Disaster Recovery Planning for Oracle Middleware Applications

Ananthan Subburaj
Technology Architect
New Jersey, United States of America

Abstract— **Disaster recovery planning is one of the most crucial component for a business but that is often ignored or inadequately planned. Business organizations must have a well-structured plan and document process for disaster recovery and business continuity, before a catastrophe occurs. The IT infrastructure disasters can be short term or long lasting failure resulting in loss of application and data, but when an organization is well planned with standby infrastructure and recover plan in place, it can quickly get the business on track. This paper will clearly outline the disaster recovery planning best practices and technical process to recover the oracle applications with zero data loss. This paper aims to provide best practices for effective disaster management planning and technical configurations to achieve faster application recovery.**

Keywords:- *Disaster Recovery; High Availability, Fusion middleware DR, DR planning.*

I. INTRODUCTION

Today's information world the IT applications have become increasingly critical for the operation of a company, the importance of ensuring the continued operation and the rapid recovery of IT applications has increased. The business organizations will be severely impacted by disaster when IT infrastructure cannot continue to function due to data loss or the application infrastructure failure, it may even go out of business. An effective disaster recovery plan ensures quick recovery of data and application infrastructure in the event of natural or technical disaster. This paper aims to provide systematic approach to plan the disaster recovery organization level process and technology implementation technique to reduce the disaster recovery time for oracle middle applications using the logical host names.

II. DISASTER RECOVERY ORGANIZATION

In the event of a disaster, the objective of the Disaster Recovery Organization (DRO) is to minimize disruption and downtime of critical business functions and data loss by rapidly recovering business critical infrastructure and application components. The DRO focuses on two metrics Recovery Point Objective (RPO) and Recovery Time Objective (RTO). As the disaster recover work cannot be planned and it is response to an unexpected event, the recovery process should be well documented and automated through tools and scripts as much as possible and the people who are assigned.

The DRO should have the well-defined procedures for conducting recovery:

- Notification and Initial workforce mobilization process
- Damage assessment process
- Disaster declaration process
- Secondary Workforce mobilization process
- DR command center establishment process
- DR support center establishment process
- Application recovery procedures

III. DR RESOURCE PLANNING AND UTILIZATION

An alternate facility should be available to function as a DR command & support center supporting the senior DR management and DR process teams along with operations, help desk, workstation support and other virtual teams. This facility should be available for use by IT operations, application and business team personnel in order to allow them to perform their respective duties during a recovery process and after the recovery process. Maintenance and testing of facilities is on an annual basis to keep pace with DR organization and recovery requirements:

A. Normal Mode

During normal day-to-day operations the D.R. hardware resources can be used for Test, Development, and QA activities with minimum allocation for the DR sync activities. All the DR servers should be on active maintenance contracts with supplying vendors, and all code fix, firmware, and Operating systems should be kept up to date same level as primary Datacenter.

B. Test Exercise Mode

The DR test can be done in two different options, one is to switch over the primary to DR and run the business operations in the DR for specific window and then switch back to original primary, the second option is testing the DR facility in isolated network non-invasive to current production environment with controlled user testing.

C. Disaster Evernt Mode

Actual DR event where all production application systems are recovered in DR location and made available for the business operations.

IV. DISASTER RECOVERY ARCHITECTURE

The fusion middleware applications production based on WebLogic server technology such as Webcenter Portal, Webcenter Content, and Imaging and Identity Management products have product software binaries, domain configuration file and application metadata schemas in the database. The Fusion Middleware application disaster recovery includes application tier file system replication and the metadata store database replication to the standby site. The recommended process for the disaster recovery is to use the shared storage for all the application tier nodes in the production site and using storage replication to synchronize the application tier file systems to the standby site, replicate the database tier using the oracle physical standby database.

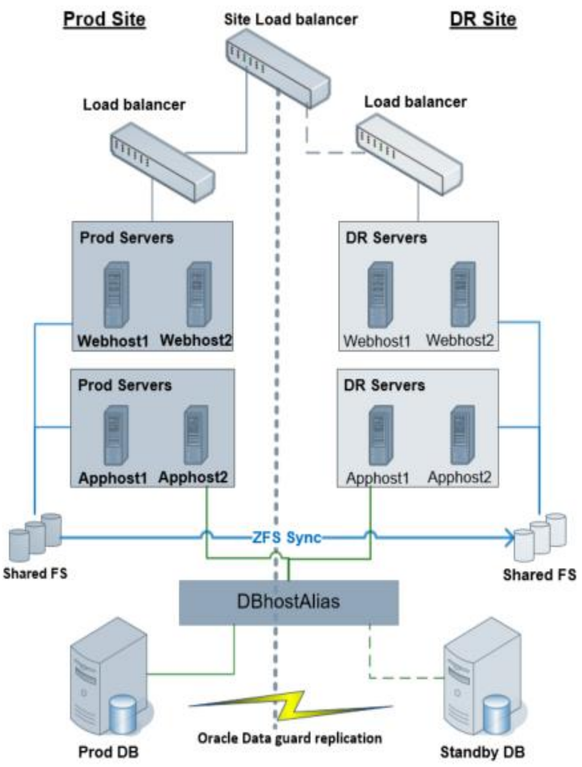

Fig. 1. Disaster Recovery Topology

The middleware applications have oracle software binaries, domain configuration files, application deployment files and application metadata files in the application tier. The application tier mount points should be created with shared storage file system on the production primary site and install the application binaries on shared file system, when using multiple mount points create all the mounts points from the same consistency group on the primary site to ensure consistent data replication across all the file systems to standby site. Setup storage replication at the project level with appropriate replication schedule. The recommended schedule for the binary and configuration files are once a day and one time on demand sync can be done whenever there is major deployment or upgrade happens on the primary site. As the file system replication and automatic scheduling of incremental sync replicates all the changes in the primary site to standby site and there is no need to install any software or update configurations in the standby site.

The Fusion middleware applications uses metadata schema and application schemas in the oracle database. Establish oracle physical standby database to replicate the data from production site to DR site. Configure the data guard broker to monitor and perform administrative tasks for the database replication.

V. NETWORK CONFIGURATIONS

The application components should be configured using the logical hostname (host name alias) instead of physical host name for the listen address on both WebLogic Admin and managed servers. The logical host name should be revolved to appropriate physical host IP address in primary site and also in the DR site. As the alias host name in production site resolves to production host physical IP address and in the DR site resolves to the DR host physical IP address there is no need to update any domain configuration when switching over to the standby site for disaster recovery operation. Configuring the application components with alias hostnames also helps with the server migration in the event of any host hardware or operating system failure with the production site, the application can be quickly started on another server by mounting the shared application file systems and updating the alias hostname to point to the new server's physical IP address.

The load balancer should be configured on both primary and standby site to route traffic to the physical IP of the servers. When the switch over happens the application DNS should be change to connect to the DR load balancer virtual IP address instead of the primary site load balancer virtual IP. The site load balancer can be used to route traffic between primary and standby site load balancers based on the health check rules or on demand when the DRO declares the disaster and the application switch over is completed.

The host name resolution can be achieved by using the local /etc/hosts file on each application and database hosts involved in the configuration or using the DNS servers. When using the /etc/hosts file maintain the entry for all the application servers with same set of entries for better manageability. When using the separate DNS servers for primary site and DR site the alias host name alias entries can be preconfigured in the DNS appropriately to point to the respective physical hosts. With global DNS server the hostname alias need to be updated to point to the DR site during the recovery process.

The host name resolution process should be decided part of the design phase of the DR process. The name resolution method can be controlled by changing the configuration order in /etc/nsswitch.conf file on each host. The entry like (hosts: files dns nis) this makes the host to use the hosts file on the server as primary resolution method.

VI. MANAGING HOST NAME

Depending on the type of host name resolution used following are the sample host name alias required for the middleware applications. When using the /etc/hosts file based resolution make sure have entries for all hosts are maintained across all the servers part of the application topology and test the name resolution by using ping command from each node to all other nodes. The naming resolution should be validated from all nodes when using the global DNS change for switch over to ensure the DNS cache is not pointing the old IP address.

Table. 1. Primary Site hostname resolution

IP Address	Physical Host Name	Alias Host name
110.24.2.101	PRIWEBHOST1.SAMPLE.COM	WEBHOST1.SAMPLE.COM WEBHOST1
110.24.2.102	PRIWEBHOST2.SAMPLE.COM	WEBHOST2.SAMPLE.COM WEBHOST2
110.24.2.103	PRIAPPHOST1.SAMPLE.COM	APPHOST1.SAMPLE.COM APPHOST1
110.24.2.104	PRIAPPHOST2.SAMPLE.COM	APPHOST2.SAMPLE.COM APPHOST2

Table. 2. Standby Site hostname resolution

IP Address	Physical Host Name	Alias Host name
110.44.2.101	DRWEBHOST1.SAMPLE.COM	WEBHOST1.SAMPLE.COM WEBHOST1
110.44.2.102	DRWEBHOST2.SAMPLE.COM	WEBHOST2.SAMPLE.COM WEBHOST2
110.44.2.103	DRAPPHOST1.SAMPLE.COM	APPHOST1.SAMPLE.COM APPHOST1
110.44.2.104	DRAPPHOST2.SAMPLE.COM	APPHOST2.SAMPLE.COM APPHOST2

VII. DISASTER RECOVERY PROCESS

To activate standby site application when there is a failure or planned outage of the production site, use the following steps to bring up the application on the standby site to assume the business operations from standby site:

- Stop the application services on the production site (for unplanned failure the applications might be already down) and stop the file system replication from the production to standby site
- Apply the last available database redo logs to the standby and execute switchover (planned maintenance of production site) or failover (unplanned failure of production site) of database using data guard.
- Mount the replicated file systems on standby servers in read write mode.
- Execute the DNS change and validate alias host name resolution to the appropriate standby site hosts.
- Start all the applications services on the standby nodes.
- Update the application URL DNS to point to the standby load balancer or use a global load balancer to route user connections to standby site.
- If original primary site is accessible, ensure to enable the replication of both database and application file systems.
- Establish appropriate database and application file system backup process to back up the files from the new standby site

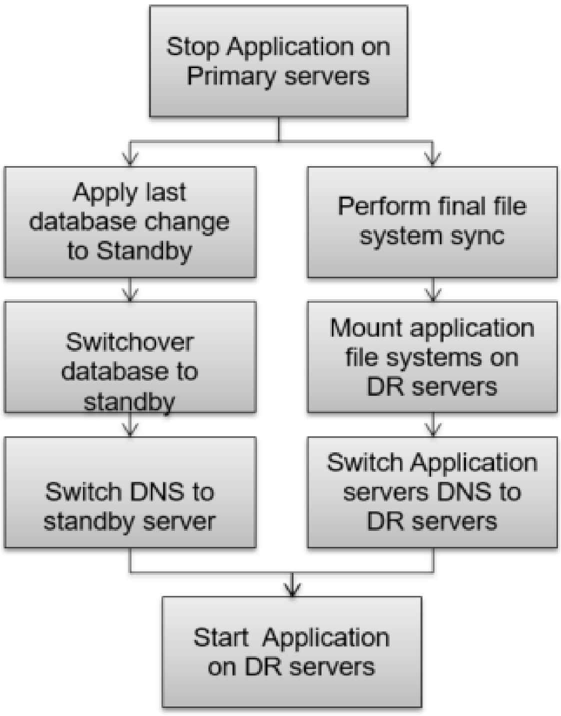
Stop Application on Primary servers
Apply last database change to Standby
Perform final file system sync
Switchover database to standby
Mount application file systems on DR servers
Switch DNS to standby server
Switch Application servers DNS to DR servers
Start Application on DR servers

Fig. 2. Disaster Recovery process flow

VIII. LOAD BALANCER SWITCHOVER

The primary site and the standby sites have independent load balancers configured to load balance application traffic across all the configured application nodes. During the DR event when the application is switched over to the standby site, client application access should be transparently redirected to the standby site which is a now configured as new primary. To redirect the client application access the DNS URL name should be updated to the DR load balancer virtual IP address. This can be automated using the site load balancer configuration to front end the local site level application load balancers. With the additional infrastructure cost using the global site load balancer provides additional capabilities to monitor and detect local load balancer failure and automate the redirection to available site based on the load balancer rule configuration. The global load balancer also avoid the impact of DNS cache issues while updating the DNS alias to DR IP address.

IX. BEST PRACTICES FOR DISASTER RECOVERY

The following are the best practices for preparing disaster recovery site and recovery procedures in readiness for a site failure.

- The Disaster recovery site should be geographically separated to ensure the site availability and avoiding possibility of losing both sites in major natural disasters.
- Use Oracle Data Guard to replicate the database changes to standby site database and use Data Guard broker to simplify the administration tasks.
- Configure Active Data Guard feature to offload read-only queries to the standby database to utilize the standby hardware resources.
- Use Oracle Flashback Database feature to reinstate the old primary database as a standby database in the event of a site failover.
- Replicate the application File Systems to the DR site using storage replication technology and establish procedure to reverse the direction of replication in the event or switchover and use cloned replica for site testing.
- Create role based database services for the application connectivity to database
- Test standby site using snapshot standby database to temporarily convert the physical standby database to updatable copy
- Create documented operational procedures to streamline the DR test process and for the actual DR event.
- To enable faster recovery and to reduce the human errors use tools or automation scripts to execute DR procedure.
- Configure Data Guard Broker to automate Data Guard operation and the database failover and switchover steps
- Create DB_ROLE_CHANGE trigger to automate the post DB switchover or failover configuration steps

X. CONCLUSIONS

Every business, large or small, in today’s information world is dependent upon their IT infrastructure servers and application data for business operations. There are many common risks such as natural disasters and internal technical failures such as hardware failure or human errors can lead to adverse effects on the information systems and hinder business operation. It is essential for a company to create well defined disaster recovery plan and test periodically. It is important for business organizations to plan for the disaster recovery IT infrastructure, create recovery procedures and test the readiness to take on the disaster challenge, also the effectively utilization of DR IT assets in the normal mode of business operations. This paper aimed to provide insight into the disaster recovery planning, infrastructure utilization, technical architecture and best practices to achieve quick recovery of business applications and reduce the business impact.

REFERENCES

[1]. https://docs.oracle.com/en/middleware/fusion-middleware/12.2.1.3/asdrg/toc.htm

[2]. https://docs.oracle.com/middleware/1212/core/ASADM.pdf

[3]. https://www.oracle.com/technetwork/database/availability/maa-site-guard-exalogic-exadata-1978799.pdf

[4]. https://www.oracle.com/technetwork/database/features/availability/wlsdatasourcefordataguard-1534212.pdf

[5]. https://www.oracle.com/technetwork/database/availability/maa-fmwsharedstoragebestpractices-402094.pdf