symmetric topology are also valid for setting up an asymmetric topology.

Section C.1, "Steps for Creating an Asymmetric Topology" describes the basic steps for creating an asymmetric topology. It does not describe in detail applicable concepts for setting up an asymmetric topology that were previously described for symmetric topologies in Chapter 2, "Implementing the Solution."

Some of the different types of supported asymmetric topologies include:

- A standby site with fewer hosts and OracleAS instances than the production site. See Section C.1.1, "Creating an Asymmetric Standby Site with Fewer Hosts and Instances" for more information about this topology.
- A standby site with fewer hosts and the same number of instances as the production site. See Section C.1.2, "Creating an Asymmetric Standby Site with Fewer Hosts and the Same Number of Instances" for more information about this topology.
- A standby site with a different database configuration than the production site. For example, a standby site could use a single instance database instead of the Real Application Clusters database used at the production site. See Section C.1.3, "Creating an Asymmetric Standby Site with a Different Database Configuration" for more information about this topology.

---

**Note:** All the host names used in this appendix are Application Server host names.

---

# C.1 Steps for Creating an Asymmetric Topology

This section describes the high level steps for creating any type of asymmetric OracleAS Disaster Recovery topology. The production site is the EDG deployment shown in Figure 2–1. The standby site will be different from the production site.

To create an asymmetric topology:

1. Design the production site and the standby site. Determine the resources that will be necessary at the standby site to ensure acceptable performance when the standby site assumes the production role.

   **Note:** The ports for the standby site instances must use the same port numbers as the peer instances at the production site. Therefore, make sure that all the port numbers that will be required at the standby site are available (not in use at the standby site).

2. Create the OracleAS Disaster Recovery production site by performing these operations:

   a. Create volumes on the production site's shared storage system for the OracleAS instances that will be installed for the production site. For more information, see Section 2.2.2.1, "Creating Volumes for the Application Server Host Clusters."

   b. Create mount points and symbolic links on the production site hosts to the Oracle home directories for the OracleAS instances on the production site's shared storage system volumes. For more information, see Section 2.2.2.2, "Creating Mount Points, Symbolic Links, and Oracle Home Directories."

   c. Create mount points and symbolic links on the production site hosts to the Oracle Central Inventory directories for the OracleAS instances on the production site's shared storage system volumes. For more information, see Section 2.2.2.3, "Creating Mount Points, Symbolic Links, and Oracle Central Inventory Directories."

   d. Create mount points and symbolic links on the production site hosts to the static HTML pages directories for the Oracle HTTP Server instances on the production site's shared storage system volumes, if applicable. For more information, see Section 2.2.2.4, "Creating Mount Points, Symbolic Links, and Static HTML Pages Directories."

   e. Install the OracleAS instances for the production site on the volumes in the production site's shared storage system. For more information, see Section 2.3, "Installing the Oracle Application Server Instances for the Production Site."

3. Create the same volumes with the same file and directory privileges on the standby site's shared storage system as you created for the OracleAS instances on the production site's shared storage system. This step is critical because it enables you to use disk replication later to create the peer OracleAS instance installations for the standby site instead of installing them using Oracle Universal Installer.

**Note:** When you configure disk replication, make sure that all the volumes you set up on the production site's shared storage system are replicated to the same volumes on the standby site's shared storage system.

Even though some of the instances and hosts at the production site may not exist at the standby site, you must configure disk replication for all the volumes set up for the production site's OracleAS instances.

4. Perform any other necessary configuration required by the shared storage vendor to enable disk replication between the production site's shared storage system and the standby site's shared storage system. Configure disk replication to asynchronously copy the volumes for the OracleAS home directories, Oracle Central Inventory directories, and Oracle HTTP Server static HTML pages directories in the production site's shared storage system to the standby site's shared storage system.
5. Create the initial baseline snapshot copy of the production site shared storage system to set up the replication between the production site and standby site shared storage systems. Create the initial baseline snapshot and subsequent snapshot copies using asynchronous replication mode. After the baseline snapshot copy is performed, validate that all the directories for the standby site volumes have the same contents as the directories for the production site volumes. Refer to the documentation for your shared storage vendor for information on creating the initial snapshot and enabled disk replication between the production site and standby site shared storage systems.
6. After the baseline snapshot has been taken, perform these steps for the OracleAS instances for the standby site hosts:
   a. Set up a mount point directory on the standby site host to the Oracle home directory for the OracleAS instance on the standby site's shared storage system. The mount point directory you set up for the peer instance on the standby site host must be the same as the mount point directory you set up for the instance on the production site host.
   b. Set up a symbolic link on the standby site host to the Oracle home directory for the OracleAS instance on the standby site's shared storage system. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link you set up for the instance on the production site host.
   c. Set up a mount point directory on the standby site host to the Oracle Central Inventory directory for the OracleAS instance on the standby site's shared storage system. The mount point directory you set up for the peer instance on the standby site host must be the same as the mount point directory you set up for the instance on the production site host.
   d. Set up a symbolic link on the standby site host to the Oracle Central Inventory directory for the OracleAS instance on the standby site's shared storage system. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link you set up for the instance on the production site host.
   e. Set up a mount point directory on the standby site host to the Oracle HTTP Server static HTML pages directory for the Oracle HTTP Server instance on the standby site's shared storage system. The mount point directory you set

up for the peer instance on the standby site host must be the same as the mount point directory you set up for the instance on the production site host.

f. Set up a symbolic link on the standby site host to the Oracle HTTP Server static HTML pages directory for the Oracle HTTP Server instance on the standby site's shared storage system. The symbolic link you set up for the peer instance on the standby site host must be the same as the symbolic link you set up for the instance on the production site host.

After completing these steps, the OracleAS instance installations for the production site have been replicated to the standby site. At the standby site, all of the following are true:

- The OracleAS instances are installed into the same Oracle home directories on the same volumes as at the production site, and the hosts use the same mount point directories and symbolic links for the Oracle home directories as at the production site.
- The Oracle Central Inventory directories are located in same directories on the same volumes as at the production site, and the hosts use the same mount point directories and symbolic links for the Oracle Central Inventory directories as at the production site.
- The Oracle HTTP Server static HTML pages directories are located in same directories on the same volumes as at the production site, and the hosts use the same mount point directories and symbolic links for the Oracle HTTP Server static HTML pages directories as at the production site.
- The same ports are used for the standby site OracleAS instances as were used for the same instances at the production site.

### C.1.1 Creating an Asymmetric Standby Site with Fewer Hosts and Instances

This section describes how to create an asymmetric standby site that has fewer hosts and OracleAS instances than the production site.

The production site for this OracleAS Disaster Recovery topology is the EDG deployment shown in Figure 2–1. Section 2.2, "Environment Preparation" through Section 2.3, "Installing the Oracle Application Server Instances for the Production Site" describe how to set up this production site and the volumes for its shared storage system, and how to create the necessary mount points and symbolic links.

Figure C–1 shows the asymmetric standby site for the production site shown in Figure 2–1.

***Figure C–1 An Asymmetric Standby Site with Fewer Hosts and Instances***

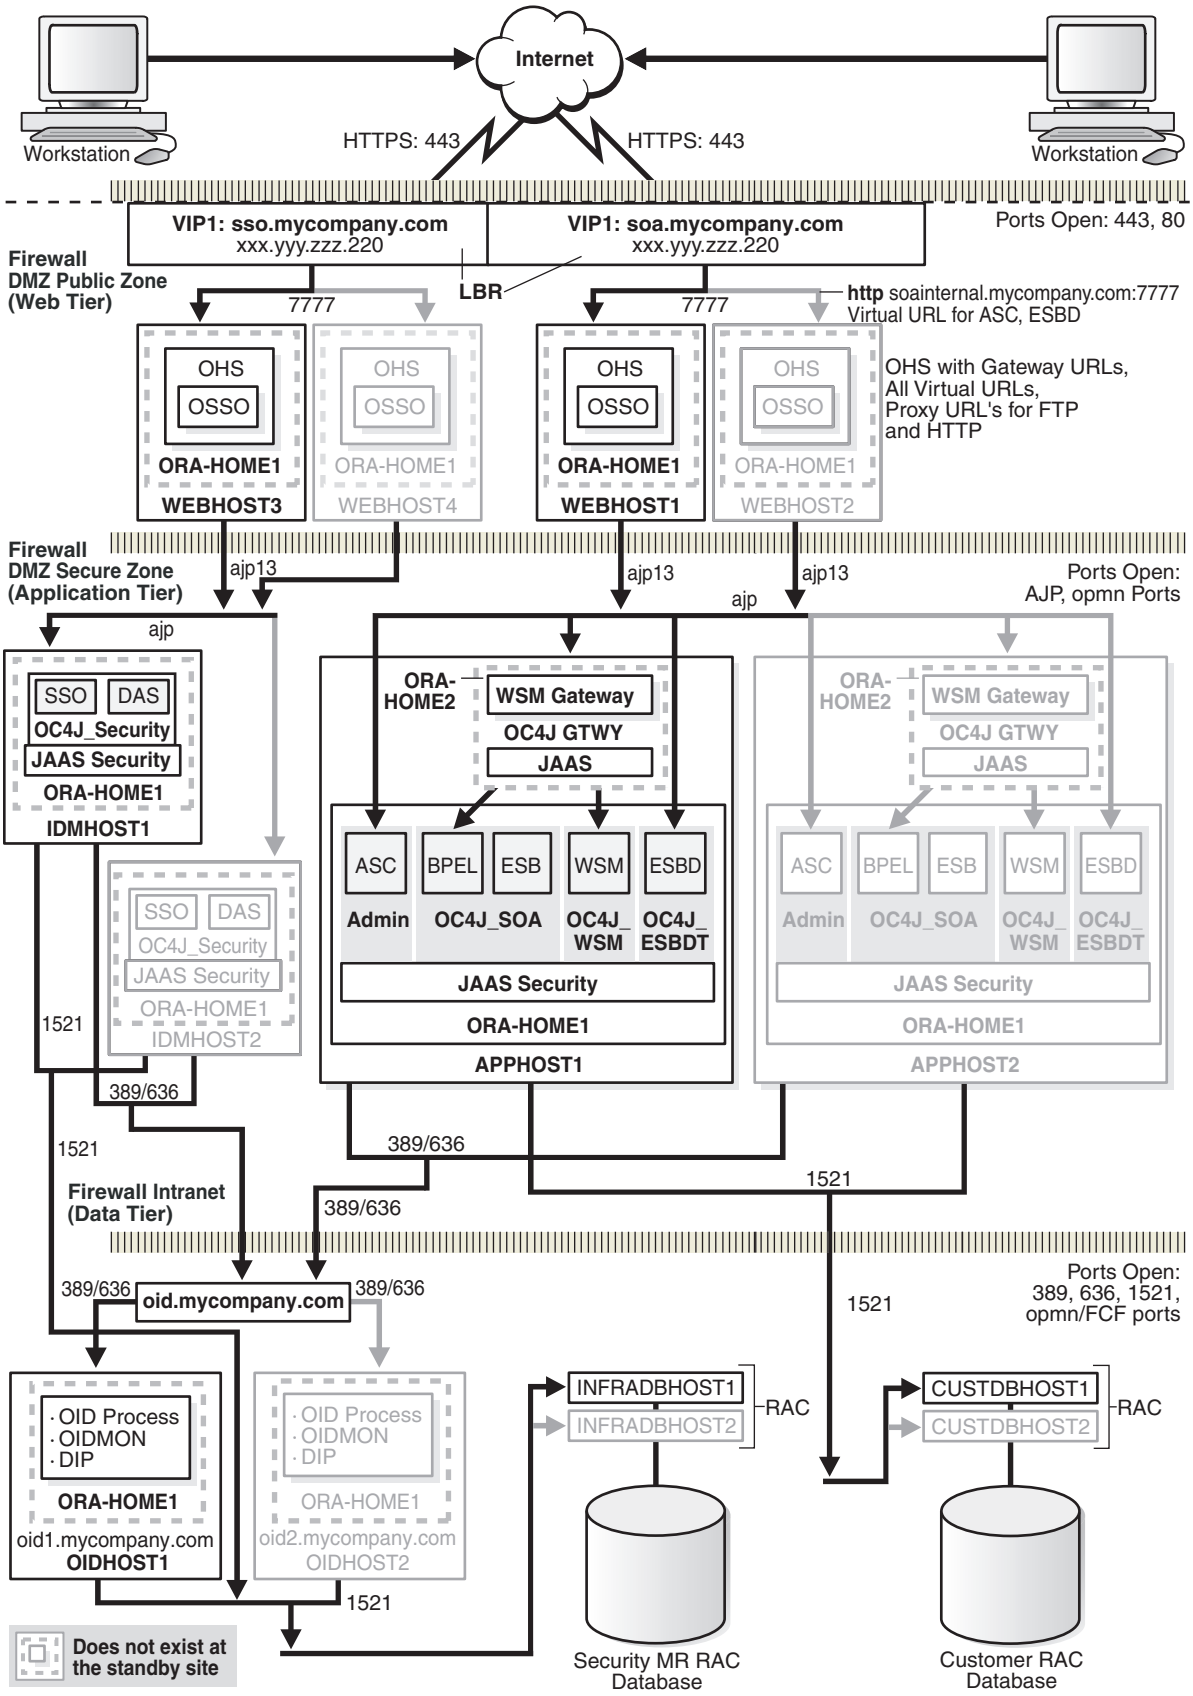
Internet
Workstation
HTTPS: 443
HTTPS: 443
Workstation
VIP1: sso.mycompany.com
xxx.yyy.zzz.220
VIP1: soa.mycompany.com
xxx.yyy.zzz.220
Ports Open: 443, 80
Firewall
DMZ Public Zone
(Web Tier)
LBR
7777
7777
http soainternal.mycompany.com:7777
Virtual URL for ASC, ESBD
OHS
OSSO
ORA-HOME1
WEBHOST3
WEBHOST4
WEBHOST1
WEBHOST2
OHS with Gateway URLs,
All Virtual URLs,
Proxy URL's for FTP
and HTTP
Firewall
DMZ Secure Zone
(Application Tier)
ajp13
ajp13
ajp13
ajp
Ports Open:
AJP, opmn Ports
ajp
SSO
DAS
OC4J_Security
JAAS Security
ORA-HOME1
IDMHOST1
IDMHOST2
ORA-
HOME2
WSM Gateway
OC4J GTWY
JAAS
ASC
BPEL
ESB
WSM
ESBD
Admin
OC4J_SOA
OC4J_
WSM
OC4J_
ESBDT
JAAS Security
ORA-HOME1
APPHOST1
APPHOST2
1521
389/636
1521
389/636
389/636
1521
Firewall Intranet
(Data Tier)
389/636
oid.mycompany.com
389/636
1521
Ports Open:
389, 636, 1521,
opmn/FCF ports
· OID Process
· OIDMON
· DIP
ORA-HOME1
oid1.mycompany.com
OIDHOST1
oid2.mycompany.com
OIDHOST2
1521
INFRADBHOST1
INFRADBHOST2
RAC
CUSTDBHOST1
CUSTDBHOST2
RAC
Security MR RAC
Database
Customer RAC
Database
Does not exist at
the standby site

The asymmetric standby site shown in Figure C–1 has fewer hosts and instances than its production site shown in Figure 2–1.

The hosts WEBHOST2, WEBHOST4, IDMHOST2, APPHOST2, and OIDHOST2 and the instances on those hosts exist at the production site in Figure 2–1, but not at the asymmetric standby site in Figure C–1. The standby site therefore has two fewer Oracle HTTP Server instances, one fewer Identity Management instance, one fewer OC4J instance and SOA Suite instance, and one fewer Oracle Internet Directory instance than the production site.

It is important to make sure that this standby site will have sufficient resources to provide adequate performance when it assumes the production role.

When you follow the steps in Section C.1, "Steps for Creating an Asymmetric Topology" to set up this asymmetric standby site, the standby site should be properly configured to assume the production role.

Table C–1 shows the standby site hosts and the volumes to mount at the standby site.

***Table C–1 Mount Points, Links, and Oracle Homes for Standby Site OracleAS Instances***

| Host | Application Server Instance Type | Host Mount Point Directory to Oracle Home Directory on Volume | Link on Host to Oracle Home Directory on Volume | Oracle Home Directory on Volume |
|---|---|---|---|---|
| OIDHOST1 | Oracle Internet Directory | /u02/voloidmount/oid1_oh | /u01/app/oracle/oid_oh | /vol/voloid/oid1_oh |
| IDMHOST1 | Oracle Identity Management | /u02/volidmmount/idm1_oh | /u01/app/oracle/idm_oh | /vol/volidm/idm1_oh |
| APPHOST1 | First OC4J Instance | /u02/volappmount/app1_oh1 | /u01/app/oracle/app_oh | /vol/volapp/app1_oh1 |
| APPHOST1 | Second OC4J Instance | /u02/volappmount/app1_oh2 | /u01/app/oracle/app_oh2 | /vol/volapp/app1_oh2 |
| WEBHOST1 | Oracle HTTP Server | /u02/volweb1mount/web1_oh | /u01/app/oracle/web_oh | /vol/volweb1/web1_oh |
| WEBHOST3 | Oracle HTTP Server | /u02/volweb2mount/web3_oh | /u01/app/oracle/web_oh | /vol/volweb2/web3_oh |

## C.1.2 Creating an Asymmetric Standby Site with Fewer Hosts and the Same Number of Instances

This section describes how to create an asymmetric standby site that has fewer hosts and the same number of instances as the production site.

For the example in this section, the production site is the EDG deployment in Figure 2–1. The asymmetric standby site does not have peer hosts for hosts IDMHOST1 and IDMHOST2 at the production site. Instead, the Identity Management instances for IDMHOST1 and IDMHOST2 at the production site are available from hosts APPHOST1 and APPHOST2 at the standby site.

At the standby site, host name resolution is set up so that Application Server host name IDMHOST1 resolves to host APPHOST1 and Application Server host name IDMHOST2 resolves to host APPHOST2.

If you plan to use this type of asymmetric standby site where the same number of production site OracleAS instances will be available on fewer hosts at the standby site, then you must follow these requirements when you design the production site and standby site to ensure that the asymmetric standby site can be set up correctly:

- Install the OracleAS instances for the separate production site hosts into unique Oracle home directories on the volumes. On the hosts, set up unique mount point directories and symbolic links to the Oracle homes. Using unique Oracle home directories, mount points, and symbolic links ensures that there will be no namespace conflicts when these instances are available from fewer hosts at the standby site.
- Set up the Oracle Central Inventory directories for the separate production site hosts in unique directories on the volumes. On the hosts, set up unique mount point directories and symbolic links to those Oracle Central Inventory directories. Using unique Oracle Central Inventory directories, mount points, and symbolic links ensures that there will be no namespace conflicts when these Oracle Central Inventory directories are accessible from fewer hosts at the standby site.
- The port numbers used for each production site OracleAS instance must be available for the instance at its standby site host.

After this asymmetric topology has been set up, Table C–2 shows the OracleAS instances, mount points, links, and Oracle home directories that will exist for hosts APPHOST1 and APPHOST2 at the standby site.

***Table C–2 Mount Points, Links, and Oracle Homes for OracleAS Instances on Standby Site Hosts APPHOST1 and APPHOST2***

| Host | Application Server Instance Type | Host Mount Point Directory to Oracle Home Directory on Volume | Link on Host to Oracle Home Directory on Volume | Oracle Home Directory on Volume |
|---|---|---|---|---|
| APPHOST1 | Oracle Identity Management[1] | /u02/volidmmount/idm1_oh | /u01/app/oracle/idm_oh | /vol/volidm/idm1_oh |
| APPHOST1 | First OC4J Instance | /u02/volappmount/app1_oh1 | /u01/app/oracle/app_oh | /vol/volapp/app1_oh1 |
| APPHOST1 | Second OC4J Instance | /u02/volappmount/app1_oh2 | /u01/app/oracle/app_oh2 | /vol/volapp/app1_oh2 |
| APPHOST2 | Oracle Identity Management[2] | /u02/volidmmount/idm2_oh | /u01/app/oracle/idm_oh | /vol/volidm/idm2_oh |
| APPHOST2 | First OC4J Instance | /u02/volappmount/app2_oh1 | /u01/app/oracle/app_oh | /vol/volapp/app2_oh1 |
| APPHOST2 | Second OC4J Instance | /u02/volappmount/app2_oh2 | /u01/app/oracle/app_oh2 | /vol/volapp/app2_oh2 |

[1] This instance was available from host IDMHOST1 at the production site.

[2] This instance was available from host IDMHOST2 at the production site.

Table C–3 shows the OracleAS instances, mount points, links, and Oracle Central Inventory directories that will exist for hosts APPHOST1 and APPHOST2 at the standby site.

*Table C–3 Mount Points and Symbolic Links to Oracle Central Inventory Directories on Volumes*

| Host | Application Server Instance Type | Host Mount Point Directory to Oracle Central Inventory Directory on Volume | Link on Host to Oracle Central Inventory Directory on Volume | Oracle Central Inventory Directory on Volume |
|---|---|---|---|---|
| APPHOST1 | Oracle Identity Management[1] | /u02/volidmmount/idm_orainv | /u01/app/oracle/idm_orainv | /vol/volidm/idm_orainv |
| APPHOST1 | First OC4J Instance | /u02/volappmount/app_orainv | /u01/app/oracle/app_orainv | /vol/volapp/app_orainv |
| APPHOST1 | Second OC4J Instance | /u02/volappmount/app_orainv | /u01/app/oracle/app_orainv | /vol/volapp/app_orainv |
| APPHOST2 | Oracle Identity Management[2] | /u02/volidmmount/idm_orainv | /u01/app/oracle/idm_orainv | /vol/volidm/idm_orainv |
| APPHOST2 | First OC4J Instance | /u02/volappmount/app_orainv | /u01/app/oracle/app_orainv | /vol/volapp/app_orainv |
| APPHOST2 | Second OC4J Instance | /u02/volappmount/app_orainv | /u01/app/oracle/app_orainv | /vol/volapp/app_orainv |

[1] This instance was available from host IDMHOST1 at the production site.

[2] This instance was available from host IDMHOST2 at the production site.

To create this asymmetric topology:

- Replicate all the volumes on the production site's shared storage system to the same volumes on the standby site's shared storage system. This copies the Oracle home directories, Oracle Central Inventory directories, and Oracle HTTP Server static HTML pages directories (if applicable) from the production site volumes to the standby site volumes.
- On host APPHOST1 at the standby site, set up the same mount points and symbolic links that you set up on host APPHOST1 at the production site. Also, on host APPHOST1 at the standby site, set up the same mount points and symbolic links that you set up for OracleAS instances on host IDMHOST1 at the production site.
- On host APPHOST2 at the standby site, set up the same mount points and symbolic links that you set up on host APPHOST2 at the production site. Also, on host APPHOST2 at the standby site, set up the same mount points and symbolic links that you set up for OracleAS instances on host IDMHOST2 at the production site.
- At the standby site, make sure that the host name IDMHOST1 resolves to host APPHOST1 and that host name IDMHOST2 resolves to host APPHOST2. Make sure that all the standby site hosts resolve host name IDMHOST1 to APPHOST1 and host name IDMHOST2 to APPHOST2. For more information on resolving an Application Server host name to a particular IP address, see Section 2.1.6, "Making /etc/hosts File Entries for Local Host Name File Resolution" if you are resolving host names using local host name file resolution or Section 2.1.7, "Resolving Host Names Using DNS Host Name Resolution" if you are resolving host names using DNS host name resolution.

## C.1.3 Creating an Asymmetric Standby Site with a Different Database Configuration

This section describes how to create an asymmetric standby site with a different database configuration than the production site. For example, the standby site could use a single instance database instead of the Real Application Clusters database used

at the production site. In Figure C–1, which shows an asymmetric standby site that can be set up for the EDG deployment production site in Figure 2–1, the RAC databases INFRADBHOST1 and INFRADBHOST2 and CUSTDBHOST1 and CUSTDBHOST2 at the production site are replaced by single instance databases INFRADBHOST1 and CUSTDBHOST1 at the standby site, respectively.

Oracle Data Guard is used to provide disaster protection and disaster recovery for the Oracle databases in an Oracle Application Server Disaster Recovery topology.

For general information on using Oracle Data Guard for databases in an Oracle Application Server Disaster Recovery topology, see Appendix A, "Using Databases in the OracleAS Disaster Recovery Solution."

For specific information on configuring Oracle Data Guard when the primary database is a RAC database and the standby database is a single instance database, see the appendix that describes Oracle Data Guard and Real Application Clusters in *Oracle Data Guard Concepts and Administration*.

# D

# Using Peer to Peer File Copy for Testing

As an alternative to using disk replication technology for disaster protection and disaster recovery of Oracle Application Server middle tier components, you can use peer to peer file copy mechanisms in test environments to replicate middle tier file system data from a production site host to a standby site peer host in an Oracle Application Server Disaster Recovery topology. An example of a peer to peer file copy mechanism is rsync (an open source utility for UNIX systems).

This appendix describes how to use rsync instead of disk replication in your Oracle Application Server Disaster Recovery topology. This appendix discusses rsync in the context of symmetric topologies. For more information about symmetric topologies, refer to Appendix C, "Creating an Asymmetric Topology." The information provided for rsync in this appendix also applies to other peer to peer file copy mechanisms.

Before you read this appendix, read the rest of this manual to ensure that you are familiar with how to use disk replication and Oracle Data Guard in an Oracle Application Server Disaster Recovery topology. There are many similarities between using disk replication and rsync for disaster protection and disaster recovery of your Oracle Application Server components.

---

**Note:** You can use rsync instead of disk replication technology to replicate middle tier file system data from the production site to the standby site. However, be aware that the following beneficial disk replication features are not available when you use rsync:

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

  With rsync, replicated production site data overwrites the standby site data, and you cannot roll back a replication.

- With disk replication, the volume you set up for each host cluster in the shared storage systems ensures data consistency for that host cluster across the production site's shared storage system and the standby site's shared storage system.

  With rsync, data consistency is not guaranteed.

Because of these deficiencies in comparison to disk replication, rsync is not supported for disaster recovery use in actual production environments.

---

# D.1 Using rsync and Oracle Data Guard for Oracle Application Server Disaster Recovery Topologies

These two basic principles apply when you use rsync and Oracle Data Guard to provide disaster protection and disaster recovery for your Oracle Application Server Disaster Recovery topology:

1. Use rsync for disaster protection of your Oracle Application Server middle tier components.
2. Use Oracle Data Guard for disaster protection of Oracle databases that are used in your Oracle Application Server topology. Appendix A, "Using Databases in the OracleAS Disaster Recovery Solution" describes how to set up Oracle Data Guard to provide disaster recovery for Oracle database.

## D.1.1 Using rsync for Oracle Application Server Middle Tier Components

Follow these steps to use rsync to provide disaster protection and disaster recovery for your Oracle Application Server middle tier components:

1. Set up rsync to enable replication of files from a production site host to its standby site peer host. See the rsync man page for instructions on installing and setting up rsync, and for syntax and usage information. Information about rsync is also available at http://rsync.samba.org.
2. For each production site host on which one or more Oracle Application Server 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 Application Server home directory and subdirectories, and all the files in them.
   - The Oracle Central Inventory directory and files for the host, which includes the Oracle Universal Installer entries for the Oracle Application Server installations.
   - If applicable, the Oracle Application Server static HTML pages directory for the Oracle HTTP Server installations on the host.
   - If applicable, the .fmb and .fmx deployment artifact files created by Oracle Forms on the host, and the .rdf deployment artifact files created by Oracle Reports on the host.

   **Note:** Run rsync as root. If you want rsync to work without prompting users for a password, set up SSH keys between the production site host and standby site host, so that SSH does not prompt for a password.
3. Set up scheduled jobs, for example, cron jobs, for the production site hosts for which you set up rsync in the previous step. These scheduled jobs enable rsync to automatically perform replication of these files from the production site hosts to the standby site hosts on a regular interval. An interval of once a day is recommended for a production site where the Oracle Application Server configuration does not change very often.
4. Whenever a change is made to the configuration of an Oracle Application Server middle tier configuration on a production site host (for example, when a new

application is deployed), you should perform a manual synchronization of that host with its standby site peer host using rsync.

5. Whenever you perform a manual rsync synchronization of an Oracle Application Server middle tier instance on a production site host to the peer standby site host, you should also manually force a synchronization of any associated database repository for the production site's Oracle Application Server instance to the standby site using Oracle Data Guard. See Section A.5, "Manually Forcing Database Synchronization with Oracle Data Guard" for more information on manually forcing a synchronization of an Oracle database using Oracle Data Guard.

## D.1.2 Performing Failover and Switchover Operations

Follow these steps to perform a failover or switchover from the production site to the standby site when you are using rsync:

1. Shut down any processes still running on the production site (if applicable).
2. Stop the rsync jobs between the production site hosts and their standby site peer hosts.
3. Use Oracle Data Guard to fail over the production site databases to the standby site.
4. Your production site hosts may include Oracle Application Server instances from different Oracle Application Server releases, such as release 10.1.2.x and 10.1.3.x. For any OC4J instances from Oracle Application Server 10.1.3.x releases, download and run the script provided with Oracle*MetaLink* patch 6785728 in the Oracle home directories for the 10.1.3.x OC4J instances to remove persistent OC4J lock files. The URL for Oracle*MetaLink* is:

   `https://metalink.oracle.com`

   **Note:** Run the script in the Oracle homes for 10.1.3.x OC4J instances only.

5. On the standby site, manually start the processes for the Oracle Application Server instances.
6. Route all user requests to the standby site (using a global DNS push or by updating the global load balancer).
7. 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).

After these steps have been performed, the former standby site has become the currrent production site. At this point, any issues that caused the original production site to become unavailable can be examined. Then work can begin on resolving those issues so that the original production site can be used again at some point in the future as either the production site or standby site, if desired.

To use the original production site as the current standby site, you need to reestablish the rsync replications between the two sites, but configure the replications so that they go in the opposite direction (from the current production site to the current standby site).

To use the original production site as the new production site, you perform the steps above again, but configure the rsync replications to go in the original direction (from the original production site to the original standby site).

# E

# Disaster Recovery for Collocated Infrastructure Deployments

In Oracle Application Server 10.1.x releases, one of the installation options was a collocated Oracle Identity Management and OracleAS Metadata Repository Infrastructure installation. Choosing this installation type caused Oracle Identity Management and the OracleAS Metadata Repository to be installed in the same Oracle home directory.

Follow the steps in this appendix to set up Disaster Recovery for an Infrastructure deployment where Oracle Identity Management and OracleAS Metadata Repository are installed (collocated) in the same Oracle home.

## E.1 Setting Up Disaster Recovery for Infrastructure Deployments with Collocated Identity Management and Metadata Repository

The steps in this section are for an Identity Management deployment which has Oracle Internet Directory and the OracleAS Metadata Repository collocated in one Oracle home and Single Sign-On Server installed in another Oracle home.

The steps below assume that the proper host names, ports, and other configuration described in Chapter 2, "Implementing the Solution" have been set up for the production site and standby site.

Follow these steps to set up Disaster Recovery for this deployment:

1. Use a remote copy mechanism such as rcp to copy the Oracle Internet Directory Oracle home from the production site host to the peer standby site host. A variation of this is to not include the data files in the copy operation to reduce the copy time. See the step below about RMAN.
2. Perform an RMAN backup of the production site Infrastructure database (which includes Oracle Internet Directory and which is in the Oracle Internet Directory Oracle home) and apply it to the standby site Infrastructure database. This is a one time operation.
3. Set up Oracle Data Guard for the Infrastructure (Oracle Internet Directory) database. See Appendix A, "Using Databases in the OracleAS Disaster Recovery Solution" for more information about setting up Oracle Data Guard.
4. Exclude the Infrastructure (Oracle Internet Directory) Oracle home from any future peer to peer file copy operations. Oracle Data Guard is synchronizing this database for the production and standby sites.
5. Use disk replication for the other Oracle homes (for this example, the Single Sign-On Server Oracle home) on a periodic basis.

# F

# Wide Area DNS Operations

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

---

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

`https://metalink.oracle.com`

---

The following topics are described in this section:

- Using a Global Load Balancer
- Manually Changing DNS Names

## F.1 Using a Global Load Balancer

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

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

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

## F.2 Manually Changing DNS Names

This method of DNS switchover involves the manual change of the name-to-IP mapping that is originally mapped to the IP address of the production site's load balancer. The mapping is changed to map to the IP address of the standby site's load balancer. Follow these instructions to perform the switchover:

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

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

# G

# Troubleshooting Disaster Recovery

This appendix describes common situations that you might encounter when deploying and managing Oracle Application Server in Disaster Recovery topologies, and explains the steps for addressing them. It contains the following topics:

- Section G.1, "Troubleshooting OracleAS Disaster Recovery Topologies"
- Section G.2, "Need More Help?"

## G.1 Troubleshooting OracleAS Disaster Recovery Topologies

This section describes common situations and steps to perform in OracleAS Disaster Recovery configurations. It contains the following topics:

- Section G.1.1, "Heartbeat Failure After Failover in Alert Logs"
- Section G.1.2, "Recommended Method of Patching an Oracle Application Server Disaster Recovery Site"
- Section G.1.3, "Changing the LockFile Directive for 10.1.2.x and 10.1.3.x Oracle HTTP Server Instances"
- Section G.1.4, "Use the DEFAULT_DMS_DIR Environment Variable for Oracle HTTP Server 10.1.3.x Instances"

### G.1.1 Heartbeat Failure After Failover in Alert Logs

A warning appears in the alert logs of the database after a failover scenario.

#### Summary

The following warning appears in the alert logs of the database after a failover scenario, where the new production site database fails to tnsping its remote database instance.

```
Errors in file c:\oracle\product\10.2.0\admin\orcl\udump\orcl1_rfs_1816.trc:
ORA-16009: remote archive log destination must be a STANDBY database
.
Fri Sep 08 09:11:13 2006
Errors in file c:\oracle\product\10.2.0\admin\orcl\bdump\orcl1_arc1_496.trc:
ORA-16009: remote archive log destination must be a STANDBY database
.
Fri Sep 08 09:11:13 2006
PING[ARC1]: Heartbeat failed to connect to standby 'orcl1_remote1'. Error is
16009.
Fri Sep 08 09:11:50 2006
Redo Shipping Client Connected as PUBLIC
-- Connected User is Valid
```

```
RFS[67]: Assigned to RFS process 628
RFS[67]: Database mount ID mismatch [0x4342404d:0x4341ffb0]
Fri Sep 08 09:11:50 2006
Errors in file c:\oracle\product\10.2.0\admin\orcl\udump\orcl1_rfs_628.trc:
ORA-16009: remote archive log destination must be a STANDBY database
.
Redo Shipping Client Connected as PUBLIC
-- Connected User is Valid
RFS[68]: Assigned to RFS process 2488
RFS[68]: Database mount ID mismatch [0x4342404d:0x4341ffb0]
Fri Sep 08 09:12:05 2006
Errors in file c:\oracle\product\10.2.0\admin\orcl\udump\orcl1_rfs_2488.trc:
ORA-16009: remote archive log destination must be a STANDBY database
.
Fri Sep 08 09:12:14 2006
Errors in file c:\oracle\product\10.2.0\admin\orcl\bdump\orcl1_arc1_496.trc:
ORA-16009: remote archive log destination must be a STANDBY database
```

#### Steps to Perform

To avoid these error messages in the alert logs, null the log_archive_dest_2 parameter using the following commands:

```
alter system set log_archive_dest_2='SERVICE=null LGWR ASYNC REOPEN=60';
alter system set log_archive_dest_state_2='defer';
```

## G.1.2 Recommended Method of Patching an Oracle Application Server Disaster Recovery Site

This section describes how to apply an Oracle Application Server patch set to upgrade the Oracle homes that participate in an Oracle Application Server Disaster Recovery site.

#### Summary

You are unsure how to apply an Oracle Application Server patch set to upgrade the Oracle homes in your Oracle Application Server Disaster Recovery production site.

#### Steps to Perform

The list in this section describes the steps for applying a patch set to upgrade the Oracle Application Server homes in an Oracle Application Server Disaster Recovery production site.

The following steps assume that the Oracle Central Inventory for any Oracle Application Server instance that you are patching is located on the production site shared storage, so that the Oracle Central Inventory for the patched instance can b e replicated to the standby site.

Use the following procedure to upgrade Oracle Application Server patch versions:

1. Perform a backup of the production site to ensure that the starting state is secured.
2. Apply the patch set to upgrade the production site instances.
3. After applying the patch set, manually force a synchronization of the production site shared storage and standby site shared storage. This replicates the production site's patched instance and Oracle Central Inventory in the standby site's shared storage.

4. After applying the patch set, use Oracle Data Guard to manually force a synchronization of the Oracle databases at the production site and standby sites. Some Oracle Application Server patch sets may make updates to repositories, so this step ensures that any changes made to production site databases are synchronized to the standby site databases.
5. The upgrade is now complete. Your Disaster Recovery topology is ready to resume processing.

---

**Note:** Patches need to be applied only at the production site for an Oracle Application Server Disaster Recovery topology. If a patch is for an Oracle Application Server instance or for the Oracle Central Inventory, the patch will be copied when the production site shared storage is replicated to the standby site shared storage. A synchronization operation should be performed when a patch is installed at the production site.

Similarly, if a patch is installed for a production site database, Oracle Data Guard will copy the patch to the standby database at the standby site when a synchronization is performed.

---

## G.1.3 Changing the LockFile Directive for 10.1.2.x and 10.1.3.x Oracle HTTP Server Instances

This manual directs you to install Oracle Application Server instances (including Oracle HTTP Server instances) on shared storage. When the installation of an Oracle home for an Oracle HTTP Server 10.1.2.x or 10.1.3.x instance is performed on shared storage (for example, NAS storage, NFS storage, or SAN storage) that does not provide reliable file locking, Oracle HTTP Server may experience performance problems.

### Summary

Some shared storage systems do not provide the reliable file locking that Oracle HTTP Server requires. In these cases, the LockFile directive in the `httpd.conf` file must be changed to point at a local file system. See the *Oracle HTTP Server Administrator's Guide* for more information about the LockFile directive.

### Steps to Perform

If any 10.1.2.x or 10.1.3.x Oracle HTTP Server instance is installed on shared storage and is experiencing performance problems, perform these steps for the Oracle HTTP Server instance to point the LockFile directive at a local file system:

1. By default, the LockFile directive is commented out. When the Oracle home for an Oracle HTTP Server installation was performed on shared storage (for example, NAS, NFS, or SAN storage), the directive will be in this format:

   ```
   #LockFile /<nfs-mounted>/Apache/Apache/logs/httpd.lock
   ```

2. Edit the `$ORACLE_HOME/Apache/Apache/conf/httpd.conf` file using the appropriate method for the version of Oracle Application Server you are using.
3. Uncomment the LockFile directive and change it to point at a local file system:

   ```
   LockFile /<local_disk>/<path>/httpd.lock
   ```

4. Restart the Oracle HTTP Server.

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

### G.1.4 Use the DEFAULT_DMS_DIR Environment Variable for Oracle HTTP Server 10.1.3.x Instances

This manual directs you to install Oracle Application Server instances (including Oracle HTTP Server instances) on shared storage. When the installation of an Oracle home for an Oracle HTTP Server 10.1.3.x instance is performed on shared storage (for example, NAS storage, NFS storage, or SAN storage) that does not provide reliable file locking, HTTP Server performance may be poor.

#### Summary

In Oracle Application Server 10.1.3.x, some new temporary files are created when Oracle HTTP Server is running. These files are created by DMS and are stored in the `$ORACLE_HOME/Apache/Apache/logs` directory.

When the Oracle home for an Oracle Application Server 10.1.3.x instance is installed on shared storage, the `dms_metrics.lock` file and other `*dms*` lock files created under the `$ORACLE_HOME/Apache/Apache/logs` directory must be moved to a local file system, otherwise Oracle HTTP Server performance may be poor.

#### Steps to Perform

If any 10.1.2.x or 10.1.3.x Oracle HTTP Server instance is installed on shared storage and is experiencing performance problems, follow these steps to define the DEFAULT_DMS_DIR environment variable and point it to a local file system directory in which to store the `*dms*` files:

1. Download and apply the Application Server 10.1.3.3.0 patch set to the Oracle home for the Oracle HTTP Server 10.1.3.x instance.
2. Define the DEFAULT_DMS_DIR environment variable in `$ORACLE_HOME/opmn/conf/opmn.xml` and point it to the local file system directory in which you want to store the `*dms*` files, for example:

   ```
   <environment>
   <variable id="TMP" value="/tmp"/>
   <variable id="DEFAULT_DMS_DIR" value="/tmp"/>
   </environment>
   ```

3. Stop all the Oracle Application Server processes and start them up again:

   ```
   opmnctl stopall
   opmnctl startall
   ```

4. Verify that the `*dms*` files have been created under the `$DEFAULT_DMS_DIR` directory and not under the `$ORACLE_HOME/Apache/Apache/logs` directory.

## G.2 Need More Help?

In case the information in the previous section is not sufficient, you can find more solutions on Oracle *MetaLink*, https://metalink.oracle.com. If you do not find a solution for your problem, log a service request.

**See Also:**

- *Oracle Application Server Release Notes*, available on the Oracle Technology Network: http://www.oracle.com/technology/documentation/index.html

# Index

## E



## F



## G



## H



## I



## L



## M

## N



## O



## P



## S

## T



## V