

Disaster Recovery for Oracle
Exalogic Elastic Cloud

WITH ORACLE EXADATA DATABASE MACHINE

Oracle Maximum Availability Architecture White Paper
May 2011

Maximum Availability Architecture

Oracle Best Practices For High Availability

Executive Overview	1
Introduction	2
Audience	4
Oracle Fusion Middleware Disaster Recovery Strategy	5
Disaster Recovery Considerations and Terminology	5
Disaster Recovery Architecture	11
Topology	11
Hardware.....	13
Software	15
Network.....	15
Load Balancers	17
Prerequisites	18
Storage Configuration.....	18
Creating Projects and Shares on the Shared Storage	19
Configuring the Storage Replication Channel	20
Configuring Remote Replication Targets	23
Host Setup.....	24
Site Setup and Configuration.....	31
Production Site Setup.....	32
Configuring Replication for the Projects and Shares.....	34
Standby Site Instantiation	36
Validate the Standby Site Setup	38
Disaster Recovery Operations.....	39

Site Switchover	39
Site Switchback	40
Site Failover	41
Oracle MAA Best Practices for Disaster Recovery	41
Appendix	44
Disaster Recovery Terminology	44
Sun ZFS Storage 7320 Operations	44
Oracle Data Guard Setup	45
Storage Scripts	46
References	49

Executive Overview

Oracle Maximum Availability Architecture (MAA) [1] is the Oracle best practices blueprint for implementing Oracle high availability technologies. Maximum Availability Architecture is one of the key requirements for any Oracle Fusion Middleware enterprise deployment. Oracle Fusion Middleware includes an extensive set of high availability features such as: process death detection and restart, server clustering, server migration, clusterware integration, GridLink, load balancing, failover, backup and recovery, rolling upgrades, and rolling configuration changes, which protect an Enterprise Deployment from unplanned down time and minimize planned downtime.

Additionally, enterprise deployments need protection from unforeseen disasters and natural calamities. The typical protection solution involves setting up a standby site at a geographically different location than the production site. The standby site may have equal or fewer services and resources compared to the production site. Application data, metadata, configuration data, and security data are replicated to the standby site on a periodic or continual basis. The standby site is normally in a passive mode; it is started when the production site is not available. This deployment model is sometimes referred to as an active/passive model.

The Oracle Fusion Middleware Disaster Recovery solution uses storage replication technology for disaster protection of Oracle Fusion Middleware middle tier components. It supports hot-pluggable deployments, and it is compatible with third-party vendor recommended solutions. Additionally, Oracle Data Guard is used to provide disaster recovery for Oracle databases that are part of Oracle Fusion Middleware deployments.

The Disaster Recovery Solution for the Oracle Exalogic Machine and Oracle Exadata Database Machine builds upon these well-established disaster protection solutions for Oracle Fusion Middleware and Oracle databases. While this paper describes the disaster

recovery solution for Oracle Exalogic and Oracle Exadata Database Machine deployments, the principles described here also apply to deployments on an Oracle Exalogic Machine with an Oracle Database and to standalone deployments on an Oracle Exalogic Machine.

Introduction

Oracle Exalogic Elastic Cloud is an integrated hardware and software system designed to provide a complete platform for a wide range of application types and widely varied workloads. Oracle Exalogic is intended for large-scale, performance-sensitive, mission-critical application deployments. It combines Oracle Fusion Middleware software and industry-standard Sun hardware to enable a high degree of isolation between concurrently deployed applications, which have varied security, reliability, and performance requirements.

The Oracle Exalogic Elastic Cloud consists of Sun Fire X4170 M2 Servers as compute nodes, a Sun ZFS Storage 7320 appliance and required InfiniBand and Ethernet networking components. The Sun ZFS Storage 7320 appliance combines multiple protocol connectivity, data services for business continuity, and ease of management into a single appliance. The appliance supports NFS, Common Internet File System (CIFS), Internet Small Computer System Interface (iSCSI), InfiniBand (IB), and Fibre Channel (FC) protocols for the data access. The appliance also supports Network Data Management Protocol (NDMP) for backing up and restoring the data.

The Storage disks on the Sun ZFS Storage 7320 appliance are allocated to a single storage pool. The storage pool is assigned to one of the server heads, also referred to as storage controllers. The server heads are configured in an active-passive cluster configuration. The compute nodes in the Oracle Exalogic Machine access the storage using the NFSv3 protocol over the InfiniBand network.

The Oracle Exadata Database Machine is an easy to deploy solution for hosting the Oracle Database, which delivers the highest levels of database performance available. The Exadata Database Machine is a “cloud in a box” composed of database servers, Oracle Exadata Storage Servers, an InfiniBand fabric for storage networking and all the other components required to host an Oracle Database.

The Oracle Exadata Database Machine provides an optimal solution for all database workloads, ranging from scan-intensive data warehouse applications to highly concurrent online transaction processing (OLTP) applications. With its combination of smart Oracle Exadata Storage Server Software, complete and intelligent Oracle Database software, and the latest industry-standard hardware components, Oracle Exadata Database Machine delivers extreme performance in a highly-available, highly-secure environment.

The goal of this technical paper is to provide:

- Oracle Fusion Middleware Disaster Recovery architecture and strategy for deployments on Oracle Exalogic with Oracle Exadata Database Machine
- Detailed deployment and configuration steps for the Oracle Fusion Middleware Disaster Recovery solution on Oracle Exalogic and Oracle Exadata Database Machine. The solution described in this paper applies to non Exadata based database deployments as well.
- Best practices for the Oracle Fusion Middleware Disaster Recovery solution with Exalogic and Exadata

Audience

This document is intended for Oracle Fusion Middleware administrators, storage-system administrators, Oracle database administrators and technical sales personnel. It is assumed that the reader is familiar with Oracle Exalogic Elastic Cloud, Oracle Exadata Database Machine, Oracle Fusion Middleware components, Oracle database concepts and Oracle Data Guard features. For details, please refer to the documents listed in the Reference section.

Oracle Fusion Middleware Disaster Recovery Strategy

Oracle Fusion Middleware product binaries, configuration and applications are deployed to the Oracle home and domain home directories. The Oracle Fusion Middleware home directories and the domain directories are stored on shared storage. The metadata and the run-time data are stored in a database repository.

The Oracle Fusion Middleware Disaster Recovery strategy facilitates data protection as follows:

- The remote replication feature of the Sun ZFS Storage 7320 appliance protects the middleware product binaries, configurations, metadata files and application data that reside on the file system.
- Oracle Data Guard protects the Oracle Database. This database contains Oracle Fusion Middleware Repository data, as well as customer data.

The clients access the production site during normal operation. During Disaster Recovery, clients access the standby site. The change is almost seamless from the client's perspective since the entire Fusion Middleware infrastructure along with the mount points and host names are configured identically on both the production and standby sites.

Disaster Recovery Considerations and Terminology

This section provides considerations for and defines the terminology for Disaster Recovery

Site Considerations

Symmetric Site:

An Oracle Fusion Middleware Disaster Recovery configuration that is completely identical across tiers on the production site and standby site is called a **symmetric** site.

A site can be **completely symmetric** or **partially symmetric**.

In a completely symmetric site the production site and standby site are identical in all respects. That is, they have identical Exalogic and Exadata hardware, load balancers, middleware instances, applications and databases.. The same port numbers are used for both sites.

In a partially symmetric site the production site and standby site are identical in topology but not hardware. That is, they have the same number of middleware instances, applications and databases on each site but the Exadata and Exalogic hardware is not identical.

It is recommended but not required to have identical Exalogic and Exadata hardware on the production and standby sites when planning a Disaster Recovery Site.

For example, you can have a full rack of Exalogic and Exadata on the production site and a half rack of Exalogic and Exadata on the standby site to create a symmetric disaster recovery topology.

This white paper describes the setup and configuration for a symmetric disaster recovery site in terms of both hardware and topology.

Asymmetric Site

An **asymmetric** topology is a disaster recovery configuration that is different across tiers at the production site and standby site.

In an asymmetric topology, the standby site has fewer resources than the production site. Typically, the standby site in an asymmetric topology has fewer hosts, load balancers, Fusion Middleware instances, and applications than the production site.

The number of database instances on the standby site must match those at the production site, but they can be Non-RAC database instances.

Many of the concepts for setting up a symmetric topology are also valid for setting up an asymmetric topology.

It is important to ensure that an asymmetric standby site has sufficient resources to provide adequate performance when it assumes the production role.

Please refer to the Oracle Fusion Middleware Disaster Recovery Guide for the steps to setup an asymmetric site.

Storage Considerations and Terminology

This section provides an overview of the terminology and storage concepts for the Sun ZFS Storage 7320 appliance. This appliance is a part of every Oracle Exalogic Machine.

Pools, Projects and Shares

Storage Pool

The storage pool (similar to a volume group) is created over a set of physical disks. File systems are then created over the storage pool. The storage pool is configured with a RAID layout such as mirrored, RAID-Z (single parity), or RAID-Z2 (dual parity).

In an Exalogic Machine, all the physical disks are **mirrored** and allocated to a **single storage pool**. This is the **default configuration** for an Exalogic Machine.

Projects

All file systems and LUNs are grouped into projects. A project can be considered a **consistency group**. A project defines a common administrative control point for managing shares. All shares within a project can share common settings, and quotas can be enforced at the project level in addition to the share level. Projects can also be used solely for grouping logically related shares together, so their common attributes (such as accumulated space) can be accessed from a single point.

Shares

Shares are file systems and LUNs that are exported over supported data protocols to clients of the appliance. Exported file systems can be accessed over CIFS, NFS, HTTP/WebDav, and FTP. LUNs export block-based volumes and can be accessed over iSCSI. The project/share is a unique identifier for a share within a pool. Multiple projects can contain shares with the same name, but a single project cannot contain shares with the same name.

Oracle strongly recommends that the compute nodes from an Exalogic Machine access the shares/projects over NFS. The shares are mounted using NFS over IPoIB (IP over InfiniBand).

ZFS Replication

ZFS replication is a method where two storage systems are replicated with a lag, such as a write is considered complete as soon as local storage acknowledges it. The remote storage is usually updated with a small lag. This has the advantage of being able to process writes much faster at the primary location, because the system does not have to wait for data to be saved at the replication site. This is usually implemented using snapshots; a snapshot of the current state of the master system is replicated to the secondary storage system. Depending on the configuration used, the process is repeated as soon as the snapshot is replicated, or it is triggered at certain times.

The main advantage of this technique is that it allows for replication over far larger distances, because the link between the storage systems can have a lower bandwidth (not every write has to be replicated; only the state of the system at certain points in time) and higher latency (because writes don't need to be confirmed at both sites at once). The obvious disadvantage is that in case of a failure on the primary system, data loss is guaranteed. The secondary system will always be missing data that has been written to the master. Performance is greatly increased, but if local storage is lost, the remote storage is not guaranteed to have a current copy of the data and most recent data may be lost.

The Sun ZFS Storage 7320 appliance in the Exalogic Machine supports snapshot-based replication of projects and shares from a source appliance to any number of target appliances in three modes: scheduled, on-demand or continuous. The replication includes both data and metadata.

- **Scheduled Replication**

In this mode, the user can define a schedule for automatic replication. If a schedule is established, then the replication occurs at the defined interval. The interval can be every half-hour, hour, day, week, or month. This mode is preferred in situations where replication during off-peak time is preferred or where backup is scheduled at the target site at specific time.

- **On-demand Replication**

In this mode, also called as a manual mode, the replication occurs only when the user requests. This is the default mode when the scheduled mode is chosen but no schedule is defined.

- **Continuous Replication**

In this mode, the replication process happens continuously without any user intervention. As soon as the package successfully arrives at the target, the subsequent replication process automatically begins. This mode is deployed where the target site is expected to be almost in sync with the source.

Project Level Replication vs. Share Level Replication

The Sun ZFS Storage 7320 appliance in the Exalogic Machine allows remote replication to be configured on both the project and share level.

By default, the shares in a project inherit the configuration of the parent project. Inheriting the configuration not only means that the share is replicated on the same schedule to the same target with the same options as its parent project, but also that the share is replicated in the same stream using the same project-level snapshots as other shares inheriting the project's configuration. This is important for applications that require data consistency among multiple shares.

Overriding the configuration means that the shares are not replicated with any project-level actions, though it may be replicated with its own share-level actions that include the project. It is not possible to override part of the project's replication configuration and inherit the rest.

More precisely, the replication configuration of a project and its shares define some number of replication groups, each of which is replicated with a single stream using snapshots taken simultaneously. All groups contain the project itself (which essentially just includes its properties). One project-level group includes all shares inheriting the replication configuration of the parent project. Any shares which override the project's configuration form a new group consisting of only the project and shares themselves.

Oracle strongly recommends that project-level and share-level replication be avoided within the same project because it can lead to surprising results (particularly when reversing the direction of replication).

Storage Replication Channel

A storage replication channel is a network channel that is dedicated specifically to replication traffic between the Sun ZFS Storage 7320 appliances at the production site and the standby site.

The Sun ZFS Storage 7320 appliance within Exalogic has four 1 Gigabit Ethernet ports (igb0, igb1, igb2 and igb3). In the current configuration two ports (igb0 and igb1) are used for managing the Sun ZFS Storage 7320 appliance.

Oracle recommends creating the replication channel by connecting the unused ports (igb2 and igb3) to the corporate network in your datacenter and creating a bonded interface using IP network MultiPathing (IPMP). This bonded interface is dedicated for replication traffic and provides high availability for the storage replication channel.

Oracle strongly recommends:

- Connecting port igb2 to the embedded Cisco Catalyst 4948 switch in the Exalogic Machine
- Connecting port igb3 directly to a corporate network drop in your data center.
- Connecting the ports two different switches provides switch level high availability.
- Configuring the two switches, to which the ports are connected, on the same VLAN.
- Disabling the management options on the two ports used for replication. This ensures the separation of the replication traffic from the management traffic. This is also an Oracle Security best practices recommendation.

Host Names

Host names play a key role in any topology. In a Disaster Recovery topology, the hostnames used for wiring intra component and inter component communication need to be same. Typically, the site where the Oracle Fusion Middleware install is done first dictates the hostname used. The standby site instantiated subsequently should be configured to resolve these hostnames to the local standby site IP addresses. Therefore, it is important to plan the host names for the production site and standby site. It is also very important that the configuration at all levels use only hostnames. That is local IP addresses must not be used.

This paper assumes that a symmetric Disaster Recovery site is being set up, where the production site and standby site have the same number of hosts. Each host at the production site has a peer host at the standby site. The peer hosts are configured the same. For example, hosts at one site use the same port numbers as their counterparts at the other site.

When configuring each component, use hostname-based configuration instead of IP-based configuration, unless the component requires you to use IP-based configuration. For example, if you are configuring the listen address of an Oracle Fusion Middleware component to a specific IP address such as 192.168.10.33, use the host name wlsvhn1.mycompany.com, which resolves to 192.168.10.33.

Oracle Data Guard

Oracle Data Guard is Oracle's disaster recovery solution prescribed by the Maximum Availability Architecture (MAA) to protect mission critical databases residing on Exadata Database Machine. Data Guard is also used to maintain availability should any outage unexpectedly impact the production database and to minimize downtime during planned maintenance. Data Guard 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. Data Guard maintains these standby databases as copies of the production database. Then, if the production database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the production role, minimizing the downtime associated with the outage. Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability.

Oracle Active Data Guard

Oracle Active Data Guard, an option built on the infrastructure of Oracle Data Guard, allows a physical standby database to be open read-only while changes are applied to it from the primary database. This enables read-only applications to use the physical standby with minimal latency between the data on the standby database and that on the primary database, even while processing very high transaction volumes at the primary database. This is sometimes referred to as real-time query.

An Oracle Active Data Guard standby database is used for automatic repair of data corruption detected by the primary database, transparent to the application. In the event of an unplanned outage on the primary database, high availability is maintained by quickly failing over to the standby database. An Active Data Guard standby database can also be used to off-load fast incremental backups from the primary database because it is a block-for-block physical replica of the primary.

Currently, Oracle Fusion Middleware does not support configuring Oracle Active Data Guard for the database repositories that are a part of the Fusion Middleware topology. However, Active Data Guard can be configured if your custom applications are designed to leverage the technology.

Data Guard Protection Modes

Maximum Availability

This protection mode provides the highest level of data protection that is possible without compromising the availability of a primary database. Transactions do not commit until all redo data needed to recover those transactions has been written to the online redo log and to the standby redo log on at least one synchronized standby database. If the primary database cannot write its redo stream to at least one synchronized standby database, it operates as if it were in

maximum performance mode to preserve primary database availability until it is again able to write its redo stream to a synchronized standby database.

This mode ensures that no data loss occurs if the primary database fails, but only if a second fault does not prevent a complete set of redo-data from being sent from the primary database to at least one standby database.

Maximum Performance

This protection mode provides the highest level of data protection that is possible without affecting the performance of a primary database. This is accomplished by allowing transactions to commit as soon as all redo data generated by those transactions has been written to the online log. Redo data is also written to one or more standby databases, but this is done asynchronously with respect to transaction commitment, so primary database performance is unaffected by delays in writing redo data to the standby database(s).

This protection mode offers slightly less data protection than maximum availability mode and has minimal impact on primary database performance.

This is the default protection mode for a database.

Maximum Protection

This protection mode ensures that no data loss occurs if the primary database fails. To provide this level of protection, the redo data needed to recover a transaction must be written to both the online redo log and to the standby redo log on at least one synchronized standby database before the transaction commits. To ensure that data loss cannot occur, the primary database shuts down, rather than continuing to process transactions, if it cannot write its redo stream to at least one synchronized standby database.

Because this data protection mode prioritizes data protection over primary database availability, Oracle recommends that a minimum of two standby databases be used to protect a primary database that runs in maximum protection mode to prevent a single standby database failure from causing the primary database to shut down.

Disaster Recovery Architecture

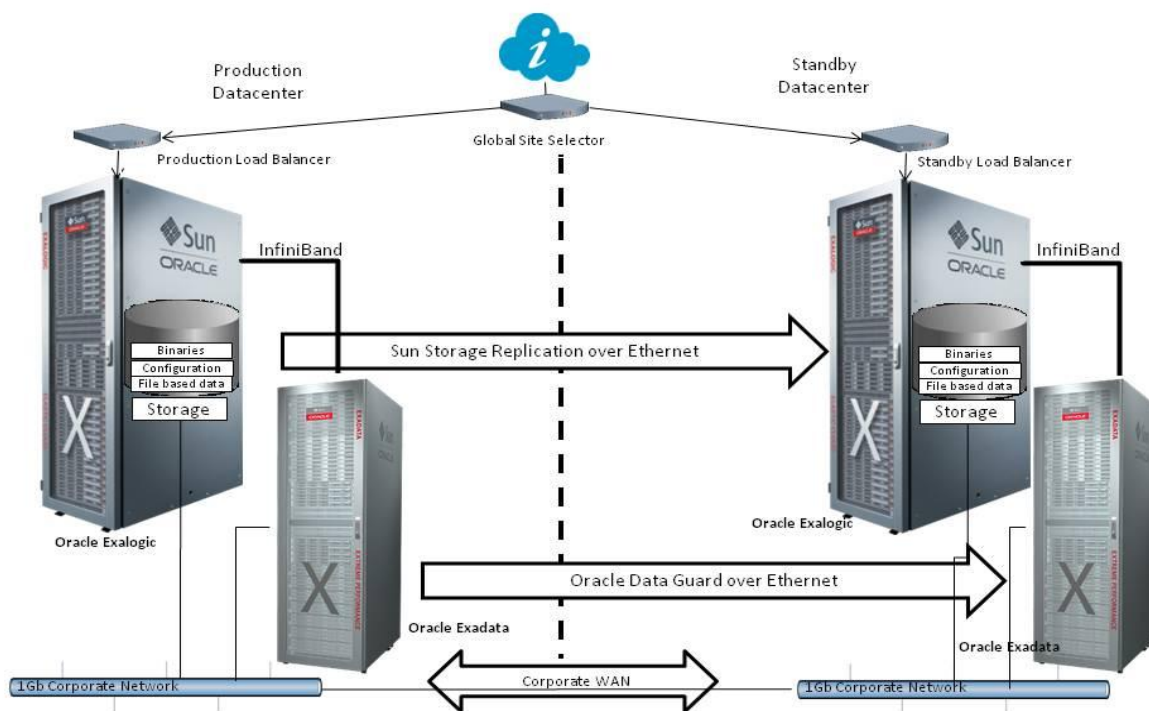
Topology

The Enterprise Deployment Topology described in the Oracle Fusion Middleware Exalogic Enterprise Deployment Guide was used as the reference architecture in this paper. An Enterprise deployment is an is a reference configuration that is designed to support large-scale, mission-critical business software applications and is an Oracle best practices blueprint based on proven Oracle high-availability and security technologies and recommendations for Oracle Exalogic. For

the purpose of testing the Oracle Fusion Middleware Disaster Recovery solution, the setup was scaled down as shown below. Each site consists of:

1. Two webhosts outside Exalogic running the Oracle HTTP Server. (webhost1, webhost2)
2. One Oracle Weblogic Administration Server running on Oracle Exalogic. (apphost1)
3. Four Weblogic Managed Servers running on Oracle Exalogic. (apphost1, apphost2)
4. Four Coherence Servers running Oracle Exalogic. (apphost1, apphost2)
5. One two node RAC database running on an Oracle Exadata Database Machine for application data. (dbhost1, dbhost2)
6. The binaries and configuration files for the Oracle Weblogic Server and Oracle HTTP server were installed on the shared storage in the Exalogic Machine.

The figure below shows a typical Disaster Recovery topology for a deployment on an Oracle Exalogic Machine and Oracle Exadata Database Machine.



Hardware

This section describes the hardware used in this paper.

Oracle Exalogic Elastic Cloud X2-2

Compute Nodes

For of this paper, two compute nodes on a Oracle Exalogic Machine at each site were used for the application server.

Production Site

COMPUTE NODE NAME	MANAGEMENT IP ADDRESS	INFINIBAND IP ADDRESS	INFINIBAND HOST NAME
apphostprod1.mycompany.com	10.204.80.40	192.168.10.1	apphost1-priv
apphostprod2.mycompany.com	10.204.80.41	192.168.10.2	apphost2-priv

Standby Site

COMPUTE NODE NAME	MANAGEMENT IP ADDRESS	INFINIBAND IP ADDRESS	INFINIBAND HOST NAME
apphoststby1.mycompany.com	10.132.251.40	192.168.10.1	apphost1-priv
apphoststby2.mycompany.com	10.132.251.41	192.168.10.2	apphost2-priv

Sun ZFS Storage 7320 Appliance

This is the storage appliance that is a part of the Exalogic Machine. The binaries and configuration files for the all middleware components are located on the shared storage and mounted by all the compute nodes on the hosts. Replication is set up between these two storage heads.

Production Site

STORAGE NODE NAME	MANAGEMENT IP ADDRESS	INFINIBAND IP ADDRESS	INFINIBAND HOST NAME
prodstor1.mycompany.com	10.204.80.100	192.168.10.15	prodstor-ib1
prodstor2.mycompany.com	10.204.80.101	192.168.10.16	prodstor-ib2

Standby Site

STORAGE NODE NAME	MANAGEMENT IP ADDRESS	INFINIBAND IP ADDRESS	INFINIBAND HOST NAME
stbystor1.mycompany.com	10.132.251.100	192.168.10.15	stbystor-ib1
stbystor2.mycompany.com	10.132.251.101	192.168.10.16	stbystor-ib2

WEBHOSTS

For this paper two standalone hosts were used at each site to run the Oracle HTTP Server. . If required, you can also run the Oracle HTTP server on two separate compute nodes in the Exalogic Machine.

Production Site

NODE NAME	IP ADDRESS	REMARKS
webhostprod1.mycompany.com	10.132.131.100	Webhost 1
webhostprod2.mycompany.com	10.132.131.101	Webhost 2

Standby Site

COMPUTE NODE NAME	MANAGEMENT IP ADDRESS	REMARKS
webhoststby1.mycompany.com	10.200.130.100	Webhost 1
webhoststby2.mycompany.com	10.200.130.101	Webhost 2

Oracle Exadata Database Machine X2-2

An Exadata Database Machine in a Quarter rack configuration on each site was used for this paper. A Quarter rack of Exadata consists of two database nodes and three storage cells.

Production Site

COMPUTE NODE NAME	MANAGEMENT IP ADDRESS	INFINIBAND IP ADDRESS	INFINIBAND HOST NAME
proddbhost01.mycompany.com	10.204.80.200	192.168.10.100	dhost01-priv
proddbhost02.mycompany.com	10.204.80.201	192.168.10.101	dbhost02-priv
prodcel01.mycompany.com	10.204.80.202	192.168.10.102	cel01-priv
prodcel02.mycompany.com	10.204.80.202	192.168.10.103	cel02-priv
prodcel03.mycompany.com	10.204.80.203	192.168.10.104	cel03-priv

Standby Site

COMPUTE NODE NAME	MANAGEMENT IP ADDRESS	INFINIBAND IP ADDRESS	INFINIBAND HOST NAME
stbydbhost01.mycompany.com	10.132.251.200	192.168.10.100	dbhost01-priv
stbydbhost02.mycompany.com	10.132.251.201	192.168.10.101	dbhost02-priv
stbycel01.mycompany.com	10.132.251.202	192.168.10.102	cel01-priv
stbycel02.mycompany.com	10.132.251.202	192.168.10.103	cel02-priv
stbycel03.mycompany.com	10.132.251.203	192.168.10.104	cel03-priv

Software

- Oracle HTTP Server 11.1.1.3
- Oracle WebLogic Server 10.3.4
- Oracle Database Enterprise Edition 11.2.0.2

Network

The Exalogic Machine has three distinct networks, i.e. the Management Network, the private InfiniBand network and the client access network (Ethernet over InfiniBand). For this paper, the following Virtual IP addresses were configured on the Private InfiniBand network and the Client Access Network.

The WebLogic servers in the topology listen on the virtual IP addresses configured on the Private InfiniBand network. All client traffic in and out of the Exalogic Machine goes over the Client Access Network

Virtual IP Addresses

Private InfiniBand Network

Production Site

VIRTUAL HOST NAME	INFINIBAND IP ADDRESS	COMMENT
adminvhn.mycompany.com	192.168.10.50	Admin Server Listen Address
wlsvhn1.mycompany.com	192.168.10.51	Managed Server 1 Listen Address
wlsvhn2.mycompany.com	192.168.10.52	Managed Server 2 Listen Address
wlsvhn3.mycompany.com	192.168.10.53	Managed Server 3 Listen Address
wlsvhn4.mycompany.com	192.168.10.54	Managed Server 4 Listen Address
dbhost01-ibvip.mycompany.com	192.168.10.110	Database Host1 VIP on InfiniBand Network
dbhost02-ibvip.mycompany.com	192.168.10.111	Database Host2 VIP on InfiniBand Network

Standby Site

VIRTUAL HOST NAME	INFINIBAND IP ADDRESS	COMMENT
adminvhn.mycompany.com	192.168.10.70	Admin Server Listen Address
wlsvhn1.mycompany.com	192.168.10.71	Managed Server 1 Listen Address
wlsvhn2.mycompany.com	192.168.10.72	Managed Server 2 Listen Address
wlsvhn3.mycompany.com	192.168.10.73	Managed Server 3 Listen Address
wlsvhn4.mycompany.com	192.168.10.74	Managed Server 4 Listen Address
dbhost01-ibvip.mycompany.com	192.168.10.110	Database Host1 VIP on InfiniBand Network

dbhost02-ibvip.mycompany.com	192.168.10.111	Database Host2 VIP on InfiniBand Network
------------------------------	----------------	--

Client Access Network

Production Site

VIRTUAL HOST NAME	CLIENT ACCESS IP ADDRESS	COMMENT
adminvhn-prod.mycompany.com	10.204.77.50	Admin Server Network Channel Address
wlsvhn1-prod.mycompany.com	10.204.77.51	Managed Server 1 Network Channel Address
wlsvhn2-prod.mycompany.com	10.204.77.52	Managed Server 2 Network Channel Address
wlsvnh3-prod.mycompany.com	10.204.77.53	Managed Server 3 Network Channel Address
wlsvhn4-prod.mycompany.com	10.204.77.54	Managed Server 4 Network Channel Address
proddbhost01-vip.mycompany.com	10.224.77.100	Database Host1 VIP
proddbhost02-vip.mycompany.com	10.224.77.101	Database Host2 VIP
proddb-scan.mycompany.com	10.224.77.101 10.224.77.102 10.224.77.103	Database SCAN address

Standby Site

VIRTUAL HOST NAME	CLIENT ACCESS IP ADDRESS	COMMENT
adminvhn-stby.mycompany.com	10.250.77.70	Admin Server Network Channel Address
wlsvhn1-stby.mycompany.com	10.250.77.71	Managed Server 1 Network Channel Address
wlsvhn2-stby.mycompany.com	10.250.77.72	Managed Server 2 Network Channel Address
wlsvnh3-stby.mycompany.com	10.250.77.73	Managed Server 3 Network Channel Address
wlsvhn4-stby.mycompany.com	10.250.77.74	Managed Server 4 Network Channel Address
stbydbhost01-vip.mycompany.com	10.250.77.100	Database Host1 VIP
stbydbhost02-vip.mycompany.com	10.250.77.101	Database Host2 VIP
stbydb-scan.mycompany.com	10.260.77.101 10.260.77.102 10.260.77.103	Database SCAN address

Storage Replication Channel

The replication channel IP address is used for replication traffic and also for mounting the shares for the Oracle HTTP Server on the webhosts at the production and standby sites.

Production Site

HOST NAME	IP ADDRESS	COMMENT
prod_repl.mycompany.com	10.204.77.120	Production Site Replication Channel

Standby Site

HOST NAME	IP ADDRESS	COMMENT
stby_repl.mycompany.com	10.200.80.120	Standby Site Replication Channel

Load Balancers

The following virtual IP addresses were configured on the load balancer for this paper

VIRTUAL HOST NAME	IP ADDRESS	COMMENT
exalogic.mycompany.com	144.25.145.54	VIP for Client Traffic
admin.mycompany.com	144.25.145.50	VIP for WLS Administration Traffic

Prerequisites

Storage Configuration

For this paper the default project and share layout described in Oracle Exalogic Enterprise Deployment Guide were followed with the following exceptions:

1. A project was created for the Oracle HTTP Server binaries and instance configuration
2. A separate share was created for the application log data
3. Redundant Middleware Homes were used for product binaries

Depending on the Oracle Fusion Middleware Components, the applications, the access policies, the replication groups and other requirements alternative layouts are possible as well. The best practice recommendations guidelines for creating projects and shares are detailed below:

Web Tier

1. Create a Project for the Oracle HTTP server product binaries and configuration. For example: **OHS**. The shares in the project do not inherit the properties from the parent project.
2. Create a share for each webhost. Each share has two files systems, one for the Oracle Home containing product binaries, the other for Oracle Instance.

Application Tier

1. Create a Project for the Oracle product binaries in the application tier. For example: **MW Binaries**. All shares in the project inherit all the properties from the project.
2. Create (n+ 1) shares, where n is the number of shares, under this project. Each of the shares will be used for a Middleware Home that contains product binaries.
3. Using two different shares for redundant Middleware Homes is a MAA best practice recommendation and provides, Maximum Availability, zero downtime rolling patching and upgrades and isolates failures on the shares.
4. Additional servers (when scaling out or up) of the same type can use either one of these two locations without requiring more installations.
5. Create a Project for the configuration files and data. For example: **Configuration**. The shares under this project do not inherit all the properties of parent project.
6. Create a share for the Administration Server domain home.

7. Create a share for each compute node in your topology. This share contains two file systems, one for the Node Manager configuration and one for the Managed Server domain home. The Managed Server domain home is shared by all the managed servers running on that compute node.
8. Create a share for all the log files.

Creating Projects and Shares on the Shared Storage

Refer to the Oracle Exalogic Enterprise Deployment for the steps to create and mount the required projects and shares. This section lists the projects, shares and file systems used for this paper along with the properties.

Production Site

Project Name: MW Binaries

PROPERTY NAME	PROPERTY VALUE	COMMENTS
Quota	100 GB	Quota for the Project, including snapshots. The quota should be allocated based on your requirements
Mount Point	/export/binaries	
All other settings	Default	
Share Name	mw_home1	Share mount point: /export/binaries/mw_home1 Mounted on apphost1
Share Name	mw_home2	Share mount point: /export/binaries/mw_home2 Mounted on apphost2

Project Name: Configuration

PROPERTY NAME	PROPERTY VALUE	COMMENTS
Quota	500 GB	Quota for the Project, including snapshots. The quota should be allocated based on your requirements
Mount Point	/export/config	
All other settings	Default	
Share Name	domains	<ul style="list-style-type: none"> • Share mount point: /export/config/domains • Mounted on apphost1, apphost2 • Contains the domain configuration for the admin server and the managed servers • A directory is created per compute node that contains the domain configuration.
Share Name	admin	<ul style="list-style-type: none"> • Share mount point: /export/config/domains

		<ul style="list-style-type: none"> • Mounted on apphost1, apphost2 • Contains node manager configuration files and start scripts • A directory is created per compute node that contains the node manager configuration
Share Name	jmsjta	<ul style="list-style-type: none"> • Share mount point: /export/config/domains • Mounted on apphost1, apphost2 • Shared location for the JMS and Transaction Log persistent stores • A directory is created per cluster in the domain and contains sub directories for the JMS and Transaction Log persistent stores.
Share Name	logs	<ul style="list-style-type: none"> • Share mount point: /export/config/logs • Mounted on apphost1, apphost2 • • Shared location for the JMS and Transaction Log persistent stores • A directory is created per cluster in the domain and contains sub directories for the JMS and Transaction Log persistent stores.

Project Name: OHS

PROPERTY NAME	PROPERTY VALUE	COMMENTS
Quota	200 GB	Quota for the Project, including snapshots
Mount Point	/export/ohs	
All other settings	Default	
Share Name	webhost1	<ul style="list-style-type: none"> • Share mount point: /export/ohs/webhost1 • Mounted on webhost1 • Contains two directories, one for Oracle HTTP Server binaries and the other for the instance configuration
Share Name	webhost2	<ul style="list-style-type: none"> • Share mount point: /export/ohs/webhost2 • Mounted on webhost2 • Contains two directories, one for Oracle HTTP Server binaries and the other for the instance configuration

Configuring the Storage Replication Channel

A **storage replication channel** is a network channel that is dedicated specifically to replication traffic between the Sun ZFS Storage 7320 appliance at the production site and the standby site. The storage replication channel must be configured on both the production site and standby site

before configuring remote replication. This section provides the steps to the configure the storage replication channel

Prerequisites

1. Connect the port igb2 from both the storage heads in the Exalogic Machine to the embedded Cisco Catalyst 4948 switch within the Exalogic Machine
2. Connect port igb3 from both the storage heads in the Exalogic Machine to a network drop in your data center. This can distribution switch in your datacenter.
3. Do this both at the production site and at the standby site
4. Ensure that the IP address assigned to the replication channel has been provisioned and is in DNS.
5. Oracle recommends provisioning the replication channel IP address on a subnet that is different from the Management IP subnet.

Configuration Steps

Follow the steps below to configure the storage replication channel. All these steps must be completed on **both the production site and the standby site**

1. Open the Browser User Interface (BUI) for the storage head.
2. Navigate to **Configuration** → **Network** to bring up the **Network** screen
3. Validate that the **Built-in Devices igb2** and **igb3** are connected and live.
4. The first step is to create **Datalinks** for these two devices.
5. Create the first datalink as follows: Click + Next to **Datalinks** table to bring up the **Network Datalink** screen. Enter the following details:
 - a. **Name:** Enter a name for the datalink. For example: **repl-1-dl**
 - b. Under **Devices**, choose **igb2**
 - c. Accept the default values for all other fields
 - d. Click Apply to apply the changes
6. Create the second datalink as follows: Click + Next to **Datalinks** table to bring up the **Network Datalink** screen. Enter the following details:
 - a. **Name:** Enter a name for the datalink. For example: **repl-2-dl**
 - b. Under **Devices**, select **igb3**
 - c. Accept the default values for all other fields

- d. Click Apply to apply the changes
7. Next create the interfaces for the two datalinks created in steps 5 and 6.
8. Create the first interface as follows: Click + Next to **Interfaces** table to bring up the **Network Interfaces** screen. Enter the following details:
 - a. **Name:** Enter a name for the interface. For example: **repl-1-interface**
 - b. Under Properties, unselect Allow Administration
 - c. Select the checkbox next to **Use IPv4 Protocol**.
 - d. Under the **Use IPv4 Protocol** Section:
 - i. Choose Static Address List
 - ii. Enter **0.0.0.0/8** for the IP address
 - e. Under the **Datalinks** section select **repl-1-dl**
 - f. Accept the default values for all other fields
 - g. Click Apply to apply the changes
9. Create the second interface as follows: Click + Next to **Interfaces** table to bring up the **Network Interfaces** screen. Enter the following details:
 - a. **Name:** Enter a name for the interface. For example: **repl-2-interface**
 - b. Under Properties, unselect Allow Administration
 - c. Select the checkbox next to **Use IPv4 Protocol**.
 - d. Under the **Use IPv4 Protocol** Section:
 - i. Choose Static Address List
 - ii. Enter **0.0.0.0/8** for the IP address
 - e. Under the **Datalinks** section select **repl-2-dl**
 - f. Accept the default values for all other fields
 - g. Click Apply to apply the changes
10. Create an **Active/Passive** bonded interface between the interfaces created in steps 8 and 9. This is accomplished by configuring an **IPMP group (IP MultiPathing)**.
11. Create the **IPMP group** as follows: Click + Next to **Interfaces** table to bring up the **Network Interfaces** screen. Enter the following details:
 - a. **Name:** Enter a name for the interface. For example: **dr-repl-interface**
 - b. Under Properties, unselect **Allow Administration**

- c. Select the checkbox next to **Use IPv4 Protocol**.
 - d. Under the **Use IPv4 Protocol** section:
 - i. Choose Static Address List
 - ii. Enter the **IP address** provisioned for the storage replication channel on the production site. Use the format: **IPv4Address/mask**. For example: **10.204.77.120/24**, where **10.204.77.120** is the IP address and the **24** is the subnet netmask.
 - iii. Use the **IP address** provisioned for the storage replication channel on the standby site when executing this step on the standby site.
 - e. Select the checkbox next to **IP MultiPathing Group**
 - f. From the list of interfaces select the interfaces created in steps 9 and 10. For example: select **repl-1-interface** and **repl-2-interface**
 - g. Accept the default values for all other fields
 - h. Click Apply to apply the changes
12. Click **Apply** on the Network page to apply the configuration changes.
13. Next create a routing table entry for the replication interface. Navigate to the **Configuration → Network → Routing** screen
14. Create a Routing Table Entry as follows: Click **+** Next to **Routing Table Entry** table to bring up the **Insert Route** screen. Enter the following details:
 - a. Family: Select **IPv4**
 - b. Kind: Default
 - c. Gateway: The **Gateway IP address** for the Replication Channel IP address. Use the appropriate Gateway IP address for the production site and standby site
 - d. Interface: Select the **IPMP group** created in step 11. For example: **dr-repl-interface**.
 - e. Click Add to add the entry
15. Validate the configuration, by pinging the replication channel IP address from one of the compute nodes

Configuring Remote Replication Targets

The Sun ZFS Storage 7320 appliance supports replication of projects and shares from a source appliance to a number of target appliances manually, on a schedule, or continuously. The

replication includes both data and metadata. This is typically one-time setup that can be through the BUI or the CLI.

This section details the steps on configuring the remote replication targets using the BUI. Follow the steps below to configure the storage replication channel. All these steps must be completed on both the production site and the standby site

1. Open the BUI for the storage head.
2. Navigate to **Configuration** → **Service** to bring up the Services screen
3. Under the **Data Services** table, click on the **Remote Replication** link to bring up the **Remote Replication** Screen
4. Setup the replication target as follows: Click the **+** next to the **Targets** table to bring up the **Add Replication Target** screen. Enter the following details:
 - a. **Name:** Enter the name for the target. For example: **dr-repl-channel**
 - b. **Hostname:** Enter the IP address for the Target appliance. This is the IP address of the storage replication channel. For example: **10.204.77.120**.
5. At this point the replication configuration has been set up between the targets

Note:

- On the production site, provide the IP address of the storage replication channel of the standby site as the target.
- On the standby site, provide the IP address of the storage replication channel of the production site as the target

- c. **Root Password:** Enter the root password for the target appliance.
- d. Click **Add** to add the replication target

Host Setup

Hostnames and Aliases

In a disaster recovery topology, the production site host names must be resolvable to the IP addresses of the corresponding peer systems at the standby site. This can be set up by creating a hostname alias in the `/etc/host` file. Create hostname aliases for all the hosts on the production and standby sites by creating the entries shown in the table below

Web Tier

For this paper, the hosts for the web tier i.e. webhost1 and webhost2 are external to the Exalogic Machine.

Edit the `/etc/hosts` file on the `webhost1` and `webhost2` at the production site and add the hostname aliases for the webhosts and the client access network VIPs.

Production Site: Web Host Alias

IP ADDRESS	NETWORK NAME	HOSTNAME ALIAS
10.132.131.100	webhostprod1.mycompany.com	webhost1.mycompany.com
10.132.131.101	webhostprod2.mycompany.com	webhost2.mycompany.com

Production Site: Client Access Network VIPs

CLIENT ACCESS IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
10.204.77.50	adminvhn-prod.mycompany.com	adminvhn-pub.mycompany.com
10.204.77.51	wlsvhn1-prod.mycompany.com	wlsvhn1-pub.mycompany.com
10.204.77.52	wlsvhn2-prod.mycompany.com	wlsvhn2-pub.mycompany.com
10.204.77.53	wlsvhn3-prod.mycompany.com	wlsvhn3-pub.mycompany.com
10.204.77.54	wlsvhn4-prod.mycompany.com	wlsvhn4-pub.mycompany.com

Also add the entries for the **storage replication channel** and load balancer VIPs to the `/etc/host` files on `webhost1` and `webhost2` at the production site.

Production Site: Replication Channel

IP ADDRESS	HOST NAME	HOSTNAME ALIAS
10.204.77.120	prodrepl.mycompany.com	None

The following Virtual IP addresses were configured on the load balancer for this paper

Production Site: Load Balancer VIPs

VIRTUAL HOST NAME	IP ADDRESS	HOSTNAME ALIAS
exalogic.mycompany.com	144.25.145.54	None
admin.mycompany.com	144.25.145.50	None

Edit the `/etc/hosts` file on the `webhost1` and `webhost2` at the standby site and add the hostname aliases for the webhosts and the client access network VIPs

Standby Site: Web Host Alias

IP ADDRESS	NETWORK NAME	HOSTNAME ALIAS
10.200.130.100	webhoststby1.mycompany.com	webhost1.mycompany.com
10.200.130.101	webhoststby2.mycompany.com	webhost2.mycompany.com

Standby Site: Client Access Network VIPs

CLIENT ACCESS IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
10.250.77.70	adminvhn-stby.mycompany.com	adminvhn-pub.mycompany.com
10.250.77.71	wlsvhn1-stby.mycompany.com	wlsvhn1-pub.mycompany.com
10.250.77.72	wlsvhn2-stby.mycompany.com	wlsvhn2-pub.mycompany.com
10.250.77.73	wlsvhn3-stby.mycompany.com	wlsvhn3-pub.mycompany.com
10.250.77.74	wlsvhn4-stby.mycompany.com	wlsvhn4-pub.mycompany.com

Add the entries for the storage replication channel and load balancer VIPs to the /etc/host files on webhost1 and webhost2 at the standby site

Standby Site: Replication Channel

IP ADDRESS	HOST NAME	ALIAS
10.200.80.120	stbyrepl.mycompany.com	None

Standby Site: Load Balancer VIPs

IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
144.25.145.54	exalogic.mycompany.com	None
144.25.145.50	admin.mycompany.com	None

Application Tier

Edit the /etc/hosts file on the apphost1 and apphost2 at the production site add the hostname aliases for the apphosts, the client access network VIPs and the private InfiniBand VIPs

Production Site: Hostname Aliases

IP ADDRESS	NETWORK NAME	HOSTNAME ALIAS
10.204.80.40	apphostprod1.mycompany.com	apphost1.mycompany.com
10.204.80.41	apphostprod2.mycompany.com	apphost2.mycompany.com
192.168.10.1	apphost1-priv.mycompany.com	None
192.168.10.2	apphost2-priv mycompany.com	None

Production Site: Client Access Network VIPs

CLIENT ACCESS IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
10.204.77.50	adminvhn-prod.mycompany.com	adminvhn-pub.mycompany.com
10.204.77.51	wlsvhn1-prod.mycompany.com	wlsvhn1-pub.mycompany.com
10.204.77.52	wlsvhn2-prod.mycompany.com	wlsvhn2-pub.mycompany.com
10.204.77.53	wlsvhn3-prod.mycompany.com	wlsvhn3-pub.mycompany.com

10.204.77.54	wlsvhn4-prod.mycompany.com	wlsvhn4-pub.mycompany.com
--------------	----------------------------	---------------------------

Production Site: InfiniBand Network VIPs

INFINIBAND IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
192.168.10.50	adminvhn.mycompany.com	None
192.168.10.51	wlsvhn1.mycompany.com	None
192.168.10.52	wlsvhn2.mycompany.com	None
192.168.10.53	wlsvhn3.mycompany.com	None
192.168.10.54	wlsvhn4.mycompany.com	None
192.168.10.110	dbhost01-ibvip.mycompany.com	None
192.168.10.111	dbhost02-ibvip.mycompany.com	None

Add the entries for the IP addresses of the storage heads and the load balancer VIPs to the /etc/host files on apphost1 and apphost2 at the standby site

Production Site: Storage Entries

MANAGEMENT IP ADDRESS	STORAGE NODE NAME	HOSTNAME ALIAS
10.204.80.100	prodstor1.mycompany.com	None
10.204.80.101	prodstor2.mycompany.com	None
192.168.10.15	prodstor-ib1.mycompany.com	None
192.168.10.16	prodstor-ib2.mycompany.com	None

Production Site: Load Balancer VIPs

IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
144.25.145.54	exalogic.mycompany.com	None
144.25.145.50	admin.mycompany.com	None

Edit the /etc/hosts file on the apphost1 and apphost2 at the production site add the hostname aliases for the apphosts, the client access network VIPs and the private InfiniBand VIPs

Standby Site: Hostname Aliases

IP ADDRESS	NETWORK NAME	HOSTNAME ALIAS
10.132.251.40	apphoststby1.mycompany.com	apphost1.mycompany.com
10.132.251.41	apphoststby2.mycompany.com	apphost2.mycompany.com
192.168.10.1	apphost1-priv	None
192.168.10.2	apphost2-priv	None

Standby Site: Client Access Network VIPs

CLIENT ACCESS IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
10.250.77.70	adminvhn-stby.mycompany.com	adminvhn-pub.mycompany.com
10.250.77.71	wlsvhn1-stby.mycompany.com	wlsvhn1-pub.mycompany.com
10.250.77.72	wlsvhn2-stby.mycompany.com	wlsvhn2-pub.mycompany.com
10.250.77.73	wlsvhn3-stby.mycompany.com	wlsvhn3-pub.mycompany.com
10.250.77.74	wlsvhn4-stby.mycompany.com	wlsvhn4-pub.mycompany.com

Standby Site: InfiniBand Network VIPs

INFINIBAND IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
192.168.10.70	adminvhn.mycompany.com	None
192.168.10.71	wlsvhn1.mycompany.com	None
192.168.10.72	wlsvhn2.mycompany.com	None
192.168.10.73	wlsvhn3.mycompany.com	None
192.168.10.74	wlsvhn4.mycompany.com	None
192.168.10.110	dbhost01-ibvip.mycompany.com	None
192.168.10.111	dbhost02-ibvip.mycompany.com	None

Add the entries for the IP addresses of the storage heads and the load balancer VIPs to the /etc/host files on apphost1 and apphost2 at the standby site

Standby Site: Storage Entries

IP ADDRESS	STORAGE NODE NAME	HOSTNAME ALIAS
10.132.251.100	stbystor1.mycompany.com	None
10.132.251.101	stbystor2.mycompany.com	None
192.168.10.15	stbystor-ib1	None
192.168.10.16	stbystor-ib2	None

Standby Site: Load Balancer VIPs

IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
144.25.145.54	exalogic.mycompany.com	None
144.25.145.50	admin.mycompany.com	None

Database Tier

Edit the /etc/hosts file on the dbhost1 and dbhost2 at the production site add the entries for the private InfiniBand VIPs. These VIP addresses will be used to configure an additional network on

the InfiniBand network to support the native SDP protocol for database access from the apphosts.

Production Site: InfiniBand Network Database VIPs

INFINIBAND IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
192.168.10.110	dbhost01-ibvip.mycompany.com	None
192.168.10.111	dbhost02-ibvip.mycompany.com	None

Standby Site: InfiniBand Network Database VIPs

INFINIBAND IP ADDRESS	VIRTUAL HOST NAME	HOSTNAME ALIAS
192.168.10.110	dbhost01-ibvip.mycompany.com	None
192.168.10.111	dbhost02-ibvip.mycompany.com	None

Mount Points

In a Disaster Recovery topology, the product binaries and component configuration is shared storage. Before the production and standby sites can be set up the hosts need to be configured with the required mount points.

Web Tier

For this paper, the hosts for the web tier i.e. webhost1 and webhost2 are external to the Exalogic Machine. However the OHS binaries and instance configuration is located on the shared storage in the Exalogic Machine. The webhosts connect to the appliance over the replication channel IP address. NFSv3 protocol is used to access these file systems from webhosts.

The table below shows the NFS mount points used on webhost1 and webhost2s at the production and standby sites

Production Site

HOSTNAME	APPLIANCE MOUNT POINT	HOST MOUNT POINT	COMMENTS
webhost1_p1.mycompany.com	prod_repl:/export/ohs/webhost1	/u01/app/oracle	OHS mount Point
webhost1_p2.mycompany.com	prod_repl:/export/ohs/webhost2	/u01/app/oracle	OHS mount Point

Standby Site

HOSTNAME	APPLIANCE MOUNT POINT	HOST MOUNT POINT	COMMENTS
webhost1_s1.mycompany.com	stby_repl:/export/ohs/webhost1	/u01/app/oracle	OHS mount Point
webhost1_s2.mycompany.com	stby_repl:/export/ohs/webhost2	/u01/app/oracle	OHS mount Point

Application Tier

The application server hosts in the Exalogic Machine connect to the shared storage over InfiniBand. NFSv3 protocol is used to access the file systems from application server hosts.

The table below shows the NFS mount points used on apphost1 and apphost2 at the production and standby sites:

Production Site

HOSTNAME	APPLIANCE MOUNT POINT	HOST MOUNT POINT	COMMENT
apphost1_p1	prodstor-ib1/export/binaries/mw_home1	/u01/app/oracle/product/fmw	MW Home
	prodstor-ib1:/ export/config/domains	/u01/app/oracle/ domains	Domain Configuration
	prodstor-ib1:/ export/config/admin	/u01/app/oracle/admin	Node Manager Configuration
	prodstor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/tlogs	T-Log Persistent Store
	prodstor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/jms	JMS Persistent Store
	prodstor-ib1:/ export/config/logs	/u01/app/oracle/logs	Application Log Directory
apphosts_p2	prodstor-ib1/export/binaries/mw_home2	/u01/app/oracle/product/fmw	MW Home
	prodstor-ib1:/ export/config/domains	/u01/app/oracle/ domains	Domain Configuration
	prodstor-ib1:/ export/config/admin	/u01/app/oracle/admin	Node Manager Configuration
	prodstor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/tlogs	T-Log Persistent Store
	prodstor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/jms	JMS Persistent Store
	prodstor-ib1:/ export/config/logs	/u01/app/oracle/logs	Application Log Directory

Standby Site

HOSTNAME	APPLIANCE MOUNT POINT	HOST MOUNT POINT	COMMENT
apphost1_s1	stbystor-ib1/export/binaries/mw_home1	/u01/app/oracle/product/fmw	MW Home
	stbystor-ib1:/ export/config/domains	/u01/app/oracle/ domains	Domain Configuration
	stbystor-ib1:/ export/config/admin	/u01/app/oracle/admin	Node Manager Configuration
	stbystor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/tlogs	T-Log Persistent

			Store
	stbystor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/jms	JMS Persistent Store
	stbystor-ib1:/ export/config/logs	/u01/app/oracle/logs	Application Log Directory
apphosts_s2	stbystor-ib1/export/binaries/mw_home2	/u01/app/oracle/product/fmw	MW Home
	stbystor-ib1:/ export/config/domains	/u01/app/oracle/ domains	Domain Configuration
	stbystor-ib1:/ export/config/admin	/u01/app/oracle/admin	Node Manager Configuration
	stbystor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/tlogs	T-Log Persistent Store
	stbystor-ib1:/ export/config/jmsjta	/u01/app/oracle/jmsjta/wlsDomain/cluster1/jms	JMS Persistent Store
	stbystor -ib1:/ export/config/logs	/u01/app/oracle/logs	Application Log Directory

Client Access Network Configuration

The client access network connects the compute nodes in an Exalogic Machine to the existing corporate network through the Sun Network QDR InfiniBand Gateway Switch. The Sun Network QDR InfiniBand Gateway Switches are connected to a 10 Gigabit Network switch to provide the Ethernet over InfiniBand (EoIB) connectivity.

Ensure that the client access network has been configured. Refer to the Oracle Fusion Middleware Exalogic Machine Owner's Guide for the configuring the Client Access Network.

Cabling the Exalogic and Exadata Database Machines over InfiniBand

Ensure that the Exalogic Machine and the Exadata Database Machine on each site are connected to each other over InfiniBand.

Refer to the Oracle Fusion Middleware Exalogic Machine Multirack Cabling Guide for the procedure to connect an Exalogic Machine with an Exadata Database Machine.

Site Setup and Configuration

This paper provides detailed steps where the configuration deviates from the steps in the **Oracle Exalogic Enterprise Deployment Guide**. Please refer to the [Oracle Exalogic Enterprise Deployment Guide](#) for the setup steps.

Production Site Setup

Web Tier

The web tier on the production site consists of two hosts called webhost1 and webhost2. Both hosts are running Oracle HTTP Server. These two webhosts are front ended by a load balancer configured to load balance traffic between webhost1 and webhost2.

Follow the [Oracle Exalogic Enterprise Deployment Guide](#) to install and configure the Oracle HTTP Server. The high level steps to configure the web tier are below:

1. Ensure that the /etc/host file is setup properly on webhost1 and webhost2
2. Ensure that the mount points are properly configured on both webhost1 and webhost2
3. On webhost1, install the Oracle HTTP Server binaries to the /u01/app/oracle/product/ohs directory.
4. Specify /u01/app/oracle/admin/ohs_instance1 as the directory for the Instance Home Location.

Application Tier

The application tier on the production site consists of two hosts called apphost1 and apphost2. Both hosts are running Oracle WebLogic Server.

For this paper, the enterprise topology described in the **Oracle Exalogic Enterprise Deployment Guide** was scaled down to consist of One Admin Server, four Weblogic Managed Servers and four coherence servers i.e. two managed servers and two coherence servers running on apphost1 and apphost2.

The table below provides a summary:

HOSTNAME	WEBLOGIC SERVER NAME	CLUSTER NAME	WEBLOGIC SERVER LISTEN ADDRESS	NETWORK CHANNEL LISTEN ADDRESS
apphost1	Admin Server	None	adminvhn.mycompany.com	adminvhn-pub.mycompany.com
	wls1	wls_cluster1	wlsvhn1.mycompany.com	wlsvhn1-pub.mycompany.com
	wls2	wls_cluster2	wlsvhn2.mycompany.com	wlsvhn2-pub.mycompany.com
	coh_server1	coh_cluster1	apphost1-priv.mycompany.com	N/A
	coh_server2	coh_cluster2	apphost1-priv.mycompany.com	N/A
apphost2	wls3	wls_cluster1	wlsvhn3.mycompany.com	wlsvhn1-pub.mycompany.com
	wls4	wls_cluster2	wlsvhn4.mycompany.com	wlsvhn2-pub.mycompany.com
	coh_server3	coh_cluster1	apphost2-priv.mycompany.com	N/A
	coh_server4	coh_cluster2	apphost2-priv.mycompany.com	N/A

Follow the [Oracle Exalogic Enterprise Deployment Guide](#) to install and configure the Oracle Weblogic Server Domain. The high level steps to configure the application tier are below:

1. Ensure that the `/etc/host` file is setup properly on `apphost1` and `apphost2`
2. Ensure that the mount points are properly configured on both `apphost1` and `apphost2`.
3. On `apphost1` and `apphost2`, install the Oracle Weblogic Server binaries to the `/u01/app/oracle/product/fmw` directory.
4. Specify **WLSDomain** for the Domain Name.
5. Specify `/u01/app/oracle/domains/apphost1` as the domain directory.
6. Configure the Weblogic Administration Server with the Listen Address specified in the table above
7. Configure the four managed servers i.e. `wls1`, `wls2`, `wls3` and `wls4` with the WLS Listen Addresses shown in the table above
8. Configure two clusters `wls_cluster1` and `wls_cluster2` with the managed servers shown in the table above.
9. Pack and unpack the domain on `apphost2` to the `/u01/app/oracle/domains/apphost2` directory.
10. Configure the HTTP and t3 Network Channels using the Network Channel Listen addresses specified in the table.
11. Specify `/u01/app/oracle/jmsjta/wlsDomain/cluster1/jms` as the directory for the JMS Persistent Store
12. Specify `/u01/app/oracle/jmsjta/wlsDomain/cluster1/tlogs` as the directory for the Transaction Log Persistent Store.
13. Change the log directory for the WebLogic Administration Server to `/u01/app/oracle/logs/apphost1/AdminServer/logs`
14. For the Managed Servers running on `apphost1` change the log directory to `/u01/app/oracle/logs/apphost1/<managedServerName>/logs`
15. For the Managed Servers running on `apphost2` change the log directory to `/u01/app/oracle/logs/apphost2/<managedServerName>/logs`
16. On `apphost1`, configure Node Manager to run under the `/u01/app/oracle/admin/apphost1/nodemanager` directory
17. On `apphost1`, configure Node Manager to run under the `/u01/app/oracle/admin/apphost2/nodemanager` directory

18. Configure the gridlink data sources using the **dbhost01-ibvip** and **dbhost02-ibvip** as the database host names
19. Specify a service name to connect to the database.

Database Tier

For this paper the customer database was running on an Exadata Database Machine in a Quarter Rack configuration. The high level steps to configure the database tier are below.

1. Follow the Oracle Exadata Database Machine Owners Guide to setup the Exadata Database Machine.

The Oracle Exadata Storage Server and Oracle Exadata Database Machine Documentation can be found on the Exadata Storage cell under the `/opt/oracle/cell/doc` directory
2. Ensure that Exalogic Machine and the Exadata Database Machine have been physically cabled as described in the [Oracle Exalogic Machine Multitrack Cabling Guide](#)
3. Ensure that the private InfiniBand networks on the Exalogic Machine and the Exadata Database Machine are configured to be on the same subnet. Refer to the [Oracle Exalogic Owners Guide](#) and the Oracle Exadata Database Machine Owners Guide for the steps to accomplish this task.
4. Follow the [Oracle Exalogic Enterprise Deployment Guide](#) to enable the SDP protocol and to configure an additional listener on the InfiniBand network.
5. For this paper all database traffic from the Exalogic Machine to the Exadata Database Machine was configured to use the private InfiniBand network
6. Configure a service name for the database hosting the application data on the production site. This service name must be configured with the **Primary role**. Refer the `srvctl` commands to configure the service name
7. Create a schema(s) according to the requirements of the application(s) deployed to the WebLogic server

Configuring Replication for the Projects and Shares

Replication must be configured on the Sun ZFS Storage 7320 appliance at the production site before instantiating the standby site. The product binaries and configuration installed/configured on the production site will be replicated to the standby site when the production site storage is replicated to the standby site storage, so no installation/configuration is required at the standby site.

Replication can be configured either at the project level or at the share level. Follow the steps below to configure replication between the production site and the standby site.

1. From the BUI, navigate to **Shares** → **Projects** screen and choose a project or share, then click the Replication.
2. Create a **Replication Target** as follows: Click + next to the Actions Table to bring up the **Add Replication Target** screen. Provide the following details.
 - a. Choose the target system from the drop-down. Note that only the targets added under the **Services** → **Remote replication** → **Targets** are listed in the drop-down. For example: **dr-repl-channel**
 - b. Select the name of the pool, by default in the Exalogic Machine, there is only one pool
 - c. Select the mode of replication. **Scheduled or Continuous**. Oracle recommends selecting the replication mode based on your requirements and the data in the project/share. Refer to the **Oracle MAA Best Practices for Disaster Recovery** section for guidelines
 - d. If using the **Scheduled** replication mode, click + next to the **Schedule table** to create a schedule.
 - e. If the replication happens within a data center, SSL can be disabled to enhance performance.
 - f. The bandwidth used for replication can be limited based on individual requirements.
 - g. If snapshots are taken at the source for the replication project or share, the user can choose to include replicating the snapshots.
3. For this paper, replication was configured as shown in the table below:

PROJECT NAME	SHARE NAME	REPLICATION LEVEL	REPLICATION TYPE	SCHEDULE
OHS	webhost1	Project	Scheduled	Scheduled at 02:00 AM
	webhost2	Project	Scheduled	Scheduled at 02:00 AM
MW Binaries	mw_home1	Project	Scheduled	Scheduled at 02:00 AM
	mw_home2	Project	Scheduled	Scheduled at 02:00 AM
Configuration	domains	Share	Scheduled	Scheduled at 02:00 AM
	admin	Share	Scheduled	Scheduled at 02:00 AM
	logs	Share	Scheduled	Scheduled at 02:00 AM
	jmsjta	Share	Continuous	N/A

4. If the target is added with **Continuous** mode of replication, the replication starts immediately. This mode is used in this exercise for maximum protection purposes.
5. If the **Scheduled** mode of replication is chosen, then it is recommended to perform a manual update one time if the schedule is expected to occur sometime in the future. This will enable a copy of the binaries and the configuration to be available at the standby site in case the production site fails before the first scheduled replication occurs.
6. Validate that the replication between the production site and the standby is configured. To validate that the packages are being received or have already been received. Use the BUI and navigate to **Shares → Projects** at the left side frame and then click **Replica**. This will list all the packages that are being received or have been received from production site.

Standby Site Instantiation

No installation and configuration is required at the standby site for Oracle Fusion Middleware components. When the production site storage is replicated to the standby site storage, the Oracle Fusion Middleware product binaries and configuration installed/configured on the production site will be replicated at the standby site.

Software installation and configuration is required for the Exadata Database Machine.

To create the initial snapshot of the product binaries and configuration on the standby site, perform a manual replication of the Projects and shares.

1. The manual update must be initiated from the storage on the production site. This can be done through the BUI or the CLI.
2. Open the BUI and navigate to the **Shares → Projects** and then click on **Replication** to bring up the Replication screen.
3. Click on the Manual replication icon next to the target to start the Manual update. The status of the update can be viewed under the status column.

On the standby site, validate that the packages have been received. Use the BUI and navigate to **Shares → Projects** at the left side frame and then click **Replica**. This will list all the packages that are being received or have been received from production site.

Web Tier

The web tier at the standby site consists of two hosts called webhost1 and webhost2. Both hosts are running Oracle HTTP Server. These two webhosts are front ended by a load balancer configured to load balance traffic between webhost1 and webhost2.

1. Ensure that the `/etc/host` file is setup properly on webhost1 and webhost2

2. Ensure that the mount points are properly configured on both webhost1 and webhost2 are correctly setup with the

Application Tier

The application tier on the standby site consists of two hosts called apphost1 and apphost2. Both hosts are running the Oracle WebLogic Server. The application tier at the standby site has the same number of hosts as the production site.

1. Ensure that the /etc/host file is setup properly on apphost1 and apphost2
2. Ensure that the mount points are properly configured on both apphost1 and apphost2

Database Tier

The database tier on the standby site consists of two hosts called dbhost1 and dbhost2. Both hosts are running Oracle database.

For this paper the standby customer database was running on an Exadata Database Machine in a Quarter Rack configuration. The high level steps to configure the database tier on the standby site are below:

1. Follow the Oracle Exadata Database Machine Owners Guide to setup the Exadata Database Machine.
2. Ensure that Exalogic Machine and the Exadata Database Machine have been physically cabled as described in the [Oracle Exalogic Machine Multitrack Cabling Guide](#).
3. Ensure that the private InfiniBand networks on the Exalogic Machine and Exadata Database Machine are configured to be on the same subnet. Refer to the [Oracle Exalogic Owners Guide](#) and the Oracle Exadata Database Machine Owners Guide for the steps to accomplish this task.
4. Follow the [Oracle Exalogic Enterprise Deployment Guide](#) to enable the SDP protocol and to configure an additional listener on the InfiniBand network.
5. For this paper all database traffic from the Exalogic Machine to the Exadata Database Machine was configured to use the private InfiniBand network.
6. Configure a service name for the database hosting the application data on the production site. This service name must be configured with **Primary role**. Refer to the `srvctl` commands to configure the service name. Use the same service name used on the production site to ensure a seamless failover.

Data Guard Setup

1. The database on the standby site can be setup as a physical standby or a logical standby database. For this paper, the database on the standby site was setup as a physical standby database.
2. Oracle Active Data Guard was not configured for this setup.
3. Oracle Active Data Guard can be configured if custom applications in your topology are designed to leverage the Oracle Active Data Guard technology.
4. Oracle Fusion Middleware does not support Oracle Active Data Guard for the database repositories that are a part of the Fusion Middleware topology.
5. The steps for setting up Oracle Data Guard are not covered in this paper. Follow the [Oracle Data Guard Concepts and Administration Guide](#) and the [Oracle Data Guard Broker Guide](#) to configure Data Guard between the databases on the production site and the standby site.
6. For Oracle Data Guard best practices for the Exadata Database Machine, please refer to the [Oracle Data Guard: Disaster Recovery Best Practices for Exadata Database Machine](#) white paper.

Validate the Standby Site Setup

It is an Oracle MAA best practice recommendation to validate the standby site once the configuration is complete. Follow the steps below to validate the site setup on the standby site

1. Shut down any processes still running on the production site.
2. Perform the **Role Reversal** procedure on the storage appliance at the standby site for the project that is replicated.
3. Use Oracle Data Guard to perform a database switchover.
4. On the standby site hosts, manually start up the processes for the application server instances.
5. Ensure that all user requests are routed to the standby site. This can be achieved through a global DNS push or similar technologies.
6. Use a browser client to perform post-failover testing to confirm that requests are being resolved and redirected to the standby site.
7. After validating that the standby site is working, make sure to switch back to make the production site.

Disaster Recovery Operations

Site Switchover

Switchovers are planned operations done for periodic validation or to perform planned maintenance on the current production site. After the switchover, the current standby site becomes the new production site, and the current production site becomes the new standby site. In other words, the roles are reversed. This paper also uses the term **switchover** to refer to **site switchover**. The section below provides the steps to perform a switchover.

1. Shut down all the Oracle Fusion Middleware components at the production site. This can be done either manually or by running a script or by using relevant management software.
2. Un- mount the NFS file systems on all the hosts at the production site.
3. Perform an Oracle Data Guard switchover for the databases on the Exadata Database Machine. This procedure converts the database on the standby site to a production database and also converts the production database to a standby database.

This step includes finishing the recovery, switching the standby database to become the primary database, shutting down and restarting the database at the new primary site. The old production site is converted to the standby site.

4. Perform a **Role Reversal** operation on the storage appliance at the standby site for the projects and shares that are being replicated.

Note: After the role-reversal, the replication is automatically set to the **manual** mode. The mode is not automatically changed to the **continuous** mode because the **old production site** may be unavailable.

5. The old standby site now becomes the new production site and the old production site now becomes the new standby site.
6. Enable replication for the projects and shares on storage appliance at the **new production site**. This will start replicating data from the **new production** to the **new standby site**.
7. Mount the file systems on the **new production site**.
8. Ensure that all user requests are routed to the **new production site**. This can be achieved through a global DNS push or similar technologies.
9. Start all the Oracle Fusion Middleware components on the **new production site** and ensure they are started successfully.

At this point, the **old standby site** has assumed the role of the **new production site** and the **old production site** is now the **new standby site**.

Site Switchback

A **site switchback** is defined as the process of reverting the current production site and the current standby site to their original roles. Site switchbacks are planned operations and are done after a switchover operation has been completed. A switchback restores the original roles of each site, the current standby site becomes the production site, and the current production site becomes the standby site. This paper also uses the term **switchback** to refer to **site switchback**. The section below provides the steps to perform a switchback.

1. Shut down all the Oracle Fusion Middleware components at the **current production site**. This can be done either manually or by running a script or by using relevant management software.
2. Un- mount the NFS file systems on all the hosts at the **current production site**.
3. Perform an Oracle Data Guard switchover for the databases on the Exadata Database Machine. This procedure converts the database on the standby site to a production database and also converts the production database to a standby database

This step includes finishing the recovery, switching the standby database to become the primary database, shutting down and restarting the database at the new primary site. The old production site is converted to the standby site.

4. Perform a **Role Reversal** operation on the storage appliance at the **current standby site** for the projects and shares that are being replicated.

Note: After the role-reversal, the replication is automatically set to the **manual** mode. The mode is not automatically changed to the **continuous** mode because the production site may be in an unavailable state.

5. Enable replication for the projects and shares on storage appliance at the **new production site**. This will start replicating data from the **new production** to the **new standby site**.
6. Mount the file systems on the **new production site**
7. Ensure that all user requests are routed to the **new production site**. This can be achieved through a global DNS push or similar technologies.
8. Start all the Oracle Fusion Middleware components on the **new production site** and ensure they are started successfully.

At this point, the **current standby site** has assumed the role of the **new production site** and the **old production site** is now the **new standby site**.

Site Failover

Site failover is the process of making the **current standby site** the **new production site** when the **current production site** becomes unavailable due to an unplanned outage. A site failover is performed when the production site is unavailable and it is faster to fail over to the standby site than fix the issues at the production site. The current standby site becomes the new production site while the issues causing the disruptions at the production site are resolved. This paper also uses the term **failover** to refer to **site failover**. The section below provides the steps to perform a failover.

1. Perform an Oracle Data Guard failover for the databases on the Exadata Database Machine. This step converts the standby database to a primary database.

On the new primary database set the **log_archive_dest_2_state** parameter to **DEFER** until the old production site is available again.
2. Perform a **Role Reversal** operation on the storage appliance at the **current standby site**. Do not enable replication at this time.
3. Mount the NFS file systems on all the hosts at the **new production site**.
4. Start the Oracle Fusion Middleware components at the **new production site**.
5. Ensure that all user requests are routed to the **new production site**. This can be achieved through a global DNS push or similar technologies.
6. At this point, the **old standby site** has assumed the role of **new production site**.
7. Once the **old production site** is back up again:
 - a. Initiate replication from the new production to the old production. The old production site now becomes the new standby site.
 - b. Use Oracle Data Guard to switch the database role at the **new standby site**.
 - c. Change the new primary database's archive log shipping to the **ENABLE** state. This starts pushing the archive logs to the new standby site.

Oracle MAA Best Practices for Disaster Recovery

1. It is recommended to test the standby site periodically. A good rule of thumb is to test the standby site after every major upgrade or once every quarter. This will help mitigate failures at both sites. Test the standby site by switching its role with the current production site.
 - a. Follow the site switchover procedure to switch over the standby site to the new production site.

- b. Once testing is complete, follow the site switchback procedures to reverse the roles.
 - c. Periodic testing validates that both the production and standby sites are completely functional and mitigates the risk of failure at both sites. It also validates the switchover and switchback procedures.
 2. Do not configure project-level and share-level replication within the same project.
 3. Use the **Scheduled** replication mode for projects and shares when:
 - a. Data does not change frequently.
 - b. Recovery Point Objective falls within your scheduled replication window.
 4. Use the **Continuous** replication mode for projects and shares when:
 - a. The standby site is required to be as close as possible to the production site.
 - b. Recovery Point Objective allows for very little data loss.
 - c. Data is of a critical nature.
 5. Follow the **Role Reversal** procedure during switchovers and failovers. This will enable much faster sync-back to the old primary during switchbacks and failbacks.
 6. Snapshots and clones can be used at the target site to offload backup, test, and development types of environment.
 7. When configuring a local standby site i.e. Disaster Recovery within the data center, consider disabling SSL on the replication channel. Removing the encryption algorithm enables a higher replication throughput.
 8. Always enable SSL when replication is across a wide-area-network.
 9. Do not perform **rollback** operations on the projects or shares either at the primary site or at the standby site. Performing a rollback operation on the Sun ZFS Storage 7320 appliance invalidates the replication configuration. It will need to be configured again.
 10. To maintain data consistency between tiers, ensure that the database and application tiers are replicated at the same time. This helps ensure that the different tiers recover to the exact point in time or as close as possible.
 11. Configure Oracle Data Guard in the Managed Recovery Mode.
 12. Configure Oracle Data Guard in the “Maximum Availability” data protection mode or in the “Maximum Protection” data protection mode
 - a. The “Maximum Availability” data protection mode enables the highest level of data protection that is possible without compromising the availability of the primary database.

- b. The “Maximum Protection” data protection mode enables the standby database to be synchronous with the primary. This mode ensures zero data loss. This data protection mode prioritizes data protection over primary database availability.
13. It is recommended to synchronize the standby database when the application tier synchronization is initiated on the storage. This synchronization occurs automatically because Oracle Data Guard is configured in Managed Recovery mode (the recommended configuration) for the database. If the standby database is not in Managed Recovery mode, then you should manually synchronize the standby database.
14. The application tier and database tier on the production site must be manually synchronized with the standby site after making configuration changes or after deploying new applications or after applying patches.
15. Oracle does not recommend synchronizing the local hard drives on the compute nodes.

Appendix

Disaster Recovery Terminology

TERM	DEFINITION
Disaster Recovery	The ability to safeguard against natural disasters or unplanned outages at a production site by having a recovery strategy for failing over applications and data to a geographically separate standby site.
Topology	The production site and standby site hardware and software components that comprise an Oracle Fusion Middleware Disaster Recovery solution
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 unplanned downtime at the production site)
Site Switchover	The process of reversing the roles of the production site and standby site. Switchovers are planned operations 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 Switchback	The process of reversing the roles of the new production site (old standby) and new standby site (old production). Switchback is applicable after a previous switchover.
Site Instantiation	The process of creating a topology at the standby site (after verifying that the primary and standby sites are valid for Oracle Fusion Middleware Disaster Recovery) and synchronizing the standby site with the primary sites so that the primary and standby sites are consistent.
Site Synchronization	The process of applying changes made to the production site at the standby site. For example, when a new application is deployed at the production site, you should perform synchronization so that the same application will be deployed at the standby site.
Recovery Point Objective (RPO)	Maximum age of the data you want the ability to restore in the event of a disaster. For example, if your RPO is six hours, you want to be able to restore the systems back to the state that they were in as of no longer than six hours ago.
Recovery Time Objective (RTO)	Time needed to recover from a disaster. This is usually determined by how long you can afford to be without your systems.

Sun ZFS Storage 7320 Operations

OPERATION	DESCRIPTION
Source	Primary (production) site of the replication.
Target	Receiving site of the replication. A target can receive one or more packages from one or more Sun ZFS Storage 7320 Appliances. In this FMW infrastructure, the target site is the standby site.
Replica/Package	The replicated copy of the project at the target site. It cannot be accessed directly. In order to access the replica, it has to be cloned and the clone is accessed for read/write operations
Snapshot	Point-in-time read-only copy of the share, used for share rollbacks and creating clones.
Clone	Read-writable copy of a snapshot. One or more clones of the share are created from a snapshot.

Export Replica	Process to access the replica at the target. A new project is created. All the shares, snapshots, clones, and so on, are all accessible under the cloned project.
Role Reversal	The direction of the replication is reversed from source → target to target → source for a package.

Oracle Data Guard Setup

Please refer to the [Oracle Data Guard documentation](#) before configuring the Oracle Data Guard. The section below only provides the high level steps required to setup Oracle Data Guard between the production and standby sites.

1. For this paper, the following database was used:

Production site: proddb, ORACLE_SID: proddb1, proddb2

Standby site: stbydb, ORACLE_SID: stbydb1, stbydb2

2. Configure the initproddb.ora parameter file at the production site and initstbydb.ora parameter file at the standby site with the parameters required by Oracle Data Guard to ship the archive logs to the other site.
3. Set the log_archive_dest_2 to point to the service at the standby site

```
SQL> alter system set log_archive_dest_2 = 'service=STBYDB
async db_unique_name=STBYDB valid_for=(primary_role,
online_logfile)'
```

4. Enable the second archive destination.

```
SQL> alter system log_archive_dest_state_2=ENABLE;
```

5. Create a backup of the database using RMAN.

```
$ backup device type disk format
'/u01/app/oracle/stage/proddb/%U' database plus archivelog;

$ backup device type disk format '
/u01/app/oracle/stage/proddb/%U' current controlfile for
standby;
```

6. From the standby host, duplicate the database for standby database.

```
SQL> startup nomount

$ rman target sys/oracle@proddb auxiliary /

RMAN> duplicate standby database for standby;
```

7. Enable the managed recovery at the standby site.

```
SQL> alter database recover managed standby database;
```

8. If standby redo logs are used, enable real-time log apply

```
SQL > alter database recover managed standby database using
current logfile disconnect;
```

9. After the Oracle Data Guard is configured, the status can be verified from v\$database and v\$instance view. With this step, the Oracle Data Guard setup is complete.

Storage Scripts

The scripts shown below are for reference only. Oracle recommends creating scripts based on the ones below

SCRIPT NAME/DESCRIPTION	SCRIPT
<p>role_reverse_at_target.aksh</p> <p>This script is invoked during switchover/switchback conditions. After the switching over is complete, the replication is set to "continuous" mode.</p> <p>The names of the scripts are self explanatory. The scripts are provided as an example.</p>	<pre>script { var myPackage; var projName='OFM-SITE1-KIT' ; run ('cd /'); run ('shares'); run ('set pool=pool-0'); try { run ('select ' + projName); run ('confirm destroy'); } catch (err) { printf("No Project to to delete.. \n"); } printf("Selecting the package to role reverse..\n"); run ('cd /'); run ('shares'); run ('set pool=pool-0'); run('replication sources select source-000'); var packages = list(); for (var i = 0; i < packages.length; i++) { run('select ' + packages[i]); var proj_name = list(); if (proj_name == projName) { myPackage = packages[i] ;</pre>

	<pre> break; } run('cd ..'); } printf("The package chosen to role reverse : %s \n", myPackage); run('cd ..'); run('select ' + myPackage); run('confirm reverse'); run('show'); printf("Source and the target roles are reversed now..\n"); printf("Setting the continuous replication.. \n"); run('cd /'); run('shares'); run('set pool=pool-0'); run('select ' + projName + ' replication'); run('select action-000'); run('set continuous=true'); run('commit'); } </pre>
<p>role_reverse_no_repl.aksh</p> <p>This script is invoked in a failover/failback conditions where the "continuous" mode of replication is not enabled at the end.</p>	<pre> script { var myPackage; var projName='OFM-SITE1-KIT' ; run ('cd /'); run ('shares'); run ('set pool=pool-0') ; try { run ('select ' + projName); run ('confirm destroy'); } catch (err) { printf("No Project to to delete.. \n"); } printf("Selecting the package to role reverse..\n"); run ('cd /'); run ('shares'); run ('set pool=pool-0') ; run('replication sources select source-000'); var packages = list(); </pre>

	<pre> for (var i = 0; i < packages.length; i++) { run('select ' + packages[i]); var proj_name = list(); if (proj_name == projName) { myPackage = packages[i] ; break; } run('cd ..'); } </pre>
<p>status_repl_src.aksh</p> <p>To check the replication status at the source site</p>	<pre> shares set pool=pool-0 select OFM-SITE1-KIT replication select action-000 show </pre>
<p>status_repl_tgt.aksh</p> <p>To check the replication status at the target site</p>	<pre> shares set pool=pool-0 replication sources select source-000 show </pre>
<p>stop_repl_at_source.aksh</p> <p>To stop the replication at the source</p>	<pre> script { var projName='OFM-SITE1-KIT'; printf("Stopping the replication for the project %s at the source \n", projName); run('cd /'); run('shares'); run('set pool=pool-0'); run('select ' + projName); run('replication select action-000'); run('set continuous=false'); run('commit'); printf("The replication is stopped.. Proceed with role reversal.. \n"); } </pre>

References

1. Oracle Maximum Availability Architecture Web site
<http://www.oracle.com/technetwork/database/features/availability/maa-090890.html>
2. Oracle Fusion Middleware Disaster Recovery Guide
http://download.oracle.com/docs/cd/E14571_01/doc.1111/e15250/toc.htm
3. Oracle Exalogic Enterprise Deployment Guide
http://download.oracle.com/docs/cd/E18476_01/doc.220/e18479/toc.htm
4. Oracle Exalogic Documentation Library
http://download.oracle.com/docs/cd/E18476_01/index.htm
5. Oracle Exadata Storage Server and Oracle Exadata Database Machine Documentation
The Oracle Exadata Storage Server and Oracle Exadata Database Machine Documentation can be found on the Exadata Storage cell in `/opt/oracle/cell/doc`
6. Oracle Data Guard: Disaster Recovery Best Practices for Exadata Database Machine
<http://www.oracle.com/technetwork/database/features/availability/maa-wp-dr-dbm-130065.pdf>
7. Oracle Database 11.2 Documentation Library
<http://www.oracle.com/pls/db112/homepage>
8. Oracle Database High Availability Library
http://www.oracle.com/pls/db112/portal.portal_db?selected=14
9. Oracle Fusion Middleware High Availability Library
http://download.oracle.com/docs/cd/E14571_01/admin.htm



Disaster Recovery for Oracle Exalogic Elastic
Cloud with Oracle Exadata Database Machine
May 2011

Author: Bharath K Reddy
Contributing Authors: Pradeep Bhat, Sridhar
Ranganathan

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



| Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.