

*Oracle Fusion Middleware Disaster Recovery
Solution using Oracle's Sun ZFS Storage
Appliance*

September 2010



Maximum Availability Architecture

Oracle Best Practices For High Availability

Introduction.....	3
About the Sun ZFS Storage Appliance.....	4
Architecture Overview.....	4
Platforms.....	4
High Availability Architecture.....	5
Storage Concept and Terminology.....	5
Data Services.....	6
Oracle Fusion Middleware Disaster Recovery Strategy.....	9
Oracle Fusion Middleware Deployment Procedure.....	14
1. Sun ZFS Storage Appliance File Systems Configuration.....	14
2. Oracle Fusion Middleware Hosts Setup.....	16
3. Oracle Fusion Middleware Deployment	17
4. Oracle Database Deployment (Primary and Standby).....	17
5. Sun ZFS Storage Appliance - Remote Replication Setup	19
6. Set Up Replication for the FMW Project.....	19
Managing Planned and Unplanned Downtime.....	20
Site Switchover Procedures.....	20
Site Switchback Procedures.....	22
Site Failover Procedures.....	22

Benefits of deploying Oracle FMW using Sun ZFS Storage Appliance.....	23
Best Practices for Oracle FMW DR Configuration and Deployment.....	27
Conclusion.....	28
Appendix.....	29
Terminology.....	29
Oracle Fusion Middleware DR Solution Test Procedure.....	32
Oracle Data Guard setup for the FMW database	33
Sample scripts	34
Sun ZFS Storage Appliance – File System Configuration via GUI (Sample)...	39
References.....	41

Introduction

Oracle Fusion Middleware (FMW) is the foundation for the Oracle application infrastructure, which enables business applications to be created and run efficiently across the enterprise. Protecting that infrastructure is critical for the business. Oracle's FMW disaster recovery solution protects the infrastructure and provides business continuity during a planned or unplanned event. In this solution, the entire Oracle FMW infrastructure is replicated to a remote site, enabling the latest business data to be made available during a site failure or site maintenance event.

At the primary site, Oracle FMW infrastructure is deployed on Oracle's Sun ZFS Storage Appliance, also referred to as *the Appliance* in this document. All the FMW binaries, application data, metadata, configuration data, logs, and security data are replicated to the remote site (also referred to as *standby site* in this document) using the *Remote Replication* feature of the Appliance. The Oracle database component of the FMW architecture is replicated to the standby site using the Oracle Data Guard feature.

The standby site is configured to be in a passive mode. It is activated only during the switchover or the failover process. This deployment model is also referred to as an *active/passive* model.

This technical document covers the following topics :

- About Sun ZFS Storage Appliance
- Oracle FMW DR architecture, strategy
- Deploying the Oracle FMW DR solution with the Appliance
- Benefits and best practices for Oracle FMW DR solution with the Appliance

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 Fusion Middleware components, and the concepts of Oracle databases and Oracle Data Guard features. For details, please refer to the documents listed in the reference section.

About the Sun ZFS Storage Appliance

This section provides Sun ZFS Storage Appliance's architecture overview, platform details, concept and terminology, and the major features.

Architecture Overview

The Sun ZFS Storage 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 Appliance is available either as single head or a clustered head for high availability. The Solaris operating system is the core of the appliance with the ZFS file system powering all the data storage, management and data services. All of these components are wrapped around with an intuitive user interface for easier management activities.

The Appliance architecture utilizes the *Hybrid Storage Pool* (HSP) model where the integrated DRAM, flash and the physical disks are seamlessly integrated for efficient data placement. Based on the user IO request and pattern, the data movement between these tiers is automatically handled by the appliance. The storage also includes a powerful performance monitoring tool called *Analytics* which provide details about the performance of various components including the network, storage, filesystems, client access and so on. There are plenty of drill-down options available. For example, the user can monitor which clients are accessing which file systems, which files, the latency, size of transfer and so on.

The Appliance also offers a variety of RAID protections to balance capacity, protection, and performance requirement of the applications.

Platforms

The Sun ZFS Storage Appliance is available in four platforms to meet our customer's requirements for price, performance, capacity, and protection capabilities. Mid-to-high-end platforms offer up to 2TBytes of read cache which enables the read response time typically in the low, single-digit milliseconds. The write flash on all the platforms provides write response time for synchronous writes with less than 1ms.

The following table illustrates the major hardware configuration information for the various platform offering.

Platform	Storage Capacity	Processor	Memory (DRAM)	Write Optimized SSD	Read Optimized SSD	Cluster Option
Sun ZFS Storage 7120	Up to 60 x 2TB SAS Disks [120TB]	1 x Quad Core Intel Westmere EP E5620 @ 2.4GHz	Up to 36GB	Up to 96GB	N/A	N
Sun ZFS Storage 7320 [details are per controller]	Up to 96 x 2TB SAS Disks [192TB]	2 x Quad Core Intel Westmere EP E5620 @ 2.4GHz	Up to 72GB	Up to 16 x 18GB	Up to 4 x 512GB	Y
Sun ZFS Storage 7420 [details are per controller]	Up to 576 x 2TB SAS Disks [1.1PB]	4 x 6C Intel Nehalem EX E7530 @ 1.86GHz [or] 4 x 8C Intel Nehalem EX X7550 @ 2GHz	Up to 512GB	Up to 96 x 18GB	Up to 4 x 512GB	Y
Sun ZFS Storage 7720	Expandable racks. Each Rack 720 TB	4 x 8C Intel Nehalem EX X7550 @ 2GHz	Up to 512GB per controller	2 x 18GB per cage	Up to 4 x 512GB per controller	Y

High Availability Architecture

The Sun ZFS Storage 7320 , Sun ZFS Storage 7420, and the Sun ZFS Storage 7720 platforms have clustering capabilities. Comprised of two appliance heads and shared storage, clustering provides improved availability when one of the heads succumbs to certain hardware or software failures. A cluster contains exactly two appliances or storage controllers, also referred as *heads*. Each head is assigned a collection of storage (one ore more storage pools), networking, and other resources from the set available to the cluster, which allows the construction of either of *active/active* or *active/passive* topologies. The terminology *active* here means the head that serves a storage pool. The other head, which is *passive* for the pool, is activated during a storage head fail-over. Two heads cannot be active for the same storage pool. It is one-to-many relationship between the the head and the pools.

Storage Concept and Terminology

Storage Pool

The storage pool (similar to a volume group) is created over a set of physical disks. Filesystems are then created over the storage pool. One or more storage pools are created over the available physical disks and flash drives are assigned. The storage pool is configured with a RAID layout such as mirrored , RAID-Z (single parity), RAID-Z2 (dual parity) and so on.

Project

All filesystems 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 filesystems and LUNs that are exported over supported data protocols to clients of the appliance. Filesystems export a file-based hierarchy and can be accessed over CIFS, NFS, HTTP/WebDay, 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. A single project can contain both filesystems and LUNs, and they share the same namespace.

Data Services

The Appliance comes with a license-free model that includes multiple data services. Remote replication, snapshots and cloning are key features for a complete disaster recovery (DR) and business continuity solution. The intention of this section is to provide an overview for these data services. For further details, please refer to the documents listed in the reference section.

Remote Replication

Data is replicated from the primary to the standby site, with data blocks asynchronously streamed to the remote storage appliance. Source data is modified at the granularity of a ZFS transaction; therefore the data is always consistent. Modified data is replicated to the target site which ensures the target side data is also consistent.

How Replication Works

Replication can be set up at a project-level or at the share-level. In project-level (consistency group) replication, all the shares (filesystems and LUNs) in the project are replicated. An implicit, consistent snapshot is performed for the share or all the shares in a project, and the data from the snapshot is streamed to the target site. The target's receipt of the replication is called *package*. After a successful receipt of the package at the target, another implicit snapshot is performed at the source and this time only the incremental data between the two snapshots is replicated. The write ordering and consistency is preserved across the shares in the project.

One source can be replicated to one or more targets. Likewise, a target system can receive packages from multiple sources. The data is optionally encrypted using SSL and then transmitted. Replication can be configured to happen over a dedicated line or over a public network. Replication is supported between different platforms of the Appliance.

Modes of Replication

The Sun ZFS Storage Appliance supports *scheduled*, *on-demand*, and *continuous* modes of replication. In all the modes, the underlying architecture is similar and the replication occurs asynchronously.

1. Scheduled Replication

In this mode, the user can define a schedule for the replication to occur automatically. If a *schedule* is established, then the replication occurs at the defined schedule. The schedule can be for every half-hour, hour, day, week, or month. This mode is preferred in situations where the replication is preferred to happen during off peak time (or) the backup is scheduled at the target site at specific time and so on.

2. On-demand Replication

Also addressed 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.

3. 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.

Role Reversal: Switchover/Failover & Switchback/Failback

In a DR setup, if a failover or switchover process is needed, the *role reversal* procedure is initiated. This process transforms the target site into a primary site and the source to become the standby. The direction of the replication is reversed. The last successfully received package data is used as the base at the target site. With role reversal, only the changes registered after the switchover/failover are replicated, thereby eliminating the need for a full replication.

Snapshots

The Sun ZFS Storage Appliance has unlimited snapshot capability. Snapshots are the read-only point-in-time copies of a filesystem, instantaneously created and with no space allocated initially. Blocks are allocated as and when changes are made to the base file system (copy-on-write). Snapshots are either initiated manually or can be automated by scheduling at specific intervals. These snapshot data can be directly accessed for any backup purposes.

Any reads to the snapshot blocks are served by the base file system's block. As the changes happen to the base file system, the older block referenced by the snapshot and the new changed block is referenced by the file system.

Project snapshots are the equivalent of performing snapshots on all shares within the project.

Clones

The Appliance supports an unlimited number of clones. A clone is an instantaneously created read-writable copy of a snapshot. One or more clones are created from a single snapshot. These clones are presented to the users as a normal file system(s). All the regular operations are allowed on the clones, including taking a snapshot from the clone. The clones are typically used in test, development, QA, and backup environments.

Similar to snapshots, when clone is created, no space is allocated. The reads to the clone are served by the base file system's blocks. Only when the blocks are changed in the clone, the changed blocks are allocated. Since the space is shared between snapshots and clones, and a snapshot can have multiple clones, a snapshot cannot be destroyed without also destroying any active clones.

From client systems, the clone file systems are accessed as though they are independent. No special requirement for accessing the clones is needed.

Oracle Fusion Middleware Disaster Recovery Strategy

Oracle Fusion Middleware Disaster Recovery Setup

Oracle Fusion Middleware Disaster Recovery solution facilitates data protection for the binaries and configuration data as well as the database contents.

- Remote replication feature of the Sun ZFS Storage Appliance is used to protect the middleware product binaries, configurations, and metadata files.
- Oracle Data Guard is used to protect the Oracle Database. This database contains the data of Oracle Fusion Middleware Repositories, as well as the customer data.
- The clients access the primary site during the normal operation. During DR conditions, they access the standby site. The change is almost seamless from the client's perspective since the entire FMW infrastructure along with the mount points and host names are configured identical for both the primary and standby sites.




Figure 1 : Oracle FMW DR setup with Sun ZFS Storage Appliance

During Failover/Switchover

Primary site is shutdown either due to a planned maintenance procedure (used for switchover) or because of an unplanned event (used for failover process) such as power failure in the data center. Then the procedure to failover to the standby sites are initiated. The procedure is explained in detail in the following sections. During the failover and switchover process, the client access to the FMW infrastructure is temporarily unavailable and resumed after the successful switchover/failover process.




Figure 2 : During the failover/switchover process

After Failover/Switchover

When the standby site is activated, the clients resume their access to the FMW infrastructure via standby site. The clients are unaware of the fact that the data is now served from a standby site. After the failover or the switchover procedure, clients can resume to access the FMW infrastructure with their requests now being served at the new primary.

When the old primary site [SITE-1] is up after the failure is fixed, it is converted to a new standby site. The Sun ZFS Storage Appliance replication and Oracle Data Guard setup are configured to replicate from SITE-2 to SITE-1. In other words, the whole infrastructure is reversed.




Figure 3 . After the Switchover/Failover process

In order to bring the setup back to the original primary and secondary, switch-back procedure is executed.

Oracle FMW Disaster Recovery Architecture Details

For the purpose of testing usage of Sun ZFS Storage Appliance for Oracle FMW Disaster Recovery solution, the setup in the lab had the following topology. Each site consists of :

1. One web host
2. One Oracle WebLogic Server Administration host
3. Two application servers with Oracle Fusion Middleware.
4. One SOA database server
5. Sun ZFS Storage clustered system

In the lab, for storage the primary site used Sun ZFS Storage 7320C with clustered head to provide high availability. The standby site's used Sun ZFS Storage 7320C clustered head. At both the sites, one of the clustered head is active for the Oracle FMW non-database components project *OFM-KIT-SITE1* and the other head is active for the SOA database server project *OFM-DR-SOADB*.

The following table shows the host information and what component is hosted in that server.

Primary Site (SITE-1) – Hardware and Software Information

Server/Storage	Alias	Operating System	Application Installed	Role
Sun Fire X4170	APPHOST1	OEL 5.4 (64-bit)	Oracle Web Logic Server 10.3	App Host 1, Admin Host
Sun Fire X4170	APPHOST2	OEL 5.4 (64-bit)	Oracle Web Logic Server 10.3	App Host 2
Sun Fire X4170	WEBHOST	OEL 5.4 (64-bit)	Oracle HTTP Server 11g	Web Host
Sun Fire X4170	SOADB	OEL 5.4 (64-bit)	Oracle 11.2.0.1	Primary Database Host
Sun ZFS Storage 7320C	aie-7320a-h1	Version 2010.Q3	N/A	Storage for the FMW primary site
Sun ZFS Storage 7320C	aie-7320a-h2	Version 2010.Q3	N/A	Storage for SOA DB Server for the primary site

Standby Site (SITE-2) – Hardware and Software Information

Server/Storage	Alias	Operating System	Application Installed	Role
Sun Fire X4170	APPHOST1	OEL 5.4 (64-bit)	Oracle Web Logic Server 10.3	App Host 1, Admin Host
Sun Fire X4170	APPHOST2	OEL 5.4 (64-bit)	Oracle Web Logic Server 10.3	App Host 2
Sun Fire X4170	WEBHOST	OEL 5.4 (64-bit)	Oracle HTTP Server 11g	Web Host
Sun Fire X4170	SOADBSBY	OEL 5.4 (64-bit)	Oracle 11.2.0.1	Standby Database Host
Sun ZFS Storage 7320C [Head-1]	aie-7320b-h1	Version 2010.Q3	N/A	Storage for FMW Standby site
Sun ZFS Storage 7320C [Head-2]	aie-7320b-h2	Version 2010.Q3	N/A	Storage for the standby database site

Figure 4 shows the architecture of the FMW DR solution using the Sun ZFS Storage Appliance.




Figure 4 : Oracle FMW DR Setup with Sun ZFS 7320C Storage

1. Four Sun Fire X4170s are used on each site for hosting the web server, application Servers, and the database server.
2. Sun ZFS Storage 7320 clustered system is deployed at both the primary and standby site. At each site, one head is active for the FMW non-database components and the other head is active for the SOA Database server.
3. All the connectivity on each site is done over 10Gb ethernet connected via Cisco 10Gb ethernet switch.
4. The DR link spans across two different labs which are approximately 200 feet apart. A dedicated 10 Gb data link is configured on each Sun ZFS Storage Appliance for the replication.

Note : Any server that is certified to run FMW and/or Oracle database can be used instead of Sun Fire X4170 servers.

Oracle Fusion Middleware Deployment Procedure

This section describes the configuration steps required for the Oracle Fusion Middleware DR solution using the Sun ZFS Storage. The steps required for the DR deployment are as follows :

1. Setting up project and file systems in the Sun ZFS Storage Appliance.
2. Installing the FMW and the database component on the primary site using the filesystems created on the Sun ZFS Storage Appliance and setting up the standby database at the standby site.
3. Setting up the Sun ZFS Storage Appliance remote replication.
4. Replicating the primary site's Oracle FMW project to the standby site.

1. Sun ZFS Storage Appliance File Systems Configuration

1.1 Clustered head setup

The Sun ZFS Storage 7320 clustered systems are used at the primary site and at the standby/target site. The clustered heads on both the sites are configured as active/passive mode for the FMW project. While one head actively serves the FMW, the other head will be activated during head failover.

1.2 Pool configuration

Storage pool (RAID set) with the name *pool-0* is created with mirrored configuration. Similar setup is done on both source and the target storage systems. If the Sun ZFS Storage is used for storing the database, then another pool *pool-1* with mirrored configuration is created so that both the heads can be active serving the database and non-database components of the Oracle FMW.

1.3 Primary site : Project and file systems setup

1. A project named "OFM-SITE1-KIT" is created on the pool *pool-0* for storing the FMW components and configuration file.

Name/Type	Value	Remarks
Quota	500 GB	For storing the Oracle FMW binaries, configuration files, logs and so on.
Mount point	/export	Can be modified to include the project name too
Record Size	128KB	Default record size.
All other setting	Default	For storing binary and configuration information, the default settings are optimal.
Filesystem	VolAdmin	Administrator server
Filesystem	VolData	Data for WLS common for both APPHOST1 and APPHOST2
Filesystem	VolOrcl1	Binaries for WLS on APPHOST1
Filesystem	VolOrcl2	Binaries for WLS on APPHOST2
Filesystem	VolWLS1	SOA Domain on APPHOST1
Filesystem	VolWLS2	SOA Domain on APPHOST2
Filesystem	VolWeb	Oracle HTTP Server (OHS)

2. If the database component is also deployed on the Appliance, then another project named “OFM-DR-SOADB” is created on the pool *pool-1*. The file systems are created as described in the following table:

Name	Value	Remarks
Quota	500GB	For storing the database binaries, and database files
Mount point	/export/	Can be modified to include the project name too
Read Access Time	OFF	Checked off
Filesystem	VolDB1_Bin	To store Oracle database binaries
Filesystem	VolDB1_Data	To store data files. Set 8KB record size
Filesystem	VolDB1_Logs	To store online redo logs . Set 128KB record size
Filesystem	VolDB1_Archive	To store archived redo logs. Set 128KB record size

3. Setup the restriction to allow only specific clients to access the project. This is done from the Project → Protocols → NFS → Add Exceptions. That procedure is to set as the NFS exception list with root squash option. Provide the subnet for the IP address used for the FMW and Database hosts. That will enable the hosts to perform *root* user operations – such as chown, mkdir and so on.

1.4 Standby site : Project and file systems setup

There is no need to set up project and filesystems for the FMW components at the standby site. When the storage-based replication is initiated, the whole OFM-SITE1-KIT project will be replicated to the standby site.

If the Sun ZFS Storage is used for storing the standby database at the standby site, a project named “OFM-DR-SOADBSBY” is created at the target Sun ZFS Storage 7320C system which is located in the standby site.

Project : *OFM-DR-SOADBSBY*

Name/Type	Value	Remarks
Quota	1 TB	Assuming 500GB of storage requirement, additional storage capacity for storing snapshots, clones etc.,
Mount point	/export/	Can be modified to include the project name too
Read Access Time	OFF	Checked off
Filesystem	VolDB2_Bin	To store Oracle database binaries.
Filesystem	VolDB2_Data	To store data files . Set 8KB record size
Filesystem	VolDB2_Logs	To store online redo logs. Set 128KB record size
Filesystem	VolDB2_Archive	To store archived redo logs. Set 128KB record size

The setup required for storing the FMW infrastructure on Sun ZFS Storage is complete. The next step is to setup the remote replication between the sites.

2. Oracle Fusion Middleware Hosts Setup

All the FMW hosts are connected to the Appliances over 10Gb ethernet infrastructure. NFSv3 protocol is used to access these filesystems from within the hosts.

2.1 Primary site: Mount points for FMW

The following table shows all the NFS mount points used in all the hosts at the primary site for FMW components :

Host Name	Appliance Mount point	Host mount point	Remarks
APPHOST1	aie-7320a-h1:/export/VolOrcl1	/u01/app/oracle/product	Binaries for WLS on APPHOST1
APPHOST1 APPHOST2	aie-7320-h1:/export/VolData	/u01/app/oracle/data	Data for WLS - Common for APPHOST1 and APPHOST2. Subdirectories apphost1 and apphost2 are created . Used to store JMS and Tlogs.
APPHOST1	aie-7320a-h1:/export/VolAdmin	/u01/app/wls/soaDomain/admin	Admin server
APPHOST1	aie-7320a-h1:/export/VolWLS1	/u01/app/wls/soaDomain/mng1	SOADomain on APPHOST1
APPHOST2	aie-7320a-h1:/export/VolOrcl2	/u01/app/oracle/product	Binaries for WLS on APPHOST2
APPHOST2	aie-7320a-h1:/export/VolWLS2	/u01/app/wls/soaDomain/mng2	SOAdomain on APPHOST2
WEBHOST	aie-7320a-h1:/export/VolWeb	/u01/app/oracle	Oracle HTTP Server

2.2 Primary site: Mount points for the primary database

The following table shows the mount points used in the standby database server, if Sun ZFS is used for storing the Oracle databases at the standby site:

Host Name	Appliance Mount point	DB Serer mount point	Remarks
SOADB	aie-7320a-h2:/export/VolDB1_bin	/u01/app/oracle	Database binaries and database at the Primary site (optional). The binaries can be stored local to the server.
SOADB	aie-7320a-h2:/export/VolDB1_data	/oradata/data	Data files for the primary database
SOADB	aie-7320a-h2:/export/VolDB1_logs	/oradata/logs	Online redo logs for the primary database
SOADB	aie-7320a-h2:/export/VolDB1_archive	/oradata/archive	Archive logs for the primary database

2.3 Standby site: Mount points for the standby database

At the standby site, if Sun ZFS Storage is used for storing the Oracle databases, the following table shows the mount points used in the standby database server :

Host Name	Appliance Mount point	DB Server mount point	Remarks
SOADBSBY	aie-7320b-h2:/export/VolDB1_bin	/u01/app/oracle	Database binaries and database at the Standby site (optional). The binaries can be stored local to the server.
SOADBSBY	aie-7320b-h2:/export/VolDB1_data	/oradata/data	Data files for the standby database
SOADBSBY	aie-7320b-h2:/export/VolDB1_logs	/oradata/logs	Online redo logs for the standby database
SOADBSBY	aie-7320b-h2:/export/VolDB1_archive	/oradata/archive	Archive logs for the standby database

2.4 NFS mount options used at the hosts

This table shows the mount options used in the hosts. The mount point and option entries can be configured in */etc/fstab* file.

Type	Mount options
For database mount points	rw,bg,hard,nointr,tcp,vers=3,timeo=600,rsize=32768,wsiz=32768,actimeo=0
For binary mount points	rw,bg,hard,nointr,tcp,vers=3,timeo=300,rsize=32768,wsiz=32768,actimeo=0,suid
For generic files	rw,bg,hard,nointr,tcp,vers=3,timeo=600,rsize=32768,wsiz=32768

3. Oracle Fusion Middleware Deployment

1. At the primary site, two Hosts (APPHOST1, APPHOST2) are used to deploy SOA over Weblogic container, with one HTTP Webserver host (WEBHOST) and one database server (SOADB) are used each site (primary/standby).
2. On each application host (APPHOST1, and APPHOST2), install binaries of Weblogic Server and SOA application on the */u01/app/oracle/product* directory.
3. On APPHOST1, one admin server and one managed SOA-Server1 (member of SOACluster) are hosted. Separate directories are created, where applications logs, configurations and application archives are stored.
4. Directory */u01/app/wls/soaDomain/admin*, is used for admin server's configuration.
5. Directory */u01/app/wls/soaDomain/mng1*, is used for SOA-Server1 configurations.
6. On APPHOST2, second member of SOACluster, SOA-Server2 is hosted and */u01/app/wls/soaDomain/mng2* directory is used for writing its logs, user configurations and so on.
7. The directories */u01/app/oracle/data/apphost1* and */u01/app/oracle/data/apphost2* are used for JMS and TLogs. APPHOST1 uses the apphost1 location and APPHOST2 uses apphost2 location.
8. Oracle HTTP Server (OHS) component is installed on */u01/app/oracle* on the WEBHOST server for its binaries as well as any user configurations.

4. Oracle Database Deployment (Primary and Standby)

Please refer to the Oracle documentation for a detailed procedure on setting up Oracle Data Guard.

At the primary site :

1. For the database deployment, the database binaries can be installed local to the server or on the NFS. Install the Oracle binaries on */u01/app/oracle*.
2. Create a database using the mount points */oradata/data*, */oradata/logs*, and */oradata/archive* locations.

3. Follow the 'Oracle Fusion Middleware Guide for Oracle SOA Suite' instructions for deploying the SOA suite on the database.
4. Create standby redo logs.

At the standby site :

1. For the database deployment, the database binaries can be installed local to the server or on the NFS. Install the Oracle binaries on */u01/app/oracle*.
2. Set up the proper *initSOADB_SBY.ora* file.
3. Using Oracle Recovery Manager (RMAN), duplicate the primary database for standby under the directories */oradata/data*, */oradata/logs*, */oradata/archive* locations.
4. Create standby redo logs.
5. Mount the standby database with managed recovery using current logfile option to catch up with the primary in real time.

At the primary site :

1. Modify the *log_archive_dest* parameters to enable shipping the Oracle changes to the standby site.
2. Enable the Oracle Data Guard and verify the SCN# of the standby catching up with the primary database.

At this point, the whole infrastructure is ready for a disaster recovery deployment.

5. Sun ZFS Storage Appliance - Remote Replication Setup

Identify the components – such as target appliance name, pool name at the target to replicate, root password for the target, schedule or continuous mode, SSL mode or not, and so on. This is typically one-time setup which could be done via GUI or using CLI along with optional scripts for faster repeated deployments. Please refer to the Sun ZFS Storage documentation for a detailed step-by-step procedure to setup the remote replication.

1. Enable the remote replication service at both source and target system. From Configuration → Services, choose Remote replication, enable the service by clicking the power button next to the service (if not already enabled).
2. Add the target Appliance information – such as name for the replication (which is typically same as the host name), Appliance host name, and the root password.
 - In the lab setup, the Name is “*aie-7320b-h1*”, target appliance host name is “*aie-7320b-h1*” along with the root password is provided.
3. After the successful addition, the target Appliance will be listed under the “Targets” defined for this Appliance.
 - The newly added *aie-7320b-h1* is listed under the targets.
4. The replication can now be enabled for any project or filesystem between these appliances.

At this point, the replication setup between *aie-7320a-h1* and *aie-7320b-h1* is complete.

6. Set Up Replication for the FMW Project

The FMW stack that is deployed under the project “*OFM-SITE1-KIT*” is configured to be replicated to the target/standby site.

1. From the GUI, choose the project *OFM-SITE1-KIT* and click the *Replication*.
2. Click (+) for adding the target appliance where this project to be replicated .
 - 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.
 - Provide the name of the pool, and the mode of replication.
 - If the replication happens within a data center, SSL can be disabled to enhance performance.
 - If the same link is used for multiple replication and also data, then the user can limit the amount of bandwidth this replication to consume.
 - If snapshots are taken at the source for the replication project or share, the user can choose to include replicating the snapshots.

3. 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. If only periodic changes happening to the FMW infrastructure, then it is recommended to use *Scheduled* mode and replicate periodically on a daily or weekly basis.
 - Note : 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 some time in the future. This will enable a copy of the FMW infrastructure available at the target site if the SITE-1 fails before the first occurrence of the replication.
4. At the target appliance (*aie-7320b-h1*), verify the package being received or already received . Click the *Shares* followed by *Projects* at the left side frame and click *Replica* , which lists the packages that are received / being received from various sources.
5. With that step, the replication is enabled, up and running between the storage systems *aie-7320a-h1* and *aie-7320b-h1* systems.

Remote replication setup of FMW stored on the Sun ZFS Storage is now complete.

Managing Planned and Unplanned Downtime

The DR configuration is setup in such a way that in the event of planned or unplanned downtime at the production site, the entire infrastructure is failed over to the standby site. This section provides the overview of the steps needed for the switchover/failover and switch-back/failback procedures.

The entire process can be scripted for a faster, more efficient DR deployment. Please refer to the Appendix for example scripts.

Site Switchover Procedures

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 [SITE-2] becomes the new primary site [SITE-1], and the current primary site becomes the new standby site. In other words, the roles are reversed.

1. Shut down all the Oracle Fusion Middleware components at the primary site. This can be done either manually or running a script [refer *RemoteShutdown.sh*] or using relevant management software.
2. Unmount the NFS file systems from all the hosts associated with the DR groups on the production site. [refer *umount_site1_ofm.sh*]
3. Perform switchover of the Oracle database using Oracle Data Guard. This procedure converts the standby site to a production database and also converts production site to a standby. [refer *switchover_db.sh* which in turns calls few *.SQL scripts].

- This step includes finishing the recovery, switching the standby to primary, shutdown and restart of the database at the new primary site. The old primary site is converted to the standby site.
4. Perform “Role reversal” procedure at the Sun ZFS Storage target site for the project that is replicated. [refer *switchover_s7000.sh*]
 - Note: After the role-reversal, the replication is automatically set to “*manual*” mode. The mode is not automatically changed to “*continuous*” - since the primary site can be in an unavailable state.
 5. Enable the *Continuous* replication for the OFM-SITE1-KIT project. This will start replicating the data from the new production to the new standby site.
 6. Mount the file systems associated with the DR groups on the standby site . [refer *mount_site2_ofm.sh*]
 7. Start all the Oracle Fusion Middleware components on standby site and ensure they are started successfully. [refer *RemoteStartup.sh* script]

At this point, the standby site has assumed the role of the production site and the old production site is now a standby site. The following block diagram show the FMW DR setup after the switchover is done. The clients are now accessing the SITE-2. SITE-1 is now the standby site.




Figure 8. After the switchover/failover

Site Switchback Procedures

The procedure is identical to the switchover procedure previously mentioned. The source and target sites are reversed. In this case, the switchback happens from the converted primary[SITE-2] to the original primary [SITE-1].

Site Failover Procedures

Failover is the process of making the current standby site the new production site after the primary production site becomes unavailable due to any unplanned events. It is necessary to understand the cause of the disruption, and that investigation may take longer. The following steps are performed assuming there is a disruption at the primary site and the only way to resume the work faster is to failover to the standby site. While the standby site is activated, the primary site is investigated for the cause of the disaster.

1. Perform Oracle Data Guard steps to convert the standby database to a primary database. This step includes finishing the recovery and switching to the primary database, shutdown and restart the database at the target site.
 - The *log_archive_dest_2_state* is set to *DEFER* until the old primary site [SITE-1] is available again.
2. Perform “Role Reversal” operation for the Oracle Fusion Middleware replicated package. Do not enable continuous replication at this time. [refer *switchover_s7000.sh* and executes *s7000_role_reverse_without_repl.aksh*]
3. Mount the NFS filesystems on all the hosts. [refer *mount_site2_ofm.sh*]
4. Start the Oracle Fusion Middleware components. [refer *RemoteStartup.sh*]
5. At this point, the standby site has assumed the role of production site.
6. Once the original primary is back up again,
 - Initiate the “*continuous replication*” for replicating from the new primary to the old primary. The old primary site will now become the new standby site. [refer *enable_repl.aksh*]
 - Perform Oracle Data Guard steps at the new standby site to convert the database from the primary to standby. [refer SQL scripts]
 - Change the new primary database's archive log shipping to “ENABLE” state. This will start pushing the archive logs to the new standby site.

At this point, the primary and standby setup is complete.

Benefits of deploying Oracle FMW using Sun ZFS Storage Appliance

The following are some of the benefits of deploying the Oracle FMW infrastructure using the Sun ZFS Storage Appliance :

Oracle tested and validated solution

The procedure described in this paper is tested and validated by Oracle. The solution description along with the script samples provided in this paper help to accelerate the deployment of the FMW DR solution.

Ease of deployment and management

With multi-protocol support, Sun ZFS Storage fits into any infrastructure. The intuitive user interface provides very easy, and convenient ways to manage the appliance. The entire solution can be scripted and executed repeatedly during the planned or unplanned events for faster and efficient DR.

Hybrid Storage Pool for quicker response

HSP model provides an optimal storage performance to the application by efficient data placement across DRAM, flash storage and the physical disks. The read-optimized flash acts as a second tier of cache in the storage which stores the recently and frequently accessed blocks, as illustrated in the following diagram :




Figure 9 : Hybrid Storage Pool components

Efficient backup, business continuity and disaster recovery

All the data services such as replication, snapshots, cloning etc., are included with the appliance. No additional costs are incurred when enabling any feature or protocol – no additional licenses to procure and manage. All these features can be used in combinations to meet specific needs for easier backup (certified with a number of leading backup applications), business continuity for almost instantaneous restores, and remote replication for disaster recovery purposes.

HA availability

The clustered storage head provides higher availability for the Oracle FMW DR infrastructure against storage head failures thus avoiding any single point of failure.

Analytics for faster resolution and planning

Analytics provides a graphical representation that illustrates the performance of the various components of the storage appliance. The following are some of these examples.

1. The dashboard provides a big picture of what is going on with the system, such as the space utilization, the IO activities going on for the different protocols, and so on as illustrated :




Figure 10. Dashboard status of Sun ZFS Storage Appliance

Figure 11 shows which files are being accessed and also which clients are accessing the files. This kind of information is extremely valuable to understand the access pattern, IO load from clients and the storage response time for those operations.




Figure 11 : Analytics screenshot for NFSv3 operations

The Analytics feature helps in

- a) understanding the overall health of the appliance
- b) observe the IOPS, latency, and throughput of the appliance
- c) understanding the client(s) access pattern for supporting SLA agreements
- d) capacity planning requirements
- e) identifying and resolving problems by comparing the performance data between the clients and the appliance
- f) exporting the historical data for analysis

and so on.

Using Snapshot and Cloning Features At The Primary Site

Snapshots taken at the primary site can be replicated along with the regular replication process. This way, multiple point-in-time copies of the production sites are available at both the source and the target.

The snapshot can be used for creating clones, which then can be used for test, development, and QA purposes.

Using Snapshot and Cloning Features at the Standby Site using Export/Clone

This is the process of making the replicated data accessible for read/write purposes at the target site. In contrast to the role-reversal operation explained earlier in this document, in this export/clone operation, the last successfully received package is cloned to be a separate project (or) filesystems for read/write purposes. This is primarily for temporary access requirements such as backup, development, testing, QA and so on. When the requirement is met, the clone can be discarded.

The clone is a point-in-time copy of the package that was last successfully received. The ongoing replication doesn't impact the cloned project. These clones and the ongoing replication are independent of each other, providing greater flexibility for creating clones at various point-in-times. If the local replica is deleted (the replication is severed) at the target, then each clone associated with the replicated package becomes an independent projects.

In summary, the export/clone operation is done for the development, testing, QA, and backup purposes while the role-reversal is to be used during switchovers and failovers.

Snapshot of the Exported Clone

A snapshot can be taken of the exported clones. This makes the target site even more flexible for various reasons—including backup, test/development, and so on, without impacting the production site.

Best Practices for Oracle FMW DR Configuration and Deployment

The following are recommendations for configuration and deployment during implementation of the Oracle Fusion Middleware Disaster Recovery solution.

Sun ZFS Storage Appliance Best Practices

- For high availability, it is strongly recommended to deploy clustered Sun ZFS Storage Appliances – such as Sun ZFS Storage 7420C, Sun ZFS Storage 7320C at each site.
- If the SLA for the FMW is same for both the primary and standby site, then use the same model of Sun ZFS Storage platform on both the sites.
- For typical deployments, 500GB to 1TB of storage quota per site may suffice. The quota can be expanded any time depending on the stored data, snapshots, and clones usage.
- NFS protocol is preferred for the FMW deployments for the ease of deployment and maintenance.
- Use *mirrored* configuration for the pool to obtain optimal performance and availability.
- Use Sun ZFS Storage Appliance to replicate all non-database content, to ensure a small RTO and RPO for the Oracle Fusion Middleware environment
- Store all the non-database FMW infrastructure within a project and replicate the project.
- Use the *Scheduled* mode of replication for a typical Oracle FMW DR deployment.
- Use the *Continuous* mode of replication if the standby site is required to be up-to-date irrespective of whatever the rate of change at the primary site.
- Follow the *Role-reversal* procedure during switchovers and failovers. This will enable much faster sync-back to the old primary during switchbacks and failbacks.
- Snapshots and clones can be used at the target site to offload backup, test, and development types of environment.
- If using the DR within the data center, consider setting the SSL set to OFF. Removing the encryption algorithm enable more data to be transferred between the sites.
- When performing replication across wide-area-network, SSL should be set to ON.
- For storing binaries on the storage, use the NFS mount options
rw,bg,hard,nointr,rsz=32768,wsz=32768,tcp,vers=3,actimeo=0,suid.
- For the Oracle database over NFS, use the mount options
rw,bg,hard,nointr,rsz=32768,wsz=32768,tcp,vers=3,actimeo=0.

- It is recommended not to perform *rollback* operation for any share at either primary or standby site, since that operation invalidates the replication and a fresh replication will need to be set up.

Oracle Fusion Middleware Best Practices

Please refer to the Oracle Fusion Middleware documentation listed in the reference section.

Oracle Database Best Practices

1. Configure Oracle Data Guard with the “Maximum Availability” option, to enable the standby database to be almost synchronized with the primary.
2. For optimal deployment, consult the Reference section where you will find links to a variety of helpful Oracle documentation.

Conclusion

Storage replication is a component of the foundation for disaster recovery protection in Oracle Fusion Middleware environments. The Sun ZFS Storage comes with a comprehensive set of features to handle all the requirements for replicating Oracle Fusion Middleware components in conjunction with Oracle Data Guard for Oracle database replication. The Sun ZFS Storage Appliance provides various modes of replication that can be used, depending on product requirements, to ensure that customer environments are protected against any unforeseen disasters.

Using Sun ZFS Storage Appliance replication in conjunction with the Oracle Data Guard allows users to obtain the maximum benefit out of their investment to protect their entire Oracle environment.

The Oracle tested and validated Oracle FMW DR configuration along with a set of example scripts and recommendations makes the Oracle FMW DR solution with Sun ZFS Storage an effective choice for the enterprise deployment.

Appendix

Terminology

Oracle FMW Disaster Recovery Terminology

Name	Description
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 (FMW)	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 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 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.
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.

Name	Description
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 Appliance Terminology and Operations

Name	Description
Project	Consistency group of related file systems and/or LUNs.
Share	A filesystem or a LUN. The file systems are exported over NFS or CIFS. LUN is exported over iSCSI or FC protocol
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 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 <i>source -> target</i> to <i>target -> source</i> for a package.
Destroy local replica	This will destroy the replica/package at the target. The clones associated with that target instantly become independent projects by themselves

Abbreviation

Abbreviation	Explanation
BUI	Browser User Interface
CLI	Command Line Interface
GUI	Graphical User Interface
DR	Disaster Recovery
GUI	Graphical User Interface
NFS	Network File System
RPO	Recovery Point Objective
RTO	Recovery Time Objective
SLA	Service Level Agreement
SSL	Secure Sockets Layer
NFS	Network File System protocol
CIFS	Common Internet File System also called Server Message Block (SMB) protocol
NDMP	Network Data Management protocol
FC	Fibre channel protocol
IB	Infiniband
SOA	Service Oriented Architecture
BPEL	Business Process Execution Language
ESB	Enterprise Server Bus
OWSM	Oracle Web Services Management

Oracle Fusion Middleware DR Solution Test Procedure

Test Criteria

The disaster recovery architecture specified in this document is tested by adhering to the following validation points:

- Show a test system architecture deploying key Oracle Fusion Middleware components and Sun ZFS Storage Appliance.
- Utilize Sun ZFS Storage Appliance's replication technology to protect all non-database content and Oracle Data Guard to protect all Oracle database content
- Document procedures to handle planned (switchover and switchback) and unplanned outages (failover and failback)
- Ensure a small RTO and RPO for the Oracle Fusion Middleware environment using Sun ZFS Storage Appliance.

Test Cases

1. Make configuration changes at production site. For example creating a new managed server, deploy a new application, and validate them at the standby site to ensure successful configuration replication between production and standby sites.
2. Keep XA transactions pending at the production site, and make sure they are committed or rolled back at the standby site after switchover/failover. This will ensure that transaction logs are replicated properly.
3. Validate the Oracle fusion middleware components on production site, such as, BPEL, ESB, and OWSM, and then re-validate them on standby site after switchover/failover.
4. Validate the oracle database replication using data guard, in the context of an Oracle fusion middleware SOA application.

Oracle Data Guard setup for the FMW database

The following are the high level overview of steps performed to setup the Oracle Standby Database and Oracle Data Guard. Refer to the Oracle documentation for further details.

- Primary site : site1db, ORACLE_SID : SOADB
- Secondary site : site2db, ORACLE_SID : SOADBSBY

The *initSOADB.ora* parameter file at the primary and *initSOADBSBY.ora* parameter file at the secondary are configured for Oracle Data Guard to ship the archive logs to the other site.

1. Set the log_archive_dest_2 to point to the service at the standby site, enable the second archive destination.

```
SQL> alter system set log_archive_dest_2 = 'service=SOADBSBY async
db_unique_name=SOADBSBY valid_for=(primary_role, online_logfile)'
SQL> alter system log_archive_dest_state_2=ENABLE;
```

2. Take backup of the database using RMAN. The output can be on a NFS share which is then accessed from both the primary and standby sites.

```
$ backup device type disk format '/opt/maa/stage/SOADB/%U' database plus
archivelog;

$ backup device type disk format '/opt/maa/stage/SOADB/%U' current
controlfile for standby;
```

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

```
$ startup nomount

$ rman target sys/oracle@SOADB auxiliary /

RMAN> duplicate standby database for standby;
```

4. Enable the managed recovery at the standby site.

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

If standby redo logs are used (11gR2 and above), then for real-time applying,

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

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.

Sample scripts

These scripts are provided for reference only. Modification required to suite any particular need.

Script Name / Description	Script
OFM Related Shell Scripts	
RemoteShutdown.sh This script shuts down the OFM infrastructure.	<pre>#!/bin/sh # set Variables; webHost=site1ofm3 appHost1=site1ofm1 appHost2=site1ofm2 password=pass user=usr #set user [lrange \$argv 1 1] #set password [lrange \$argv 3 3]; #Shutdown OHS running on web server; shutOHS='/u01/app/oracle/product/11.1.1/ohs_1/instances/ohs_instance1/bin/opm nctl stopall' ssh \$user@\$webHost \$shutOHS; #shutdown any managed cluster/server ssh \$user@\$appHost1 '/u01/app/oracle/product/fmw/wlserver_10.3/common/bin/wlst.sh ~/scripts/stopClusters.py'; #Kill any Node manager and shutdown admin server ssh \$user@\$appHost1 sh -c '~/scripts/shutNM.sh'; #Kill any java component ssh \$user@\$appHost2 sh -c '~/scripts/shutNM.sh';</pre>
RemoteStartup.sh This script bring all the OFM infrastructure up.	<pre>#!/bin/sh webHost=site1ofm3 appHost1=site1ofm1 appHost2=site1ofm2 password=pass user=usr #set user [lrange \$argv 1 1] #set password [lrange \$argv 2 2]; #start Node manager on AppHost2 ssh \$user@\$appHost2 sh -c '/u01/app/oracle/product/fmw/wlserver_10.3/server/bin/startNodeManager.sh > ~/nm.out 2> nm.out &'; #start Node manger and admin server on AppHost1 ssh \$user@\$appHost1 sh -c '/u01/app/wls/soaDomain/admin/user_projects/domains/soa_domain/bin/startWeb Logic.sh > ~/admin.out 2>&1 &'; ssh \$user@\$appHost1 sh -c '/u01/app/oracle/product/fmw/wlserver_10.3/server/bin/startNodeManager.sh > ~/nm.out 2>&1 &'; #Startup OHS running on web server; startOHS='/u01/app/oracle/product/11.1.1/ohs_1/instances/ohs_instance1/bin/opm nctl startall' ssh \$user@\$webHost \$startOHS; echo "Waiting 3 mins for Admin Server to be up" sleep 180 #seconds #startup SOA cluster amd TLogServer ssh \$user@\$appHost1 '/u01/app/oracle/product/fmw/wlserver_10.3/common/bin/wlst.sh ~/scripts/startClusters.py'</pre>
ShutNM.sh Shuts down the node manager	<pre>#!/bin/bash ps -ef grep weblogic.NodeManager grep -v grep awk '{print \$2}' xargs -i kill -9 {} }</pre>

Script Name / Description	Script
	/u01/app/wls/soaDomain/admin/user_projects/domains/soa_domain/bin/stopWebLogic.sh &
StartNM.sh Starts the node manager	#!/bin/bash /usr/bin/nohup /u01/app/oracle/product/fmw/wlserver_10.3/server/bin/startNodeManager.sh &
StartClusters.py Starts the SOA server and TlogServer	connect('weblogic','welcome1','t3://site2ofm1:7001'); start('soa_Cluster','Cluster'); start('TLogServer','Server');
StopClusters.py Shuts down the SOA server and TlogServer	connect('weblogic','welcome1','t3://site2ofm1:7001'); shutdown('soa_Cluster','Cluster',force=true); shutdown('TLogServer','Server',force=true);
Data Switchover/Switchback Shell scripts	
switchover_db.sh / switchback_db.sh This shell script calls for few SQL scripts which are used to convert physical standby database to a primary database and vice versa. Note that during failover/failback conditions, the failed site won't be accessible. So, converting the primary site to a standby is to be done at a later time once the site is accessible again.	#!/bin/bash ## Interchange the name for switchbacks CURRENT_STANDBY_HOST='site2db' CURRENT_PRIMARY_HOST='site1db' # #CURRENT_PRIMARY_S7000='aie-7320a-h1' #CURRENT_STANDBY_S7000='aie-7320b-h1' # # Get the confirmation/check_instances.sh echo "Going to convert \$CURRENT_PRIMARY_HOST to Standby site and \$CURRENT_STANDBY_HOST to primary" echo "Press <Enter> to continue, <ctrl+c> to abort.." read a ## Convert primary to standby echo "Converting \$CURRENT_PRIMARY_HOST to standby .." ssh \$CURRENT_PRIMARY_HOST ". .bash_profile; sqlplus -s / as sysdba @/home/oracle/scripts/convert_to_standby_1" echo "Mounting the standby database at \$CURRENT_PRIMARY_HOST.." ssh \$CURRENT_PRIMARY_HOST ". .bash_profile; sqlplus -s / as sysdba @/home/oracle/scripts/convert_to_standby_2" ## ## Convert standby to primary echo "Converting standby database at \$CURRENT_STANDBY_HOST to primary database.." ssh \$CURRENT_STANDBY_HOST ". .bash_profile; sqlplus -s / as sysdba @/home/oracle/scripts/convert_to_primary_1" echo "Starting the primary database at \$CURRENT_STANDBY_HOST.." ssh \$CURRENT_STANDBY_HOST ". .bash_profile; sqlplus -s / as sysdba @/home/oracle/scripts/convert_to_primary_2" # echo "Waiting for 60 seconds before enabling automatic recovery at the standby" sleep 60 ## Recover automatic standby.. echo "Setting the automatic recovery at the standby database" ssh \$CURRENT_PRIMARY_HOST ". .bash_profile; sqlplus -s / as sysdba @/home/oracle/scripts/convert_to_standby_3" # # Run the Query against primary and standby and wait for user confirmation before proceeding.. ## ./check_instances.sh # # Failover complete # echo "Dataguard failover Complete.. Proceed with S7000 role reversal"
Database SQL Scripts	
check_instance.sql To check the role for the database. This can be run against the primary database or the standby database.	select substr(host_name,1,40) "Host", instance_name "Instance", status "Status" from v\$instance; select DATABASE_ROLE,current_scn,protection_mode from v\$database; exit
convert_to_primary_1.sql This is the first to be executed against the	-- Steps to conver from standby to primary alter database recover managed standby database finish;

Script Name / Description	Script
standby site to convert to a primary site.	alter database commit to switchover to primary with session shutdown; shutdown immediate; exit
convert_to_primary_2.sql This is the second step to be executed against the standby site to convert to a primary site.	<pre>-- Restarting the database startup alter system register; alter system set log_archive_dest_2='service="(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=site2db)(PORT =1521))(CONNECT_DATA=(SERVICE_NAME=SOADB) (SERVER=DEDICATED)))";LGWR SYNC AFFIRM delay=0 optional co mpression=disable max_failure=0 max_connections=1 reopen=300 db_unique_name="SOADBSBY" net_timeout=30 ','valid_for=(all_logfiles,primary_role)' scope=both; alter system set log_archive_dest_state_2=ENABLE scope=both; alter database set standby database to maximize availability; show parameter log_archive_dest exit</pre>
convert_to_standby_1.sql This is the first step to be executed against the primary site to convert to a standby database..	<pre>-- To convert to a standby database alter database commit to switchover to physical standby with session shutdown; shutdown immediate; exit</pre>
convert_to_standby_2.sql This is the second step to be executed against the primary site to convert to a standby database..	<pre>-- Part 2 startup mount ; alter system set log_archive_dest_state_2=DEFER scope=both; alter database set standby database to maximize availability; alter system register; exit</pre>
convert_to_standby_3.sql This is the last step to be executed against the primary site to convert to a standby database..	<pre>-- Part 3 alter database recover managed standby database using current logfile disconnect from session; alter database set standby database to maximize availability; exit</pre>
Shell Script for invoking S7000 scripts	
switchover_s7000.sh This script performs the role-reversal for the project OFM-SITE1-KIT where the source becomes the target and the target site becomes the source. After this step, the shares are mounted in the various hosts at the standby site. The shell script executes the appliance kit shell script for automated role-reversal.	<pre>[oracle@site2db site1db_scripts]\$ cat switchover_s7000.sh # This script executes the failover/switchover of S7000 by doing role-reversal # Interchange the names for switchback CURRENT_PRIMARY_HOST='site1db' CURRENT_STANDBY_HOST='site2db' # CURRENT_PRIMARY_S7000='aie-7320a-h1' CURRENT_STANDBY_S7000='aie-7320b-h1' ## ## Failover the S7000 ## Get the confirmation.. echo "Status at the source [\$CURRENT_PRIMARY_S7000] " ssh -T \$CURRENT_PRIMARY_S7000 < s7000_status_repl.aksh # echo "Status at the target [\$CURRENT_STANDBY_S7000] " ssh -T \$CURRENT_STANDBY_S7000 < s7000_status_repl_tgt.aksh # echo " " echo " " echo "Going to failover from the source : \$CURRENT_PRIMARY_S7000 to the target : \$CURRENT_STANDBY_S7000 ." echo "Press <Enter> to proceed.. <Ctrl+C> to abort.." read a # Stop the replication echo "Suspending the continuous replication at the source : \$CURRENT_PRIMARY_S7000 " ssh -T \$CURRENT_PRIMARY_S7000 < s7000_stop_repl_at_source.aksh echo "Performing the role reversal at the target : \$CURRENT_STANDBY_S7000" ssh -T \$CURRENT_STANDBY_S7000 < s7000_role_reverse_at_target.aksh # echo "Status at the new source [\$CURRENT_STANDBY_S7000] " ssh -T \$CURRENT_STANDBY_S7000 < s7000_status_repl_src.aksh # echo "Status at the new target [\$CURRENT_PRIMARY_S7000] " ssh -T \$CURRENT_PRIMARY_S7000 < s7000_status_repl_tgt.aksh # echo "Now setting the mount point at the new source \$CURRENT_STANDBY_S7000 .. "</pre>

Script Name / Description	Script
	ssh -T \$CURRENT_STANDBY_S7000 < s7000_set_mount.aksh echo "Role reversal Complete.. Proceeding with umount / mount in the new Primary \$CURRENT_STANDBY_HOST .." ssh \$CURRENT_STANDBY_HOST /home/oracle/scripts/umount_site2_ofm.sh ssh \$CURRENT_STANDBY_HOST /home/oracle/scripts/mount_site2_ofm.sh
S7000 – aksh scripts	
s7000_role_reverse_at_target.aksh This script is invoked during switchover/switchback conditions. After the switching over is complete, the replication is set to "continuous" mode. The names of the scripts are self explanatory. The scripts are provided as an example.	<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] ; 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>
s7000_role_reverse_no_repl..aksh This script is invoked in a failover/failback conditions where the "continuous" mode of replication is not enabled at the end.	<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] ; break; } } run('cd ..'); } </pre>

Script Name / Description	Script
s7000_status_repl_src.aksh To check the replication status at the source site	<pre>shares set pool=pool-0 select OFM-SITE1-KIT replication select action-000 show</pre>
s7000_status_repl_tgt.aksh To check the replication status at the target site	<pre>shares set pool=pool-0 replication sources select source-000 show</pre>
s7000_stop_repl_at_source.aksh To stop the replication at the source.	<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>

Sun ZFS Storage Appliance – File System Configuration via GUI (Sample)

1. Project with name “OFM-SITE1-KIT” is created with the quota of 900GB. The mount point can be left to the default “/export” or optionally a subdirectory name can be entered. The database record size is set to 8kB. Deduplication and compressions are off.

OFM-SITE1-KIT... Shares **General** Protocols Access Snapshots Replication

REVERT APPLY

Space Usage

DATA

Quota ☒ 900 G

Reservation ☐ 0 G

USERS & GROUPS

User or Group Show All

Usage none

Inherited Properties

Mountpoint /export

Read only ☐

Update access time on read ☐

Non-blocking mandatory locking ☐

Data deduplication ☐

Data compression Off

Checksum Fletcher4 (Standard)

Cache device usage All data and metadata

Synchronous write bias Latency

Database record size 8k

Additional replication Normal (Single Copy)

Virus scan ☐

Prevent destruction ☐

2. Setup the restriction to allow only specific clients to access the project. That is set as the NFS exception list with root squash option. In this step, the network with the subnet 172.20.67 is allowed to be the root user for the project. That will enable the hosts to perform root operations – such as chown, mkdir etc.,

NFS

isv-7310b.central:/export

Share Mode None

Disable setuid/setgid file creation ☐

Prevent clients from mounting subdirectories ☐

Anonymous user mapping nobody

Character encoding default

Security mode Default (AUTH_SYS)

NFS Exceptions

TYPE	ENTITY	ACCESS MODE	CHARSET	ROOT ACCESS
Network	172.20.67.0/24	Read/write	default	☒

3. Filesystem shares are created under the project OFM-SITE1-KIT. In order to preserve consistency across the filesystems that are mounted from the NAS appliance for hosting the Oracle FM stack, all the filesystems were created within this project. The mount point is inherited. Optionally, the user and group Ids from the hosts can be provided here for more security.

Create Filesystem CANCEL APPLY

Project

OFM-SITE1-KIT

Name

VolAdmin

Data migration source

None

User

nobody

Group

other

Permissions

R

W

X

User

R

W

X

Group

R

W

X

Other

Use Windows default permissions

Inherit mountpoint

☒

Mountpoint

Reject non UTF-8

☒

Case sensitivity

Mixed

Normalization

None

4. The following is the list of filesystems created under the project “OFM-SITE1-KIT”.

OFM-SITE1-KI...		Shares	General	Protocols	Access	Snapshots	Replication
+ Filesystems		LUNs	7 Total				
NAME	SIZE	MOUNTPOINT					
VolAdmin	8.91G	/export/VolAdmin					
VolData	1.17M	/export/VolData					
VolOrcl1	3.79G	/export/VolOrcl1					
VolOrcl2	7.79G	/export/VolOrcl2					
VolWLS1	368M	/export/VolWLS1					
VolWLS2	3.99G	/export/VolWLS2					
VolWeb	3.27G	/export/VolWeb					

References

Oracle MAA site :

- [Oracle Maximum Availability Architecture Web site](#)

Oracle Fusion Middleware related documents :

- [Fusion Middleware Disaster Recovery Guide](#)
- [Fusion Middleware High Availability Guide](#)
- [Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite](#)

Oracle Data Guard related documents :

- [Oracle Data Guard with Oracle Database 11g Release 2](#)
- [Oracle DataGuard 11g – Installation and Configuration on Oracle RAC Systems](#)
- [Oracle Data Guard with Oracle Database 11g Release 2 Technical Information](#)

Oracle's Sun ZFS Storage Appliance related documents :

- [Sun ZFS Storage Appliance Documentation](#)
- [Sun download site](#)

Oracle Databases and Sun ZFS Storage Appliance documents :

- [Deploying Oracle Database over Sun ZFS Storage using NFS protocol](#)
- [Configuring Sun ZFS Storage Appliance for Oracle Databases](#)



Oracle White Paper Title:
September 2010

Author: Sridhar Ranganathan, Anuj Sahni
Contributing Authors: Pradeep Bhat, Sunita
Sharma, David Krenik

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 © 2010, 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.

0109