Oracle Fusion Middleware Disaster Recovery Solution Using HP EVA Storage

Oracle Maximum Availability Architecture White Paper
September 2009

Maximum Availability Architecture

Oracle Best Practices For High Availability

ORACLE

Oracle Fusion Middleware Disaster Recovery Solution Using HP EVA Storage

Table of Contents

Oracle Fusion Middleware Disaster Recovery Solution Using HP EVA Storage

EXECUTIVE OVERVIEW

The Oracle Fusion Middleware Disaster Recovery solution uses disk replication technology provided by storage vendors for disaster protection of key information on file systems. In addition, Oracle Data Guard is used to replicate database content.

This document describes how to achieve disaster recovery for an Oracle Fusion Middleware 10*g* environment using HP StorageWorks EVA for data storage, HP StorageWorks Continuous Access EVA for data replication, and Oracle Data Guard 10*g* to keep the Oracle primary and standby databases synchronized.

INTRODUCTION

Enterprise deployments need protection from unforeseen disasters and natural calamities. One protection solution involves setting up a standby site at a geographically different location than the production site. All data including application data, metadata, configuration data, and security data is replicated to the standby site. The Oracle Fusion Middleware environment on a 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. Oracle Fusion Middleware Disaster Recovery solution is dependent on storage replication techniques to keep middle-tiers across production and standby sites synchronized in addition to using Oracle Data Guard 10*g* to keep the Oracle primary and standby databases synchronized.

ORACLE FUSION MIDDLEWARE DISASTER RECOVERY STRATEGY

Oracle Fusion Middleware (OFM) Disaster Recovery solution facilitates data protection for Oracle Fusion Middleware data and database content, as follows:

- To protect middleware product binaries, configuration, and metadata files, use disk replication technologies that are offered by storage vendors.
- To protect Oracle Database content, use Oracle Data Guard for disaster protection of Oracle Databases because of its superior level of protection and high availability. Oracle Data Guard protects the databases used for Oracle Fusion Middleware Repositories, as well as the databases for customer data.

- To protect non Oracle database content, use vendor-recommended solutions.

If a failure or a planned outage occurs on the production site, synchronization to the standby site stops. The services and applications are subsequently started on the standby site. The network traffic should then be routed to the standby site. As a result, the standby site becomes the new production site.

HP StorageWorks EVA Components

The Oracle Fusion Middleware Disaster Recovery solution depends on the HP StorageWorks components listed in the following table:

Type of Component	HP StorageWorks Components
Storage array	HP StorageWorks Enterprise Virtual Arrays (EVAs)
Fabric	The network of Fibre Channel switches that connects the arrays
Array management software	Replication Solutions Manager (RSM) and Command View EVA
Replication software	RSM and Command View EVA

HP StorageWorks Continuous Access EVA

HP StorageWorks Continuous Access EVA software is the critical component of the OFM Disaster Recovery Solution because it:

- Provides an array-based application that uses advanced replication technologies to replicate data over distances between HP StorageWorks Enterprise Virtual Arrays.
- Uses the graphical user interface (GUI) provided by Replication Solutions Manager (RSM) software to create, manage and configure remote replication on the entire HP StorageWorks EVA family of storage arrays.

Benefits

The HP Continuous Access EVA software provides the following benefits:

- Continuous replication of local virtual disks to remote virtual disks
- Synchronous and asynchronous replication modes
- Automated failover when used with other solution software
- Failsafe data protection
- Ability to suspend and resume replication
- Bidirectional replication
- Graphical and command line user interfaces
- Automatic suspension of replication if the links between arrays are down

- Support for array-to-array fan-in and fan-out

For more information, see the *HP StorageWorks Continuous Access EVA Implementation Guide* [3]

Hardware Configuration

For the HP StorageWorks Continuous Access EVA Software to work properly, some hardware configuration is required. Specifically, the EVA's at two different sites must be connected through Intersite links (ISL). Figure 1 shows the relationship between the hardware and software components for a typical remote replication configuration (described in the "Remote Replication" section).

In Figure 1, arrays on source and destination sites are connected by two linked fabrics, and two ISLs connect the fabrics between the source and destination sites.

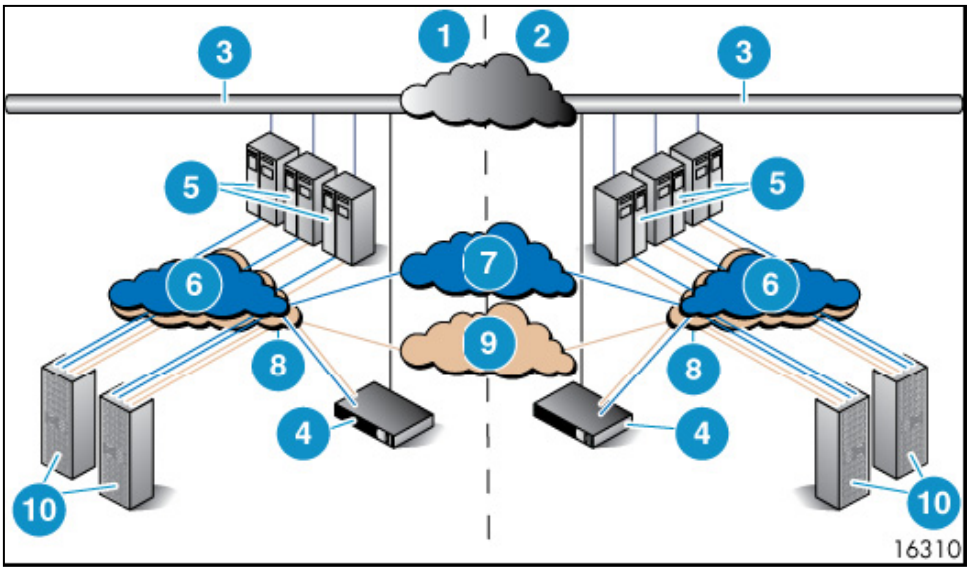

Figure 1: Basic HP Continuous Access EVA Configuration

1. Source site	6. Source/destination fabric—dark-colored line
2. Destination site	7. Intersite link—dark-colored line
3. LAN connection	8. Source/destination fabric—light-colored line
4. Management server	9. Intersite link— light-colored line
5. Hosts	10. Enterprise Virtual Arrays

Remote Replication

With HP StorageWorks Continuous Access EVA, **remote replication** is the continuous copying of data from selected virtual disks on a source (local) array to replica virtual disks on a destination (remote) array. Applications continue to run while data is replicated in the background. Remote replication requires a fabric

connection between the source and destination arrays and a software connection (DR group) between source virtual disks and destination virtual disks.

Remote replication operates in one of the following write modes:

- **Asynchronous**—The array acknowledges I/O completion before data is replicated on the destination array. Asynchronous write mode can be standard or enhanced, depending on the software version of the controller.
- **Synchronous**—The array acknowledges I/O completion only after the data is cached on both the source and destination arrays.

For Oracle Fusion Middleware Disaster Recovery, the best practice is to choose either the enhanced asynchronous replication mode, or synchronous replication write mode. For more information, see the "Choosing Replication Write Modes" section.

DR (Data Replication) Groups

A **DR group** is a logical group of virtual disks in a remote replication relationship between two arrays. Hosts write data to the virtual disks in the source array, and the array copies the data to the corresponding virtual disks in the destination array. I/O ordering is maintained across the virtual disks in a DR group, ensuring I/O consistency on the destination array in the event of a failure of the source array.

The virtual disks in a DR group:

- Fail over together
- Share a write history log (DR group log)
- Preserve the order of write I/Os within the DR group

A pair of source and destination virtual disks is called a **copy set**.

Figure 2 shows the replication of one DR group between a source array and a destination array. For more information, see the "Planning DR groups" section.

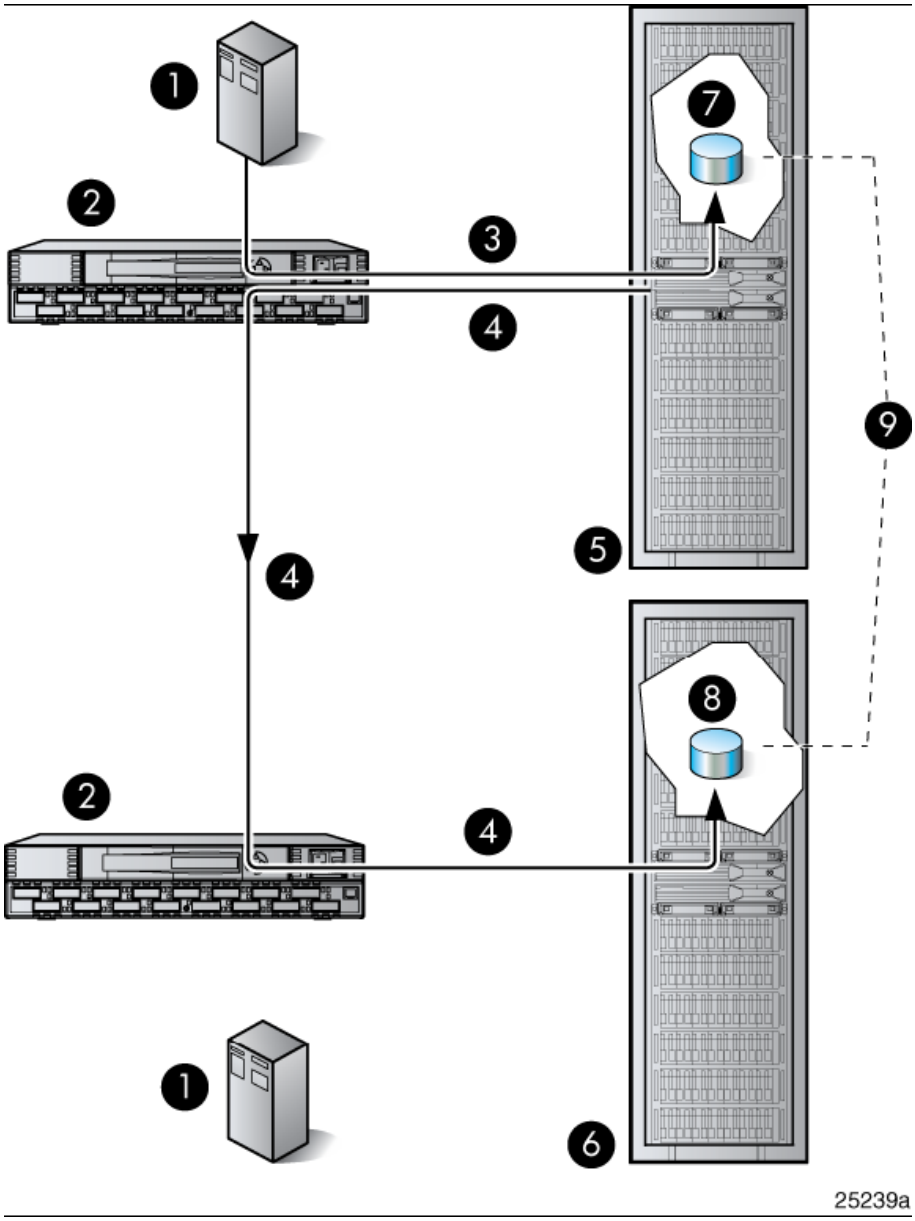

Figure 2: DR Group Replication

1. Host server
2. Fibre Channel switch
3. Host I/O
4. Replication writes
5. Source array
6. Destination array
7. Source virtual disk
8. Destination virtual disk
9. DR group

DR Group Write History Log

The DR group write history log is a virtual disk that stores a DR group's host write data. The log is created when you create the DR group. Once the log is created, it cannot be moved. For more information, see "Choosing Replication Write Modes" in the Best Practices section.

Managed Sets

Managed sets are a feature of HP Replication Solutions Manager. A managed set is a named collection of resources banded together for convenient management. Although managed sets can be used for a variety of different types of resources, our implementation used them only to aggregate DR groups. By performing an action on a managed set, in effect you are performing the action on all members of the set.

> **Note:** Managed sets enable you to manage multiple resources easily. They do not contribute to the data consistency of a DR group. Write order consistency is maintained at the DR group level.

In managed sets:

- You must ensure that all resources, or members, in a single managed set are of the same type (for example, all virtual disks).
- You can add a specific resource to more than one managed set.
- You can add resources on more than one array to a managed set.
- You should create separate managed sets for DR groups so that if a failover occurs, you can perform the actions that correspond to the changed source or destination role of the managed set members.

Failover

In HP Continuous Access EVA replication, failover reverses replication direction for a DR group. The destination array assumes the role of the source, and the source array assumes the role of the destination. For example, if a DR group on array A were replicating to array B, a failover would cause data for the DR group to be replicated from array B to array A. You can fail over a single DR group or you can failover multiple DR groups with a single command using a managed set. When you specify a failover action for a specific managed set, the failover occurs for all DR groups contained in the specified managed set. Without managed sets, you must fail over each DR group individually. For more information about failover settings, see the "Planning DR groups" section.

Best Practices for DR Configuration and Deployment

This white paper provides recommendations for configuration and deployment best practices to implement the Oracle Fusion Middleware Disaster Recovery solution. The DR solution uses the following technologies:

- Use HP StorageWorks Continuous Access EVA to protect all non-database content and to ensure a small RTO and RPO for the Oracle Fusion Middleware environment.
- Use Oracle Data Guard to protect all Oracle Database content.

Disaster Recovery Architecture

This section describes the recommended system configuration and operational best practices for deploying disaster-recovery architecture.

We tested these scenarios by using the disaster recovery architecture specified in this section and by adhering to the following validation points:

(1) Make configuration changes at the production site. For example, create a new WebLogic managed server and deploy a new application and validate them at the standby site to ensure successful configuration replication between the production and standby sites.

(2) Keep XA transactions pending at the production site and ensure the transactions are committed or rolled back at the standby site after performing a switchover or failover to ensure that transaction logs get replicated properly.

(3) Validate the Oracle Fusion Middleware components on the production site—such as, BPEL, ESB, and OWSM—and then revalidate the components on the standby site after performing a switchover or failover.

(4) Validate the Oracle Database replication, using Oracle Data Guard redo transport services, in the context of an Oracle Fusion Middleware SOA application.

System Configuration

Figure 3 depicts the system architecture to deploy key Oracle Fusion Middleware components and HP replication technologies. Each site consists of two Web hosts, an Oracle WebLogic Server Administration host, and two clustered application servers with Oracle Fusion Middleware. All non-database objects, including the Oracle binaries and Oracle Fusion Middleware configuration, are protected using HP Continuous Access EVA. Also, the Oracle Real Application Clusters (Oracle RAC) database is protected with Oracle Data Guard.

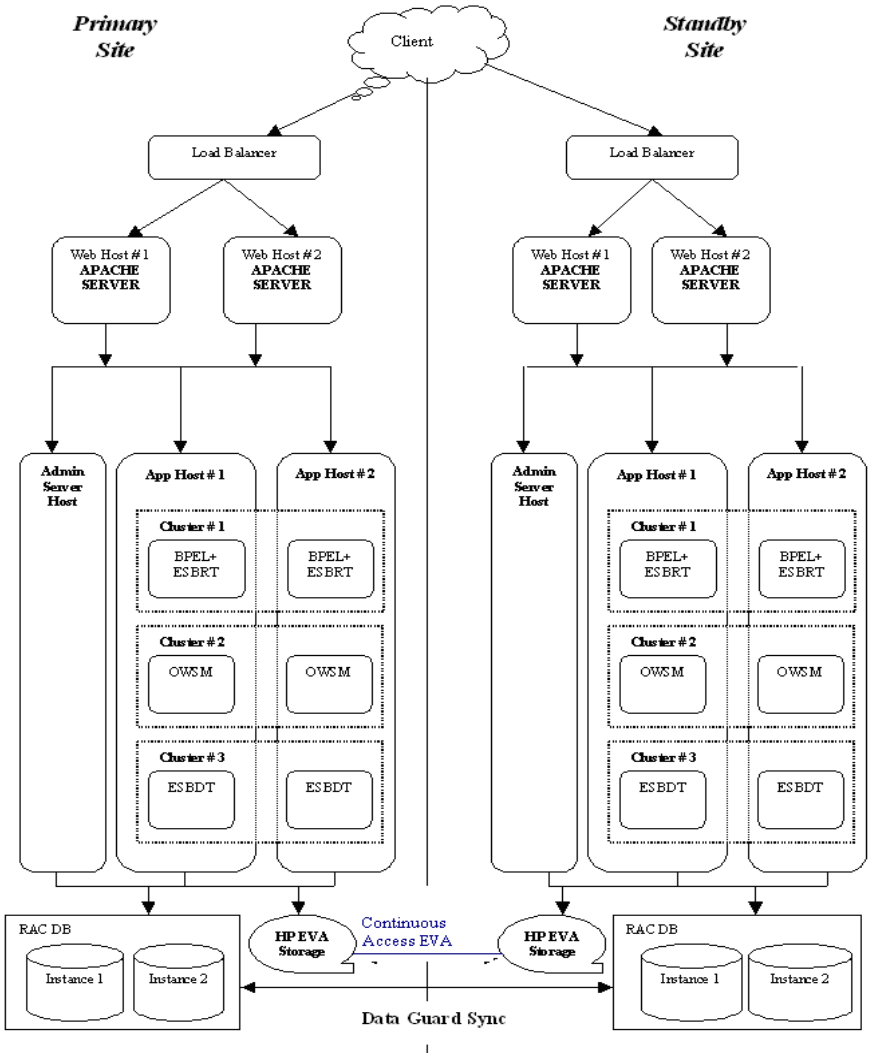

Figure 3: Reference Topology Diagram

Figure 4 shows the hardware system configuration for the disaster-recovery architecture.

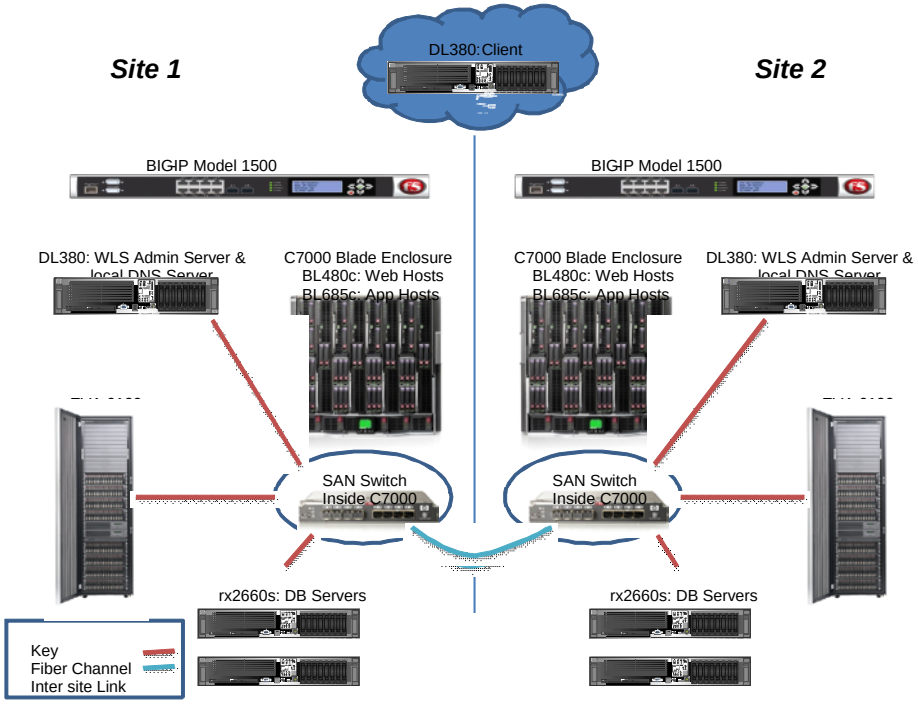

Figure 4: Hardware Diagram

The configuration in Figure 4 includes the following hardware and operating system software:

- Client Host Configuration (used for validation testing)

  Proliant DL380 (8 core, 2.8 GHz x86, 4 GB RAM)
  Microsoft Windows 2003 Version 5.2
  Selenium Remote Control v1.0

- Load Balancer (one per site)

  Site 1: BigIP Model 1500
  Site 2: BigIP Model 3400
  Software Version: BIG-IP 9.4.6 Build 401.0 Final

- Administration Server Host Configuration (one per site)

  Proliant DL380 (8 core, 2.8 GHz x86, 4 GB RAM)
  Oracle Enterprise Linux 5.1 (32 bit)
  Oracle Web Logic Server 9.2

- Web Host Configuration (two per site)

  Proliant BL480c (8 core, 3.2 GHz x86, 8 GB RAM)
  Oracle Enterprise Linux 5.1 (64 bit)

Apache Server 2.2.4

- App Host Configuration (two per site)

  Proliant BL685c (8 core, 2.6 GHz ADM, 16 GB RAM)
  Oracle Enterprise Linux 5.1 (64 bit)
  Oracle Web Logic Server 9.2
  Oracle SOA Suite 10.1.3.4

- Database Server Host (2 per site)

  HP Integrity rx2660 (4 core, 1.66 GHz Itanium 2, 32 GB RAM)
  HPUX 11iV3 March '08 HA OE
  Oracle Database release 10.2.0.4

- Storage Management Host

  Microsoft Windows 2003 Version 5.2
  Command View EVA 9.0
  Replication Solutions Manager 5.0.118

- Storage Array (one per site)

  HP StorageWorks EVA 6100
  2 x HSV Controllers
  2.25 TB raw capacity
  Firmware: XCS 6200

Best Practices and Recommendations

The following list summarizes the recommendations for implementing an Oracle Fusion Middleware Disaster Recover solution and provides links to the sections for additional information:

- Use disk replication to copy Oracle Fusion Middleware file systems and data
- Use Oracle Data Guard to copy database data
- Perform failover for unplanned downtime and switchover for planned downtime
- Assign virtual disks to DR Groups
- Set up the network configuration to reach services through a single URL
- Use the HP StorageWorks Command View utility to configure storage software
- Use RSM to manage DR Groups and other replication tasks
- Choose replication write modes
- Choose the Size of the Write History Log
- Use a persistent file-based store for high availability of JMS and TLogs

Use Disk Replication to Copy Oracle Fusion Middleware File Systems and Data

Use disk replication to copy Oracle Fusion Middleware file systems and data from the production site shared storage to the standby site shared storage. During normal operations, the production site is active and the standby site is passive. When the production site is active, the only write operations made to the standby site shared storage are the disk replication operations from the production site shared storage to the standby site shared storage.

Use Oracle Data Guard to Transport Database Redo Data

The Oracle Data Guard *redo transport services* transmit database redo data from the production database to the databases at standby site. On the standby site, the Data Guard *apply services* automatically apply redo data to standby databases to maintain synchronization with the production database and allow transactionally consistent access to the data. When the production site is active, the only write operations made to the standby databases are the database synchronization operations performed by Oracle Data Guard.

Perform Failover for Unplanned Downtime and Switchover for Planned Downtime

When the production site becomes unavailable, you enable the standby site to take over the production role by performing a failover or a switchover:

- If the current production site becomes unavailable unexpectedly, then perform a site failover operation to enable the standby site to assume the production role.
- If the current production site is taken down intentionally (for example, for planned maintenance), then perform a site switchover operation to enable the standby site to assume the production role.

Assign Virtual Disks to DR Groups

The general rule for assigning virtual disks to DR groups is that virtual disks associated with the same *application* should be configured in the same DR group. In this context, the term "application" refers to a distinct software installation.

In the Oracle Fusion Middleware environment, there are three distinct software installations:

- Apache Server installation on the Web Hosts.

  For the Apache installation, there is a single file system (virtual disk) for each host. Thus, the associated DR group consists of two virtual disks, with one file system configured for each Web Host.

- Oracle SOA Suite release 10.1.3.4 on Oracle WebLogic Server release 9.2 installation on the Managed Server hosts.

  For the Managed Servers, there are two file systems per host (home and configuration). Thus, the associated DR group consists of four virtual disks, with two virtual disks configured for each Managed Server.

- Oracle WebLogic Server release 9.2 management installation on the Administration Server host.

  For the Administration Server, there are two file systems (home and configuration), plus there is a data file system for the Java Message Services (JMS) and Transaction Logs (TLogs) that the Administration Server shares with the Managed Servers. (See the section, "Remote Replication," for a description about how this file system is a candidate for synchronous replication to prevent any data loss.)

  Therefore, the Administration Server must have its own DR group. This configuration requires you configure two DR groups: one DR group for the home and configuration file systems, and another DR group for the data file system that supports the JMS and TLogs because of their different synchronization requirement. Thus, there is a total of four DR groups, as shown in the following table.

Oracle Fusion Middleware DR Groups			
DR Group Name	**Host(s)**	**File systems**	**# of virtual disks**
Admin	Administration Server	/oracle/home, /oracle/config	2
Web Host	WebHost1, WebHost2	/oracle (Apache install)	2
Managed	AppHost1, AppHost2	/oracle/home, /oracle/config	4
JMS-TLOGS	Admin, AppHost1, AppHost2	/oracle/data	3

Examples Showing How to Mount File Systems

```
/usr/bin/rsh `hostname` -l root mount /oracle
/usr/bin/rsh `hostname` -l root mount /oracle/config
/usr/bin/rsh `hostname` -l root mount /oracle/home
/usr/bin/rsh `hostname` -l root mount /oracle/data
```

Set Up the Network Configuration to Reach Services Through a Single URL

Use a single URL to reach the services provided by Oracle Fusion Middleware.

When a site switchover or site failover occurs, the hostnames and IP addresses associated with the primary site will continue to work on the standby site, because

the DNS server on the secondary site has aliases for all of the hosts on the primary site.

For example:

- Primary site: Web Host webhost1_1; IP address 10.10.11.101
- Secondary site: Web Host webhost1_2; IP address 10.10.11.102.

To enable the same name to be used on the secondary site, the DNS server on the secondary site would contain an alias (canonical name) for `webhost1_1`, associating it with `webhost1_2`. The `BigIP` load balancer would use the actual names and IP addresses of the hosts to do its load balancing, but the Oracle Fusion Middleware configuration would still be able to make use of the primary site names.

In our test environment, the client host can see both sites because the primary site's DNS server was its primary DNS server and the secondary site's DNS server was its alternate DNS server. However, similar to any client out on the Web, the client software only knows how to reach the services provided by Oracle Fusion Middleware through a single URL. For example:

`http://bigiplb.mycompany.com:<port number>/console`

Using a single URL works well, because both the primary and secondary site DNS servers have aliases for `bigiplb`, even though the aliases point to different BigIP load balancers. When a failover occurs, the client host is able to detect that its primary DNS server was down, and begin to use its alternate DNS server (on the secondary) site, yet the client software is unaware that anything has changed.

Use the HP StorageWorks Command View Utility to Configure Storage Software

Use the HP StorageWorks Command View utility to manage HP StorageWorks. The Command View Utility runs a Management Server host that is connected via Fiber Channel to the SAN associated with the EVA. The following screenshot shows a typical EVA command view screen:

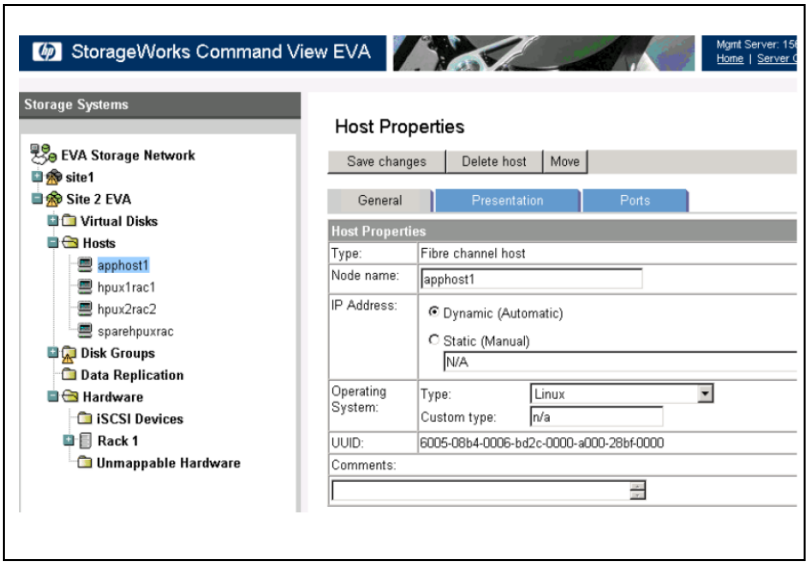

Through the Command View GUI (graphical user interface), you can perform the following array management tasks:

- Install licenses
  Ensures that the EVA licenses are updated, including the EVA command view utility, and Continuous Access. The command view license needs to be version 9.0 (or a later version), which is the minimum required release that supports asynchronous failover.
- Upgrade Firmware (XCS)
  Ensure that the latest firmware is installed on the EVA controllers.
- Create and Manage Disk Groups
- Define Hosts
- Create Virtual Disks
- Present Virtual Disks to Hosts
- Create DR groups
- Add and Remove members from DR groups

Use Replication Solutions Manager to Manage DR Group and Other Replication Tasks

HP provides Replication Solutions Manager (RSM), which is a special utility designed specifically for data replication. This tool provides a robust set of capabilities for managing DR groups and other features associated with data replication including creating and failing over Managed Sets.

You can use RSM to:

- Automatically discover array, virtual disk, host, and application resources
- Present virtual disks for host access
- Copy virtual disks and host volumes using snapshot, snap clone, and mirror clone technology
- Remotely replicate and fail over virtual disks
- Group and replicate resources as a unit called a managed set
- Dynamically mount virtual disks on enabled hosts
- Automate replication tasks using replication manager jobs and schedule jobs for future purposes
- Monitor replication status by array
- Back up and restore replication manager configuration and jobs
- Replicate application resources on enabled hosts (for example, Oracle tablespaces)
- Visually manage replication resources with the topology viewer

- Configure security using operating system authentication and user administration with audit capabilities
- Integrate with a variety of backup and recovery solutions
- Create round-robin snapshots and snap clones of host volumes using the host volume replication wizard. The oldest replica is deleted automatically, allowing the space it occupied to be reused
- Create virtual disks, containers and DR groups
- Manually or automatically perform dynamic capacity management

Note: You cannot use the RSM utility to perform functions such as installing licenses or upgrading firmware. You should manage these types of functions with the HP StorageWorks Command View Utility.

Choose Replication Write Modes

For Oracle Fusion Middleware Disaster Recovery, Enhanced Asynchronous and Synchronous replication write modes are recommended. The replication write mode can be set on specific DR groups depending upon frequency of data change. For example, Synchronous mode should be set for DR group pertaining to Tlogs due to its dynamic nature.

The base Asynchronous replication write mode is not recommended

- **Synchronous replication** prevents any loss of data, however it also requires each write I/O to be completed on the destination array before it is considered completed for the source array. Therefore, in an environment where there are a lot of write I/O's, synchronous replication is a potential drag on performance.
- Enhanced asynchronous replication. Fortunately, there is a solution that is nearly as robust as synchronous replication and that is **enhanced asynchronous replication**. In enhanced asynchronous replication, write I/Os do not have to complete on the destination array before they are marked as completed locally.

At the same time, there is protection against data loss, because each write I/O is written both to the source array and to the DR group write history log before it is considered to be complete. The write history log is written in the same order that the write I/O's are written to the source array. As the write I/O's are propagated to the destination array, they are removed from the write history log, so the write history log is a sequential record of all write I/Os written to the source array that have not yet been acknowledged to be completed on the destination array.

In the event of a failure while enhanced asynchronous write mode is being used, all pending write I/O's are preserved in the write history log. In this scenario, one option is to simply wait until the source array can be brought back up. If the failure is only temporary and can be corrected in a short period of time, this is probably the best option, because it ensures that no data will be lost.

In the case where the failure is not temporary or the production environment needs to be brought back online quickly, the customer will have to fail over the production site to the standby site. In enhanced asynchronous write mode, this means that all pending write I/O's in the write history log will be lost. The number of writes lost can be minimized if the writes are being processed quickly, and the therefore the number of pending writes is low. The rate of write processing should be estimated by customers when they are setting their RPO. The RPO is dependent on the bandwidth of the inter site link, which is in turn dependent on the distance between the arrays, the type of interconnect, and other factors. Careful analysis of the application's write profile and the replication link speed can determine what the worst case RPO will be for the solution. For complete details on RPO's, bandwidth, and inter site links, see the *HP StorageWorks Continuous Access EVA Implementation Guide* [3].

While it is possible that a failover using enhanced asynchronous write mode could result in zero data loss if the write log is empty, enhanced asynchronous replication is can never be guaranteed to achieve that objective. Synchronous replication is the only way to guarantee zero data loss. While an Oracle Fusion Middleware environment does not generally require synchronous replication, there is one part of the environment for which customers may want synchronous replication, the file systems used for the Java Message Services (JMS) and Transaction Logs (TLogs). For these logs to be replicated synchronously when the rest of environment is using enhanced asynchronous replication, they must be in a separate file system and a different DR group. For more details, see the "Planning DR Groups" section.

The third write mode option, asynchronous without the write history log, is not recommended. It is the only asynchronous option available for older versions of firmware (pre XCS 6.xxx versions). To use enhanced asynchronous mode, both source and destination arrays must be running XCS 6.xxx firmware or greater. For arrays that are not capable for running XCS 6.xxx or greater, we would recommend either upgrading to new storage arrays or running in synchronous mode.

Choose the Size of the Write History Log

For enhanced asynchronous write mode[1] to work properly, the write history log must be large enough to hold all write I/Os for a system that is under peak load. This is important because:

- A full write log results in a process called **normalization**, which will force a synchronization of the source and destination arrays. Under peak load, a forced normalization would have a very negative impact on performance.

[1] Enhanced asynchronous write mode is the preferred write mode for most parts of the Oracle Fusion Middleware environment

- Setting the size of the write history log correctly from the point when the DR group is created avoids the need to:
    - Switch to synchronous (which is required when changing the size of the write history log)
    - Drain the write log
    - Switch back to enhanced asynchronous mode

Use a Persistent File-Based Store for High Availability of JMS and TLogs

The WebLogic application servers are usually clustered for high-availability. For high availability of the SOA Suite within a site, a persistent file-based store is used for the Java Message Services (JMS) and Transaction Logs (TLogs). This file store needs to reside on shared disk, which is accessible by all members of the cluster.

MANAGING PLANNED AND UNPLANNED DOWNTIME

This section documents the step-by-step procedures used to handle planned (switchover and switchback) and unplanned outages (failover and failback).

Site Switchover Procedures

Switchovers are planned operations done for periodic validation or to perform planned maintenance on the current production site. During a switchover, the current standby site becomes the new production site, and the current production site becomes the new standby site. The following steps need to be performed in sequence –

1. Shut down all the Oracle Fusion Middleware components on the production site either manually or using relevant management software.
2. If the replication write mode is not already set to synchronous, then change the write mode to synchronous and wait for any pending write I/Os to complete.
3. Unmount the file systems associated with the DR groups on the production site.
4. Switch over the database to the standby site using Oracle Data Guard and ensure the standby database comes up successfully.
5. Fail over the DR groups associated with Oracle Fusion Middleware.

The following screenshots demonstrate the failover process for one DR group; this process needs to be repeated for all the DR groups.

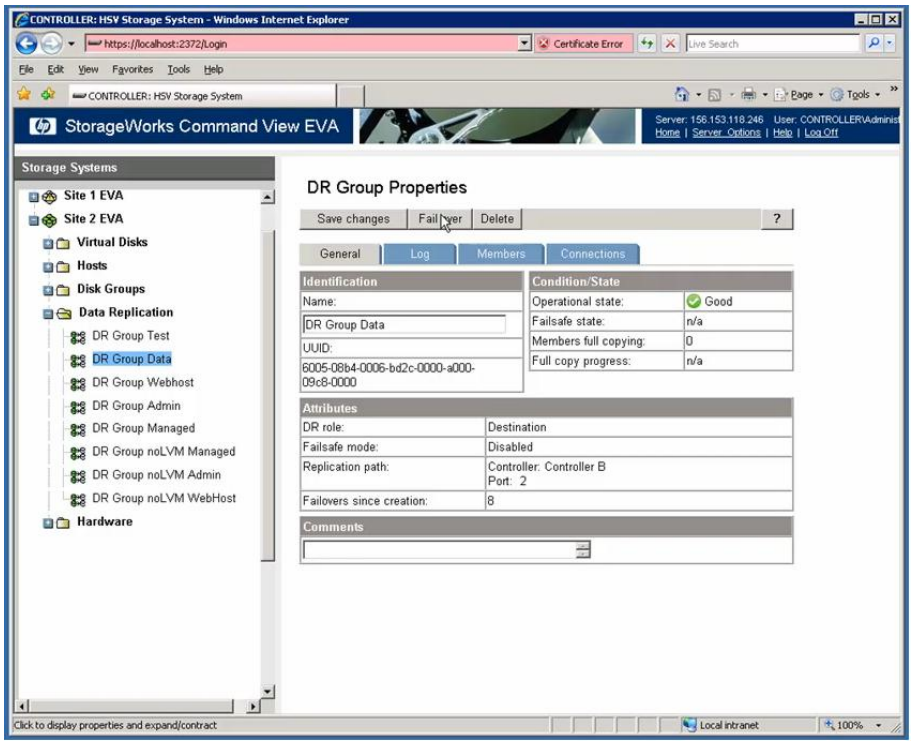

Figure 6 Select Failover for a DR Group

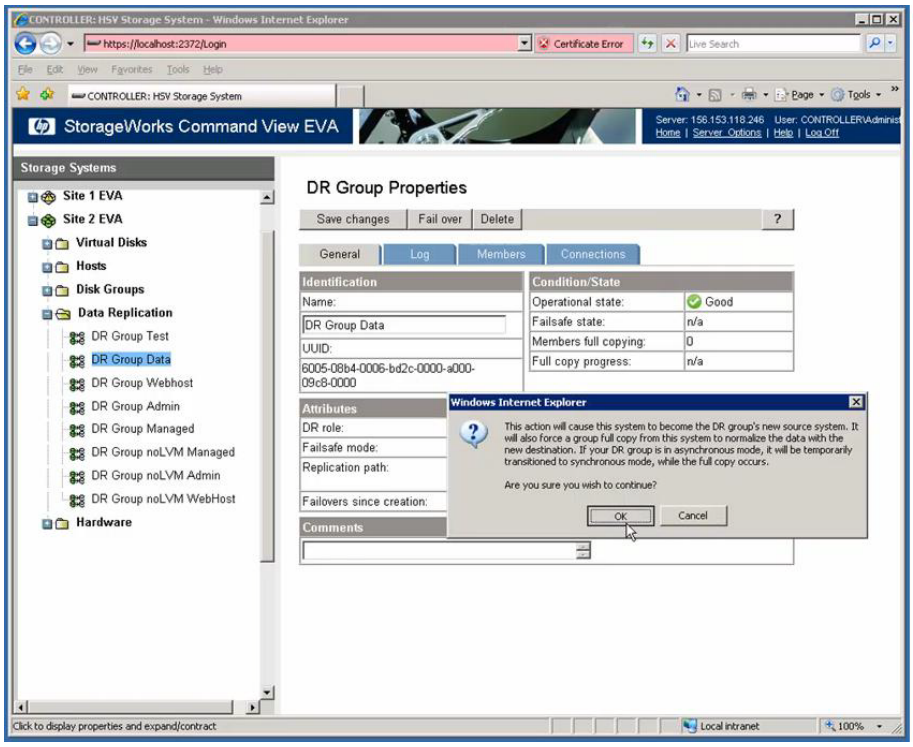

Figure 7 DR Group Failover Warning

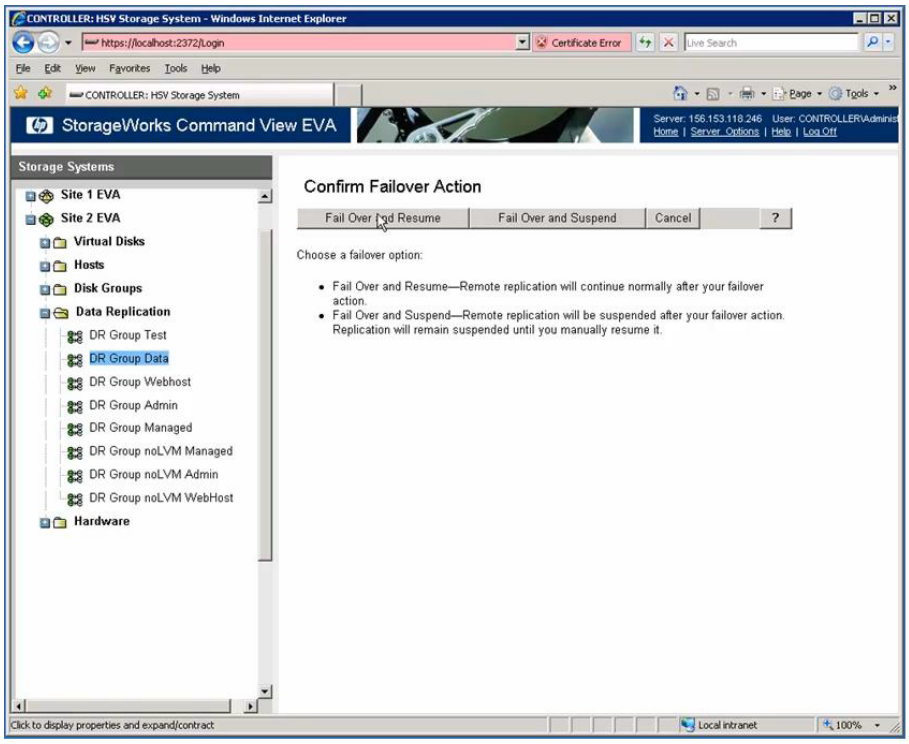

Figure 8 Confirm Failover for DR Group

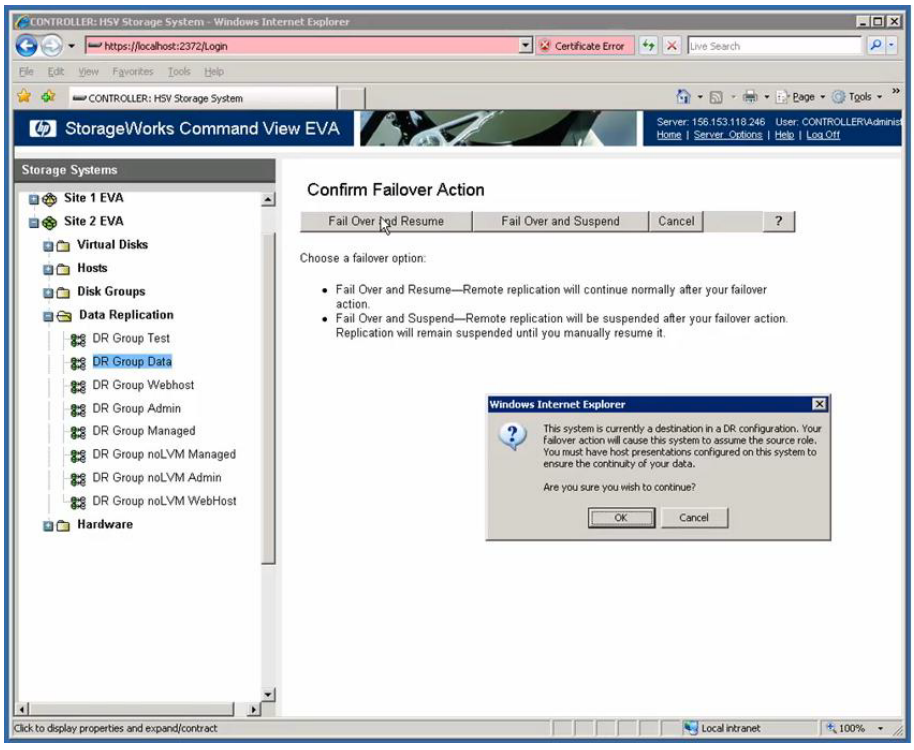

Figure 9 Confirm Failover Action for DR Group

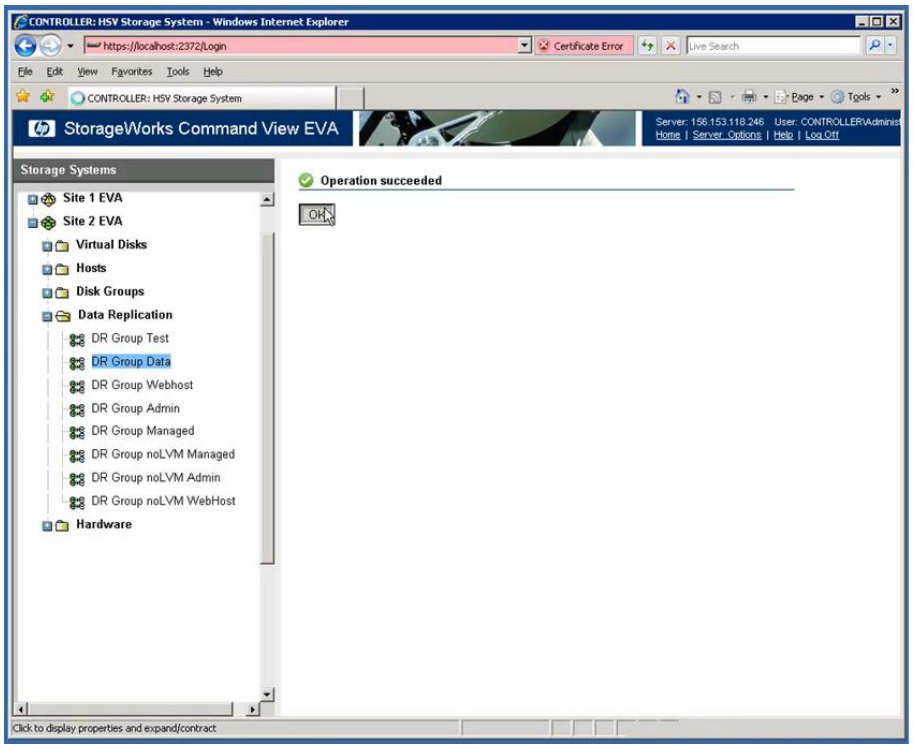

Figure 10 DR Group Failover Completion

6. Reconfigure the replication write mode to its original asynchronous setting, if any of the DR groups settings were modified in Step 2
7. Run device scan utility on each of the standby site's middle-tier hosts to ensure that failed over storage is now visible. This is a process in which the host scans all the devices presented to it. It may be necessary to perform this step on each server node to make sure all the LUNs are visible
8. Mount the file systems associated with the DR groups on the standby site
9. Start all the Oracle Fusion Middleware components on standby site and ensure they are started successfully
10. Ensure all user requests are routed to the standby (new production) site. This can be achieved through a global DNS push or something similar
11. The standby site has assumed the role of the production site and vice versa
12. Use a browser client to perform post-switchover testing to confirm that requests are being resolved and redirected to the new production site

Site Switchback Procedures

Repeat all the steps in the "Site Switchover Procedures" section to switch back to the original production site.

Site Failover Procedures

Failover is the process of making the current standby site the new production site after the production site becomes unexpectedly unavailable (for example, due to a disaster at the production site). Perform the following steps in sequence:

1. Detect that the production site no longer available by pinging the host.
2. Determine the actual status of production site. For example, has a real disaster occurred, or it something simpler like a network node failure? Can it be corrected without performing a site failover?
3. If you determine that a site failover is required, then fail over the DR groups associated Oracle middleware to the standby site using HP Command View EVA, as shown in 'Site Switchover Procedures" section.
4. Run device scan utility on each of the standby site's middle-tier hosts to ensure that the failed over storage is now visible. This is process in which host runs a scan on all the devices presented to it. It may be required to be done on each server node to make sure all the LUNs are visible.
5. Mount the file systems on standby site hosts.
6. Fail over the database using Oracle data guard and ensure it is started successfully.
7. Start Oracle Fusion Middleware components on the standby site and ensure they are started successfully.
8. Ensure that all user requests are routed to the standby site by performing a global DNS push.
9. At this point, the standby site has assumed the role of production site.
10. Use a browser client to perform post-failover testing to confirm that requests are being resolved and redirected to the new production site.
11. When the original production site is back up, resynchronize the DR groups on the new production site (old standby site) with the DR groups on the original productions site.
12. If the DR groups on original production site are not recoverable, then re-create them and resynchronize the DR groups on the new production site (old standby site) with the DR groups on the original productions site.

Site Failback Procedures

Failback involves re-instantiating both database and middle tiers.

CONCLUSION

Storage replication is a key requirement when providing disaster recovery protection for Oracle Fusion Middleware environments. HP StorageWorks Continuous Access EVA provides comprehensive features to handle all the unique requirements for replicating Oracle Fusion Middleware components in conjunction with Oracle Data Guard for Oracle database replication. HP Continuous Access EVA provides different modes of replication techniques that could be used depending on product requirements to ensure that customer environments are protected against any unforeseen disasters.

Using HP StorageWorks Continuous Access EVA with Oracle Data Guard provides users with the maximum benefit out of their investment to protect their entire Oracle environment.

APPENDIX A: TERMINOLOGY

Oracle Disaster Recovery Terminology

- **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.
- **Oracle Fusion Middleware (OFM)**—A collection of standards-based software products that spans a range of tools and services, from Java EE and developer tools, to integration services, business intelligence, and collaboration.
- **SOA (Service Oriented Architecture) suite**—An architecture with infrastructure components, such as, BPEL, ESB, and OWSM.
- **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 done for periodic validation or to perform planned maintenance on the current production site. During a switchover, the current standby site becomes the new production site, and the current production site becomes the new standby.
- **Site 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.
- **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 site 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.
- **Weblogic Server Transaction Logs**—Each Weblogic Server instance has a transaction log that captures information about committed transactions that may not have completed. The transaction logs enable Weblogic Server to recover transactions that could not be completed before the server failed.

- **Recovery Point Objective (RPO)**—The 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 systems back to the state they were in as of no longer than six hours ago.
- **Recovery Time Objective (RTO)**—The time needed to recover from a disaster, which is usually determined by how long you could afford to be without your systems.

HP Continuous Access EVA Terminology

- **Array**—See virtual array and storage system.
- **Asynchronous**—A descriptive term for computing models that eliminate timing dependencies between sequential processes. In asynchronous write mode, the array controller acknowledges that data has been written at the source before the data is copied at the destination. Asynchronous mode is an optional DR group property. See also synchronous.
- **Bandwidth**—The transmission capacity of a link or system, usually measured in bits per second.
- **Copy Set**—A pair of source and destination virtual disks.
- **Disk Group**—A named group of disks selected from all available disks in an array. One or more virtual disks can be created from a disk group.
- **DR Group**—Data replication group. A logical group of virtual disks in a remote replication relationship with a corresponding group on another array.
- **Enhanced Asynchronous**—A write mode in which all host write I/Os are added to write history log. The controller then acknowledges that data has been written at the source before the data is copied at the destination.
- **Enterprise Virtual Array (EVA)**—An HP StorageWorks product that consists of one or more virtual arrays. See also virtual arrays.
- **Fabric**—A network of Fiber Channel switches or hubs and other devices.
- **HP Continuous Access EVA**—An HP StorageWorks product consisting of two or more arrays performing disk-to-disk replication, along with the management user interfaces that facilitate configuring, monitoring, and maintaining the replicating capabilities of the arrays.
- **Intersite Link**—A connection from an E-port on a local switch to an E-port on a remote switch.
- **LUN (Logical unit number)** —LUNs are the components within SCSI targets that execute I/O commands. Virtual disks that are presented to hosts correspond to logical units and are identified by LUN IDs.

- **Managed Set**—Selected resources that are grouped for convenient management. For example, you can create a managed set to manage all DR groups whose sources reside in the same rack.
- **Management Server**—A server, on which HP StorageWorks Enterprise Virtual Array (EVA) management software is installed, including HP StorageWorks Command View EVA and HP StorageWorks Replication Solutions Manager, if used. A dedicated management server runs EVA management software exclusively.
- **Mount Point**—The file system path and directory for a host volume.
- **Normalization**—The initial copy that occurs between source and destination virtual disks or any complete re-synchronization that occurs after the initial copy.
- **Present LUN**—The process in which the Management Console of the storage makes the LUN or virtual disk to be presented (made visible) to the World Wide ID (WWID) of the host (database or Middleware) server QLogic HBA.
- **Remote Copy**—A virtual disk on the destination array that is a replica of a virtual disk in the source array.
- **Source**—The virtual disk, DR group, or virtual array where I/O is stored before replication. See also Source-Destination Pair.
- **Storage Area Network (SAN)**—A network of storage devices and the initiators that store and retrieve information on those devices, including the communication infrastructure.
- **Storage System**—Synonymous with virtual array. The HP StorageWorks Enterprise Virtual Array consists of one or more storage systems. See also virtual array.
- **Synchronous**—A descriptive term for computing models that perform tasks in chronological order without interruption. In synchronous write mode, the source waits for data to be copied at the destination before acknowledging that it has been written at the source. See asynchronous.
- **Virtual Array**—Synonymous with disk array and storage system, a group of disks in one or more disk enclosures combined with control software that presents disk storage capacity as one or more virtual disks.
- **Virtual Disk**—Variable disk capacity that is defined and managed by the array controller and presentable to hosts as a disk.
- **XCS**—The HP Enterprise Virtual Array software on specific EVA controller models. Controller software manages all aspects of array operation, including communication with HP StorageWorks Command View EVA.

REFERENCES

1. *Oracle's Middleware Disaster Recovery Guide*
2. *HP StorageWorks Command View EVA User Guide*
3. *HP StorageWorks Continuous Access EVA Implementation Guide*
4. HP StorageWorks Continuous Access EVA Software
5. Configure DR Solution using HP EVA Continuous Access
6. *HP StorageWorks Enterprise Virtual Array Compatibility Reference*
7. "HP StorageWorks Enterprise Virtual Array Configuration Best Practices" white paper
8. HP StorageWorks Replication Solutions Manager help and user guide
9. HP licenses installation

Oracle Fusion Middleware Disaster Recovery Solution Using HP EVA Storage

September, 2009
Author: Anuj Sahni and Sunita Sharma from Oracle, Bill Cortright and Sathya Krishnaswamy from HP
Contributing Authors: Pradeep Bhat
Editor: Viv Schupmann

Oracle USA, Inc.
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