Oracle® Fusion Middleware

Continuous Availability for Oracle WebLogic Server

12*c* (12.2.1)

E65845-02

December 2015

Documentation that describes the features and benefits of Oracle Continuous Availability, the architectures it supports, and how you can how you can use the continuous availability features in the supported architectures.

ORACLE®

Oracle Fusion Middleware Continuous Availability for Oracle WebLogic Server, 12*c* (12.2.1)

E65845-02

Contents

Preface

Oracle WebLogic Server Continuous Availability provides an integrated solution for building maximum availability architectures that span data centers across distributed geographical locations. This guide describes the features and benefits of Oracle WebLogic Server Continuous Availability, the architectures it supports, and how you can use the continuous availability features in the supported architectures.

Audience

This document is intended for administrators, developers, and others whose role is to configure and manage Oracle WebLogic Server with continuous availability requirements.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc`.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

Related Documents

For more information, see the following Fusion Middleware documents:

- *Administering Zero Downtime Patching Workflows*
- *Using WebLogic Server Multitenant*
- *Developing JTA Applications for Oracle WebLogic Server*
- *Administering JDBC Data Sources for Oracle WebLogic Server*
- *Administering Oracle Coherence*

Conventions

The following text conventions are used in this document:

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

1

Introduction to WebLogic Continuous Availability

Oracle WebLogic Server Continuous Availability provides an integrated solution for building maximum availability architectures that span data centers across distributed geographical locations. This chapter provides an introduction to continuous availability and its key features.

Topics include:

- What is Continuous Availability?
- Continuous Availability Key Features
- Value-Added WebLogic Server and Coherence High Availability Features
- Disaster Recovery for the Oracle Database
- Continuous Availability Terminology

1.1 What is Continuous Availability?

Continuous availability is the ability of a system to provide maximum availability by employing both high availability and disaster recovery solutions to ensure that applications are available when they are needed. Typically, a high availability solution provides redundancy in one data center. Disaster recovery solutions provide 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.

Oracle WebLogic Server Continuous Availability provides an integrated solution for building maximum availability architectures (MAA) that span data centers across distributed geographical locations. Integrated components include Oracle WebLogic Server, Oracle Coherence, Oracle Traffic Director, and Oracle SiteGuard. The major benefits of this integrated solution are faster failover or switchover, increased overall application availability, data integrity, reduced human error and risk, recovery of work, and local access of real-time data.

The key features of Continuous Availability are described in the next section. Common terminology used in this document is described in "Continuous Availability Terminology" on page 1-6.

1.2 Continuous Availability Key Features

Continuous Availability provides maximum availability, reliability and application stability during planned upgrades or unexpected failures. It builds on the existing

high availability features in Oracle WebLogic Server, Oracle Coherence, and Oracle Fusion Middleware, and supports the key features described in the following sections.

- Automated Cross-Domain Transaction Recovery
- WebLogic Server Zero Downtime Patching
- WebLogic Server Multitenant Live Resource Group Migration
- Coherence Federated Caching
- Coherence GoldenGate HotCache
- Oracle Traffic Director
- Oracle Site Guard

1.2.1 Automated Cross-Domain Transaction Recovery

Automated cross-domain transaction recovery provides automatic recovery of XA transactions across an entire domain, or across an entire site with servers running in a different domain or at a different site. In active/active architectures, transactions can be recovered when an entire domain or site fails by having an active server running in a different domain either collocated at the same site or at a different site. In active/passive architectures, the server at the passive (standby) site at a different location can be started when the production site is no longer available. For more information, see "Transaction Recovery Spanning Multiple Sites or Data Centers" in *Developing JTA Applications for Oracle WebLogic Server*.

Automated Cross Domain Transaction Recovery also takes advantage of the WebLogic Server high availability features described in "WebLogic Server High Availability Transaction and Data Source Features" on page 1-4.

1.2.2 WebLogic Server Zero Downtime Patching

WebLogic Server Zero Downtime Patching provides an automated mechanism to orchestrate the rollout of patches while avoiding downtime or loss of sessions. It reduces risks and downtime of mission-critical applications that require availability and predictability while applying patches.

Using workflows that you define, you can patch or update any number of nodes in a domain with little or no manual intervention. Changes are rolled out to one node at a time, allowing a load balancer such as Oracle Traffic Director to redirect incoming traffic to the remaining nodes until the node has been updated.

For more information, see *Administering Zero Downtime Patching Workflows*.

1.2.3 WebLogic Server Multitenant Live Resource Group Migration

In WebLogic Server Multitenant environments, you can migrate partition resource groups that are running from one cluster/server to another within a domain without impacting the application users. A key benefit of migrating the resource groups is that it eliminates application downtime for planned events.

Resource groups are a collection of (typically) related deployable resources, such as Java EE applications and the data sources, JMS artifacts, and other resources that the applications use. When you migrate a resource group, you change the virtual target used by the resource group from one physical target (cluster/server) to another. After migration, the virtual target will point to the new physical target (cluster/server).

For more information about resource groups and migration, see the following topics in *Using WebLogic Server Multitenant*:

- Configuring Resource Groups
- Migrating a Resource Group: Main Steps and WLST Example

1.2.4 Coherence Federated Caching

The Oracle Coherence federated caching feature replicates cache data asynchronously across multiple geographically distributed clusters. It supports multiple replication topologies including active-active, active-passive, hub-spoke, and central-replication. Cached data is replicated across clusters to provide redundancy, off-site backup, and multiple points of access for application users in different geographical locations.

Federated caching supports multiple replication topologies. These include:

- Active-passive—Replicates data from an active cluster to a passive cluster. The passive site supports read-only operations and off-site backup.
- Active-active—Replicates data between active clusters. Data that is put into one active cluster, is replicated at the other active clusters. Applications in different sites have access to a local cluster instance.
- Hub and spoke—Replicates data from a single hub cluster to multiple spoke clusters. The hub cluster can only send data and spoke clusters can only receive data. This topology requires multiple geographically dispersed copies of a cluster. Each spoke cluster can be used by local applications to perform read-only operations.

For more information about federated caching, see "Replicating Caches Across Clusters" in *Administering Oracle Coherence*.

1.2.5 Coherence GoldenGate HotCache

The Oracle Coherence GoldenGate HotCache feature detects and reflects database changes in cache in real time. Third-party updates to the database can cause Coherence applications to work with data that can be stale and out-of-date. Coherence GoldenGate HotCache solves this problem by monitoring the database and pushing any changes into the Coherence cache in real time. It employs an efficient push model that processes only stale data. Low latency is assured because the data is pushed when the change occurs in the database.

In Maximum Availability Architectures, when the database is replicated to a secondary site during failover, the database changes are reflected to the cache using GoldenGate HotCache.

For more information, see "Integrating with Oracle Coherence GoldenGate HotCache" in *Integrating Oracle Coherence*.

1.2.6 Oracle Traffic Director

Oracle Traffic Director is a fast, reliable, and scalable software load balancer that routes HTTP, HTTPS, and TCP traffic to application servers and web servers on the network. It distributes the requests that it receives from clients to available servers based on the specified load-balancing method, routes the requests based on specified rules, caches frequently accessed data, prioritizes traffic, and controls the quality of service.

The architecture of Oracle Traffic Director enables it to handle large volumes of application traffic with low latency. For high availability, you can set up pairs of Oracle

Traffic Director instances for either active-passive or active-active failover. As the volume of traffic to your network grows, you can easily scale the environment by reconfiguring Oracle Traffic Director with additional back-end servers to which it can route requests.

For more information, see *Administering Oracle Traffic Director*.

1.2.7 Oracle Site Guard

Oracle Site Guard, a component of Oracle Enterprise Manager Cloud Control, is a disaster-recovery solution that enables administrators to automate complete site switchover or failover, thereby minimizing downtime for enterprise deployments. Because Oracle Site Guard operates at the site level, it eliminates the need to tediously perform manual disaster recovery for individual site components like applications, middleware, databases, and so on. The traffic of an entire production site can be redirected to a standby site in a single operation.

Administrators do not require any special skills or domain expertise in areas like databases, applications, and storage replication. Oracle Site Guard can continuously monitor disaster-recovery readiness and it can do this without disrupting the production site.

You can manage an Oracle Site Guard configuration by using either the Enterprise Manager Command-Line Interface (EMCLI), or a compatible version of Oracle Enterprise Manager Cloud Control (Cloud Control).

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

1.3 Value-Added WebLogic Server and Coherence High Availability Features

In addition to the features described in "Continuous Availability Key Features" on page 1-1, Oracle Continuous Availability also takes advantage of the high availability features provided with WebLogic Server and Coherence as described in the following sections.

1.3.1 WebLogic Server High Availability Transaction and Data Source Features

The following high availability transaction and data source features can be used with Automated Cross Domain Transaction Recovery for Continuous Availability:

- Active GridLink data sources that use Fast Connection Failover to provide rapid failure detection of Oracle RAC nodes, and failover to remaining nodes for continuous connectivity. For more information, see "Using Active GridLink Data Sources" in *Administering JDBC Data Sources for Oracle WebLogic Server*.
- Transaction logs in the database (JDBC Tlogs) that store information about committed transactions coordinated by the server that may not have been completed. WebLogic Server uses the TLogs when recovering from system crashes or network failures. For more information, see "Using Transaction Log Files to Recover Transactions" in *Developing JTA Applications for Oracle WebLogic Server*.
- No transaction TLog writes (No TLOG) where you eliminate writes of the transaction checkpoints to the TLog store. For more information, see "XA Transactions without Transaction TLog Write" in *Developing JTA Applications for Oracle WebLogic Server*.
- Logging Last Resource (LLR) transaction optimization which is a performance enhancement option that enables one non-XA resource to participate in a global

transaction with the same ACID (atomicity, consistency, isolation, durability) guarantee as XA. For more information, see "Logging Last Resource Transaction Optimization" in *Developing JTA Applications for Oracle WebLogic Server*.

These features work with Oracle Data Guard which replicates databases to make transaction logs needed for recovery to be highly available. For more information about Oracle Data Guard, see *Data Guard Concepts and Administration*.

1.3.2 Coherence Persistence and Clusters

Coherence persistence is a set of tools and technologies that manage the persistence and recovery of Coherence distributed caches. Cached data is persisted so that it can be quickly recovered after a catastrophic failure or after a cluster restart due to planned maintenance. Persistence and federated caching can be used together as required. For more information about Coherence persistence, see "Persisting Caches" in *Administering Oracle Coherence*.

Coherence clusters consist of multiple Coherence server instances that distribute data in-memory to increase application scalability, availability, and performance. Application data is automatically and transparently distributed and backed-up across cluster members. For more information about Coherence clusters, see "Configuring and Managing Coherence Clusters" in *Administering Clusters for Oracle WebLogic Server*.

1.4 Disaster Recovery for the Oracle Database

Oracle Continuous Availability also takes advantage of existing database failover and switchover capabilities using Oracle Data Guard, Oracle Data Guard Broker, and Oracle Clusterware. All of these components contribute to managing and orchestrating the failover and switchover of the Oracle Database as follows:

- Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. It provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions. Oracle Data Guard maintains these standby databases as transactionally consistent copies of the primary database. If the primary database becomes unavailable because of a planned or an unplanned outage, Oracle Data Guard enables you to switch any standby database to the production role, thus minimizing the downtime associated with the outage.
- Oracle Data Guard broker logically groups these primary and standby databases into a broker configuration that allows the broker to manage and monitor them together as an integrated unit. It sends notifications to WebLogic Active GridLink which then makes new connections to the database in the failover site, and coordinates with Oracle Clusterware to fail over role-based services.
- Oracle Clusterware manages the availability of instances of an Oracle RAC database. It works to rapidly recover failed instances to keep the primary database available. If Oracle Clusterware cannot recover a failed instance, the broker continues to run automatically with one less instance. If the last instance of the primary database fails, the broker provides a way to fail over to a specified standby database. If the last instance of the primary database fails, and fast-start failover is enabled, the broker can continue to provide high availability by automatically failing over to a pre-determined standby database.

Oracle Site Guard uses Data Guard Broker to perform failover/switchover of the Databases. The integration of all these products makes the failover of the database fast and automatic.

1.5 Continuous Availability Terminology

The following list describes the common terminology that applies to continuous availability:

- Active-active—An active-active solution deploys two or more active servers to improve scalability and provide high availability. In active-active deployments, all instances handle requests concurrently. When an entire domain or site fails, transactions can be recovered by an active server in a different domain either collocated in the same site or on a different site.
- Active-passive—Active-passive solutions involve setting up and pairing a standby site at a geographically different location with an active (production) site. The standby site may have equal or fewer services and resources compared to the production site. Application data, metadata, configuration data, and security data are replicated periodically to the standby site. The standby site is normally in a passive mode; it is started when the production site is not available. This model is usually adopted when the two sites are connected over a WAN and network latency does not allow clustering across the two sites.
- WebLogic Server cluster—A collection of WebLogic Server server instances running simultaneously and working together to provide increased scalability and reliability. In a cluster, most resources and services are deployed identically to each Managed Server, enabling failover and load balancing.
- Coherence cluster—A collection of JVM processes, called Coherence servers, that run Coherence. A Coherence cluster consists of multiple Coherence server instances that distribute data in-memory to increase application scalability, availability, and performance. Application data is automatically and transparently distributed and backed-up across cluster members.
- Stretch cluster—A cluster in which nodes can span datacenters within a proximate geographical range, usually with guaranteed, relatively low-latency networking between the sites. Stretch clusters are also referred to as extended clusters.
- High availability—The ability of a system or device to be available when it is needed. A high availability architecture ensures that users can access a system without loss of service. Deploying a high availability system minimizes the time when the system is down, or unavailable, and maximizes the time when it is running, or available.
- 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.
- 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.
- 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).

2

Supported MAA Architectures

This chapter describes the supported maximum availability architecture (MAA) multi-datacenter solutions that can be used to provide continuous availability to protect an Oracle WebLogic Server system against downtime across multiple data centers. It also describes how the Continuous Availability features can be used with each architecture.

The supported MAA solutions include:

- Active-Active Application Tier with Active-Passive Database Tier
- Active-Passive Application Tier with Active-Passive Database Tier
- Active-Active Stretch Cluster with Active-Passive Database Tier

2.1 Active-Active Application Tier with Active-Passive Database Tier

The following figure shows a recommended continuous availability solution using an Active-Active application infrastructure tier with an Active-Passive database tier.

Figure 2–1 Topology for Active-Active Application infrastructure Tier with Active-Passive Database Tier

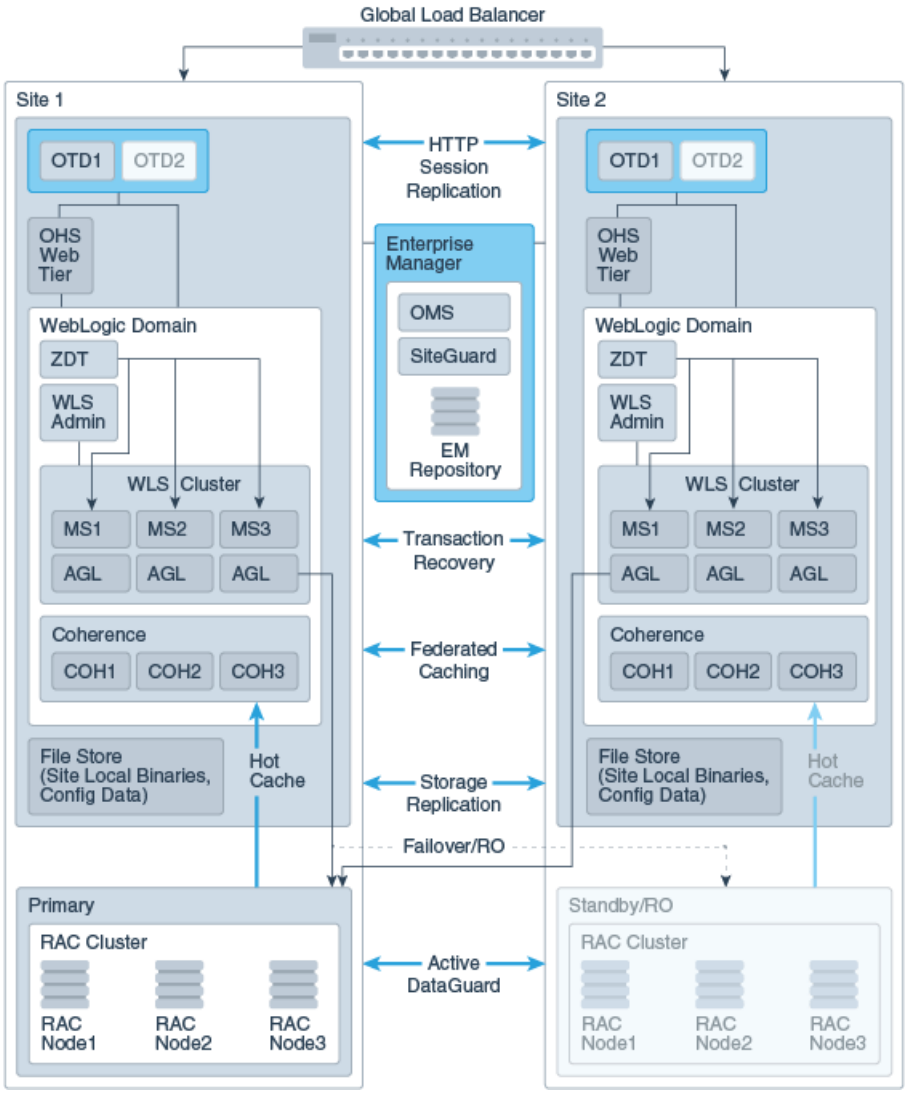

The key aspects of this sample topology include:

- Two separate WebLogic domains configured in two different data centers, Site 1 and Site 2. The domains at both sites are active. The domains include:
  - A collection of Managed Servers (MS1, MS2, and MS3) in a WebLogic Server cluster, managed by the WebLogic Server Admin Server in the domain. In this sample, Active Gridlink (AG) is being used to connect the Managed Servers to the primary database. (Although generic DataSource or MultiDataSource can be used, Active Gridlink is preferable because it offers high-availabity and improved performance). The Zero Downtime Patching (ZDT) arrows represent patching the Managed Servers in a rolling fashion.
  - A Coherence cluster (COH1, COH2, and COH3) managed by the WebLogic Server Admin Server in the domain.
- A global load balancer.
- WebLogic Server HTTP session replication across clusters.

- Two instances of Oracle Traffic Director (OTD) at each site, one active and one passive. OTD can balance requests to the web tier or to the WebLogic Server cluster.
- Oracle HTTP Server (OHS) Web Tier. (Optional component based on the type of environment.)
- A filestore for the configuration data, local binaries, logs, and so on that is replicated across the two sites using any replication technology.
- Oracle Site Guard, a component of Oracle Enterprise Manager Cloud Control, that orchestrates failover and switchover of sites.
- Two separate Oracle RAC database clusters in two different data centers. The primary active Oracle RAC database cluster is at Site 1. Site 2 contains an Oracle RAC database cluster in standby (passive) read-only mode. The clusters can contain transaction logs, JMS stores, and application data. Data is replicated using Oracle Active DataGuard. (Although Oracle recommends using Oracle RAC database clusters because they provide the best level of high availability, they are not required. A single database or multitenant database can also be used.)

This topology uses the Continuous Availability features as follows:

- Automated cross-domain transaction recovery—Because both domains are active, you can use the full capabilities of this feature as described in "Automated Cross-Domain Transaction Recovery" on page 1-2. In this architecture, transactions can be recovered automatically, without any manual intervention.
- WebLogic Zero Downtime Patching—Because both domains are active, you can orchestrate the roll out of updates separately on each site. See "WebLogic Server Zero Downtime Patching" on page 1-2.
- Coherence Federated Caching—You can use the full capabilities of this feature as described in "Coherence Federated Caching" on page 1-3. Cached data is always replicated between clusters. Applications in different sites have access to a local cluster instance.
- Coherence GoldenGate HotCache—Updates the Coherence cache in real-time for any updates that are made on the active database. See "Coherence GoldenGate HotCache" on page 1-3.
- Oracle Traffic Director—Adjusts traffic routing to application servers depending on server availability. You can achieve high-availability with Oracle Traffic Director by having a pair of instances configured at each site, either active-active or active-passive. See "Oracle Traffic Director" on page 1-3.
- Oracle Site Guard—Because only the database is in standby mode in this architecture, Oracle Site Guard will control database failover only. It does not apply to the application architecture tier because both domains are active. For more information, see "Oracle Site Guard" on page 1-4.

2.2 Active-Passive Application Tier with Active-Passive Database Tier

The following figure shows a recommended continuous availability topology using an Active-Passive application infrastructure tier with an Active-Passive database tier.

Figure 2–2 Topology for Active-Passive Application infrastructure Tier with Active-Passive Database Tier

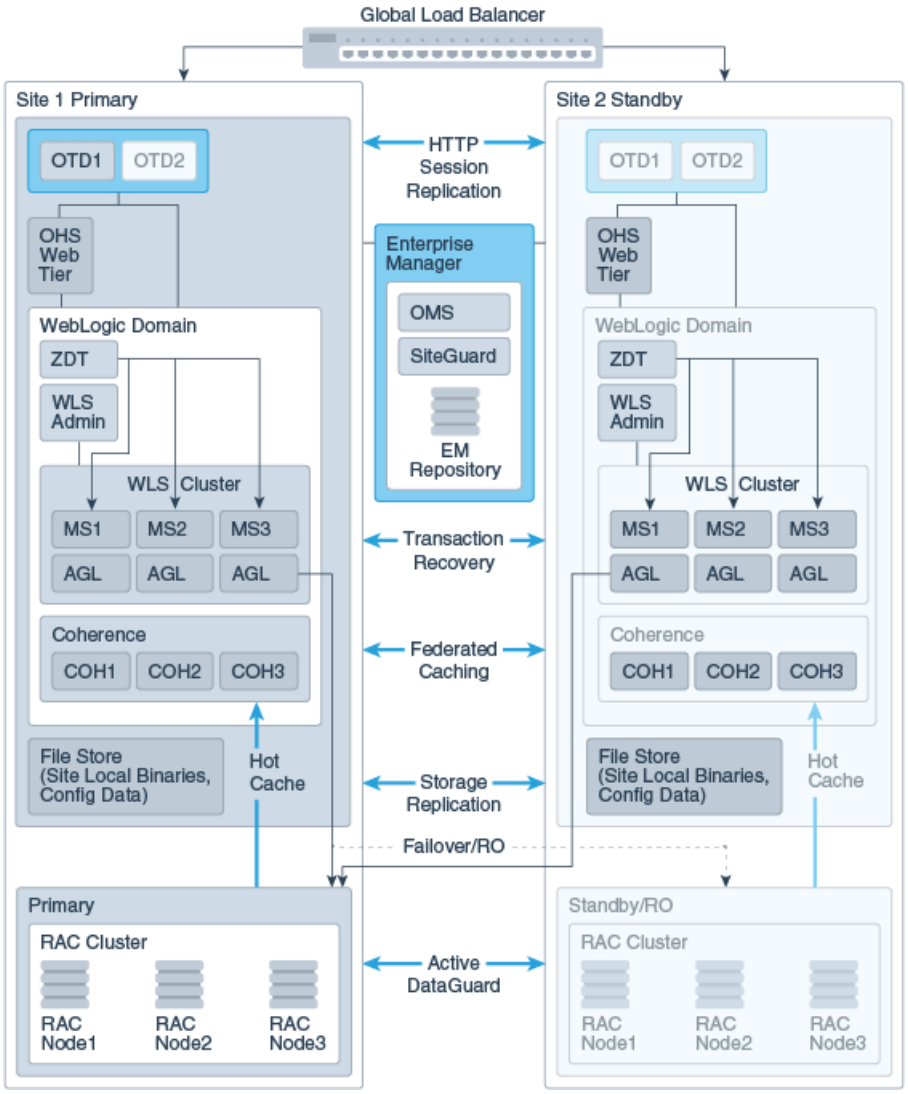

The key aspects of this topology include:

- Two separate WebLogic domains configured in two different data centers, Site 1 and Site 2. The domain at Site 1 is active and the domain at Site 2 is in standby (passive) mode.

  All active-passive domain pairs must be configured with symmetric topology; they must be identical and use the same domain configurations such as directory names and paths, port numbers, user accounts, load balancers and virtual server names, and the same versions of the software. Hostnames (not static IPs) must be used to specify the listen address of the Managed Servers.

  The domains include:

  – A collection of Managed Servers (MS1, MS2, and MS3) in a WebLogic Server cluster, managed by the WebLogic Server Admin Server in the domain. In this sample, Active Gridlink (AG) is being used to connect the Managed Servers to the primary database. (Although generic DataSource or MultiDataSource can be used, Active Gridlink is preferable because it offers high-availabity and

improved performance). The Zero Downtime Patching (ZDT) arrows represent patching the Managed Servers in a rolling fashion.

  - A Coherence cluster (COH1, COH2, and COH3) managed by the WebLogic Server Admin Server in the domain.

- A global load balancer.
- WebLogic Server HTTP session replication across clusters.
- Two instances of Oracle Traffic Director (OTD) at each site. OTD can balance requests to the web tier or to the WebLogic Server cluster. At Site 1, one instance is active and one passive. On Site 2 they are both on standby. When Site 2 becomes active, the OTD instances on that site will start routing the requests.
- Oracle HTTP Server (OHS) Web Tier. (Optional component based on the type of environment.)
- A filestore for the configuration data, local binaries, logs, and so on that is replicated across the two sites using any replication technology.
- Oracle Site Guard, a component of Oracle Enterprise Manager Cloud Control, that orchestrates failover and switchover of sites.
- Two separate Oracle RAC database clusters in two different data centers. The primary active Oracle RAC database cluster is at Site 1. Site 2 contains an Oracle RAC database cluster in standby (passive) read-only mode. The clusters can contain transaction logs, JMS stores, and application data. Data is replicated using Oracle Active DataGuard. (Although Oracle recommends using Oracle RAC database clusters because they provide the best level of high availability, they are not required. A single database or multitenant database can also be used.)

This architecture uses the Continuous Availability features as follows:

- Automated cross-domain transaction recovery—Because the domain at Site 2 is in standby mode, during failover you must first start the domain at Site 2. Once you do so, you can use the cross-domain transaction recovery features as described in "Automated Cross-Domain Transaction Recovery" on page 1-2.
- WebLogic Server Zero Downtime Patching—In this architecture, you can use Zero Downtime Patching feature on the active domain in Site 1 as described in "WebLogic Server Zero Downtime Patching" on page 1-2. Because the servers are not running in the standby (passive) domain at Site 2, you can use OPatch. When the servers become active they will point to the patched Oracle home. For more information about OPatch, see *Patching with OPatch*.
- Coherence Federated Caching—In this architecture, the passive site supports read-only operations and off-site backup. See "Coherence Federated Caching" on page 1-3.
- Coherence GoldenGate HotCache—In this architecture, updates on the active database at Site 1 update the Coherence cache in real time and the database updates are replicated to Site 2. When the data replication occurs on Site 2, GoldenGate HotCache updates the cache in real time. See "Coherence GoldenGate HotCache" on page 1-3.
- Oracle Traffic Director—In this architecture, OTD is in standby mode on the standby (passive) domain at Site 2. When Site 2 becomes active after failover, the OTD instance on the standby site is activated (by SiteGuard using a script that is run before or after failover) and the OTD instance will start routing requests to the recently started servers.

- Oracle Site Guard—In this architecture, when Site1 fails, SiteGuard will initiate the failover using scripts that specify what should occur and in what order. For example, it can start WebLogic Server, Coherence, web applications, and other site components in a specific order. You can execute the scripts either pre- or post-failover. SiteGuard integrates with Data Guard Broker to failover the database. See "Oracle Site Guard" on page 1-4.

2.3 Active-Active Stretch Cluster with Active-Passive Database Tier

The following figure shows a recommended continuous availability solution using an Active-Active stretch cluster application infrastructure tier with an Active-Passive database tier.

Figure 2–3 Topology for Active-Active Stretch Cluster Application Infrastructure Tier and Active-Passive Database Tier

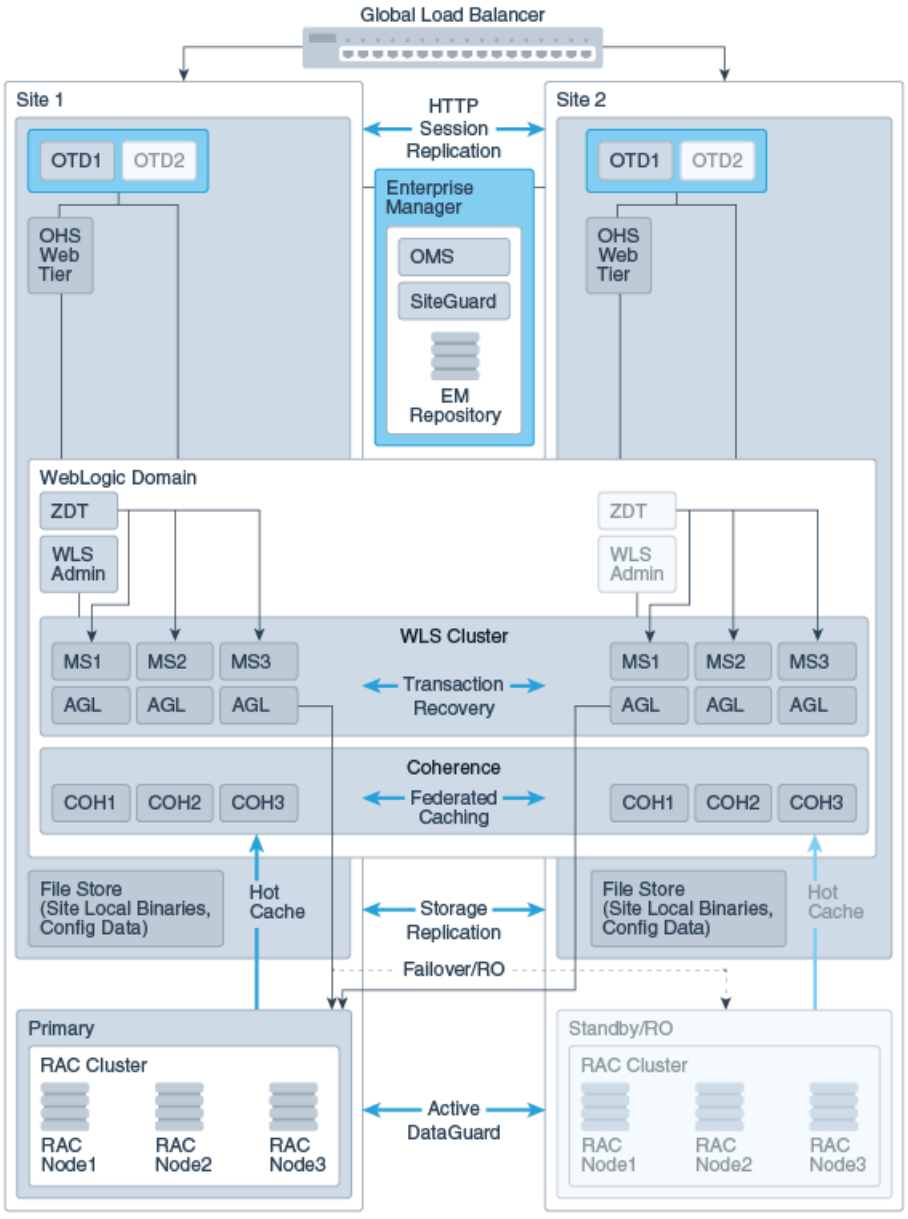

The key aspects of this topology include:

- WebLogic Server configured as a cluster that stretches across two different data centers, Site 1 and Site 2. All servers in the cluster are active.
- The domain includes:
  - A WebLogic Server cluster that consists of a group of Managed Servers (MS1, MS2, and MS3) at Site 1 and another group of Managed Servers (MS4, MS5, and MS6) at Site 2. The Managed Servers are managed by the WebLogic Server Admin Server at Site 1. In this sample, Active Gridlink (AG) is being used to connect the Managed Servers to the primary database. (Although generic DataSource or MultiDataSource can be used, Active Gridlink is preferable because it offers high-availabity and improved performance). The Zero Downtime Patching (ZDT) arrows represent patching the Managed Servers in a rolling fashion.
  - A Coherence cluster that consists of a group of Coherence instances (COH1, COH2, and COH3) at Site 1, and another group (COH4, COH5, and COH6) at Site 2, all managed by the WebLogic Server Admin Server at Site 1.
- A global load balancer.
- WebLogic Server HTTP session replication across clusters.
- Two instances of Oracle Traffic Director (OTD) at each site, one active and one passive. OTD can balance requests to the web tier or to the WebLogic Server cluster.
- Oracle HTTP Server (OHS) Web Tier. (Optional component based on the type of environment.)
- A filestore for the configuration data, local binaries, logs, and so on that is replicated across the two sites using any replication technology.
- Oracle Site Guard, a component of Oracle Enterprise Manager Cloud Control, that orchestrates failover and switchover of sites.
- Two separate Oracle RAC database clusters in two different data centers. The primary active Oracle RAC database cluster is at Site 1. Site 2 contains an Oracle RAC database cluster in standby (passive) read-only mode. The clusters can contain transaction logs, JMS stores, and application data. Data is replicated using Oracle Active DataGuard. (Although Oracle recommends using Oracle RAC database clusters because they provide the best level of high availability, they are not required. A single database or multitenant database can also be used.)

This architecture uses the Continuous Availability features as follows:

- Automated cross-domain transaction recovery—Because all of the servers are in the same cluster, you can use existing WebLogic Server high availability features, server and service migration, to recover transactions.
  - In whole server migration, a migratable server instance, and all of its services, is migrated to a different physical machine upon failure. For more information, see "Whole Server Migration" in *Administering Clusters for Oracle WebLogic Server*.
  - In service migration, in the event of failure, services are moved to a different server instance within the cluster. For more information, see "Service Migration" in *Administering Clusters for Oracle WebLogic Server*.

- WebLogic Zero Downtime Patching—Because all of the servers in the cluster are active, you can use the full capabilities of this feature as described in "WebLogic Server Zero Downtime Patching" on page 1-2.
- Coherence Federated Caching—You can use the full capabilities of this feature as described in "Coherence Federated Caching" on page 1-3.
- Coherence GoldenGate HotCache—You can use the full capabilities of this feature as described in "Coherence GoldenGate HotCache" on page 1-3.
- Oracle Traffic Director—Adjusts traffic routing to application servers depending on server availability. You can use the full capabilities of this feature as described in "Oracle Traffic Director" on page 1-3.
- Oracle Site Guard—Because only the database is in standby mode in this architecture, Oracle Site Guard will control database failover only. It does not apply to the application architecture tier because all servers in the cluster are active. For more information, see "Oracle Site Guard" on page 1-4.