Oracle Maximum
Availability Architecture

Oracle Commerce MAA

With a Case Study on Exadata and Exalogic

ORACLE WHITE PAPER | APRIL 2015

Table of Contents

Executive Overview

Oracle Maximum Availability Architecture (MAA) is Oracle's best practices blueprint based on proven Oracle high availability technologies and recommendations. The goal of MAA is to achieve the optimal high availability architecture at the lowest cost and complexity. Papers are published on the Oracle Technology Network (OTN) - http://www.oracle.com/goto/maa.

This paper describes the architecture along with installation, configuration, and operational best practices for deploying Oracle Commerce with MAA best practices. For the purpose of this paper, Oracle Commerce is comprised of Oracle Commerce Platform, Oracle Guided Search, and Oracle Commerce Merchandising.

Oracle Commerce MAA is implemented in two different MAA configurations: 1) fully Active/Passive and 2) Active/Active application with Active/Passive databases (also known as Active/Active/Passive). Each is tested to validate MAA best practices and to measure and observe application impact in various outage scenarios. The results are presented in this paper. A third configuration exists, which is fully Active/Active. Due to the many business rules required for this type of configuration to be constructed, it is not in the scope of this document and not formally an MAA supported configuration. It will be touched upon briefly to show that the option exists.

Introduction to Engineered Systems

Oracle's engineered systems combine best-of-breed hardware and software components with game changing technical innovations. Designed, engineered, and tested to work best together, Oracle's engineered systems can power the cloud or streamline data center operations to make traditional deployments even more efficient. The components of Oracle's engineered systems are preassembled for targeted functionality and then, as a complete system, optimized for extreme performance. By taking the guesswork out of these highly available, purpose-built solutions, Oracle delivers a solution that is integrated across every layer of the technology stack—a simplicity that translates into less risk and lower costs for your business. Only Oracle can innovate and optimize at every layer of the stack to simplify data center operations, drive down costs, and accelerate business innovation.

Oracle Exalogic

Oracle Exalogic is an engineered system on which enterprises deploy Oracle business applications, Oracle Fusion Middleware, or third-party software products. Exalogic comes pre-built with compute nodes, memory, flash storage, and centralized storage, all connected using InfiniBand in a highly available architecture with fault tolerance and zero downtime maintenance. Exalogic is optimized for Oracle Commerce, and Oracle Commerce is optimized for Exalogic through enhancements such as full InfiniBand between Commerce Platform and Oracle Guided Search.

Oracle Exadata Database Machine

Oracle's Exadata Database Machine is Oracle's database platform delivering extreme performance for database applications including online transaction processing (OLTP), data warehousing, reporting, batch processing, and consolidation of mixed database workloads. Exadata is a pre-configured, pre-tuned and pre-tested integrated system of servers, networking, and storage all optimized around Oracle Database. Exadata is an excellent platform for Oracle Commerce because of the many enhancements contained within it, such as write-back flash cache.

Oracle Commerce MAA

Oracle Commerce MAA is a high availability (HA) architecture – providing local HA – layered on top of the MAA best practices for Oracle Database and Oracle Fusion Middleware. Operating out of multiple data centers is increasingly important to maintain some form of high availability. The full MAA architecture includes two sites to provide business continuity in the event of a single site failure. The Configuration Types section covers the different configurations.

Commerce Platform MAA

Oracle Commerce MAA can be achieved by configuring applications in two separate sites or data centers. A typical MAA configuration is shown below Figure 1: Oracle Commerce MAA.

Figure 1: Commerce Platform MAA

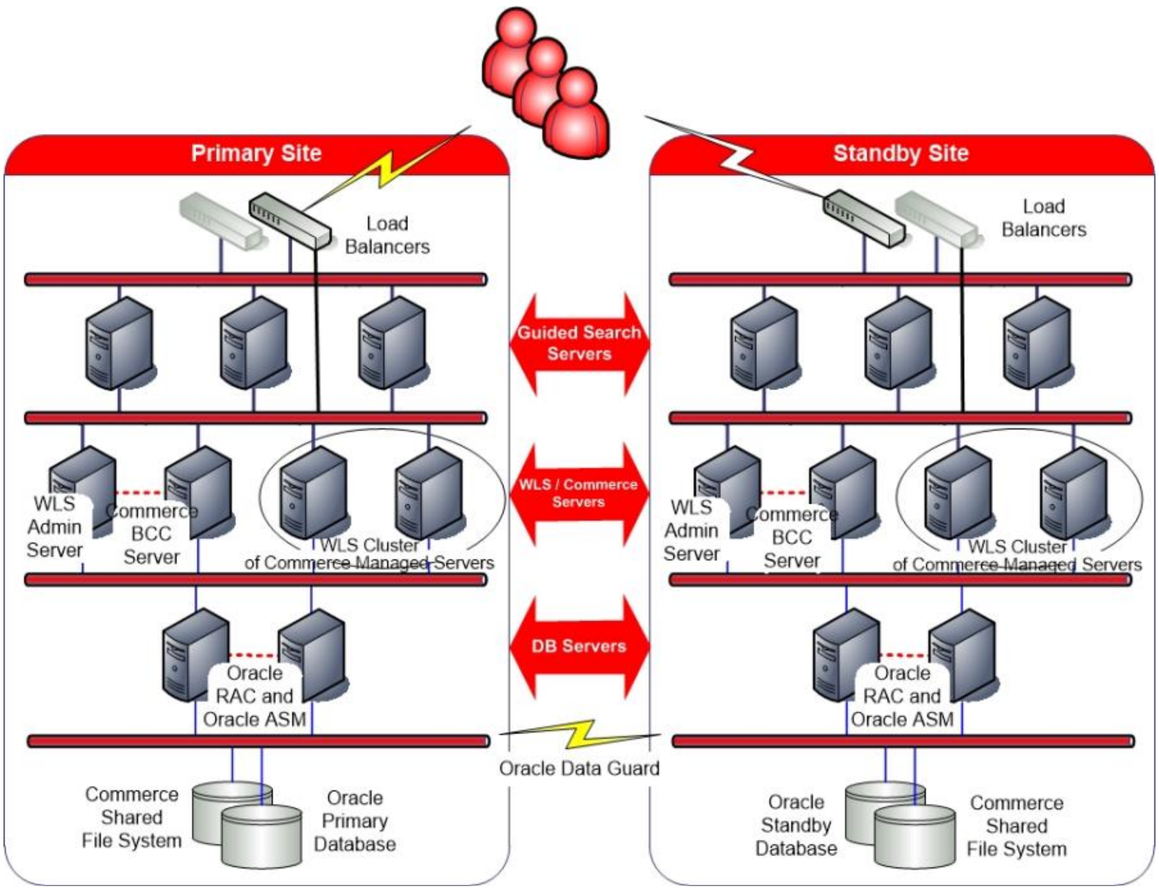

Oracle Database MAA

To achieve maximum Commerce Platform application availability, Oracle recommends deploying it on an Oracle Database MAA foundation as depicted in Figure 2: Oracle Database MAA, which includes the following database technologies. MAA includes high availability within a data center and across data centers.

- *Oracle Real Application Clusters*
- *Oracle Clusterware*
- *Oracle Data Guard and Online Upgrade*
- *Oracle Flashback Database*
- *Oracle Automatic Storage Management*
- *Oracle Recovery Manager and Oracle Secure Backup*
- *Oracle GoldenGate*

See also: "*Oracle Database High Availability Overview*" for a thorough introduction to Oracle Database high availability products, features, and best practices.

Oracle Database MAA is achieved with several configurations, some of which are shown below. These configurations make up a data structure, which remains highly available.

Figure 2: Oracle Database MAA

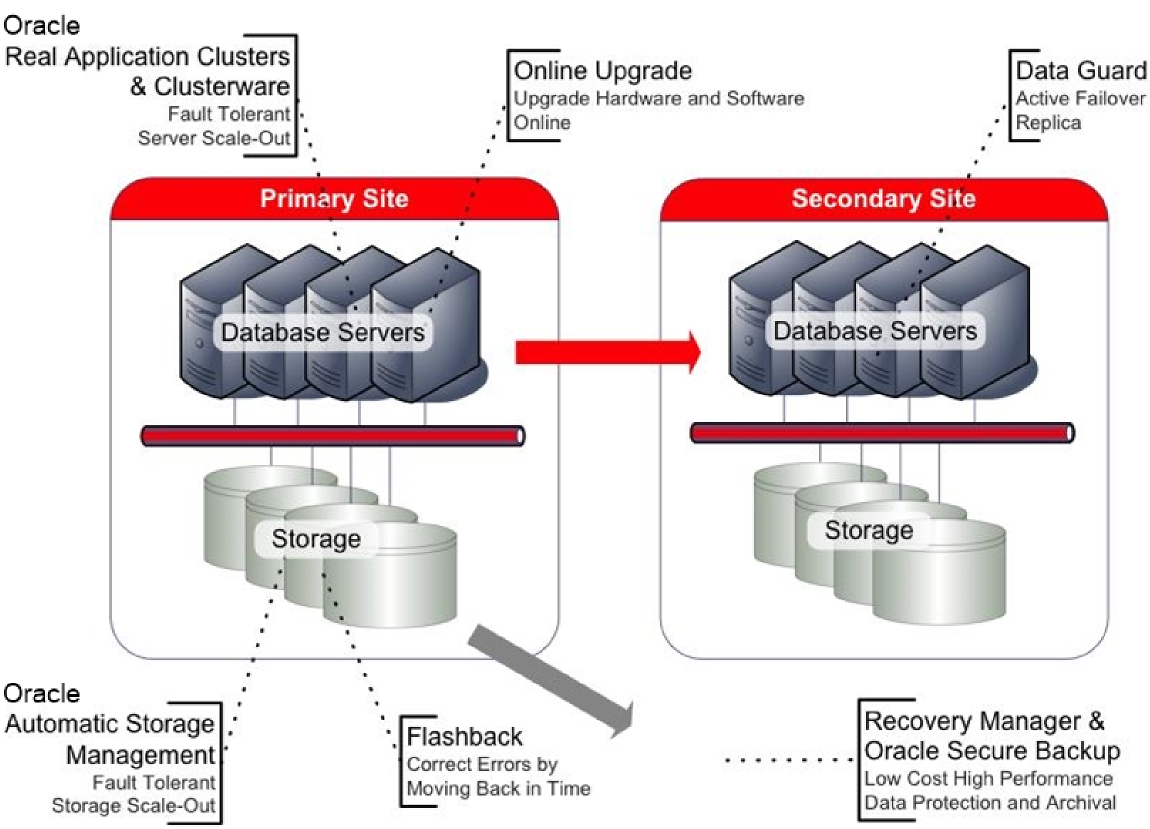

Oracle Real Application Clusters

Oracle Real Application Clusters (Oracle RAC) allows Oracle Database to run any packaged or custom application unchanged across a set of clustered nodes. This capability provides the highest levels of availability and the most flexible scalability. If a clustered node fails, the Oracle database continues running on the surviving nodes. When more processing power is needed, another node can be added without interrupting user access to data.

The Oracle Commerce JDBC client integrates with the Oracle RAC Fast Application Notification (FAN) and Fast Connection Failover (FCF) features to provide a more aware connection pool.

See also: "*Oracle Real Application Clusters Administration and Deployment Guide*".

Oracle Clusterware

Oracle Clusterware is a general purpose clustering solution originally designed for the Oracle Real Application Clusters Active/Active multi-instance database and which has been extended to support clustering of all applications. Oracle Clusterware provides traditional HA failover support in addition to online management of protected resources, such as online relocation of applications for planned maintenance. Oracle Clusterware is a policy engine providing a rich dependency model for start and stop dependencies, ordered startup and shutdown of applications, and defined placement of resources for affinity, dispersion, or exclusion. Oracle Clusterware provides a suite of integrated stand-alone or bundled agents for Oracle Application Express high availability and application resource management.

See also: "*Oracle Clusterware Administration and Deployment Guide*".

Oracle Data Guard and Online Upgrade

Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable Oracle production databases to survive failures, disasters, user errors, and data corruption. Oracle Data Guard maintains these standby databases as transactionally consistent copies of the production database. If the production database becomes unavailable due to a planned or an unplanned outage, Oracle Data Guard can switch any standby database to the production role, thus greatly reducing the application downtime caused by the outage. Oracle Data Guard can be used with traditional backup, restore, and clustering solutions to provide a high level of data protection and data availability.

See also: "*Oracle Data Guard Concepts and Administration*".

Oracle Data Guard allows different data transport methods, depending on the type of data integrity and transport speed required:

» Synchronous – allows for zero data loss. This transport method sacrifices speed for better data integrity.

» Asynchronous – allows for near-zero data loss. This transport method may sacrifice some data integrity for transport speed.

Oracle Commerce supports both physical and logical standby databases. A physical standby database provides a physically identical copy of the primary database with on-disk database structures that are identical to the primary database on a block-for-block basis. A physical standby database is kept synchronized with the primary database using Redo Apply, which recovers the redo data received from the primary database and applies the redo to the physical standby database.

With Oracle Active Data Guard, a physical standby database can receive and apply redo while it is open for read-only access and so may be used for other purposes as well as disaster recovery.

A logical standby database contains the same logical information as the production database, although the physical organization and structure of the data can be different. The logical standby database is kept synchronized with the primary database using SQL Apply, which transforms the data in the redo received from the primary database into SQL statements and then executes the SQL statements on the standby database. A logical standby database can be used for disaster recovery and reporting requirements, and can also be used to upgrade the database software and apply patch sets while the application is online and with almost no downtime.

See also: MAA white paper "*Database Rolling Upgrades Made Easy by Using a Data Guard Physical Standby Database*".

With a single command, a physical standby database can be converted into a snapshot standby database and become an independent database open read/write, which is ideal for QA and other testing scenarios. The snapshot standby database continues to receive and archive redo data from the primary database while it is open read/write, thus protecting primary data at all times. When testing is complete, a single command converts the snapshot back into a standby database and automatically resynchronizes it with the primary database.

A physical standby database can be used for rolling database upgrades using the SQL Apply (logical standby) process, and return to its function as a physical standby database once the upgrade is complete.

See also: MAA white paper "Database Rolling Upgrades Made Easy *by Using a Data Guard Physical Standby Database*".

It is possible to deploy a local standby database at the primary site as well as a remote standby at a secondary site. A local standby offers the advantage that a failover can be performed while the Oracle Commerce Servers continue running, and it can be done almost transparently to the end users. The local standby also offers the ability to perform an online database upgrade without the need to switch to another site. We recommend that both a local and remote standby be deployed for maximum availability.

Oracle Flashback Database

Oracle Flashback Database provides a more efficient alternative to database point-in-time recovery. With Oracle Flashback Database, current data files can be reverted to their contents of a past version. The result is much like restoring data from data file backups and executing point-in-time database recovery. However, Oracle Flashback Database skips the data file restoration and most of the application of redo data. Typically, flashback is used when a speedy recovery to a point-in-time is required. Also, flashback is limited to the amount of time that the undo data is retained, so Oracle Recovery Manager (RMAN) restore may be able to recover further into the past than flashback may be able to.

Enabling Oracle Flashback Database provides the following benefits:

» Eliminates the time to restore a backup when fixing human error that has a database-wide impact.
» Because human errors can be quickly undone, it allows standby databases to use real-time apply to synchronize with the primary database.
» Allows quick standby database re-instantiation after a database failover.

See My Oracle Support note, "Flashback Database Best Practices & Performance (MOS article 565535.1)," which includes the use of pre-allocated flashback logs (see _DB_FLASHBACK_LOG_MIN_TOTAL_SPACE*) to improve performance. See also:* Using Oracle Flashback Technology in the "*Oracle Database Development Guide*".

Oracle Automatic Storage Management

Oracle Automatic Storage Management (Oracle ASM) provides a vertically integrated file system and volume manager directly in the Exadata kernel, resulting in:

- Significantly less work to provision database storage.
- Higher levels of availability.
- Elimination of the expense, installation, and maintenance of specialized storage products.
- Unique capabilities for database applications.

For optimal performance Oracle ASM spreads files across all available storage. To protect against data loss Oracle ASM extends the concept of the stripe and mirror everything (SAME) methodology and adds more flexibility in that it can mirror at the database file level rather than the entire disk level.

Oracle Recovery Manager and Oracle Secure Backup

Oracle Recovery Manager (RMAN) is an Oracle Database tool that can back up, restore, and recover database files. It is a feature of Oracle Database and does not require separate installation. RMAN integrates with sessions running on an Oracle database to perform a range of backup and recovery activities, including maintaining a repository of historical data about backups.

See also: "*Oracle Database Backup and Recovery User's Guide*".

Oracle Secure Backup (OSB) is a centralized tape backup management solution providing performant, heterogeneous data protection in distributed UNIX, Linux, Windows, and Network Attached Storage (NAS) environments. By protecting file system and Oracle Database data, OSB provides a complete tape backup solution for your IT environment. OSB is tightly integrated with RMAN to provide the media management layer for RMAN.

Oracle GoldenGate

Oracle GoldenGate provides bidirectional replication of data between sites to allow for each site to act independently. It is normally used in compulsory systems, where the highest availability is required. Oracle GoldenGate is normally used in fully Active/Active implementations and can provide more flexibility in site planning. The key features of Oracle GoldenGate include:

- Log-based replication, which does not have a distance limitation. This allows for greater distance between sites.
- Two or more sites have full read/write capabilities and act independently of each other.
- Maintains data integrity and allows for conflict detection and resolution.

Create Role-Based Database Services

A database service provides a simple named access point to the database. A service can physically span multiple instances in Oracle RAC and can be simply moved from one instance to another. By requiring Commerce to connect only through a service, we are able to relocate or reconfigure the service without reconfiguring Commerce. Oracle Commerce on Oracle WebLogic Server uses JDBC and, more specifically, Oracle WebLogic Server Active GridLink for Oracle RAC connections to configure its connections. Oracle WebLogic Server Active GridLink for Oracle RAC data sources provides connectivity directly to the database services. Together with the Oracle Notification Service (ONS), the data source can respond to state changes of the Oracle RAC instances, which are contained in the database service. Role-based database services should be created for each database role and purpose as summarized in Table 1: Role-Based Database Services.

TABLE 1: ROLE-BASED DATABASE SERVICES

Database Role	Purpose
Primary database	Production access to primary
Standby database	Production offload of queries on standby
Snapshot standby database	Standby site testing

A database service can be created and configured using Oracle Enterprise Manager or by using the `srvctl` command line tool. Commerce uses JDBC and, therefore, does not use transparent application failover (TAF). Here is an example of how a service can be created with the `srvctl` command:

```
srvctl add service -d atgmaa -s atgsvc -r atgmaa1,atgmaa2 -l PRIMARY -q FALSE -e NONE -m NONE -w 0 -z 0
```

Role-based database services are also an important part of client failover best practices as detailed in the recommended MAA white paper "*Client Failover Best Practices for Highly Available Oracle Databases: Oracle Database 11g Release 2*". It applies as much to Oracle Database 12*c* as it does to Oracle Database 11*g*.

Configure HugePages (Linux Database and Application Servers Only)

Web Commerce typically runs with many database connections and a large System Global Area (SGA), or a group of shared memory areas, supporting an Oracle RAC instance. Therefore, configuring HugePages for the Commerce database instances is essential. HugePages are not configured by default within an Exadata system, so it is necessary to manually configure sufficient HugePages for the Oracle ASM instance and all database instances on each Linux database server node. This results in more efficient page table memory usage, which is critically important with a large SGA or when there are high numbers of concurrent database connections. HugePages can only be used for SGA memory space and should not be configured for more than is required.

MOS ID 361468.1, "*HugePages on Oracle Linux 64-bit*" describes how to configure HugePages. Automatic shared memory management (ASMM) can be used with HugePages , so use the `SGA_MAX_SIZE` parameter to set the SGA size for each instance.

Automatic memory management cannot be used in conjunction with HugePages, so the `MEMORY_TARGET` and `MEMORY_MAX_TARGET` parameters should be unset for each database instance. See MOS ID 749851.1 "*HugePages and Oracle Database 11g Automatic Memory Management (AMM) on Linux*" for details.

Set the parameter `USE_LARGE_PAGES='only'` for each instance so that the instance only starts if sufficient HugePages are available. See MOS ID 1392497.1 "*USE_LARGE_PAGES To Enable HugePages*" for details.

It may be necessary to restart the database server to bring the new HugePages system configuration into effect. Check to ensure that you have sufficient HugePages by starting all of the database instances at the same time.

Starting with Oracle Database 11*g* Release 2 (11.2.0.2), a message is logged to the database alert log when HugePages are being used. For example:

```
********************** Huge Pages Information ***********************
Huge Pages memory pool detected (total: 18482 free: 17994)
DFLT Huge Pages allocation successful (allocated: 4609)
***********************************************************************
```

HugePages have also been proven to be a more performant method of running applications on Oracle WebLogic Server. It is a best practice to set up the operating system of a server running Oracle WebLogic Server to use HugePages to support the memory usage of Oracle WebLogic Server managed servers.

Note: *When configuring HugePages for database or Oracle WebLogic Server managed server use, it is important to calculate the amount of memory required for either application and allocate slightly more than required to HugePages in the operating system. Once memory is allocated to HugePage, it is reserved only for HugePages. If the database or Oracle WebLogic Server is not configured to use HugePages, and the operating system has been configured to allocate memory for HugePages, the database or application server may run out of memory or possibly not start at all.*

Handle Database Password Expiration

The default behavior of Oracle Database has changed as of release 11*g* such that database user passwords expire after 180 days. Processes should be put in place to refresh passwords regularly, or expiration should be extended or disabled. Commerce application availability is impacted if passwords are allowed to expire. Password expiration for the default user profile can be disabled with the following command:

```
alter profile default limit password_life_time unlimited;
```

Configure Dead Connection Detection

When a Commerce Server node fails suddenly, there may not be time for the operating system to reset the TCP connections and, as a result, the connections on the database server remain open. To clean up the "dead" connections, it is recommended that Dead Connection Detection is configured. See MOS ID 151972.1 "*Dead Connection Detection (DCD) Explained*" for details. Making these configuration changes may have an adverse effect on network utilization. Therefore, all changes should be tested and monitored carefully.

Reduce Timeout on Oracle RAC Node Failure (Exadata Only)

On Exadata, it is possible to failover more quickly in the event of an Oracle RAC node failure by reducing the `misscount` parameter. The parameter defines how long to wait after a node becomes unresponsive before evicting the node from the cluster. The parameter should not be set to less than `30` seconds. To update the CSS `misscount` setting, log in as the `root` user on one of the database servers and run the following command:

```
$GRID_HOME/bin/crsctl set css misscount 30
```

Oracle Fusion Middleware MAA

To achieve maximum Commerce application availability, deploy Commerce on an Oracle Fusion Middleware WebLogic Server MAA foundation that uses the following WebLogic Server high availability features:

» WebLogic Server cluster
» Oracle WebLogic Server Active GridLink data sources
» Shared storage and storage replication

WebLogic Server Cluster

A WebLogic Server cluster consists of multiple WebLogic Server instances running simultaneously and working together to provide increased scalability and availability. A cluster appears to clients to be a single WebLogic Server instance. The server instances that constitute a cluster can run on the same system or be located on different systems. You can increase the capacity of a cluster by adding additional server instances to the cluster on an existing system or you can deploy additional systems on which the cluster can host incremental server instances. In a WebLogic Server cluster, application processing can continue when a server instance fails. You can cluster application components by deploying them on multiple server instances in the cluster. Then, if a server instance on which a component is running fails, another server instance on which that component is deployed can continue application processing. Clusters are created for ease of configuration and application deployment, because a single configuration or deployment can be performed on the cluster. This single configuration or deployment is propagated to all of the cluster members.

Configuring Commerce session backup (see "Enabling Component Backup" in the "*Oracle Commerce Platform Installation and Configuration Guide*" for details) enables Commerce to utilize the application failover and server migration aspects of WebLogic Server clustering. Depending on which methods are used for replication, it may be possible to replicate session states and data caching across sites. Starting with Commerce version 11.1, Oracle Coherence 12c can be used for HTTP session replication caching.

See also: "Using Oracle Coherence*Web for Session Failover" in the "*Oracle Commerce Platform Installation and Configuration Guide*".

Oracle WebLogic Server Active GridLink for Oracle RAC Data Sources

A single Oracle WebLogic Server Active GridLink for Oracle RAC data source provides connectivity between WebLogic Server and an Oracle RAC database. It uses ONS to respond adaptively to state changes in an Oracle RAC database and responds to FAN events to provide Fast Connection Failover (FCF), Runtime Connection Load-Balancing (RCLB), and Oracle RAC instance graceful shutdown. It also provides capabilities of Affinities.

See also: "Using GridLink Data Sources" in the "*Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*" guide for more information about Oracle WebLogic Server Active GridLink for Oracle RAC data sources.

Storage Replication

Storage replication technology is used to replicate the middle tier file systems and other data from the shared storage of the production site to the shared storage of the standby site. After storage replication is enabled, application deployment, configuration, metadata, data, and product binary information is replicated from the production site to the standby site. It is not necessary to perform any Oracle software installations at the standby site hosts. This method is best used with an Active/Passive configuration, because an Active/Active middle tier application requires independent configurations.

When the production site storage is replicated to the standby site storage, the equivalent Oracle home directories and data are written to the standby site storage. The recommended interval for scheduling incremental replications is once a day for the production deployment because 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 Commerce application updates at the production site).

Additional References

See these Oracle documentation references for more background on WebLogic Server MAA:

» *"Oracle Fusion Middleware 12c High Availability Guide"*
» *"Oracle Fusion Middleware Using Clusters for Oracle WebLogic Server"*
» *"Oracle Fusion Middleware High Availability Guide, Release 11g"*
» *"Oracle Fusion Middleware Disaster Recovery Guide"*

Refer to these MAA white papers and additional documentation:

» *"Oracle WebLogic Server and Highly Available Oracle Databases: Oracle Integrated Maximum Availability Solutions"*
» *"Oracle WebLogic on Shared Storage: Best Practices"*
» *"Oracle Fusion Middleware 11g Backup and Recovery"*
» *"Oracle Fusion Middleware Disaster Recovery Guide"*

Oracle Commerce Platform and Merchandising MAA

To achieve Oracle Commerce Platform and Merchandising MAA, Oracle recommends deploying Oracle Commerce on a foundation that includes the architecture depicted in Figure 3: Oracle Commerce High Availability on both the primary site and secondary site shown in Figure 1: Oracle Commerce MAA.

Figure 3: Oracle Commerce High Availability

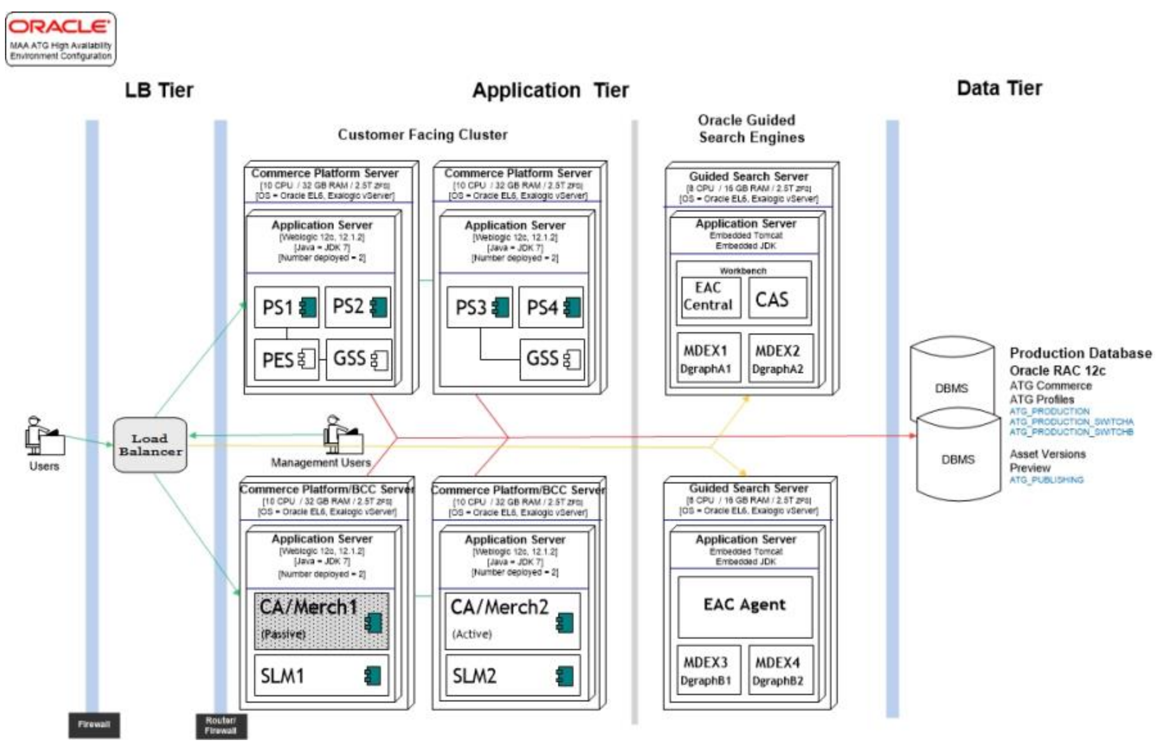

Oracle Commerce Platform high availability is built on the foundation of Oracle Database MAA and Oracle Fusion Middleware MAA as described previously in this white paper.

In Figure 3: Oracle Commerce High Availability, the only customer-facing cluster created (as shown in Figure 4: Oracle Commerce Customer Facing Cluster) is the Commerce web store application. The components of the Oracle Commerce Platform cluster are as follows:

» **Page Servers (PS)**: At least two Page Servers are required for HA. Ideally, there should be at least one on each application server for MAA. These serve the actual page content to the end users. The best practice is not to combine a PS with any other server type. This is the only Commerce component which is load balanced and, therefore, can be configured as part of a WebLogic Server cluster.

  For this particular environment, the PES/GSS have been combined into two of the four PS applications, because there is no need for global scenarios or for editing of scenarios.

» **Process Editor Server (PES)**: PES is required for the editing of Oracle Commerce Platform Scenarios. Only one PES can be running in a platform environment at any given time. This component does not require High Availability, because the main store application does not require PES in order to run.

» **Global Scenario Server (GSS)**: GSS is required for applications that use long running or Commerce Platform scenarios that are globally-scoped. This type of scenario is uncommon. There can be many GSSs in a given environment and the PES is configured as a GSS by default.

» **Server Lock Manager (SLM)**: Use one primary SLM instance and one backup SLM instance specified only by run order. Either SLM instance can become the primary, simply by being the first SLM to start. The SLMs are used by any repository whose item-descriptor is set with locked cache mode and by any custom code that uses the SLM API. The cache mode is used when more than one SLM could possibly write data to the same object, in the same repository, at the same time, across multiple threads. The best example of this is user profile data. There are also several out-of-the-box OOTB and possibly custom code calls to this component.

» **Commerce Merchandising**: Commerce Merchandising is the management tier. The management tier is not usually set up as HA, because it only needs to be available for internal users for the management of the content for the customer-facing site. The base component of the management tier is the Commerce Business Control Center (BCC). Sub-components of the BCC are:

  » **Content Administration (CA)**: CA is the means by which business users create new content for the customer-facing site and deploy it either to the database or to the file systems.

  » **Merchandising**: Merchandising is a more granular method of creating content for the customer-facing site. If configured, Merchandising includes preview functionality to view how content will look based on certain criteria.

  » **Site Management**: The BCC has the ability to maintain the data of multiple sites in a single configuration set. This is also known as Commerce Platform Multisite. The default installation of the Commerce Platform Commerce Reference Store (CRS) includes the CRS Store and the CRS Home sites.

While the management tier is not required to be set up with high availability, redundancy is required in order to allow component access in the event of an outage. Therefore, the management tier is set up to run in an Active/Passive capacity with the use of shared storage.

Data Tier

The data tier has already been covered in detail in the Oracle Database MAA section.

WebLogic Server / Commerce Platform Clustering

Commerce Platform uses WebLogic Server clustering described earlier in this white paper. Because Commerce applications run within a WebLogic Server application server, clustering is essential to maintain session state for certain aspects of a user's HTTP session. For this, Commerce uses the combination of Coherence*Web clustering and Commerce session component backup to the WebLogic Server cluster. Commerce places session states into the HTTP session object as configured in the `/atg/dynamo/Configuration.sessionBackupPropertyList` property. Session data is replicated to all other members of the cluster using this component. That way, if a Commerce managed server should become unavailable while a user has an active session, the user's full session is failed over to another active cluster member.

See "WebLogic Server Clustering" in the "*Oracle Fusion Middleware 12c Understanding Oracle WebLogic Server*" for more information on the setup of a WebLogic Server cluster.

See: Enabling Component Backup in the "*Oracle ATG Installation and Configuration Guide*" for information on Commerce session backup configuration.

Note: *as mentioned earlier, clustering of WebLogic Server applications is optional. Commerce Page Servers are often clustered for ease of configuration and application deployment.*

Commerce Platform, consisting of several page servers and two server lock managers, are typically deployed to multiple hosts. The Commerce Content Administration and Merchandising applications can also be incorporated into the same hosts. Figure 4: Oracle Commerce Customer-Facing and Management-Facing Clusters shows a sample of a small environment deployment.

Figure 1: Oracle Commerce Customer-Facing and Management-Facing Clusters

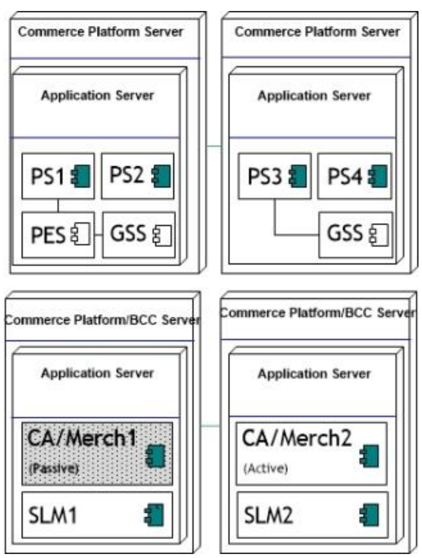

Oracle Commerce Experience Manager MAA

The Platform/Experience Manager integration enables customers of Oracle Commerce Platform and Oracle Experience Manager to index the product catalog data in MDEX engines, the query engine that is the core of Oracle Guided Search, where it can then be queried and the results can be displayed on commerce sites. The MDEX engine consists of the Indexer (Dgidx), the Dgraph, and the Agraph. The MDEX engine loads the indices generated by the indexing component of the Endeca Information Transformation Layer. The components of the Oracle Guided Search architecture are documented in "*Oracle Commerce Guided Search Common Documentation*".

» **Experience Manager Application Controller (EAC)**: This application is used to control any Oracle Guided Search application. There is a central EAC which handles the overall control of Experience Manager subcomponents including distributing content to the MDEX engines. There are also EAC agents that are required to be running on each host that runs a standalone MDEX. See also: "*Oracle Commerce Guided Search Common Documentation*".

» **Content Acquisition System (CAS)**: CAS is the component where data is stored after Commerce exports it from its repositories. See also: "*Oracle Commerce Guided Search CAS Documentation*".

» **Tools & Frameworks w/ Experience Manager**: Tools and Frameworks is the installation for Workbench, Experience Manager, and a number of reference applications. Within the Tools and Frameworks installation is Workbench and Experience Manager, where merchandisers can configure custom landing pages, make modifications to search results, control dimensions, establish thesaurus entries, etc. See also: "*Oracle Commerce Guided Search Tools and Frameworks Documentation*".

» **MDEX:** The MDEX engine (sometimes known as the Dgraph, which is the actual process name) is the workhorse application that processes the search and navigation queries. See also: "*Oracle Commerce Guided Search MDEX Engine Documentation*".

To achieve maximum availability for the Platform-Experience Manager integration, Oracle recommends load balancing MDEX engines in order to ensure consistent display of Oracle Guided Search driven content for the Commerce Platform applications. While Experience Manager is made up of several components, the Dgraph engines are the components that Platform and Merchandising access in order to view this content. The other components are used either by business users or are only required during startup of Platform page servers, therefore, HA is not a concern. However, redundancy for these components is required.

Figure 5: Commerce Guided Search Deployment is a typical Oracle Guided Search small environment deployment. The MDEX engines need to be configured as highly available, because this is the component used for serving the customer-facing search content.

Figure 2: Commerce Guided Search Deployment

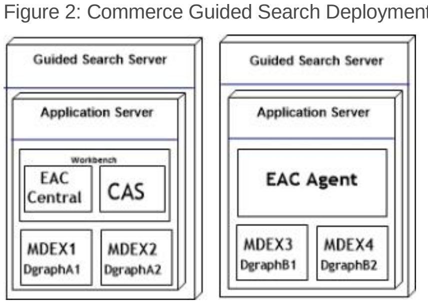

Load Balancing

Commerce Platform components are installed and deployed on multiple servers and run in an Active/Active configuration for high availability and scalability purposes. Client initiated workload is distributed across multiple component instances running on multiple servers through load balancing. Web server load is distributed by an HTTP load balancer directly to WebLogic Server or to the Oracle HTTP Server (OHS). OHS is optional and not utilized or depicted in Figure 3: Oracle Commerce High Availability. Refer to the '*F5 Oracle ATG Deployment Guide*" for a detailed example.

The Endeca MDEX engine load balancing is covered in Part II of the "*Oracle Commerce Guided Search Performance Tuning Guide*".

Oracle Commerce Site State Model and State Transitions

In Figure 6: MAA Site State Model and State Transitions, we illustrate the states that a deployment goes through as it progresses from the initial single site implementation through the setup, testing, and eventual dual site MAA deployment. The systems have a specific configuration in each state and there is a set of documented steps to move from one state to the next.

Figure 3: MAA Site State Model and State Transitions

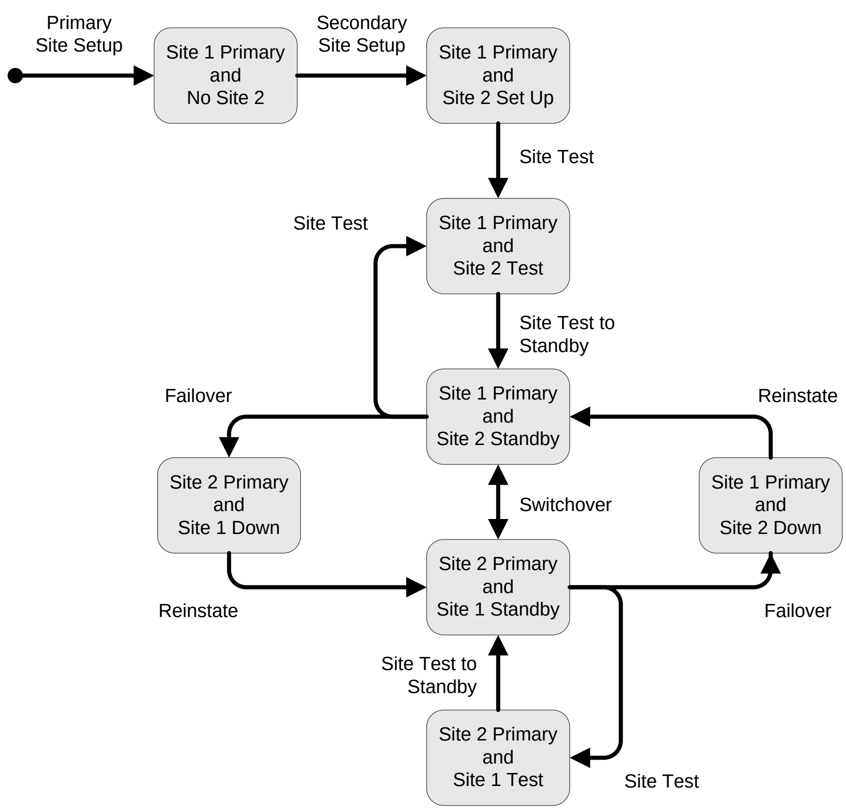

Table 2: State Transition Descriptions contains the steps taken for different portions of switchover or failover operations mentioned in the section titled Oracle Commerce MAA on Exalogic and Exadata later in this white paper.

TABLE 2: STATE TRANSITION DESCRIPTIONS

Transition	Description
Primary Site Setup	Install and configure the primary site.
Secondary Site Setup	Establish the secondary site.
Site Test	Prepare the standby site for a site test.
Site Test to Standby	Convert the site performing a site test back to standby mode.
Switchover	Switch the roles so that the current standby site becomes the primary site and the current primary site becomes the standby site.
Failover	Switch the current standby to primary mode. The current primary mode is assumed to be down or unavailable.
Reinstate Standby	Reinstate Standby

Table 3: State Summary for Database and File System shows the progression of Site 1 as a single site without HA, through the setting up of Site 2 as a standby and switching over or failing over from Site 1 to Site 2.

TABLE 3: STATE SUMMARY FOR DATABASE AND FILE SYSTEM

Site State	Commerce Database Data Guard	Commerce File System - Replication
Site 1 Primary and No Site 2	Not configured	Not configured
Site 1 Primary and Site 2 Set Up	Site 1 primary and Site 2 physical standby. Snapshot standby during setup.	Site 1 primary and Site 2 physical standby. Snapshot standby during setup.
Site 1 Primary and Site 2 Test	Site 1 primary and Site 2 snapshot standby.	Site 1 primary with continuous replication to Site 2. Site 2 snapshot created for test.
Site 1 Primary and Site 2 Standby	Site 1 primary and Site 2 physical standby.	Site 1 primary with continuous replication to Site 2.
Site 2 Primary and Site 1 Down	Site 2 primary through failover and Site 1 down.	Site 2 primary established from replica and Site 1 down.
Site 2 Primary and Site 1 Standby	Site 2 primary and Site 1 physical standby.	Site 2 primary and continuous replication to Site 1.
Site 1 Primary and Site 2 Down	Site 1 primary through failover and Site 2 down.	Site 1 primary established from replica and Site 2 down.
Site 2 Primary and Site 1 Test	Site 2 primary and Site 1 snapshot standby.	Site 2 primary with continuous replication to Site 1. Site 1 snapshot created for test.

Oracle Commerce MAA on Exalogic and Exadata

Oracle Commerce 11.1 is deployed into an Oracle Commerce Reference Store application and Merchandising MAA configuration on Exalogic and Exadata machines. It is deployed in an Active/Passive configuration.

Exadata MAA is a very mature pre-optimized, pre-configured, integrated system of software, servers, storage, and MAA configuration best practices that comes ready-built to implement the highest database and application availability and performance. Additionally, the MAA Best Practices white paper, *"Deploying Oracle Maximum Availability Architecture with Exadata Database Machine*," includes an Exadata MAA outage and solution matrix which includes Exadata specific components in addition to the aforementioned reference to "*Oracle Database High Availability Best Practices, 12c Release 1"*. Most outages should incur zero database downtime and a minimal application impact for any affected connections.

Figure 7: Oracle Commerce Exadata and Exalogic Disaster Recovery Topology is a high level overview of the configured system.

Figure 4: Oracle Commerce Exadata and Exalogic Disaster Recovery Topology

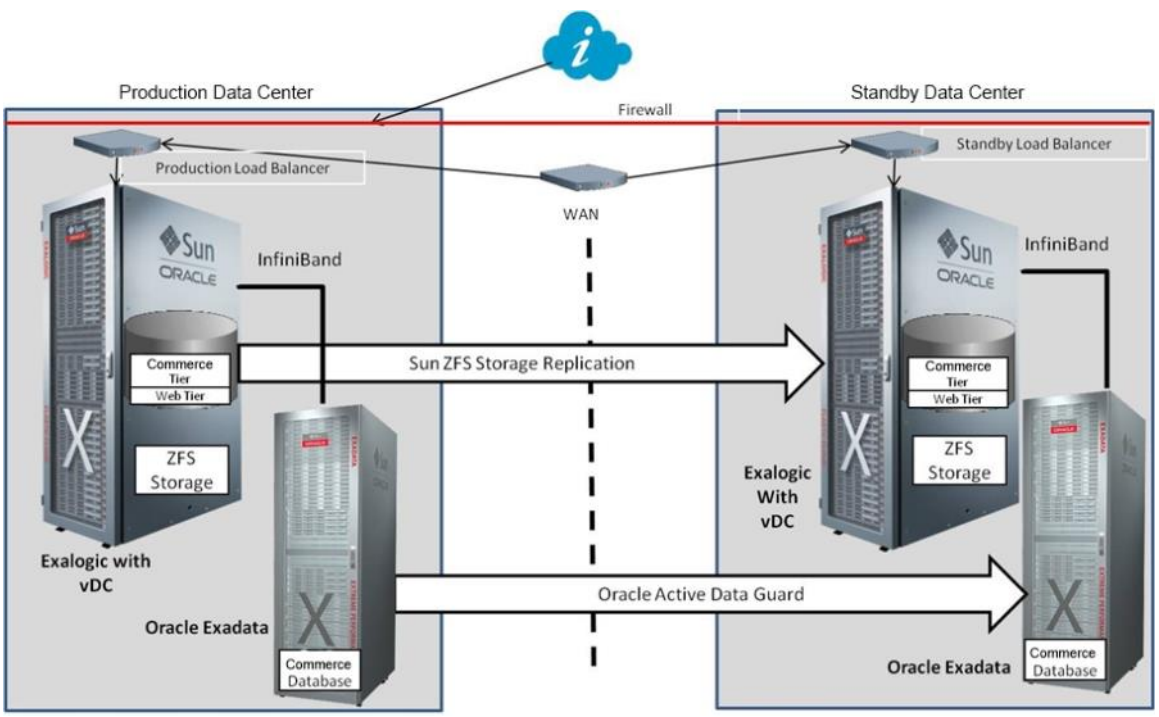

Server Deployment

The production data center is comprised of an Oracle Exadata Database Machine quarter rack X3-2 with the database services listed in Table 4: Database Services. Additionally, bare metal X3 Exalogic compute nodes are used for the Commerce Platform and Merchandising application tier components as detailed in Table 5: Primary Site Servers, and their function details are pictured in Figure 3: Oracle Commerce High Availability, Figure 4: Oracle Commerce Customer-Facing Cluster and Figure 5: Commerce Experience Manager Deployment*.* The steps to configure the application tier are detailed in the Oracle Commerce MAA Configuration Best Practices white paper.

TABLE 4: DATABASE SERVICES

Service Name	Database Role	Purpose
ATGSVC	Primary database	Production
ATGSVC_TST	Snapshot standby database	Site test (to test application and database switching)
ATGSVC_STBY	Standby database	Production offload of read-only workloads

TABLE 5: PRIMARY SITE SERVERS

Host Name	Function
scan04cn21	WebLogic Server administration server Commerce Platform server 1 Oracle Guided Search server 1
scan04cn22	Commerce Platform server 2 Commerce Merchandising Guided Search server 1
scan0cn23	Commerce Platform server 3
scan04cn24	Commerce Platform server 4

The standby data center is comprised of an Oracle Exadata Database Machine quarter rack X4-2 with the database services listed in Table 4: Database Services. Additionally, X3 Exalogic virtual machine (VM) servers are configured inside of the Exalogic virtual data center (vDC) for the Commerce Platform and Merchandising application tier components as detailed in Table 6: Standby Host Names, and their functional details are described in "*Oracle Commerce Platform High Availability*" and pictured in Figure 3: Oracle Commerce High Availability, Figure 4: Oracle Commerce Customer-Facing Cluster, and Figure 5: Commerce Experience Manager Deployment.

TABLE 6: STANDBY SITE SERVERS

Host Name	Function
scan03vm0019-eoib2	WebLogic Server administration server Commerce Platform server 1 Oracle Guided Search server 1
scan03vm0020-eoib2	Commerce Platform server 2 Commerce Merchandising Oracle Guided Search server 1
scan03vm0021-eoib2	Commerce Platform server 3
scan03vm0021-eoib2	Commerce Platform server 4

F5 4200v Local Traffic Managers load balancers are used for Oracle Commerce Platform Server load balancing.

The Commerce Platform and Experience Manager file systems and all other shared file systems are hosted on the ZFS Storage Appliances located on the Exalogic machines.

Workload

The "*Oracle Application Testing Suite*" (Oracle ATS) is used to capture and create a repeatable workload against the Commerce Reference Store. The behavior of this workload is monitored during outage testing and any anomalies are noted. The workload is comprised of a single script with 2000 virtual users executing each scenario in a loop. The script consisted of:

- Browse and search, which utilizes Oracle Guided Search.
- Adding to shopping cart.
- Full anonymous user checkout.

The primary metric is page views/second and averages 434 or 28.9 page views per second per CPU in our baseline test as illustrated in Figure 8: Baseline Workload Chart. The workload utilizes fifteen Commerce Page Servers.

Figure 8: Baseline Workload Chart

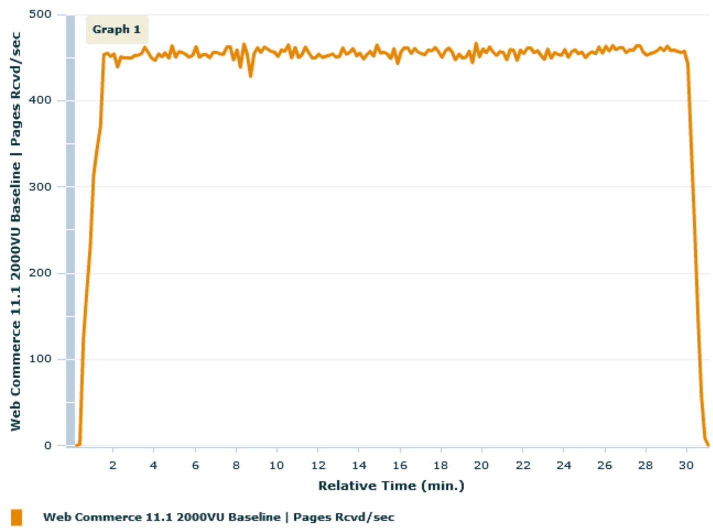

Configuration Types

The following sections give a brief explanation of the different ways in which two sites or data centers can be configured to provide MAA for Oracle Commerce applications. Each configuration lists its benefits and drawbacks.

Fully Active/Passive Configuration

When a 2-site configuration is set up where the applications of one site are live and the other applications are on standby, it is known as an Active/Passive configuration. This is a common configuration for most clients. As time and technology progresses, this configuration is losing popularity. Figure 9: Fully Active/Passive Configuration shows the typical Commerce Active/Passive configuration.

Figure 9: Fully Active/Passive Configuration

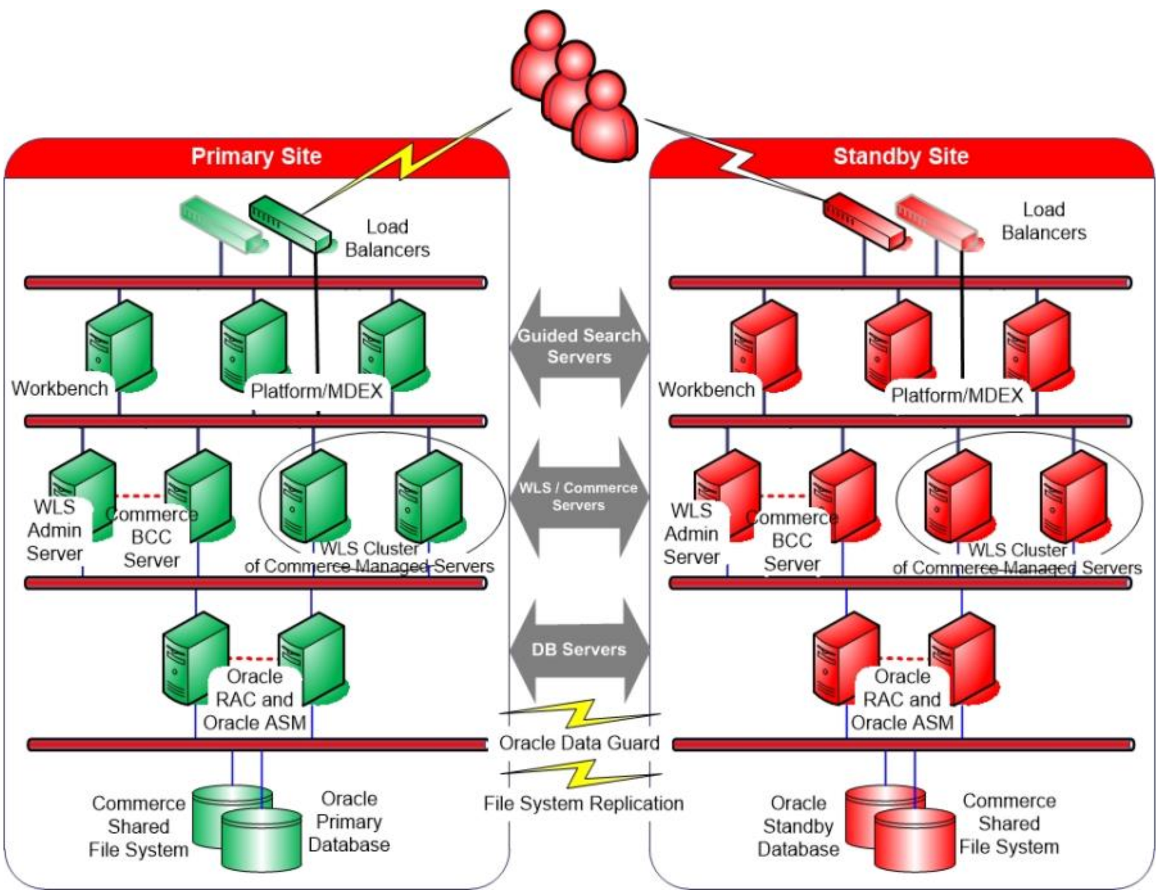

Benefits of Active/Passive Configuration

The benefits of an Active/Passive configuration are:

» Easiest to configure.

» Supported by Oracle support agreements in all configurations.

» Configurations are limited to individual data centers and each is independent.

» Only one data center functions at a time.

» Supports out-of-the-box Commerce Platform configuration types.

» ZFS Replication can be used in Exalogic implementations because the standby site can be configured in an identical manner.

» Easier application deployments between sites because replication handles the standby site deployment.

» Two completely independent sites. Switching of sites to maintain equal usage should be observed, such as at each code or data deployment.

Drawbacks to Active/Passive Configuration

The drawbacks of an Active/Passive configuration are:

» Less traffic can be run through the environments because only half of the servers are in use at a given time.

» Time is needed to reconfigure business-facing applications in alternative data centers (storefront can be passive).

» Peak load may cause strain on applications.

» Does not support zero-downtime if the entire primary site goes down.

Active/Active/Passive Configuration

The Active/Active/Passive site configuration has become increasingly popular among Commerce clients. In this configuration, certain customer-facing applications are set up as Active/Active and other internal-facing applications may not be depending on how they are written and deployed. The database is set up in the Active/Passive configuration with a one-way replication from the primary to the standby sites. Because of the combination of Active/Passive and Active/Active, this type of configuration is also known as Active/Active/Passive (A/A/P). Figure 10: Active/Active/Passive Configuration shows a typical Active/Active/Passive configuration.

Figure 10: Active/Active/Passive Configuration

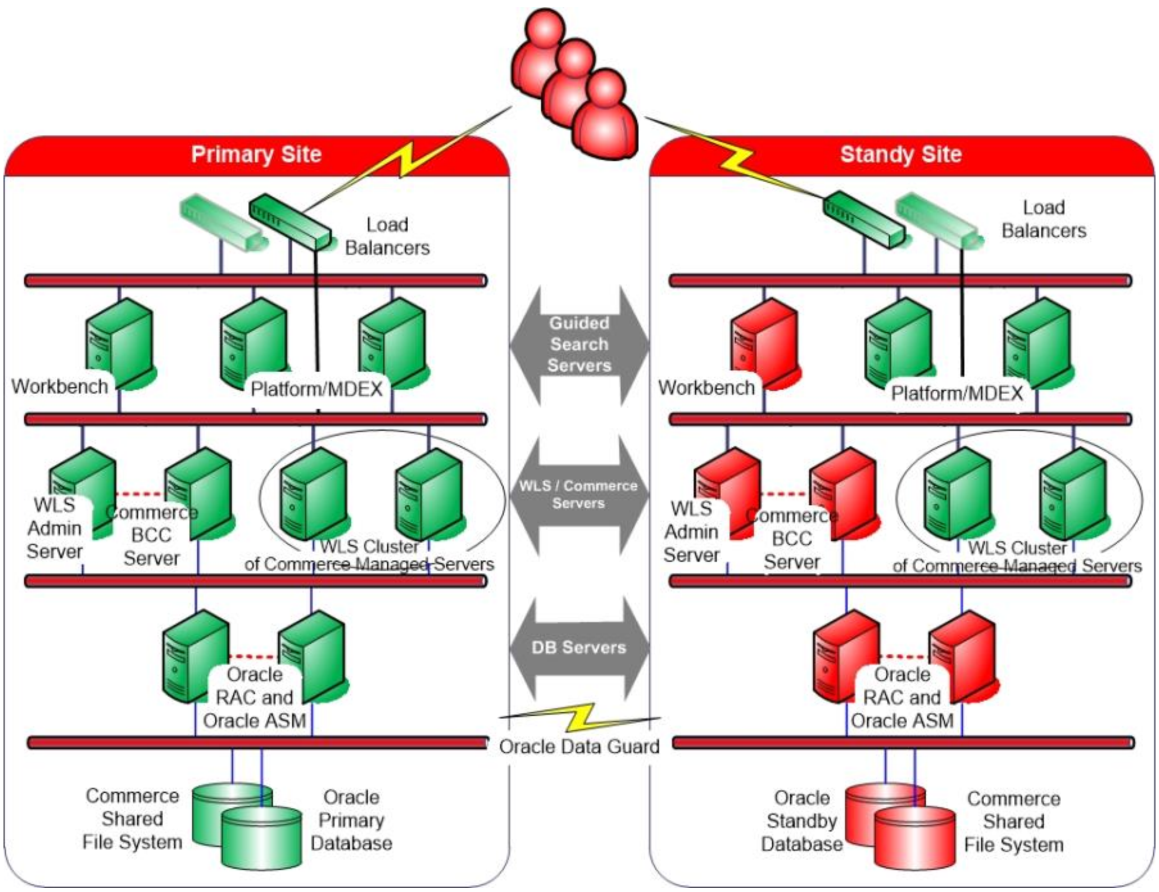

Benefits of Active/Active/Passive Configuration

The benefits of an Active/Active/Passive configuration are:

» Supports more traffic because both sites serve traffic for customer-facing applications.

» Minimal downtime if there is an application failure on a single site location.

Drawbacks to Active/Active/Passive Configuration

The drawbacks of an Active/Active/Passive configuration are:

» More complicated to set up because separate configurations are required for each site. Also, several different failover configurations options can be considered.

» Data centers cannot be very far apart because latency plays a large factor when the site with the passive database has its applications connecting to the site with the active database. The maximum distance should be within 100 miles.

» Data center applications do not act independently of each other because they both connect to the same database nodes.

» May not support out-of-the-box Commerce configuration types and custom configuration may be required depending on the BCC deployment type selected.

» Oracle support agreements may not support custom code or configurations if a client develops a new component which does not comply to support guidelines.

» Downtime to swap to a passive database if there is a database failure.

» Only the Server Lock Managers from the data center, where the active database nodes reside, should be used for all production JVMs in both data centers. This is to prevent additional latency of having to jump back and forth unnecessarily between sites.

» Only certain applications can be set up as Active/Active due to how they are designed.

    » Only the Platform Page Servers and Oracle Guided Search MDEX processes can be Active/Active and load balanced.

    » Commerce Merchandising and the Experience Manager Workbench cannot be Active/Active, so plans need to be put in place to fail these applications over if the primary site fails.

Fully Active/Active Configuration

This configuration allows for the highest availability. However, it is the most complex to set up and some of the configuration is not available for out-of-the-box. Fully Active/Active configuration also incorporates the use of Oracle GoldenGate to perform asynchronous replication between both data centers. Items replicated include the new or changed catalog data for the Commerce sites, as well as order and profile data. Figure 11: Fully Active/Active Configuration shows a typical Active/Active configuration.

Figure 11: Fully Active/Active Configuration

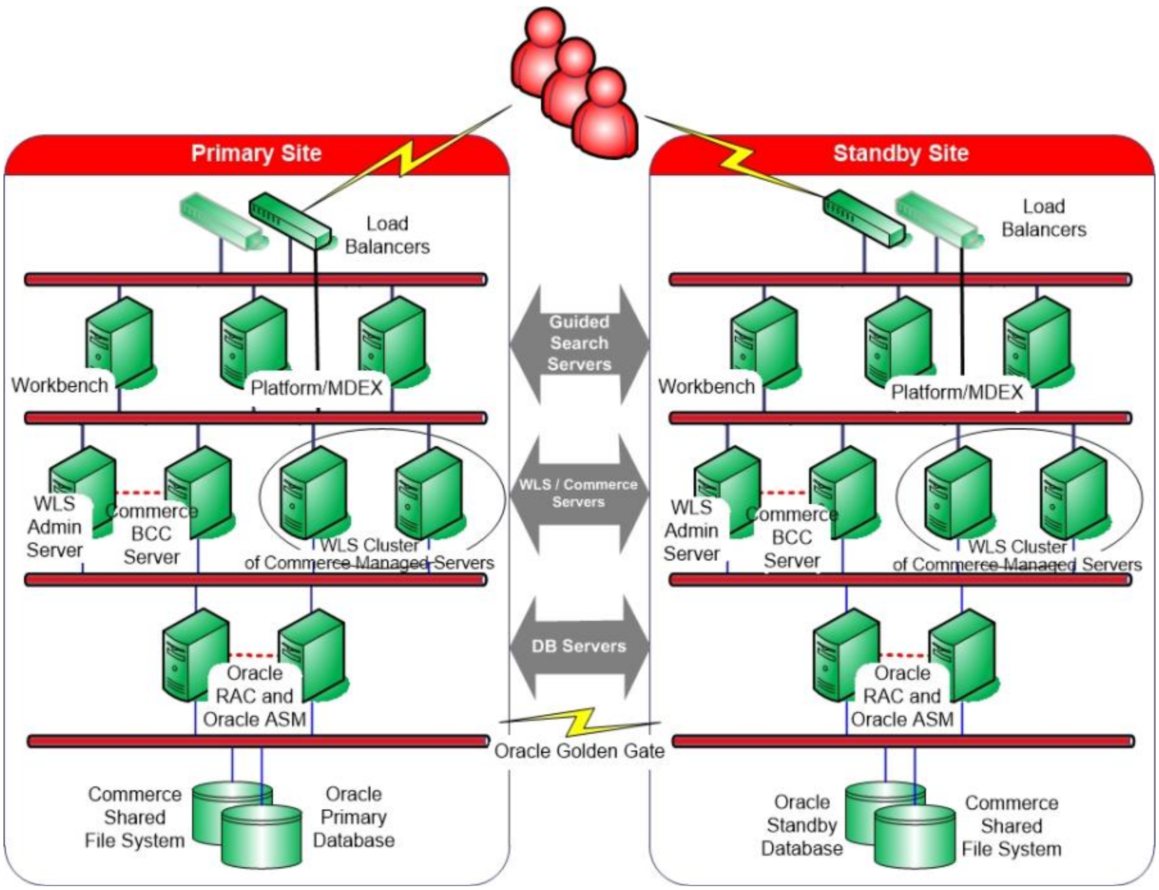

Benefits of Active/Active Configuration

The benefits of an Active/Active configuration are:

» Supports more traffic than an Active/Passive database configuration because more traffic can be routed to each database.

» Zero downtime if there is an application or database failure in either site.

» Each data center acts completely independent of the other, meaning all applications can be Active/Active.

» Data centers can be long distances apart, such as several thousand miles.

» Latency for Oracle GoldenGate replication is not such a large factor, even over great distances.

Drawbacks to Active/Active Configuration

The drawbacks to an Active/Active configuration are:

» Most complicated to set up because it incorporates Oracle GoldenGate.

» Not an out-of-the-box Commerce configuration type.

» Oracle support agreements may not support custom code or configurations if a client develops a new component which does not comply with support guidelines.

» Time is needed to reconfigure business-facing applications in the alternative data center.

» Significant design time and effort may need to be spent to account for the various profile transition use cases that can arise to eliminate data conflicts across the two databases.

» Not specifically endorsed by the MAA Applications team due to the differences in business rules required for the setup of this type of configuration.

Outage Testing and Results

Most Oracle Commerce customers have provisions in place for planned outages. Therefore, the scope of this white paper covers unplanned outages only.

Unplanned Outage Procedures

Unplanned outages are those which occur unexpectedly. Because unplanned outages occur when an Oracle Commerce customer least expects it, provisions should be put in place to avoid any unexpected database or application downtime.

Database Server Outages

With Commerce deployed and layered on the MAA database architecture as shown earlier in Figure 2: Oracle Database MAA, many database related outages have no impact. A summary of outage times with recovery time objectives (RTOs) and their detailed recovery steps are described in the "*Maximize Availability with Oracle Database 12c*" white paper.

Additionally, as mentioned in Oracle Commerce MAA on Exadata and Exalogic earlier in this white paper, there is a discussion of the integration of MAA operational and configuration best practices with Oracle Exadata Database Machine (Exadata MAA) and how it provides the most comprehensive high availability solution available for Oracle Database. The MAA Best Practices white paper Deploying Oracle Maximum Availability Architecture with Exadata Database Machine includes an Exadata MAA Outage and Solution Matrix which includes Exadata specific components that is not be repeated here.

The workload is ramped up and the following outages are simulated:

- Database instance outage – A database instance stops abruptly using the SQL*Plus `shutdown abort` command.
- Database node outage – A database server node stops abruptly using a power reset from the Exadata Integrated Lights Out Manager (ILOM).

Commerce Platform Application Server Outages

The workload is ramped up and the following outage is simulated:

- Instance: One or more Commerce Platform JVMs stop abruptly using the `kill -9` command on the server process.
- Node: This environment contains both VM servers and bare metal servers. A VM stops abruptly using the `xm destroy` command issued from the VM Server (DOM0) or the bare metal servers are brought down with the `poweroff` command.

Commerce Platform Application Server Lock Manager (SLM) Outage

This outage test is performed on one and then on both SLM instances. It shows that there is a brief interruption until the SLM fails over.

- Instance: One or more Commerce Platform JVMs stop abruptly using the `kill -9` command on the server process.

- Node: This environment contains both VM servers and bare metal servers. A VM stops abruptly using the `xm destroy` command issued from the VM Server (DOM0) or the bare metal servers are brought down with the `poweroff` command.

Commerce Oracle Guided Search Server Outages

The workload is ramped up and the following outage is simulated:

- Instance: An MDEX engine process (Dgraph) stops abruptly using the `kill -9` command on the server process.
- Node: This environment contains both VM servers and bare metal servers. A VM stops abruptly using the `xm destroy` command issued from the VM Server (DOM0) or the bare metal servers are brought down with the `poweroff` command.

WebLogic Server Administration Server Outage

This outage is performed on the WebLogic Server administration server. It shows that there is no interruption in the Oracle Commerce application when the WebLogic Server administration server is brought down.

- Instance: The WebLogic Server administration process stops abruptly using the `kill -9` command.
- Node: This environment contains both VM servers and bare metal servers. A VM stops abruptly using the `xm destroy` command issued from the VM Server (DOM0) or the bare metal servers are brought down with the `poweroff` command.

Coherence Server Outage

This outage is performed on one or more Coherence cache servers. It shows how there may be a brief interruption in service as the caching is moved from one cache server to another.

- Instance: One or more Coherence Server JVMs stop abruptly using the `kill -9` command on the server process.
- Node: This environment contains both VM servers and bare metal servers. A VM stops abruptly using the `xm destroy` command issued from the VM Server (DOM0) or the bare metal servers are brought down with the `poweroff` command.

Site Outage

The workload is ramped up and the following outage is simulated:

- Site: All servers on the primary site stop abruptly and site failover is performed.

Unplanned Outage Test Results

The test results for the various outages are summarized in Table 7: Summary of Unplanned Outage Results.

TABLE 6: SUMMARY OF UNPLANNED OUTAGE RESULTS

Component	Outage	Outage (seconds)	User Errors	Performance Degradation	Time to Failover (seconds)
Commerce Platform CRS Server	Instance	0	49%	30%	0
	Node	0	37%	54%	0
WebLogic Server Administration Server	Instance	0	0%	0%	30
	Node	0	0%	0%	30
Coherence Server	Instance	0	0%	0%	30
	Node	0	0%	0%	30
Oracle Guided Search Server	Instance	0	0%	0%	30
	Node	0	0%	0%	100
Database Server	Instance	10 (partial)	1.9%	0%	0
	Node	45 (complete)	2.1%	0%	0
Site	Site-wide	86 (complete)	100%	100%	86
Commerce Platform SLM Server	Instance	0	0%	0%	0

Key: CRS=Commerce Reference Store | JVM=Java Virtual Machine | SLM=Server Lock Manager

Observations on Database Server Outages

The following two outage simulations show what occurs when a database instance is brought down abruptly and when an entire database node is abruptly powered down.

Database Instance Failure

When a database instance fails, connections on that instance fail over to the surviving node and the Commerce Platform continues processing. Figure 12: Database Instance Failure Graph shows that there is only a very slight and short reduction in performance after the outage.

Figure 12: Database Instance Failure Graph

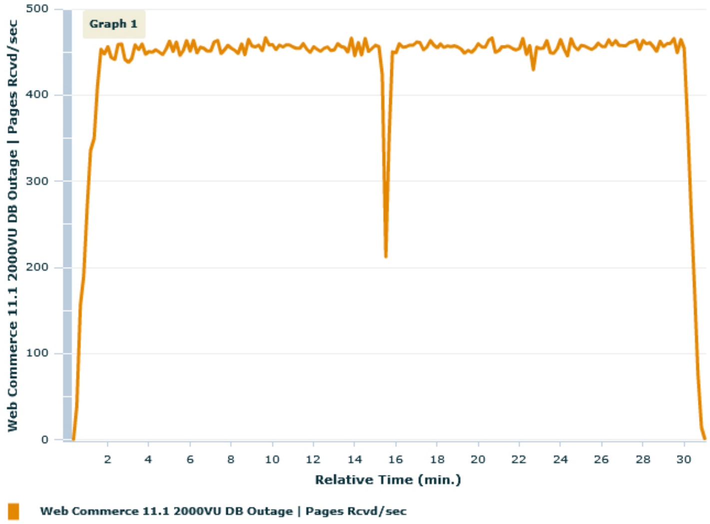

Database Node Failure

Similar to the database instance outage, when the database node fails, connections to that instance fail over to the surviving node and Commerce Platform continues processing. Figure 13: Database Node Failure Chart shows that there is only a very slight and short reduction in performance after the outage.

Figure 13: Database Node Failure Chart

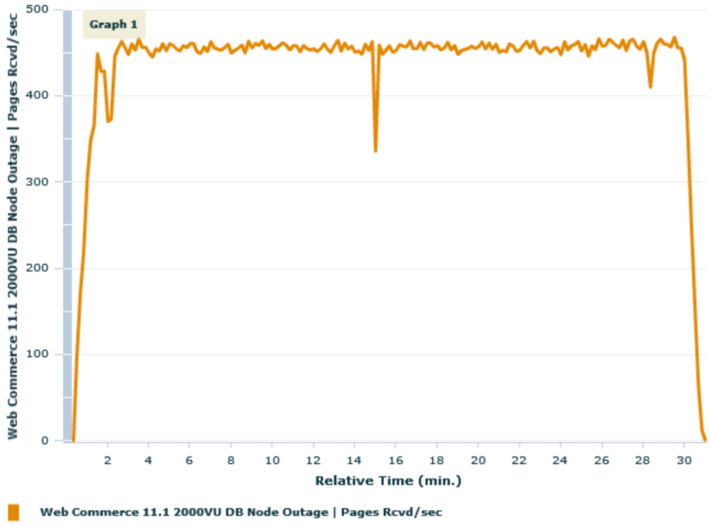

Observations on Unplanned Commerce Platform Server Outages

Failures are minimal because Commerce Platform employs WebLogic Server clustering and Coherence http session clustering to handle session failures. The bulk of failures are due to a quick increase in volume. Figure 14: Commerce Platform CRS Instance Outage Graph shows the impact of killing half of the JVM instances in the cluster. Figure 15: Commerce Platform CRS Node Outage Graph shows the impact of running an `xm destroy` or `poweroff` on two of four server nodes in the Commerce Platform cluster. Figure 16: Commerce Platform SLM Instance Outage Graph shows the impact of killing one and then both of the Commerce Platform SLM instances.

Figure 14: Commerce Platform CRS Instance Outage Graph

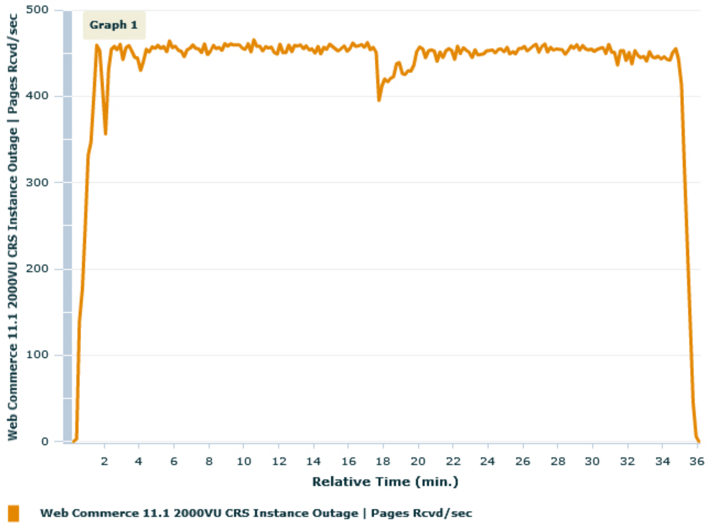

Figure 15: Commerce Platform CRS Node Outage Graph

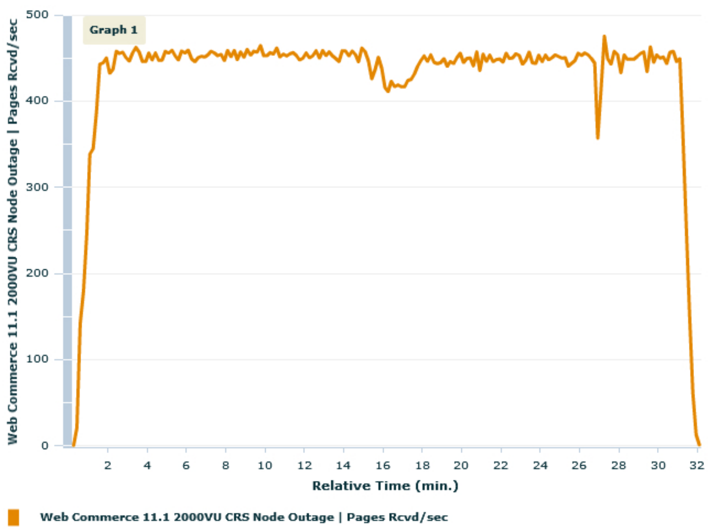

Figure 16: Commerce Platform SLM Instance Outage Graph

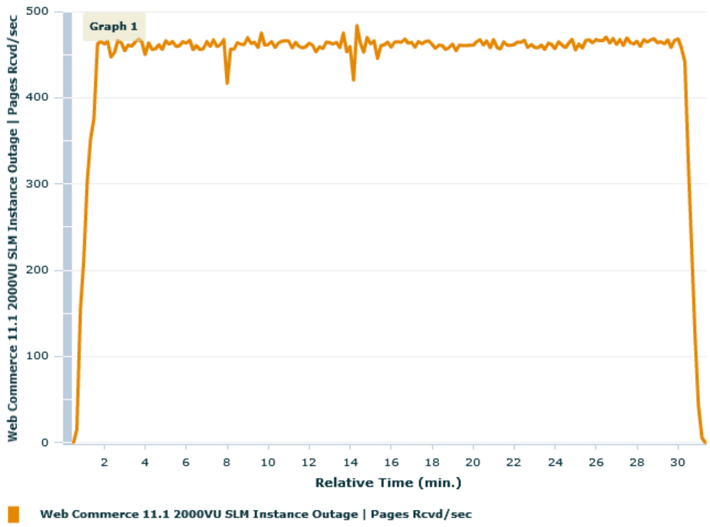

Coherence Server Failure

Similar to the Commerce CRS Outages, two of four Coherence Cache Servers are stopped with the `kill -9` command. Connections to the failed server instances fail over to the surviving nodes and Commerce Platform continues processing. Figure 17: Coherence Server Instance Outage Graph shows that there is no service disruption. HTTP sessions fail over to the other cache servers as expected.

Figure 17: Coherence Server Instance Outage Graph

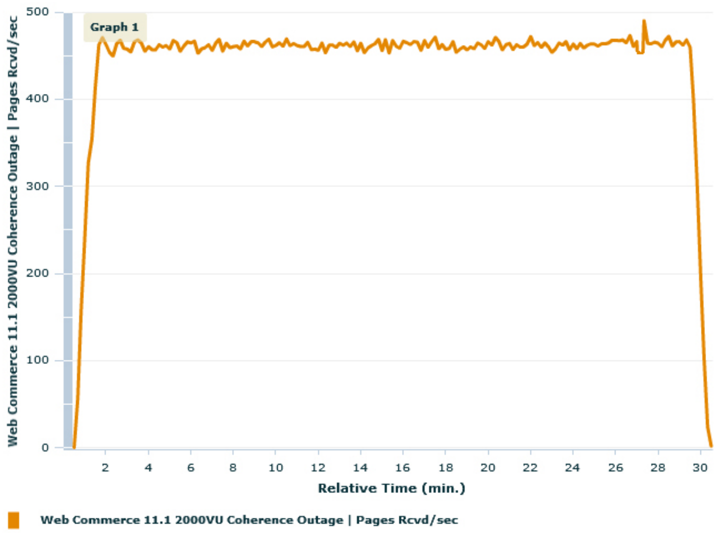

Observations on Oracle Guided Search Outages

The performance remains steady when one of two Oracle Guided Search instances (Dgraph) are killed and when an MDEX node fails. Figure 18: Commerce Oracle Guided Search Outage Graph shows that there is no impact.

Figure 18: Commerce Oracle Guided Search Outage Graph

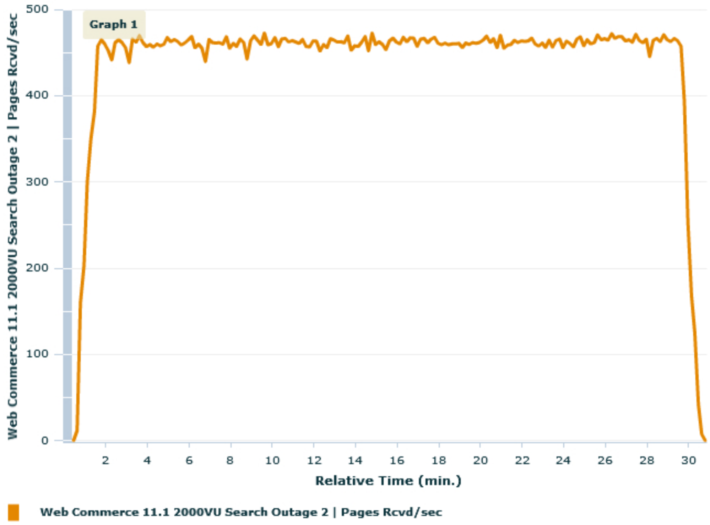

WebLogic Server Administration Server Outage

The failure of the administration server does not have any impact. When the administration server becomes unavailable, the server instances remain up and running. The clustered server instances, load balancing, and failover capabilities supported by the domain configuration remain available even if the administration server fails.

Because a virtual IP is used for the WebLogic Server administration server listener address, the WebLogic Server administration server can be started on any of the WebLogic Servers that have been configured to use the virtual IP address.

Observations on Site Outage

All users are lost on site outage and the site failover procedure is followed to restore service on the standby site. The failover procedure is performed manually and the timings for each step are shown in Table 8: Recovery Times for Site Outage.

TABLE 7: RECOVERY TIMES FOR SITE OUTAGE

Recovery Step	Elapsed Time (secs)	Cumulative Time (secs)	Notes
Database Data Guard Failover	35	35	These two steps performed in parallel.
Commerce Platform File System Replication Role Reversal & Mount	32		
WebLogic Server Administration Server Startup	6	39	N/A
Coherence Server Startup	10	50	Performed in parallel across servers.
Commerce Platform CRS Server Startup	3	42	Performed in parallel across servers.
Experience Manager Server Startup	5	47	Performed in parallel across servers.
First Commerce Platform Login	35	86	Login and select Contacts View.

Each step in the process could have been executed by script with no delay between steps, and in that case the total time to perform recovery, excluding the first user login, is estimated to be 47 seconds.

Summary of Best Practices

The following sections are a summary of the best practices that are presented in this paper, providing a checklist for a Commerce Platform and Merchandising MAA implementation.

Best Practices Commerce Platform Database HA

The following are the Commerce Platform database best practices that should be applied to the primary and secondary site to achieve highest availability:

» Deploy Oracle Commerce on an Oracle RAC database for the highest availability and scalability.

» Use Oracle Automatic Storage Management to simplify the provisioning and management of database storage.

» Enable Oracle Flashback Database to provide the ability to “rewind” the database in the event of user errors.

» Use Oracle Recovery Manager to regularly backup the Oracle Commerce database.

» Always use HugePages for Oracle Commerce databases on Linux. Monitor memory usage and adjust the workload and parameters accordingly.

» Configure database Dead Connection Detection to actively remove dead connections in the event of Oracle Commerce Platform Server node failure.

» For Exadata deployments, configure `cluster misscount` to `30` seconds to reduce downtime in the event of database node failure.

» Revalidate the configuration regularly and especially after changes are made. `exachk` can be used to assist in the validation process for when deployed on Exadata.

Best Practices for Oracle Commerce Application HA

The following are the Web Oracle Commerce Platform application best practices that should be applied to the primary and secondary site to achieve the highest availability:

» Always use HugePages for Web Oracle Commerce servers on Linux. Monitor memory usage and adjust the workload and parameters accordingly.

» Deploy multiple Web Oracle Commerce servers and deploy all critical Oracle Commerce Platform, Merchandising, and Experience Manager components in a load balanced, distributed service, or clustered configuration, so that work can continue in the event of a server node failure.

» Deploy a load balancer in a redundant configuration and load balance server load using the recommended logic.

» Deploy the Oracle Commerce File System on a fault tolerant filer.

» Take regular backups of the Oracle Commerce Servers and Oracle Commerce File System.

» Connect to the database through the role based services using Oracle WebLogic Server Active GridLink for Oracle RAC data sources.

» Reduce TCP Keepalive Timeout on Oracle Commerce Servers.

Best Practices for Disaster Readiness and Recovery

The following are the best practices for deploying a secondary site and recovery procedures in readiness for a site outage:

» Deploy a second geographically separated site that can run the Oracle Commerce workload in the event the primary site is down. The environment configuration type determines the distance between the two sites.

» Use Data Guard to replicate all database changes to a standby database located on the secondary site.

» Take advantage of Oracle Active Data Guard to offload read-only queries to the standby database.

» Enable Oracle Flashback Database so that the old primary database can be quickly reinstated as a standby database in the event of a site failover.

» Replicate the Oracle Commerce File System to the secondary site if using an Active/Passive environment configuration. Develop procedures to reverse the direction of replication in the event of failover or switchover and to clone the replica for site testing.

» Export the Oracle Commerce File System primary, standby replica, and clones with different names to avoid mounting the incorrect one.

» Create different role-based database services for the Oracle Commerce database in primary, standby and snapshot standby mode.

» Develop and document operational procedures in line with the Oracle Commerce MAA state model and state transitions.

» Use Oracle Data Guard Broker to simplify Data Guard administration.

» Use the snapshot standby to provide an updatable replica of the primary database for temporary site testing.

» Use Oracle Golden Gate when replicating between two active sites or data centers.

Glossary

Commerce Platform – formerly known as ATG Web Commerce, Commerce Platform is a framework with which clients can build and develop large-scale B2C or B2B web sites. Commerce Platform offers a complete commerce software platform that enables you to deliver a personalized customer buying experience across all customer touch points, including the web, contact center, mobile devices, social media, physical stores, and more.

Commerce Experience Manager – formerly known as Endeca Commerce, Commerce Experience Manager gives clients the flexibility to set up the selling experience as the clients see fit. The application suite adds search capabilities when used with Oracle Commerce.

Commerce Merchandising – part of the Oracle Commerce Business Control Center, Commerce Merchandising allows a client to create and deploy content directly to its commerce web site in a manner that suits its business rules. Commerce Merchandising uses Oracle Commerce Content Administration to deploy both data-based content to the commerce database instance(s) and file-based content to the file systems with which the commerce applications are configured.

Commerce Service Center – is the commerce customer service application. This is a fully integrated system, which allows CSRs to view and edit orders and profile information. Commerce Service Center is not in the scope of this paper; however, its implementation is similar to that of the base commerce web site.

References

1. Oracle Maximum Availability Architecture Web site
2. Exalogic & Exadata: The Optimal Platform for Oracle Commerce
3. Deploying Oracle Maximum Availability Architecture with Exadata Database Machine (technical white paper)
4. Oracle Exalogic Elastic Cloud Administrator's Guide, Release EL X2-2, X3-2, X4-2
5. Oracle Fusion Middleware High Availability Guide, 12c Release 1 (12.1.2)
6. Oracle Fusion Middleware Administering Clusters for Oracle WebLogic Server, 12c Release 1 (12.1.2)
7. Oracle Commerce Platform Installation and Configuration Guide, Release 11.1
8. Oracle Commerce Reference Store Installation and Configuration Guide, Release 11.1
9. Oracle Commerce Reference Store Overview, Release 11.1
10. Oracle Commerce Guided Search Common Documentation, Release 1.1
11. Oracle Commerce Guided Search MDEX Engine, Release 6.5.1
12. Oracle Commerce Guided Search Content Acquisition System (CAS), Release 11.1
13. Oracle Commerce Guided Search with Tools and Frameworks, Release 11.1
14. Oracle Commerce Guided Search Platform Services, Release 11.1
15. MOS 1908576.1 - Oracle Commerce Supported Environments Matrix
16. MOS 1461463.1 – Oracle ATG Web Commerce Reference Architecture
17. Oracle WebLogic Server and Highly Available Oracle Databases: Oracle Integrated Maximum Availability Solutions
18. Oracle Fusion Middleware Administering JDBC Data Sources for Oracle WebLogic Server, 12c Release 1 (12.1.2)
19. Oracle Fusion Middleware Disaster Recovery Solution using Oracle's Sun ZFS Storage Appliance
20. Disaster Recovery for Oracle Exalogic Elastic Cloud with Oracle Exadata Database Machine
21. Automating Disaster Recovery using Oracle Site Guard for Oracle Exalogic with Oracle Exadata Database Machine

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/oracle

facebook.com/oracle

twitter.com/oracle

oracle.com

Hardware and Software, Engineered to Work Together