Disaster Recovery Solution for Oracle Traffic Director

Oracle Maximum Availability Architecture White Paper
September 2013

Maximum Availability Architecture

Oracle Best Practices for High Availability

ORACLE

Executive Overview

With the advent of Oracle Engineered Systems, enterprises have immensely enhanced their business value by leveraging the systems' ability to support high transaction volumes, consolidation, and simplified manageability. One of the key Engineered Systems is the **Oracle Exalogic Elastic Cloud**, which has been widely recognized for delivering breakthrough performance for packaged and custom applications. The responsiveness from complex, distributed applications running on Oracle Exalogic Elastic Cloud is beyond the capabilities of typical servers used in data centers today, primarily due to two major elements:

- **Exalogic Elastic Cloud Hardware:** This high-performance hardware system, assembled by Oracle, integrates storage and computing resources using a high-performance I/O subsystem called Oracle Exabus, which is built on the Oracle Quad Data Rate (QDR) Infiniband.

- **Exalogic Elastic Cloud Software**: This essential package of Oracle Exalogic software, device drivers, and firmware is pre-integrated with Oracle Linux and Oracle Solaris, enabling the Exalogic advanced performance and Infrastructure as a Service (IaaS) capabilities, server and network virtualization, and storage and cloud-management capabilities.

One of the key components of the Exalogic Elastic Cloud Software is **Oracle Traffic Director,** a built-in application delivery controller (ADC, also defined as a software load balancer). It is a fast, reliable, and scalable layer-7 software load balancer.

This white paper describes the disaster recovery solution for Oracle Traffic Director based on Oracle Maximum Availability Architecture (MAA) principles. Oracle Maximum Availability Architecture [1] is the Oracle best-practices blueprint for implementing Oracle high-availability technologies.

Oracle Traffic Director can serve as an entry point for all HTTP, HTTPS, and TCP traffic to application servers and web servers in an Exalogic deployment while distributing the requests based on the specified load-balancing method, routing the requests based on

specified rules, caching frequently accessed data, prioritizing traffic, and controlling the quality of service. Oracle Traffic Director can also be used for distributing initial context requests that come from Java clients over WebLogic T3 protocol.

Along with features like high performance, flexible routing, load control, and quality of service, Oracle Traffic Director provides many required high-availability features, such as health checks for the back end, failover for load balancing, and dynamic reconfiguration.

While high-availability features typically protect Oracle Traffic Director deployments against local outages, such as application failures or system-level problems, the disaster tolerance solution for Oracle Traffic Director protects against larger outages, such as a catastrophic data center failure. For maximum availability, the loss of a site cannot be the cause for an outage of the load-balancing web tier that routes traffic to the various origin servers--such as web servers, application servers, or Lightweight Directory Application Protocol (LDAP) servers--hosting various enterprise applications.

The emphasis of this white paper is on two different options to achieve protection of Oracle Traffic Director across geographically spread-out sites.

- Active-Passive Mode: This protection solution involves setting up a standby site at a location that is geographically different from the primary site. The standby site can have equal or fewer resources compared to the primary site. Installation binaries, configuration, and security data are replicated to the standby site on a periodic or continual basis. The standby site is normally in a passive mode; it is started in the event of any planned or unplanned outages of the primary site.

- Active-Active mode: This protection solution involves setting up each site with its own Oracle Traffic Director implementation, with separate installation binaries, and keeping the Oracle Traffic Director instance homes synchronized between both sites. Traffic from external clients, however, is routed to only one site at any given time.

Audience

This document is intended for Oracle Fusion Middleware administrators and web applications administrators. The reader should be familiar with Oracle Exalogic Elastic Cloud, Oracle Fusion Middleware components, storage replication techniques, Oracle Enterprise Manager, and Oracle Site Guard. For additional details, refer to the documents listed in the "References" section.

Introduction

Oracle Traffic Director, one of the key components of the Exalogic Elastic Cloud Software, provides hardware-level application acceleration, including SSL encryption. It supports fault tolerance through built-in load balancing, connecting components with the Exabus fabric. As traffic volume on the Exalogic system varies, Oracle Traffic Director can easily be scaled in lockstep with the required application computing resources. With Oracle Traffic Director, enterprises give applications administrators the ability to build service levels for applications and to shape traffic into and out of them. It enables administrators to set service levels for different workloads and to change rules dynamically, without the need to coordinate with other IT constituencies. There's no need for a separate, software-based ADC because Oracle Traffic Director is built in.

Enterprise deployments can take advantage of key features of Oracle Traffic Director on a high-performing Oracle Exalogic platform. Oracle Traffic Director is fully integrated with the Oracle Exabus I/O subsystem and can support both extremely high throughput and low-latency application-traffic workloads. Varying volumes of application traffic can easily and dynamically be scaled through the use of Oracle Traffic Director in the enterprise architecture. Oracle Traffic Director can easily be configured to apply multiple, declarative rules when distributing requests to the back-end servers and when forwarding responses to clients. Oracle Traffic Director is not limited to routing inbound traffic from the external client network. It can also be used to route traffic between processes running on the same Infiniband fabric. Internal traffic like service calls between various Oracle SOA Suite components never have to leave the IB Fabric. Maximum availability of various deployments can be achieved using high-availability features of Oracle Traffic Director, such as active-passive or active-active failover.

The goal of this technical paper is to provide the following information:

- Disaster recovery deployment options
- Configuration flow and steps
- Disaster recovery operations

This paper describes the disaster recovery scenarios for Oracle Traffic Director with Oracle HTTP Server instances as origin servers; however, the concepts presented herein can also be adapted for use in other deployments supported by Oracle, such as these:

- Oracle Traffic Director deployed in Active-Active high-availability mode at each site
- Oracle Weblogic Server instances serving as origin servers for the Oracle Traffic Director setup
- Oracle Fusion Middleware product services such as Oracle WebCenter Content listening over TCP-based protocols, like the Remote Intradoc Client (RIDC) socket-based protocol.
- Oracle Identity Management suite services listening over an LDAP authentication provider, like Oracle Internet Directory.

This paper does not describe any disaster recovery procedures for the Exalogic compute nodes or the Exalogic Control vServers, switches, storage appliances, and guest vServers used for Oracle Traffic Director setup.

This document does not describe any disaster recovery procedures for configuration of the Exalogic infrastructure components or for backing up and restoring the management components (the Exalogic Control Stack) of the cloud infrastructure used for creating the virtual machines on which the Oracle Traffic Director instances are installed. For information about the backup and recovery procedures for these components not described in this paper, refer to the "Oracle Exalogic Backup and Recovery Best Practices" white paper.

In this white paper, OTD is used as an abbreviation for Oracle Traffic Director.

Overview of Oracle Traffic Director Deployment

Depending on the user requirements and to maintain production service-level agreements (SLAs), Oracle Traffic Director can be deployed in different modes, such as a single instance for a development or test setup, Active-Active or Active-Passive mode for a highly available deployment, or Disaster Recovery mode for protecting the primary production site. A typical Oracle Traffic Director deployment would comprise the following components:

- **Administration Server:** A specially configured Oracle Traffic Director instance that hosts the user interfaces--the Administration Console and the command-line interface--through which you can create OTD configurations, deploy them as instances on Administration Nodes, and manage the instances. The Oracle Traffic Director Administration Server is not created automatically when you install the product, so it should be created after installation on the Administration Server node.

- **Administration Node:** A physical host on which Oracle Traffic Director instances are deployed. Note that on an Administration Node only one instance of a particular configuration can be created.

- **Oracle Traffic Director configuration:** A collection of configurable elements (metadata) that determine the run-time behavior of an Oracle Traffic Director instance. A typical OTD configuration contains definitions for the listeners (IP address and port combinations) on which OTD should listen for requests and information about the servers in the back end to which the requests should be sent. OTD reads the configuration when an Oracle Traffic Director instance starts and while processing client requests. All of the configurable elements of an Oracle Traffic Director instance are stored as a configuration, which is a set of files created in the config-store directory under OTD *INSTANCE_HOME.*

- **Oracle Traffic Director instance:** A server process that is instantiated from an Oracle Traffic Director configuration and deployed on an Administration Node or the Administration Server.

- **Oracle Traffic Director Failover Group:** This ensures the high availability of OTD instances by combining two OTD instances, using one or two virtual IP (VIP) addresses.

  Whenever two OTD instances are grouped by a virtual IP address (VIP) to provide high availability, they are known to be in **Active-Passive** failover mode. Requests are received at the VIP and routed to the OTD instance that is designated as the primary instance. If the primary instance is not reachable, requests are routed to the backup instance.

  For **Active-Active** failover mode, two failover groups are required, each with a unique VIP, but both consisting of the same nodes with the primary and backup roles reversed. Each instance in the failover group is designated as the primary instance for one VIP and the backup for the other VIP.

Deployment Scenarios -- Single Site

Single-site deployment scenarios include Oracle Traffic Director setups in single-instance mode and in high-availability mode.

Single-Instance Mode

In the simplest implementation, you can have a single Oracle Traffic Director instance running on a dedicated compute node distributing client requests to a pool of servers in the back end. In this topology, however, the Oracle Traffic Director instance becomes a single point of failure. Development and test environments may have such a single-instance mode setup.

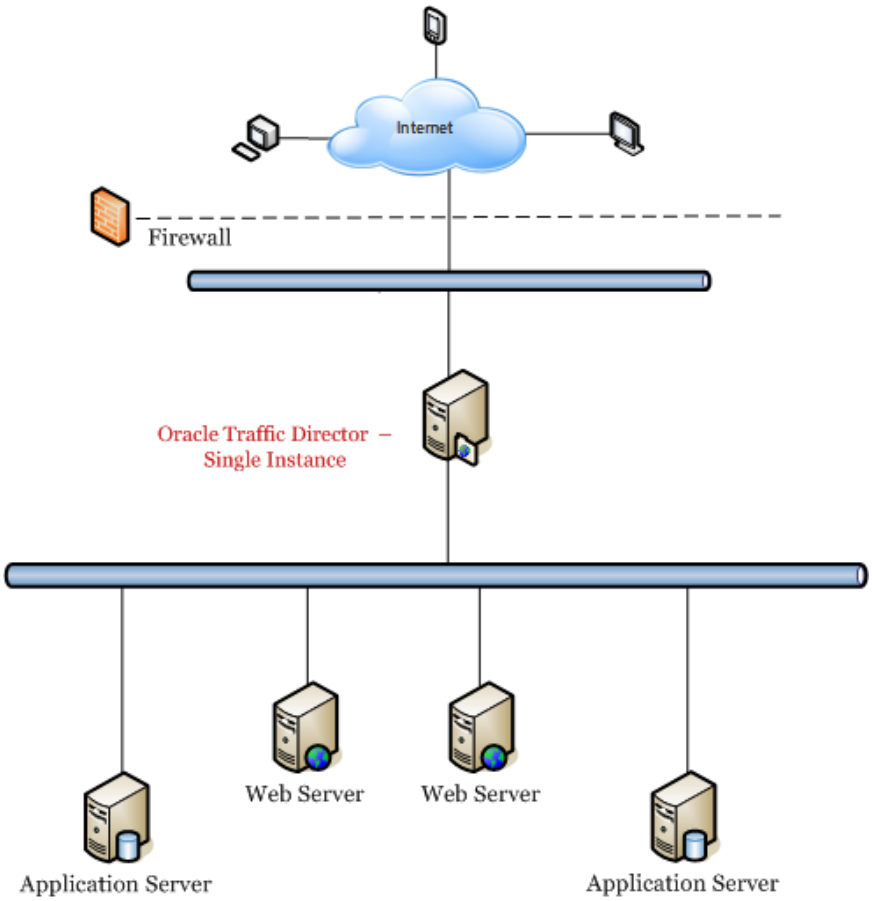

Figure 1-1. Oracle Traffic Director Network Topology: Single Site with Single-Instance Mode

High-Availability Mode

To ensure that the node on which an Oracle Traffic Director instance runs does not become the single point of failure in the topology, and to establish a highly available traffic routing and load-balancing service for the enterprise applications and services, you can configure two Oracle Traffic Director instances to provide **active-active** or **active-passive** failover. The high availability of Oracle Traffic Director instances is achieved by combining two Oracle Traffic Director instances in a failover group represented by one or two virtual IP (VIP) addresses.

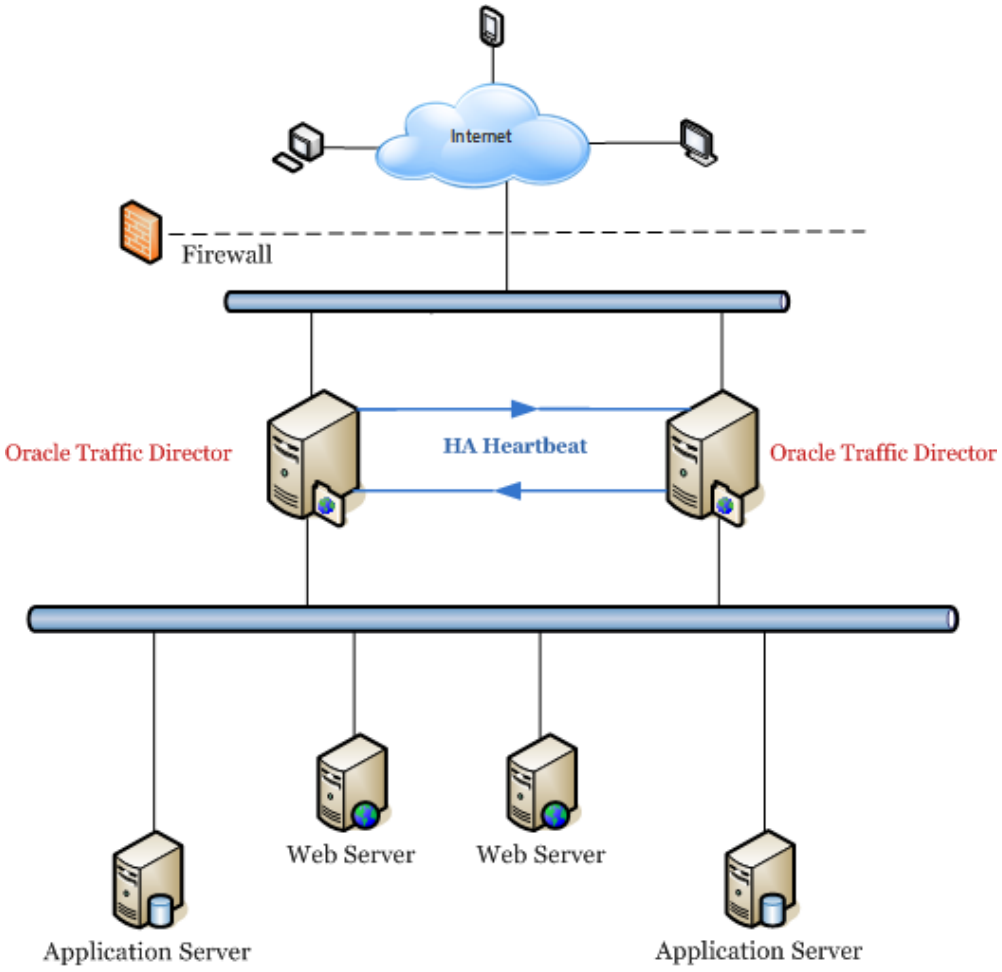

Figure 1-2. Oracle Traffic Director Network Topology: Single Site with High-Availability Mode

The failover group can be configured to work in the following high-availability modes:

- Active-Passive: In this mode, a single VIP address is used. One instance in the failover group is designated as the primary node. If the primary node fails, the requests are routed through the same VIP to the other instance.

- Active-Active: This mode requires two VIP addresses. Each instance in the failover group is designated as the primary instance for one VIP address and the backup for the other VIP address. Both instances receive requests concurrently through the two different VIP addresses instead of one VIP, as is the case in the Active-Passive mode. This mode is used mainly for load balancing end-user access traffic between the two failover group VIPs, front-ended with a single load-balancer virtual IP.

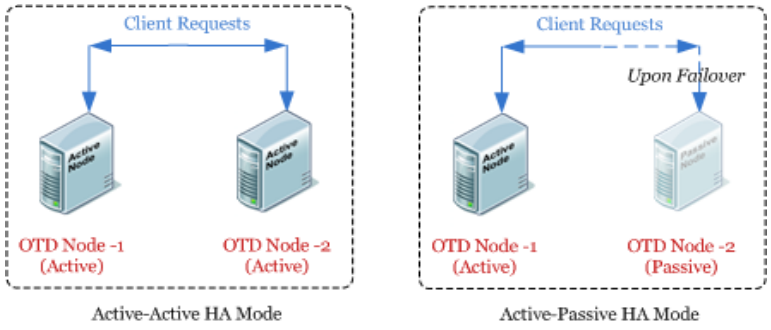

Figure 1-3. Oracle Traffic Director High-Availability Modes

For additional details on configuring high availability for Oracle Traffic Director, refer to the "Configuring Oracle Traffic Director for High Availability" chapter of the *Oracle Traffic Director Administrator's Guide.*

Because the focus of this white paper is to describe the disaster recovery solution and its steps, a detailed explanation of how Oracle Traffic Director functions in a highly available mode is not included. The solution described in this paper can be applied to single-instance mode, active-passive mode, and active-active mode. However, the disaster recovery operation and tests were carried out on a deployment where Oracle Traffic Director instances were configured in Active-Passive mode at both sites.

Deployment Scenarios -- Disaster Recovery Setup

While high availability of Oracle Traffic Director instances at a particular data center is ensured by combining two OTD instances in a failover group represented by one or two virtual IP (VIP) addresses, the protection of OTD deployment at a site is achieved by having an equivalent OTD deployment at a standby site. The recovery site gets activated in the event of any disaster at the primary site.

This white paper covers the disaster recovery options only for OTD setup and not for origin servers (like Oracle HTTP Server) used in the setup. However, for recovering a complete site (comprising OTD and origin servers), disaster recovery for the origin servers must also be configured.

The two most viable disaster recovery scenarios for Oracle Traffic Director follow:

- **ZFS storage replication-based standby**
- **OTD instance synchronization-based standby**

The topology at each site for these validated disaster recovery scenarios consists of two Oracle Traffic Director instances: OTD admin node-1 and OTD admin node-2, forming an **Active-Passive failover pair** and providing a **single virtual IP** address for client requests.

Specific details of each disaster recovery scenario for Oracle Traffic Director are covered later in this document.

In this Active-Passive high-availability mode, one node in the failover group is redundant at any point in time. Each instance caters to requests received on one virtual IP address and backs up the other instance. When the active instance (OTD admin node-1 in this example) receives a request, it determines the server pool to which the request should be sent and forwards the request to one of the servers in the pool based on the load-distribution method defined for that pool.

For the exercises in this paper, two symmetric sites were built:

- Primary Data Center, referred to as **OTD_A**
- Standby Data Center, referred to as **OTD_B**

In each site's topology, two Oracle HTTP Server instances are configured as a pool of origin servers in the back end. Although it was not performed in the white paper exercise, Oracle Traffic Director in general can also be configured to route requests to servers in multiple server pools that include Oracle WebLogic Server instances or LDAP servers.

At both sites, the hosts for Oracle Traffic Director and Oracle HTTP Server are provisioned as vServers from a Virtual Data Center (vDC) hosted on the Oracle Exalogic machines. Each vServer has interfaces on the following networks:

- A public (Ethernet-over-Infiniband) network used for data center connectivity
- A private (Infiniband-based) network used for communication between vServers hosted on the Exalogic machine
- A private (Infiniband-based) network, known as the IPoIB-vserver-shared-storage network in the Exalogic virtualized environment, used for access to the ZFS Storage Appliance in the Exalogic machine

The following toplogy is used when site OTD_A operates as the primary site.

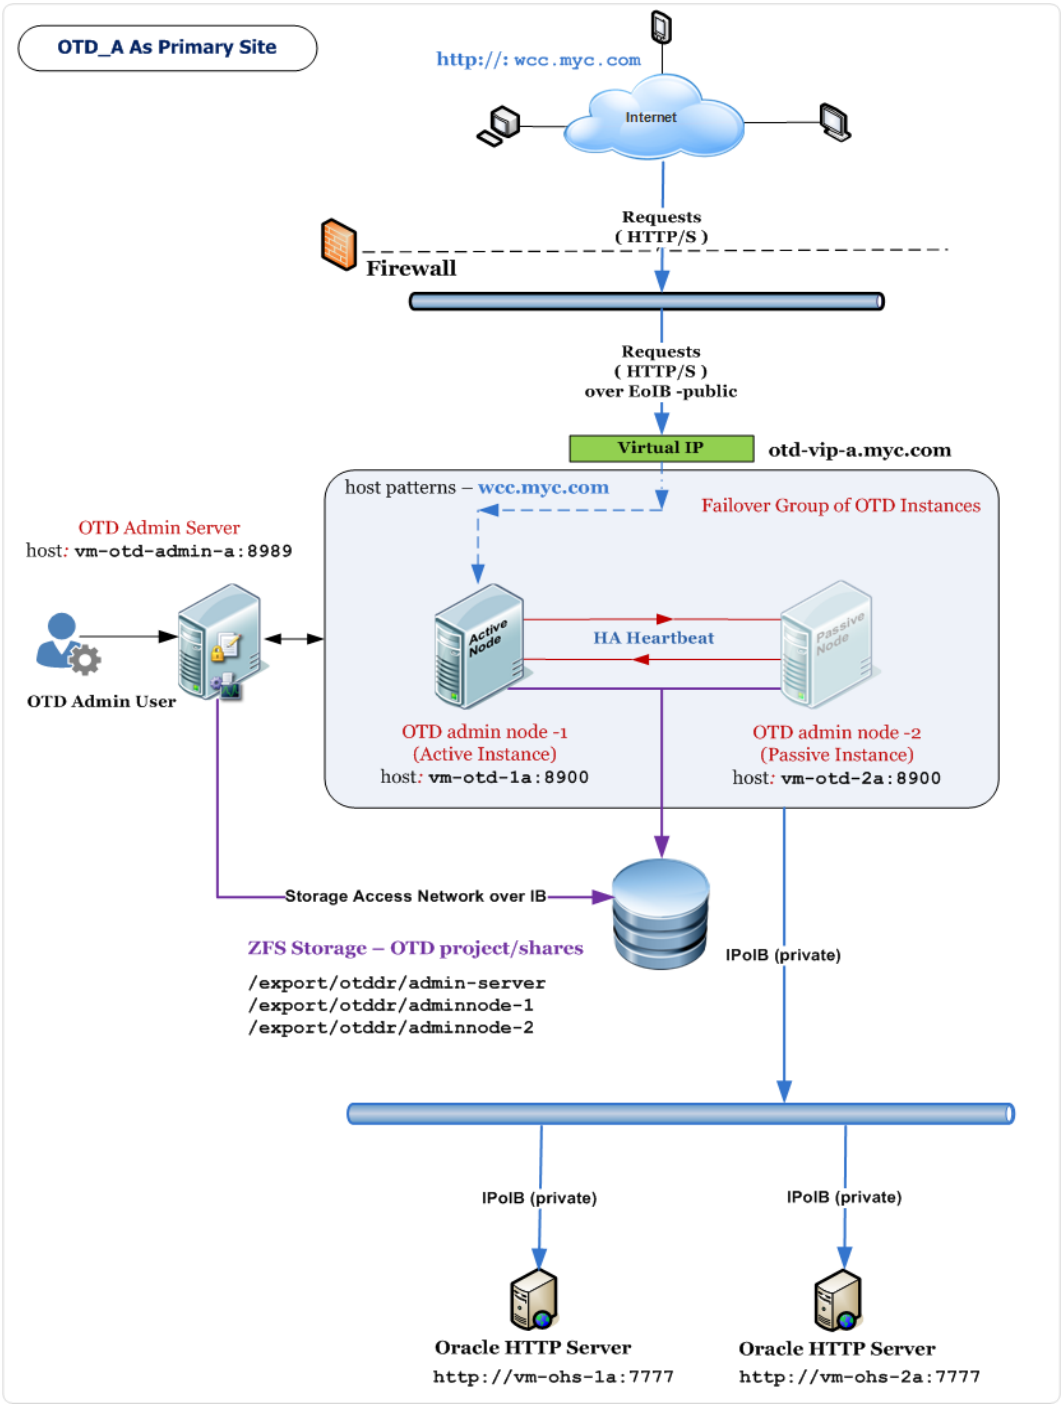

Figure 1-4. Oracle Traffic Director Network Topology: OTD_A Operating As Primary Site

The following toplogy is used when site OTD_B operates as the primary site.

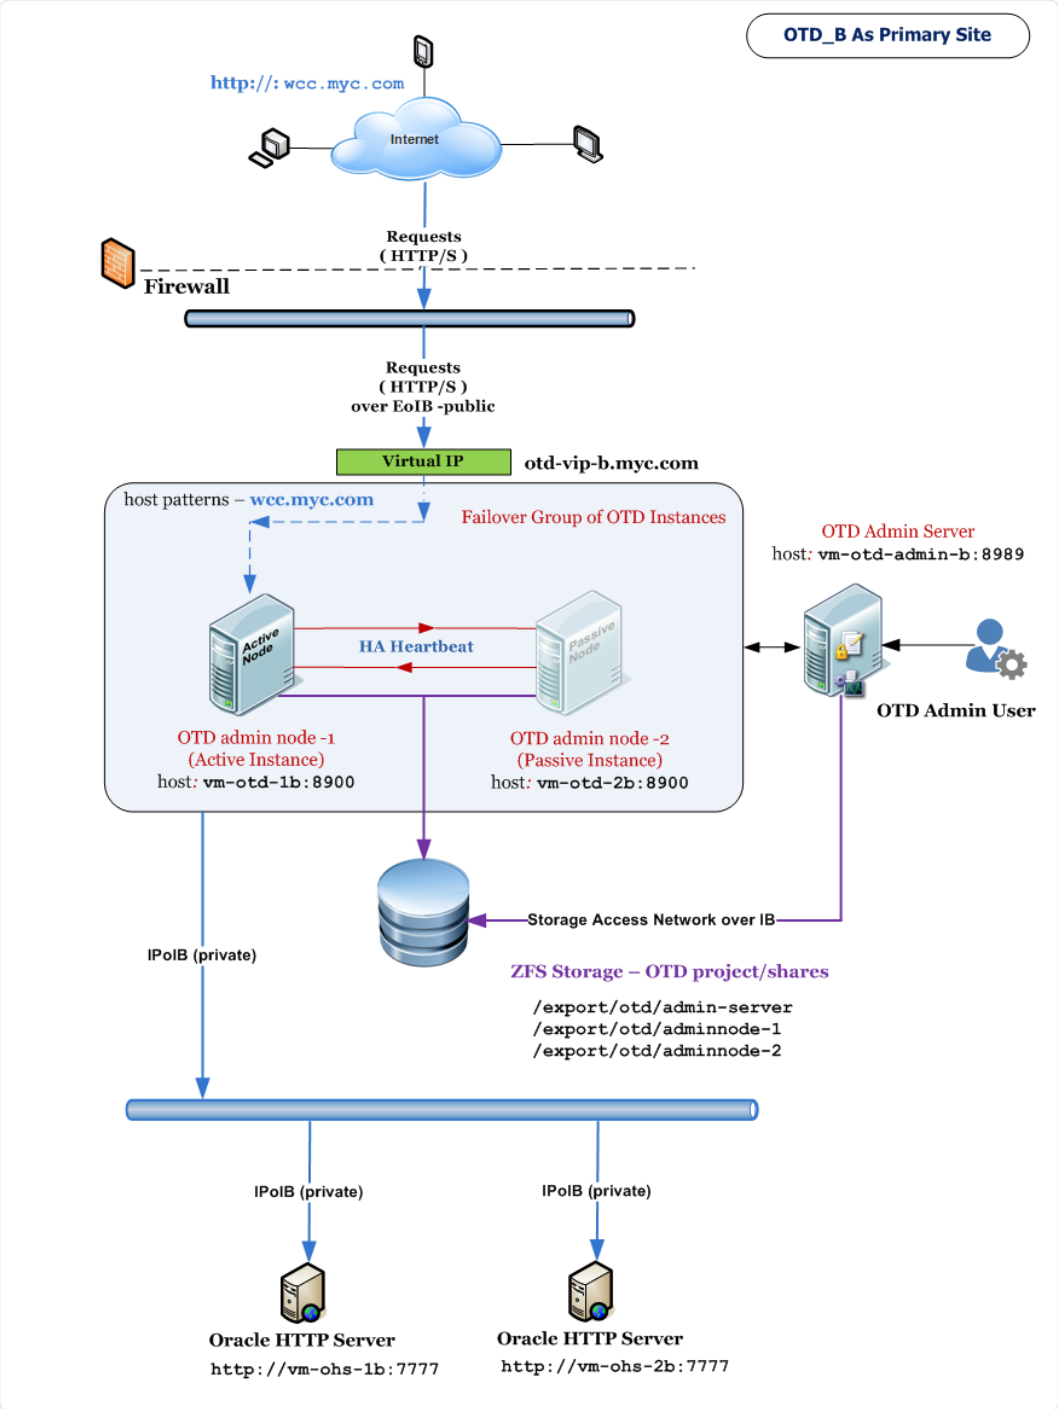

Figure 1-5. Oracle Traffic Director Network Topology: OTD_B Operating As Primary Site

Hardware Details

Hosts at Primary Site [OTD_A]

Host Name *	EoIB IP (public)	IPoIB (private)	Comments
vm-otd-admin-a	10.133.235.20	192.158.1.40	OTD admin server
vm-otd-1a	10.133.235.21	192.158.1.41	OTD admin node-1
vm-otd-2a	10.133.235.22	192.158.1.42	OTD admin node-2
vm-ohs-1a	10.133.235.23	192.158.1.43	OHS-node-1
vm-ohs-2a	10.133.235.24	192.158.1.44	OHS-node-2

Hosts at Standby Site [OTD_B]

Host Name *	EoIB IP (public)	IPoIB (private)	Comments
vm-otd-admin-b	10.143.245.30	192.168.2.50	OTD admin server
vm-otd-1b	10.143.245.31	192.168.2.51	OTD admin node-1
vm-otd-2b	10.143.245.32	192.168.2.52	OTD admin node-2
vm-ohs-1b	10.143.245.33	192.168.2.53	OHS-node-1
vm-ohs-2b	10.143.245.34	192.168.2.54	OHS-node-2

Storage Appliance at Primary Site [OTD_A]

Host Name *	Net0(IGB0) IP	IPoIB-vserver-shared-storage IP	Comments
el-prim-sn01	10.133.41.80	172.47.1.1	Active storage head
el-prim-sn02	10.133.41.81		Passive storage head

Storage Appliance at Standby Site [OTD_B]

Host Name *	Net0(IGB0) IP	IPoIB-vserver-shared-storage IP	Comments
el-stby-sn01	10.143.47.78	172.27.2.2	Active storage head
el-stby-sn02	10.143.47.79		Passive storage head

* The domain name for all hosts is **myc.com**.

Administrative access to the storage appliance is through this URL:

https://*ipaddress*:215

The IP address, or host name, is the one assigned to the NET0 port of either storage head.

ZFS Storage Project and Shares at Primary Site [OTD_A]

Project	Shares	Mount Points	Mounted on Host
OTDDR	/export/otddr/admin-server	/u01/otd_base	vm-otd-admin-a
	/export/otddr/adminnode-1	/u01/otd_base	vm-otd-1a
	/export/otddr/adminnode-2	/u01/otd_base	vm-otd-2a

ZFS Storage Project and Shares at Standby Site [OTD_B]

Project	Shares	Mount Points	Mounted on Host
OTDDR	/export/otddr/admin-server	/u01/otd_base	vm-otd-admin-b
	/export/otddr/adminnode-1	/u01/otd_base	vm-otd-1b
	/export/otddr/adminnode-2	/u01/otd_base	vm-otd-2b

All the NFS mounts used for the OTD setup were **NFSv4**. The prerequisites for using NFSv4, like the NIS setup, are not covered in this document. For details, refer to the section "Configuring NFS Version 4 (NFSv4) on Exalogic" of *Oracle Exalogic Elastic Cloud Machine Owner's Guide.*

The project and shares are available at the standby storage appliance and ready to mount on the standby site servers after storage reversal during a switchover or failover from the primary site. During the storage replication, the project and shares are available as the **Replica** project and shares at the standby storage appliance.

Software Details

The following products were used to test the deployment in this paper. No additional patches were required.

- Oracle Traffic Director 11.1.1.7
- Oracle HTTP Server 11.1.1.7
- Oracle Enterprise Manager Cloud Control 12.1.0.3
- Oracle Enterprise Manager Plug-in for Oracle Fusion Middleware 12.1.0.4 (includes the Oracle Site Guard Plug-in)

  Oracle Site Guard is used for automating disaster recovery operations.

Network Details

The network topologies in the deployments tested for this white paper take advantage of the high bandwidth and performance of the Exalogic internal Infiniband (IPoIB) network fabric. The Exalogic default IPoIB network is used for all internal communications between the Traffic Director Instances and the origin servers deployed on Exalogic compute nodes (physical) or the vServers in a virtualized environment.

Virtual Servers

Virtual Server	Virtual IP *	Site
otd-vip-a.myc.com	10.133.235.50	Primary Site [OTD_A]
otd-vip-b.myc.com	10.143.245.60	Standby Site [OTD_B]

* EoIB (public) IP assigned to the Oracle Traffic Director failover groups. These virtual IPs are managed by OTD/VRRP and do not need to be explicitly enabled with ifconfig.

Global Site Selector

In the event of a primary-site disaster and after the standby site has assumed the production role, a **global site selector** is used to reroute user requests to the standby site. Global site selectors, like F5 – BigIP Global Traffic Manager (GTM) and Cisco – Global Site Selector (GSS), also handle DNS server resolution (by offloading the resolution process from the traditional DNS servers).

As described in each site's topology, **'wcc.myc.com'** is the end-user access URL from the Internet, and it front-ends both the active site and the standby site. The active site uses the DNS alias **'wcc.myc.com'** to point to **'otd-vip-a.myc.com'**, which resolves to the IP address allocated to the failover group of the OTD_A setup. The standby site (OTD_B) has the virtual IP **'otd-vip-b.myc.com'**, which is allocated for the failover group of the OTD_B setup.

Without a global site selector, failover or switchover would rely on the manual process of updating DNS to point 'wcc.myc.com' to the standby site virtual IP 'otd-vip-b.myc.com'. With a global site selector, 'wcc.myc.com' resolves automatically to the site that is currently active without any need for manual intervention. The global site selector accomplishes this task automatically, by being aware of the current state of the OTD setup at each site. DNS resolution is then instantly changed to reroute traffic to the new primary site.

The configuration of the global site selector is not described in this white paper. However, for the configuration of a global site selector for selection of the Oracle Traffic Director deployments, the network administrator needs to follow these guidelines:

- Configure corporate-wide DNS to have Address records for the virtual IPs assigned to the Oracle Traffic Director failover group.

- Configure corporate-wide DNS to hand off a request to the global site selector used for Oracle Traffic Director site selection by adding it as an authoritative child domain.

Prerequisites

The following prerequisites are required for setting up the Disaster Recovery Solution for Oracle Traffic Director:

- **Hosts Preparation**
  - All the hosts in the Oracle Traffic Director deployment at each site must run the same operating system version, using identical patches and service packs. Also at each site, all the hosts must be available on a network with the same subnet.
  - For this MAA exercise, all the hosts have Oracle Linux as their operating system. Oracle Traffic Director should be the only consumer of the *Keepalived* process in all the hosts used for configuring the failover group, by not having any other applications requiring that the *Keepalived* process run on these hosts.
  - In the case where NFSv4 mounts of the ZFS storage volumes are used for the Oracle Traffic Director setup, as in this white paper exercise, ensure that the hosts have the proper NIS settings. Also make sure that the NIS server used for the NIS settings on the hosts is the same as the NIS server used in the ZFS Storage Appliance NIS settings.
  - For the OTD instance synchronization-based standby disaster recovery option, there must a remote sync tool and a time-based scheduler application on the Administration Server host at each site for transferring the OTD instance changes between sites. For the white paper exercise, rsync and cron utilities available with Oracle Linux were used.
- **Storage Configuration**
  - Ensure that the NIS settings are configured and the NIS service is started on the ZFS Storage Appliance at both sites. This is a requirement for using NFSv4 mounts on the hosts used in the overall setup.
  - All the shares, as mentioned in the “Hardware Details” section under the ”ZFS Storage Projects and Shares” table, must be created on the ZFS Storage Appliance at the primary site, prior to Oracle Traffic Director deployment at the primary site (OTD_A). For a disaster recovery scenario where the standby site is locally installed and configured, create the same storage project and shares on the ZFS Storage Appliance at the standby site.

- **Software Configuration**
  - Oracle Enterprise Manager Cloud Control with the Oracle Fusion Middleware plug-in must be installed and configured at some location in the corporate wide area network (WAN) in such a way that both sites in the deployment are accessible from it. While it is recommended to use Oracle Enterprise Manager for monitoring an Oracle Traffic Director setup, this is not compulsory if Oracle Site Guard is not used for automating the disaster recovery operations.
  - Install the Oracle Enterprise Manager Management Agents on all hosts in the overall deployment, and make sure that the storage locations where the Management Agents are installed on each host are not replicated to the standby site.
- **Network Preparation**
  - Allocate one virtual IP address for each site to be used for the failover group of Oracle Traffic Director. The addresses must belong to the same subnet as that of the nodes in the failover group. They should be DNS resolvable and accessible over the EoIB network. Ensure that for the network interface on which the failover-group virtual IP is created is the same on all the Administration Node hosts. For example, if the failover group is created on the bond0 EoIB interface at the primary site, make sure that bond0 is available as the EoIB interface at the standby site. This is a requirement for smooth migration of the failover group from the primary site to the standby site.
  - At the standby site, ensure that the primary site's host names and the primary site's virtual IP, 'otd-vip-a.myc.com', resolve to the IP addresses of the corresponding peer systems. This can be set up by creating aliases for host names in the /etc/hosts file. For both disaster recovery deployment options, make sure aliases for all the systems and the virtual IP names exist.

ZFS Storage Replication-Based Standby

In this solution, the production site (also referred to as the primary site) is in active mode, while the second site is serving as a standby site, in passive mode. Oracle Traffic Director is installed and configured only at the primary site. The virtual IP for the failover group in the Oracle Traffic Director setup is enabled only at the primary site, allowing each external client access only to traffic routed to it. The Oracle Traffic Director setup resides on shared storage that gets replicated to the remote site, making the Oracle Traffic Director binaries and latest configuration data available at the standby site during a site failure or site maintenance event. All the Oracle Traffic Director binaries, configuration data, logs, and security data are replicated to the remote site using the Remote Replication feature of the Sun ZFS Storage Appliance.

Using a remotely replicated standby setup has these advantages:

- Installation, configuration, patching, and upgrade need to be performed at only one site.

- Keeping primary and standby sites synchronized is greatly simplified.

In this disaster recovery scenario, where only one site can be active at any given time, the application traffic is directed to the appropriate site by a DNS entry that is updated when the site is considered as primary. The standby site must be similar to the primary site in terms of hardware and network resources to ensure there is no loss of performance when failover happens. There must be sufficient network bandwidth between the primary and standby sites to handle remote storage replication.

For monitoring the deployment at both sites, Oracle Enterprise Manager Cloud Control is used. Disaster recovery operations in this deployment are automated through Oracle Site Guard, a component of Oracle Enterprise Manager Cloud Control.

Deployment Topology

The following topology is for the solution in which the hosts at a standby site are in standby mode with OTD projects and shares being remotely replicated across the corporate WAN.

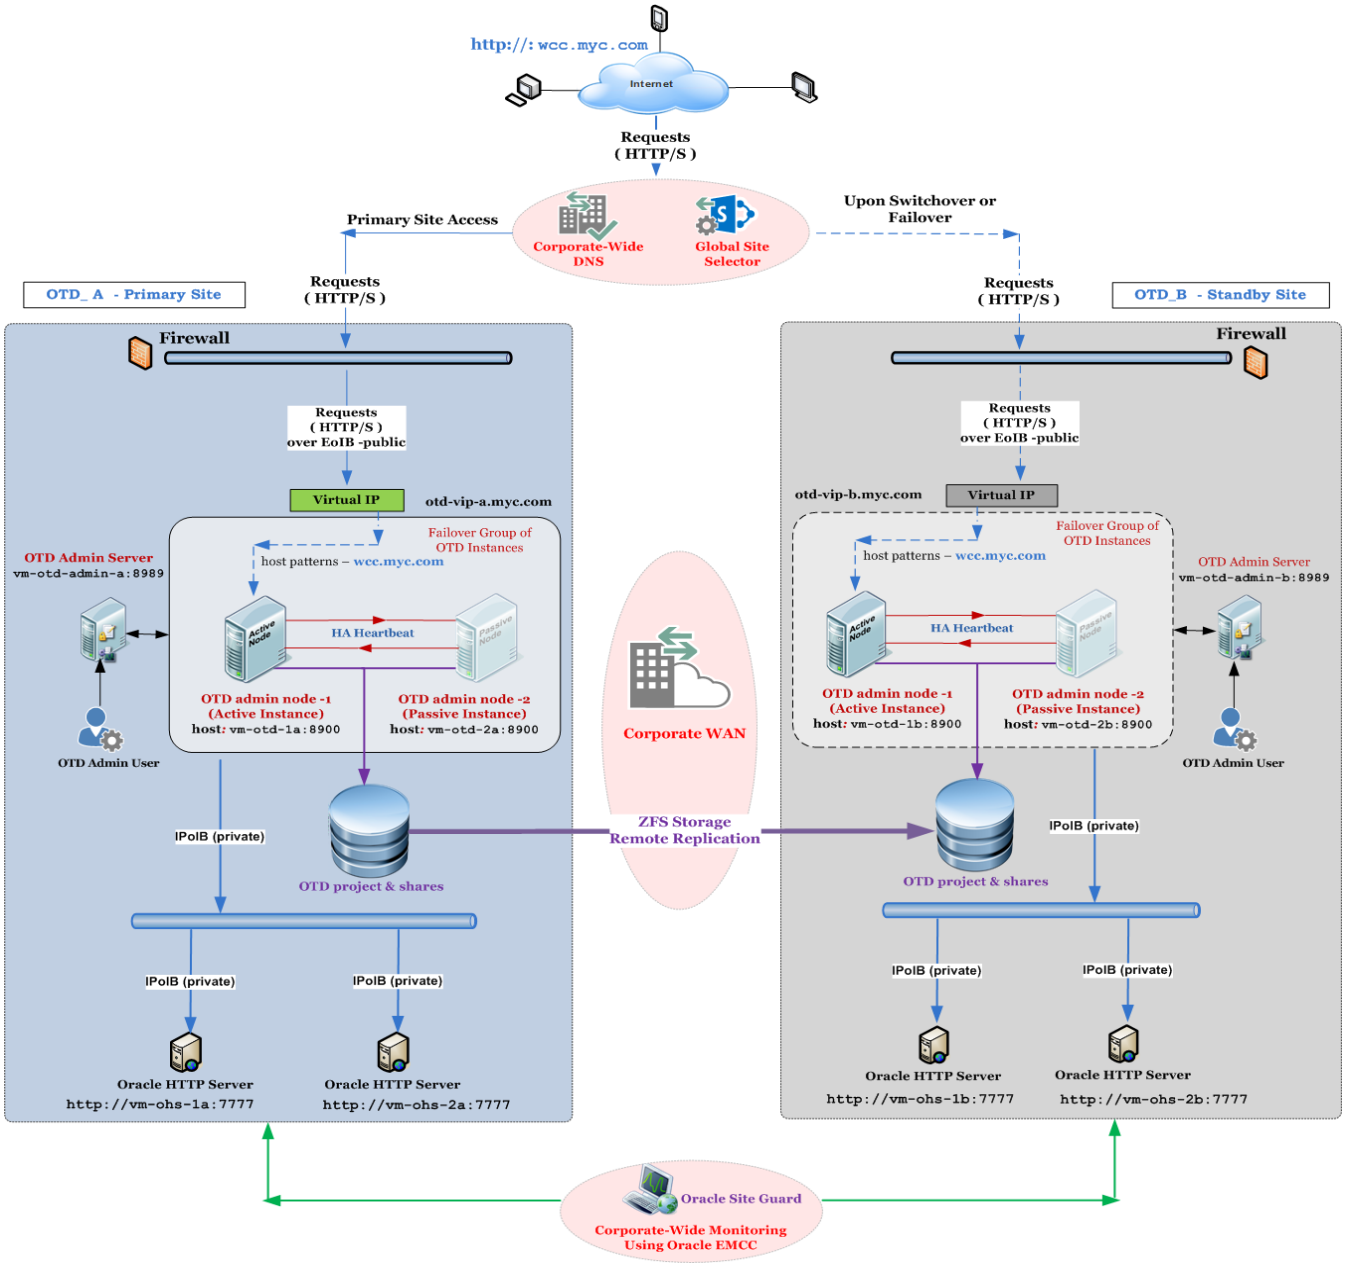

Figure 1-6. Oracle Traffic Director Disaster Recovery Topology – ZFS Storage Replication-Based

Configuration Flow

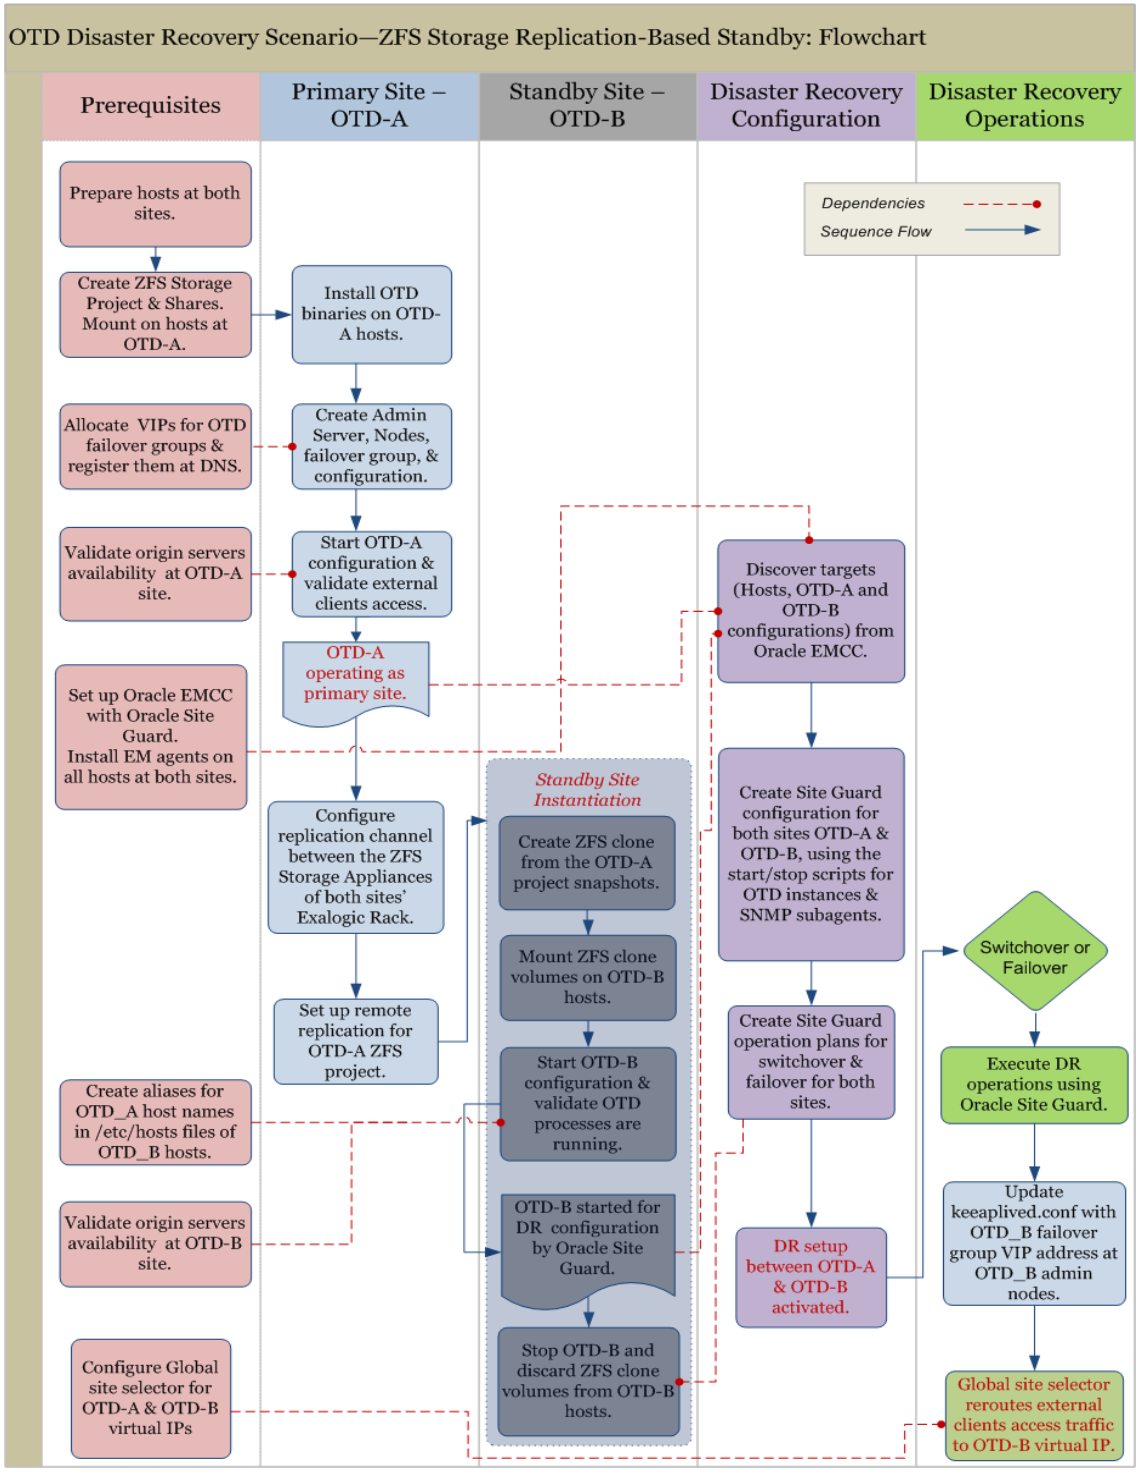

Detailed Procedures

As outlined in the flowchart, the steps followed for the Oracle Traffic Director Disaster Recovery setup are described in this section, with a remotely replicated standby site using ZFS storage replication.

Primary Site (OTD_A) Setup

1. Install binaries for Oracle Traffic Director on the ZFS storage volumes mounted on the respective hosts at the OTD_A site, which is the primary site.

2. At site OTD_A, create and start the Administration Server and the Administration Nodes on their respective hosts. On the OTD instances, deploy a newly created configuration that has the virtual server with the required listener and host patterns. The OTD server instances can be started at this stage.

   Details of the OTD configuration used in the MAA exercise follow.

	Values	Comments
Configuration	**elv-wcc-config**	
Virtual Server	**elv-wcc-config**	
Hosts	**wcc.myc.com**	This entry is for the list of host names or URL patterns that will be served by the virtual server for access by the external clients.
HTTP listener port	**80**	
HTTP listener address	*	Setting the IP address to * enables the HTTP listener in the configuration to listen on all IP addresses of the hosts and on that port. The failover virtual IP is also one of the IPs available on the active host.

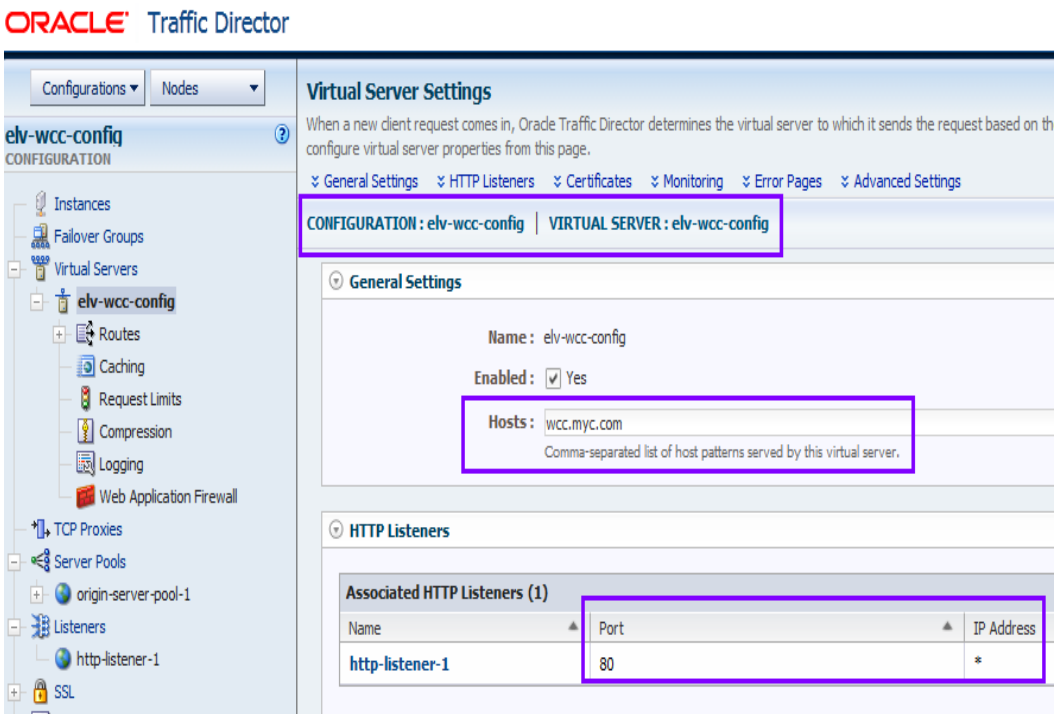

3. Create a failover group of the Oracle Traffic Director server instances in Active-Passive high-availability mode with the virtual IP 'otd-vip-a.myc.com' assigned to the failover group. The Router ID is unique across failover groups at the primary site, and upon a switchover or failover, it should also be unique at the standby site.

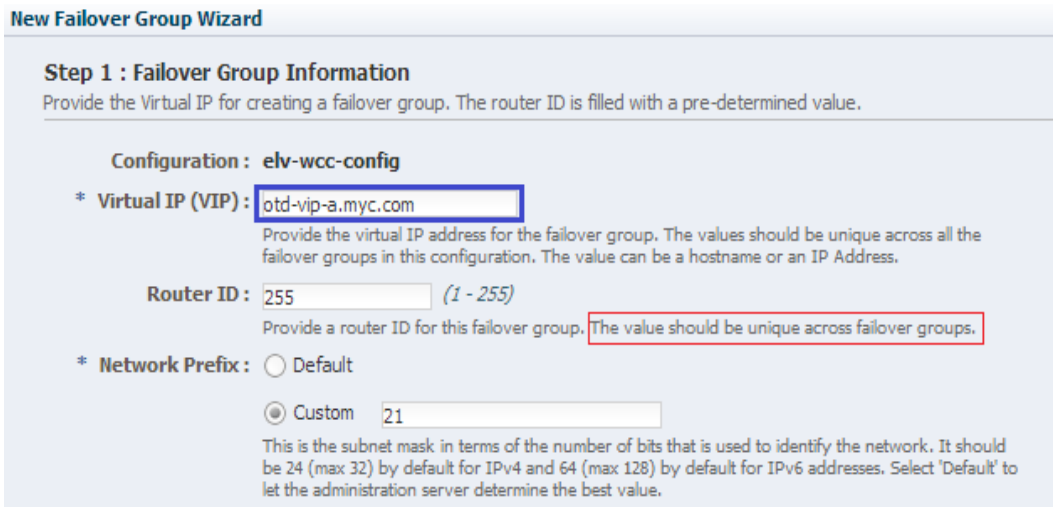

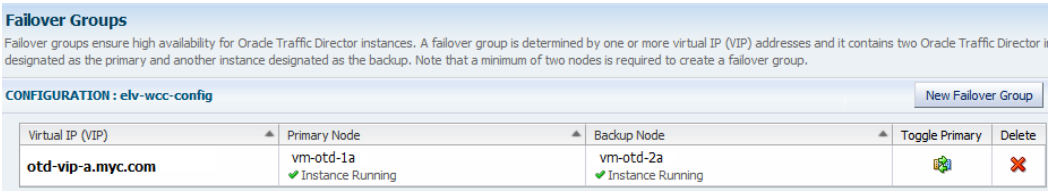

4. Ensure that the instances at site OTD_A are available. Validate external-client access to the URL **wcc.myc.com,** and check for the corresponding web pages fetched from the origin servers (Oracle HTTP Server instances). This confirms that OTD_A is operational as the primary site.

   Note: At this stage, prepare the start and stop scripts for all components of the OTD_A setup. The scripts will be used later for the Oracle Site Guard configuration for Oracle Traffic Director.

   For details on installing and configuring Oracle Traffic Director, refer to its "Installation Guide" and "Administration Guide".

Storage Replication Configuration

5. Configure a storage replication channel for replication traffic between the ZFS Storage Appliances in the Oracle Exalogic racks at the primary site (OTD_A) and the standby site (OTD_B).

6. On the ZFS Storage Appliance at the primary site (OTD_A), configure and enable remote replication for the project **OTDDR** in continuous-replication mode. Ensure that snapshots are included in the replication setup. Validate that the replica and snapshots of this project get created on the ZFS Storage Appliance at the standby site. The tasks in this step can be performed using the ZFS Storage Appliance Administration Console or the storage ILOM console.

   For details on ZFS storage replication, refer to the "Replication" chapter in the *Sun ZFS Storage System Administration Guide.*

The following ZFS remote replication mode can be used:

a) Continuous replication

b) On-demand replication after any OTD administration operation

c) Scheduled replication at regular intervals

d) Combination of (b) and (c)

Option d is suggested because it covers most of the cases.

Standby Site (OTD_B) Instantiation

7. The standby site (OTD_B) must be instantiated for configuring the disaster recovery setup and operations by Oracle Site Guard. For OTD_B instantiation, create a ZFS clone, named '**OTDDRClone',** of the OTD project OTDDR on the ZFS Storage Appliance at the OTD_B site, using the snapshot received from the OTD_A site. This clone is basically a local ZFS project on a standby storage appliance with the volumes containing the OTD_A binaries and configuration. This task is performed only once during the disaster recovery configuration step.

8. Mount the local ZFS project **OTDDRClone** onto the hosts at the OTD_B site so that the installation binaries, configurations, and security data from the OTD_A deployment are now available to the hosts at OTD_B.

9. Ensure that the primary site host names and the primary site virtual IP '**otd-vip-a.myc.com'** resolve to the IP addresses of the corresponding peer systems at the standby site. This can be set up by creating **aliases for host names** in the **/etc/hosts** file.

10. Start the Administration Server, Nodes, and instances at OTD_B. Ensure that the origin servers (Oracle HTTP Server instances) for OTD_B deployment are also available. Validate that the OTD_B deployment is running without any external client access. Access to the OTD_B Administration Console is done from the internal network using host-name aliasing from the client machine for validating OTD_B.

Note: At this stage, prepare the start and stop scripts for all components of the OTD_B setup. The scripts will be used later for the Oracle Site Guard configuration for Oracle Traffic Director. Refer to "Appendix" for the start and stop scripts

Oracle Site Guard Configuration for Oracle Traffic Director

11. Before configuring Oracle Site Guard for Oracle Traffic Director, ensure that Oracle Enterprise Manager Cloud Control is set up at some location in the corporate WAN and has access to the hosts at both of the OTD sites. Install Oracle Enterprise Manager Agents on each host at both sites, and manually add all the hosts as targets.

12. Start the SNMP subagent available in the Oracle Traffic Director instance home on all the hosts at each site. This is a requirement for Oracle EMCC monitoring of the OTD instances through SNMP and also for configuring the availability status of the instances during disaster recovery operations by Oracle Site Guard. SNMP subagents can be started either from the Oracle Traffic Director Administration Console or from a command-line script.

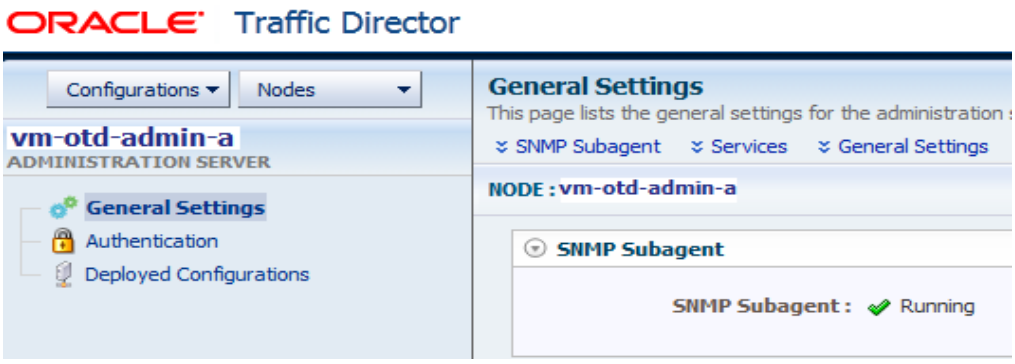

The SNMP subagent can also be started from the command-line interface as follows. Starting the SNMP subagent from the command-line interface does not require OTD administrator user credentials.

```
/u01/otd_base/otd_home/bin/tadm start-snmp-subagent --instance-
home=/u01/otd_base/otd_instance
```

13. Add the Oracle Traffic Director configuration targets in Oracle Enterprise Manager Cloud Control by selecting **'Traffic Director'** in the **Target Types** field, as shown in the following screenshots.

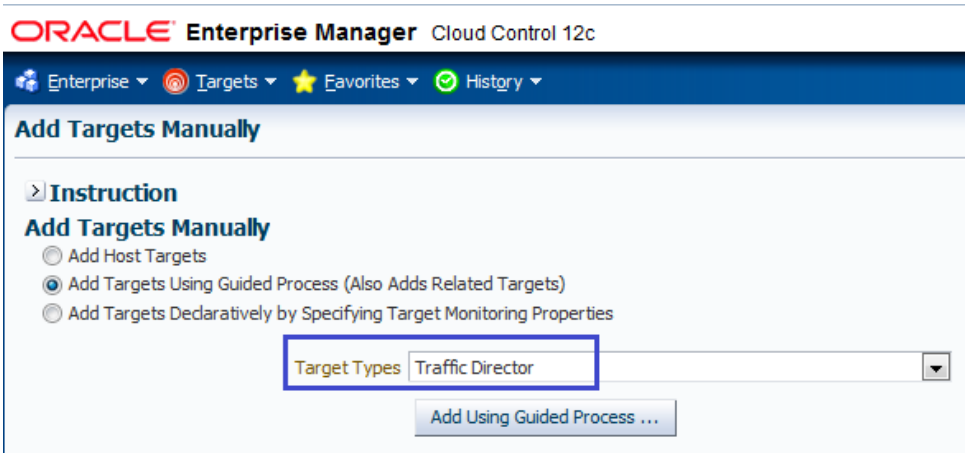

The targets of the Oracle Traffic Director deployment, which includes configuration and instances, are added by the discovery done on the Administration Server host for each site, as shown in the following figure. No separate additions of targets are required for the Administration Node hosts.

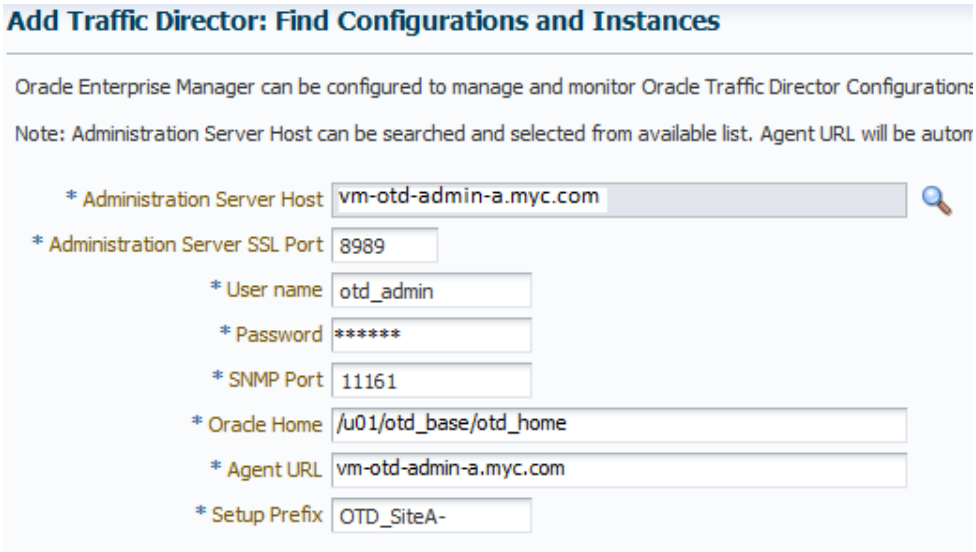

14. Create a generic system for each site with Oracle Site Guard. In this exercise, the sites are named OTD_A and OTD_B. The following screenshots of site OTD_B show the generic system-creation step.

    Log in to Oracle Enterprise Manager Cloud Control. Select **Targets** and then select **Systems**.

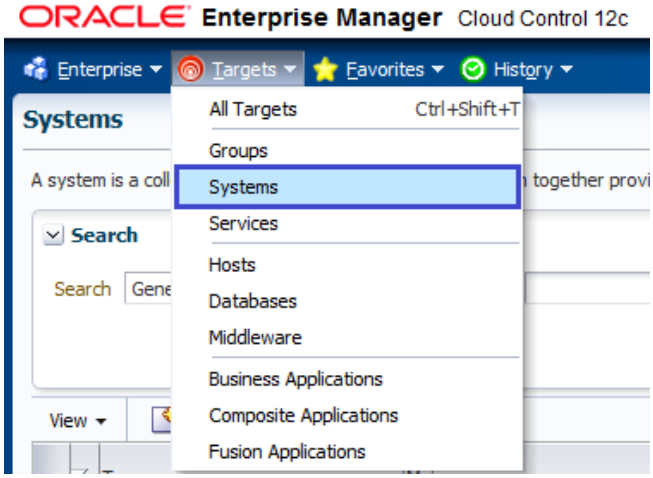

    Once the Systems page is displayed, select **Generic System** from the drop-down menu and click **Add**.

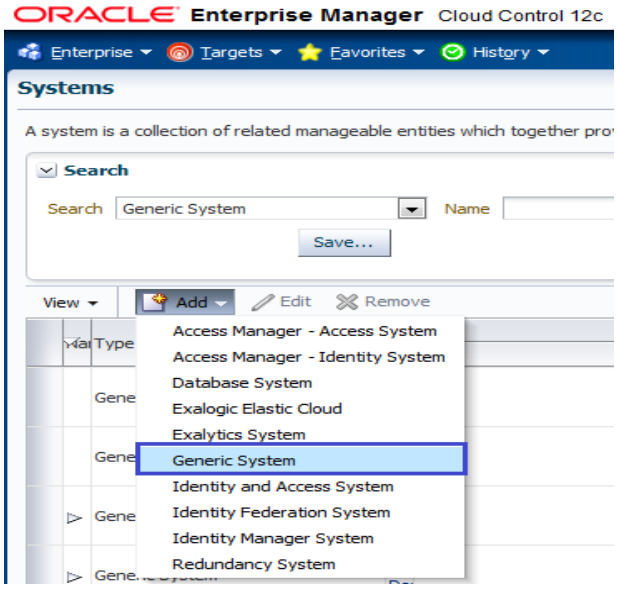

Ensure that in the **Members** section, the correct target of Oracle Traffic Director configuration is added for the intended site.

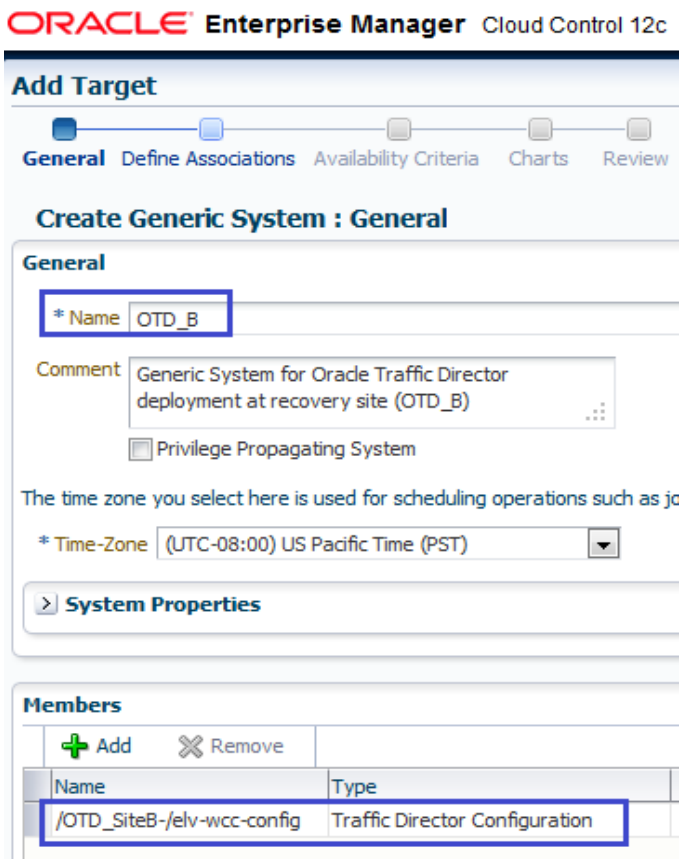

After creation of generic systems for both sites, validate that the number of members shown for each generic system in Cloud Control is the same as that of the actual deployment, to ensure that all site components are included for the disaster recovery operations.

Name	Type	Members
OTD_A	Generic System	Host(3), Traffic Director Instance(2)
OTD_B	Generic System	Host(3), Traffic Director Instance(2)

15. Create and associate credentials for both sites. Create and associate Pre and Post scripts and storage scripts for the sites. Using Oracle Enterprise Manager, create Oracle Site Guard operations for switchover and failover operations for both sites. For the list of storage scripts and the Pre and Post scripts, refer to "Appendix" of this white paper.

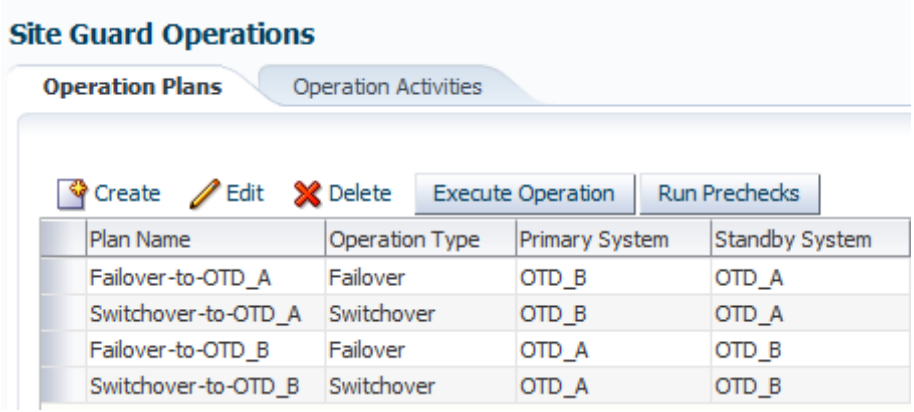

Refer to the "Configuring Oracle Site Guard Operations for Disaster Recovery" section of the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide* for instructions on creating Oracle Site Guard operations. Also refer to the steps described in the MAA white paper "Automating Disaster Recovery using Oracle Site Guard for Oracle Exalogic."

16. After completing the Oracle Site Guard configuration for both sites for Oracle Traffic Director deployment, stop the OTD_B site completely, and discard the ZFS clone volume OTDDRClone. At this stage, ensure that the storage replication for the project OTDDR is active on the primary site storage appliance and that its replica is available at the standby site storage replication. This confirms that the disaster recovery setup between the OTD_A and OTD_B sites is now complete and activated.

Disaster Recovery Operations and Testing

The following operations were validated in the MAA exercise.

SITE GUARD OPERATION	DESCRIPTION
Switchover-to-OTD_B	Switch over operations from primary to standby site
Switchback-to-OTD_A	Switch operations back to primary site from standby site
Failover-to-OTD_B	Fail over operations from primary to standby site
Fallback-to-OTD_A	Fail over operations back to primary site from standby site

For executing any disaster recovery operation (switchover or failover), submit an operation plan using Oracle Enterprise Manager Cloud Control.

1. Log in to Oracle Enterprise Manager Cloud Control. From the **Targets** menu, click **Systems**.

2. On the Systems page, click the name of the system that you want to update.

3. On the system's home page, from the Generic System menu, select **Site Guard** and then **Operations**. The Site Guard Operations page is displayed.

4. Click the **plan** listed in the Plan Name column.

5. Click **Execute Operation**. The Confirmation screen appears.

6. Select **Run PreChecks** on the Confirmation screen. Click **Yes** to submit the operation plan.

The Oracle Site Guard operation plans submitted for execution can be monitored from Oracle Enterprise Manager Cloud Control.

1. From the **Targets** menu, click **Systems**.

2. On the Systems page, click the name of the system that you want to update.

3. On the system's home page, from the **Generic System** menu, select **Site Guard** and then **Operations**. The Site Guard Operations page is displayed.

4. Click **Operation Activities**. A table listing all of the submitted operation plan executions is displayed.

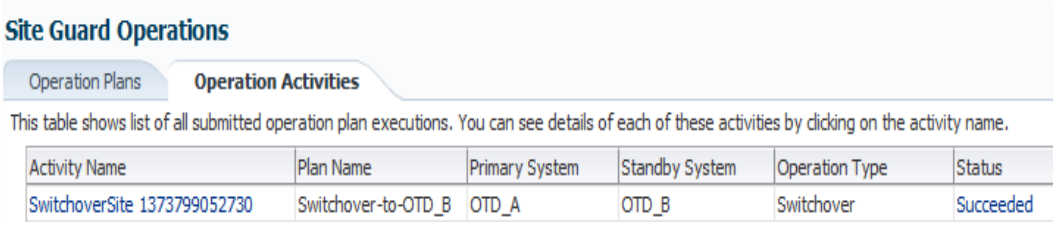

Detailed procedure steps for an operation-plan activity and their status can be monitored by clicking the activity name.

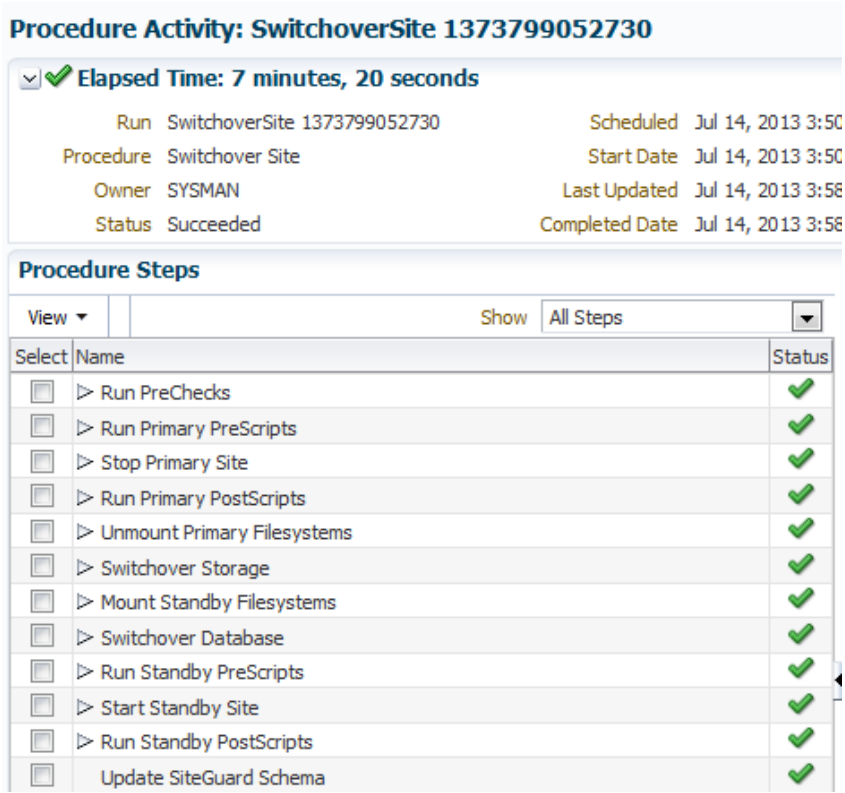

In the MAA exercise, the topology of OTD at each site has a failover group of OTD Administration Nodes. After a switchover or failover to the standby site, for ensuring that the failover group gets activated and the high availability of the OTD instances is maintained, perform the following steps at the standby site before re-routing of the external traffic to the OTD_B failover-group virtual IP:

- Edit the ***keepalived.conf*** file located in the OTD *INSTANCE_HOME*/net-**elv-wcc-config**/config/ directory at each Administration Node server (vm-otd-1b and vm-otd-2b) to replace the failover group VIP address with the IP address corresponding to **otd-vip-b.myc.com**.

  Note: For the version of Oracle Traffic Director used in this MAA exercise, it is fine to edit the keepalived.conf file, even though there is a comment in the file for not modifying it.

- As the ***keepalived.conf*** file on each Administration Node instance is modified, in the preceding step, it is important to keep the configuration store on the Administration Server in sync with the latest changes done at the Administration Nodes. An alert, **Instance Configuration Modified,** is displayed until the configuration stored on the Administration Server is synchronized with that of all its instances by the following commands, executed from the OTD_B Administration Server:

  tadm> pull-config --config=**elv-wcc-config vm-otd-1b**

  tadm> deploy-config -f **elv-wcc-config**

  Since both the instances are identical, the configuration can be pulled from either of the Administration Nodes.

  Details of these synchronization steps can be found in the section "Synchronizing Configurations Between the Administration Server and Nodes" of the *Oracle Traffic Director Administrator's Guide.*

Also refer to the section "Accessing the Command-Line Interface of Oracle Traffic Director" in the appendix for obtaining the **tadm** shell.

- After the replacement tasks for the failover-group VIP address, log in to the OTD Administration Console to verify the status of all components at OTD. As the primary site's host names resolve to the IP addresses of the corresponding peer systems at the standby site, the self-signed Administration Server certificates are valid on the standby site hosts. Secure access to the Administration Server through a browser or command-line interface can be performed without any re-creation of self-signed certificates at the standby site due to the aliasing of host names.

  The Administration Console at the OTD_B site is accessed through the URL of the Administration Server host, as follows:

  https://vm-otd-admin-b:8989

  Note: Due to the replicated content of the OTD Instances, the names for various OTD components as seen through the OTD Administration Console at OTD_B will be the same as those of the OTD_A setup.

- During an event of switchover or failover to standby site and upon the service availability of the OTD_B VIP address **otd-vip-b.myc.com** from the global site selector, the client traffic gets rerouted to the standby site OTB_B, which is now fully operational as the new primary site. Validate the following end-user URL as per the host pattern set in the example configuration in this white paper:

  http://wcc.myc.com

- Subsequent remote replication of the ZFS storage projects can be set up from site OTD_B to OTD_B. The steps for switchback or failback to site OTD_A are not described in this white paper.

OTD Instance Synchronization-Based Standby

In this solution, Oracle Traffic Director is installed and configured at both sites (primary and standby). The virtual IP for the failover group at each site is enabled. However, when the primary site is active, all external-client access traffic is routed only to the virtual IP of the failover group at the primary site, even though the virtual IP of the failover group at the standby site is active and enabled. This site selection is controlled by the corporate DNS and the global site selector. When failover or switchover is required, the global site selector seamlessly directs traffic to the already configured and enabled Oracle Traffic Director setup at the standby site (OTD_B). This task is automatically accomplished by the global site selector, which is aware of the current state of the Oracle Traffic Director instance at the primary site (OTD_A).

Any changes to the OTD configuration at the primary site are kept in sync at the standby site by synchronizing both sites' OTD *INSTANCE_HOME* directories of the Administration Server hosts. This can be achieved by using any secure and reliable remote sync application, like **rSync,** which is available on most UNIX-based operating systems. The synchronization process is scheduled to run at intervals by using a time-based job scheduler, like the **Cron** utility.

At the standby site, during a switchover or failover, the updated configurations are deployed to the required instances running on the Administration Nodes. The configuration deployment is one of the administrator-driven tasks during disaster recovery operations. Prior to any new deployment, the previous local configurations get saved as backed-up configurations.

Using an OTD instance synchronized-based standby has these advantages:

- Time to recover from a disaster is fast. RTO gets short.
- The Oracle Traffic Director setup can be shared as a primary site by other deployments with different host patterns in the OTD configurations.
- A dedicated ZFS storage-replication channel is not required.

Deployment Topology

The following topology is for the solution in which each site runs an independent Oracle Traffic Director setup, and the standby site is kept updated with any OTD configuration changes done at the primary site.

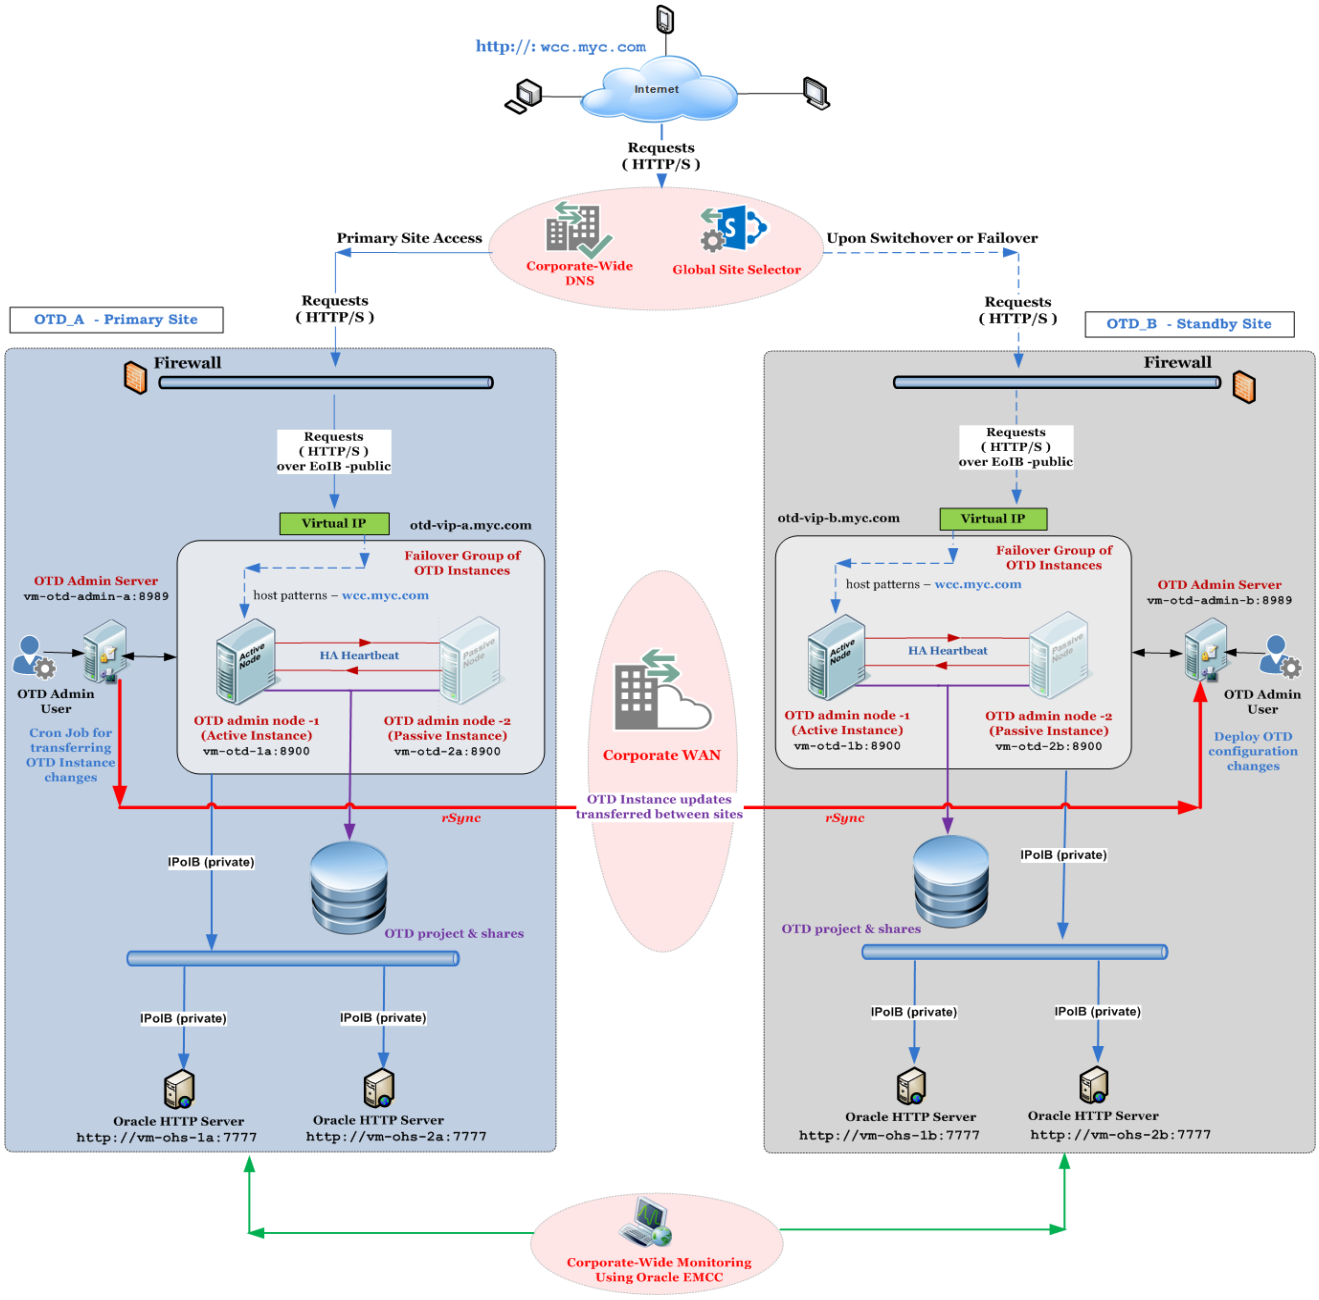

Figure 1-7. Oracle Traffic Director Disaster Recovery Topology – OTD Instance Synchronization Based

Configuration Flow

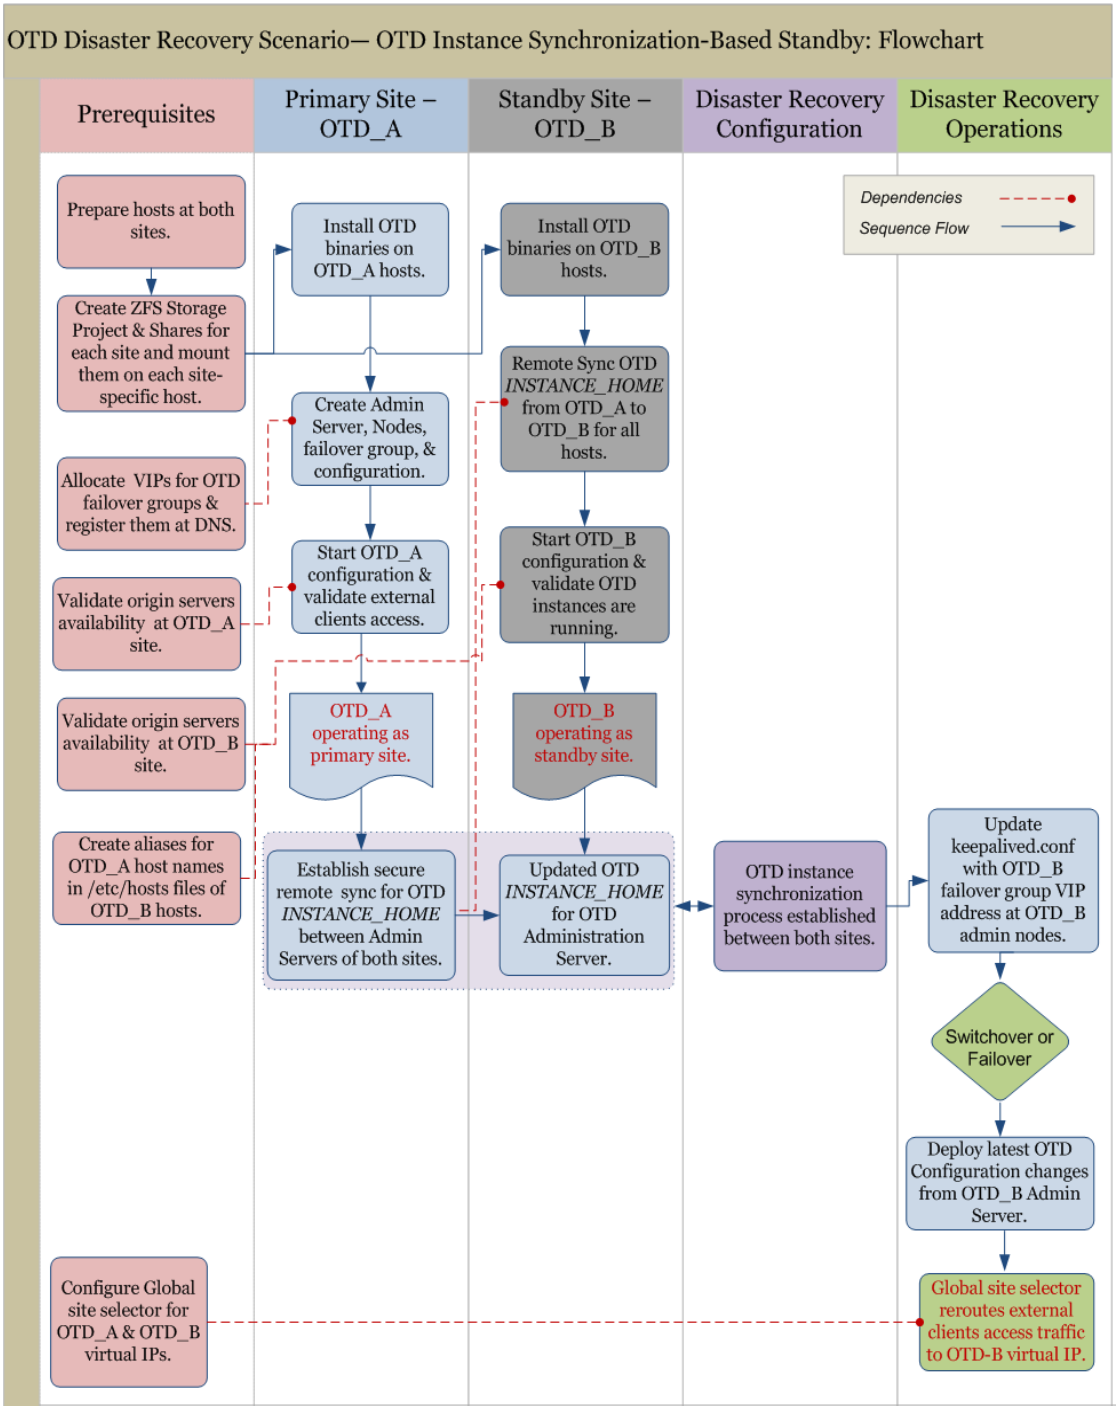

Detailed Procedures

As outlined in the flowchart, this section describes the steps followed for the Oracle Traffic Director Disaster Recovery scenario where the standby site is locally configured and updated manually with the configuration changes received from OTD_A.

Primary Site (OTD_A) Setup

1. Install binaries for Oracle Traffic Director on the ZFS storage volumes mounted on the respective hosts at the OTD_A site (primary site).

2. At the OTD_A site, create the Administration Server and the Administration Nodes on the respective hosts. Create and configure Oracle Traffic Director instances in Active-Passive high-availability mode with the virtual IP '**otd-vip-a.myc.com'** assigned to the failover groups at OTD_A. On the OTD instances, deploy a newly created configuration that has the required listener and host patterns (for example, '**wcc.myc.com**', as used in the exercise for this white paper).

3. Start all components of the Oracle Traffic Director setup at the OTD_A site.

4. For OTD_A, validate external-client access to the URL **wcc.myc.com,** and check for the corresponding web page fetched from the origin servers (Oracle HTTP Server instances). This confirms that OTD_A is operational as the primary site. The external-client access is controlled at the corporate DNS and the global site selector.

Standby Site (OTD_B) Instantiation

1. Install binaries for Oracle Traffic Director on the ZFS storage volumes mounted on the respective hosts at the OTD_B site (primary site).

2. Ensure that the primary site's host names resolve to the IP addresses of the corresponding standby site peer systems, including the virtual IP host name.

3. Remote sync the OTD *INSTANCE_HOME* of the primary site's Administration Server and Administration Node to the respective systems at the standby site. Ensure that the same directory structure is maintained for the *INSTANCE_HOME* at the standby site. The synchronization of *INSTANCE_HOME* for Administration Nodes is done only once during the standby site instantiation. Any subsequent changes at OTD_A are transferred to OTD_B by continuous synchronization of the OTD *INSTANCE_HOME* of only the Administration Server, as described in the steps later in this section.

4. At site OTD_B start the Administration Server, Administration Nodes, and server instances. Ensure that the virtual IP for the failover group at the OTD_B site, **'otd-vip-b.myc.com',** is accessible.

   Note: Ensure the networking prerequisites stated earlier in this paper are addressed for the failover group. The unique router ID and the network interface for EoIB IP for the failover groups at each site need to be the same.

Synchronization of OTD_A and OTD_B Sites

5. For the OTD_B setup to serve as a standby for OTD_A, the active configurations from OTD_A need to be in sync with OTD_B. This is achieved by synchronizing the OTD ***INSTANCE_HOME*** of the Administration Server at OTD_A with that of OTD_B. You can use any reliable and secure remote sync tool and time-based scheduler available on the host's platform.

   For the MAA exercise, this is achieved by using **rsync** over **ssh** from the operating system of the Administration Server host. Also, for the exercise, the remote sync was set unidirectional and was scheduled to execute at certain intervals by the scheduler **cron,** available on a Linux operating system. As this would be a scheduled job, set up the user equivalence using **ssh** to enable remote sync with a prompt for the password.

The following command is executed as a **cron** job on the Administration Server host at OTD_A:

rsync -avz -e ssh /u01/otd_base/**otd_instance/** root@vm-otd-admin-b:/u01/otd_base/**otd_instance/**

6. With the preceding synchronization step, the changes transferred from site OTD_A instances are ready to be published to the OTD_B server instances running on the Administration Nodes, in the event of a switchover or failover. The deployment of these configuration changes is done using the CLI **deploy-config** command from the Administration Server host, as described in the next section, “Disaster Recovery Operations and Testing”.

Disaster Recovery Operations and Testing

Similar to the ZFS storage replication DR option, in this OTD *INSTANCE_HOME* sync DR option, the topology of OTD at each site has a failover group of OTD Administration Nodes. For the failover group to become active and to achieve high availability of the OTD instances at the standby site, perform the following steps before rerouting the external traffic to the OTD_B failover group virtual IP:

- Edit the ***keepalived.conf*** file located at OTD *INSTANCE_HOME*/net-**elv-wcc-config**/config directory on each Administration Node server (vm-otd-1b and vm-otd-2b) to replace the failover-group VIP address with the IP address corresponding to **otd-vip-b.myc.com**.

  Note: For the version of Oracle Traffic Director used in this MAA exercise, it is fine to edit the keepalived.conf file, even though there is a comment in the file for not modifying it.

- As the ***keepalived.conf*** file in each Administration Node instance is modified, in the preceding step, it is important to keep the configuration store on the Administration Server in sync with the latest changes done at the Administration Nodes. An alert, **Instance Configuration Modified,** is displayed until the configuration stored on the Administration Server is synchronized with that of all its instances by the following commands, executed from the OTD_B Administration Server:

  tadm> pull-config --config=**elv-wcc-config vm-otd-1b**

  tadm> deploy-config -f **elv-wcc-config**

  Since both the instances are identical, the configuration can be pulled from either of the Administration Nodes.

  Details of these synchronization steps can be found in the section "Synchronizing Configurations Between the Administration Server and Nodes" of the *Oracle Traffic Director Administrator's Guide.*

Also refer to the section "Accessing the Command-Line Interface of Oracle Traffic Director" in the appendix for obtaining the **tadm** shell.

- During an event of switchover or failover to standby site and upon the service availability of the OTD_B VIP address otd-vip-b.myc.com from the global site selector, the client traffic gets rerouted to the standby site OTB_B, which is now fully operational as the new primary site. The mechanism for rerouting traffic is the same as that described for the global site selector in the "Network Details" section of this white paper and in the disaster recovery scenario of a remotely replicated standby site.

  The Administration Console at the OTD_B site is accessed through the URL of the Administration Server host, as follows:

  https://vm-otd-admin-b:8989

  Validate the following end-user URL as per the host pattern set in the example configuration in this white paper.

  http://wcc.myc.com

- Publish any subsequent configuration changes received from site OTD_A to all the OTD server instances running on the Administration Nodes at OTD_B by executing the following deploy-config command from the OTD_B Administration Server.

  tadm> deploy-config -f elv-wcc-config

- Validate the latest configuration changes at OTD_B from the OTD_B Administration Console.

Oracle MAA Best Practices for Disaster Recovery

1. Oracle recommends that you test the standby site periodically. This will help mitigate failures at both sites. Test the standby site by switching its role with the current primary site:
    a. Follow the site switchover procedure to switch over the standby site to the new primary site.
    b. Once testing is complete, follow the site switchback procedure to reverse the roles.

    Periodic testing validates that both the primary and standby sites are completely functional and mitigates the risk of failure at both sites. It also validates the switchover and switchback procedures.
2. Do not configure project-level and share-level replication within the same project.
3. Use the Scheduled replication mode for projects and shares in these cases:
    a. Data does not change frequently.
    b. The Recovery Point Objective value falls within your scheduled replication window.
4. Use the Continuous replication mode for projects and shares in these cases:
    a. The standby site is required to be as close as possible to the primary site.
    b. The Recovery Point Objective value is a range of a few seconds, and the allowance is for very little data loss. Data is of a critical nature.
5. Snapshots and clones can be used at the target site to offload backup, test, and development types of environments.
6. When configuring a local standby site (disaster recovery within the data center), consider disabling SSL on the replication channel. Removing the encryption algorithm enables a higher replication throughput.
7. Always enable SSL when replication is across a wide area network.

Conclusion

As validated in this MAA exercise, Oracle Traffic Director provides high-end layer-7 load-balancing features with built-in high availability, leverages the Infiniband fabric for efficient throughput, and extends to provide multisite availability. The disaster recovery options stated in this white paper are essential to maximize the Oracle Traffic Director-based web tier and load-balancing capabilities of geographically spread-out enterprise deployments. Based on the enterprise deployment requirements, either of the options can be used. This technical white paper has highlighted the procedures for configuring various disaster recovery operations.

Appendix

ZFS Storage Scripts Used

All the storage scripts in this deployment come bundled with Oracle Site Guard and were used with no modifications. The following table lists the bundled storage scripts used.

SCRIPT NAME	PURPOSE
zfs_storage_role_reversal.sh	Top-level shell script that triggers the storage role reversal. This script invokes the other AKSH action scripts as required.
retrieve_actionid_and_validate_source.aksh	Retrieves replication action and validates that the project and source pool exist.
retrieve_sourceid_and_validate_target.aksh	Retrieves replication source and validates that the package and target pool exist.
validate_last_sync_source.aksh	Validates that at least one sync was performed on the source appliance.
validate_last_sync_target.aksh	Validates that at least one sync was performed on the target appliance.
retrieve_replication_properties_source.aksh	Retrieves all replication properties from the source appliance.
sync_project_source.aksh	Performs a sync before storage-role reversal.
break_replication_source.aksh	Breaks replication before performing role reversal.
role_reverse_storage_target.aksh	Performs the actual reversal on the target appliance.

For the various options that get passed in the storage scripts and for their descriptions, refer to the "Creating Storage Scripts" section of the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide.*

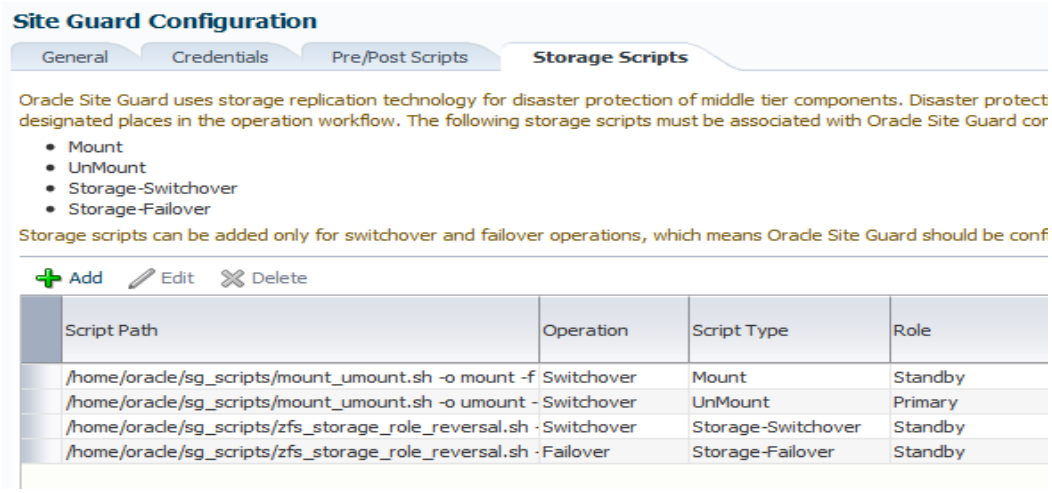

Mounting of Shares

The following script mounts all the NFS shares. It is executed from one node for all nodes in the setup at each site.

```
/home/oracle/sg_scripts/mount_umount.sh -o mount -f /u01/app/ora_base
```

Unmounting of Shares

The following script unmounts all the NFS shares. It is executed from one node for all nodes in the setup at each site.

```
/home/oracle/sg_scripts/mount_umount.sh -o umount -f /u01/app/ora_base
```

Custom Pre and Post Scripts for Oracle Site Guard Configuration

The following scripts, known as Pre and Post scripts, were used in the Oracle Site Guard configuration for OTD. These are the start and stop scripts for the OTD Administration Server, Administration Node, instances, and SNMP subagents. These scripts are located on each of the hosts and are executed by Oracle Site Guard by logging in to these hosts using the credential framework set in Oracle Enterprise Manager.

In the exercise for this white paper, user **root** on each host was used for starting and stopping the OTD components; however, any user other than **root** with proper privileges can also be used to install and start the OTD components. The credentials for the ZFS Storage Appliance set in the credential framework in Oracle Enterprise Manager are used for executing storage scripts.

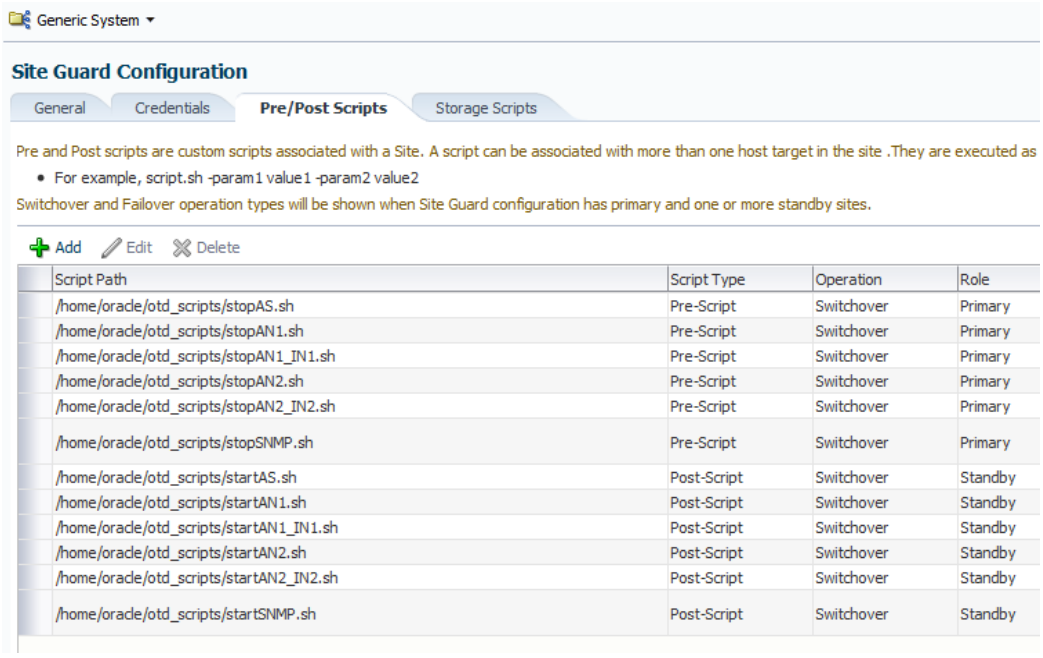

Script Path	Script Type	Operation	Role
/home/oracle/otd_scripts/stopAS.sh	Pre-Script	Switchover	Primary
/home/oracle/otd_scripts/stopAN1.sh	Pre-Script	Switchover	Primary
/home/oracle/otd_scripts/stopAN1_IN1.sh	Pre-Script	Switchover	Primary
/home/oracle/otd_scripts/stopAN2.sh	Pre-Script	Switchover	Primary
/home/oracle/otd_scripts/stopAN2_IN2.sh	Pre-Script	Switchover	Primary
/home/oracle/otd_scripts/stopSNMP.sh	Pre-Script	Switchover	Primary
/home/oracle/otd_scripts/startAS.sh	Post-Script	Switchover	Standby
/home/oracle/otd_scripts/startAN1.sh	Post-Script	Switchover	Standby
/home/oracle/otd_scripts/startAN1_IN1.sh	Post-Script	Switchover	Standby
/home/oracle/otd_scripts/startAN2.sh	Post-Script	Switchover	Standby
/home/oracle/otd_scripts/startAN2_IN2.sh	Post-Script	Switchover	Standby
/home/oracle/otd_scripts/startSNMP.sh	Post-Script	Switchover	Standby

startAS.sh

This script is used to start the Administration Server on hosts vm-otd-admin-a and vm-otd-admin-b at site OTD_A and site OTD_B, respectively.

```
# startAS.sh
/u01/otd_base/otd_instance/admin-server/bin/startserv
```

stopAS.sh

This script is used to stop the Administration Server on hosts vm-otd-admin-a and vm-otd-admin-b at site OTD_A and site OTD_B, respectively.

```
# stopAS.sh
/u01/otd_base/otd_instance/admin-server/bin/stopserv
```

startAN1.sh

This script is used to start the first Administration Node on hosts vm-otd-1a and vm-otd-1b at site OTD_A and site OTD_B, respectively.

```
# startAN1.sh
/u01/otd_base/otd_instance/admin-server/bin/startserv
```

stopAN1.sh

This script is used to stop the first Administration Node on hosts vm-otd-1a and vm-otd-1b at site OTD_A and site OTD_B, respectively.

```
# stopAS.sh
/u01/otd_base/otd_instance/admin-server/bin/stopserv
```

startAN2.sh

This script is used to start the second Administration Node on hosts vm-otd-2a and vm-otd-2b at site OTD_A and site OTD_B, respectively.

```
# startAN2.sh
/u01/otd_base/otd_instance/admin-server/bin/startserv
```

stopAN2.sh

This script is used to stop the second Administration Node on hosts vm-otd-2a and vm-otd-2b at site OTD_A and site OTD_B, respectively.

```
# stopAN2.sh
/u01/otd_base/otd_instance/admin-server/bin/stopserv
```

startAN1_IN1.sh

This script is used to start the OTD instance on the first Administration Node on hosts vm-otd-1a and vm-otd-1b at site OTD_A and site OTD_B, respectively.

```
# startAN1_IN1.sh
/u01/otd_base/otd_instance/net-elv-wcc-config/bin/startserv
```

stopAN1_IN1.sh

This script is used to stop the OTD instance on the first Administration Node on hosts vm-otd-1a and vm-otd-1b at site OTD_A and site OTD_B, respectively.

```
# stopAN1_IN1.sh
/u01/otd_base/otd_instance/net-elv-wcc-config/bin/stopserv
```

startAN2_IN2.sh

This script is used to start the OTD instance on the second Administration Node on hosts vm-otd-2a and vm-otd-2b at site OTD_A and site OTD_B, respectively.

```
# startAN2_IN2.sh
/u01/otd_base/otd_instance/net-elv-wcc-config/bin/startserv
```

stopAN2_IN2.sh

This script is used to stop the OTD instance on the second Administration Node on hosts vm-otd-2a and vm-otd-2b at site OTD_A and site OTD_B, respectively.

```
# stopAN2_IN2.sh
/u01/otd_base/otd_instance/net-elv-wcc-config/bin/stopserv
```

startSNMP.sh

This script is used to start the SNMP subagent on all the hosts at site OTD_A and site OTD_B. This is required for the Oracle Enterprise Manager Agent to monitor the Oracle Traffic Director configuration on each host.

```
# startSNMP.sh
/u01/otd_base/otd_home/bin/tadm start-snmp-subagent --instance-
home=/u01/otd_base/otd_instance
```

stopSNMP.sh

This script is used to stop the SNMP subagent on all the hosts at site OTD_A and site OTD_B.

```
# stopSNMP.sh
/u01/otd_base/otd_home/bin/tadm stop-snmp-subagent --instance-
home=/u01/otd_base/otd_instance
```

Accessing the Command-Line Interface of Oracle Traffic Director

The command-line interface (CLI) of Oracle Traffic Director can be accessed by running the tadm command from the ORACLE_HOME/bin directory, as follows:

```
./tadm [subcommand] --user=admin_user --host=adminserver_host [--
password-file=path_to_file] --port=adminserver_port
```

The CLI uses password-based authentication to allow access to the Administration Server. If the `--password-file` option is not specified, a prompt to enter the administrator user password is displayed. After the password is entered, the specified subcommand is executed.

The `tadm` command supports a comprehensive set of subcommands for creating, viewing, updating, and managing settings for all of the features of Oracle Traffic Director.

If the tadm command is used without specifying the subcommand, the shell mode of the CLI is obtained. In the shell mode, the options to connect to the Administration Server—`user`, `host`, `port`, `and password`—have already been specified, so individual subcommands can be executed without specifying the connection options each time.

Terminology

The following table provides terminology for Oracle Traffic Director.

Term	Description
Configuration	A collection of configurable elements (metadata) that determine the run-time behavior of an Oracle Traffic Director instance. A typical configuration contains definitions for the listeners (IP address and port combinations) on which Oracle Traffic Director should listen for requests and information about the servers in the back end to which the requests should be sent. Oracle Traffic Director reads the configuration when an Oracle Traffic Director instance starts and while processing client requests.
Instance	An Oracle Traffic Director server that is instantiated from a configuration and deployed on an Administration Node.
Failover group	Two Oracle Traffic Director instances grouped by a virtual IP address (VIP) to provide high availability in Active-Passive mode. Requests are received at the VIP and routed to the Oracle Traffic Director instance that is designated as the primary instance. If the primary instance is not reachable, requests are routed to the backup instance. For Active-Active failover, two failover groups are required, each with a unique VIP, but both consisting of the same nodes with the primary and backup roles reversed. Each instance in the failover group is designated as the primary instance for one VIP and the backup for the other VIP.
Administration Server	A specially configured Oracle Traffic Director instance that hosts the Administration Console and command-line interfaces, which you can use to create and manage Oracle Traffic Director configurations, deploy instances on Administration Nodes, and manage the lifecycle of these instances. Note that you can deploy instances of the Oracle Traffic Director configuration on the Administration Server. In this sense, the Administration Server can function as an Administration Node as well.
Administration Node	A specially configured Oracle Traffic Director instance that is registered with the remote Administration Server. The Administration Node running on a host acts as the

	agent of the remote Administration Server and assists the Administration Server in managing the instances running on the host. Note that on a given node you can deploy only one instance of a configuration.
INSTANCE_HOME	A directory of your choice, on the Administration Server or an Administration Node, in which the configuration data and binary files pertaining to Oracle Traffic Director instances are stored.
ORACLE_HOME	A directory of your choice in which you install the Oracle Traffic Director binaries.
Administration Console	A web-based graphical user interface, on the Administration Server, used for creating, deploying, and managing Oracle Traffic Director instances.
Client	Any agent—a browser or an application, for example—that sends HTTP, HTTPS, and TCP requests to Oracle Traffic Director instances.
Origin server	A server in the back end to which Oracle Traffic Director forwards the HTTP, HTTPS, and TCP requests that it receives from clients, and from which it receives responses to client requests. Origin servers can be application servers like Oracle WebLogic Server Managed Server instances, web servers, and so on.
Origin-server pool	A collection of origin servers that host the same application or service that you can load-balance by using Oracle Traffic Director. Oracle Traffic Director distributes client requests to servers in the origin-server pool based on the load-distribution method that is specified for the pool.
Virtual server	A virtual entity within an Oracle Traffic Director server instance that provides a unique IP address (or host name) and port combination through which Oracle Traffic Director can serve requests for one or more domains. An Oracle Traffic Director instance on a node can contain multiple virtual servers. Administrators can configure settings, such as the maximum number of incoming connections, specifically for each virtual server. They can also customize how each virtual server handles requests.

The following table provides terminology for disaster recovery.

Term	Description
Disaster recovery	The ability to safeguard against natural disasters or unplanned outages at a primary site by having a recovery strategy for failing over applications and data to a geographically separate standby site.
Topology	The primary site and standby site hardware and software components that compose an Oracle Fusion Middleware disaster recovery solution.
Site failover	The process of making the current standby site the new primary site after the primary site becomes unexpectedly unavailable (for example, due to unplanned downtime at the primary site).
Site switchover	The process of reversing the roles of the primary site and standby site. Switchovers are planned operations on the current primary site. During a switchover, the current standby site becomes the new primary site, and the current primary site becomes the new standby.
Site switchback	The process of reversing the roles of the new primary site (old standby) and new standby site (old primary). Switchback is applicable after a previous switchover.
Site instantiation	The process of creating a topology at the standby site (after verifying that the primary and standby sites are valid for Oracle Traffic Director 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 primary site at the standby site. For example, when a new application is deployed at the primary site, you should perform synchronization so that the same application will be deployed at the standby site.
Recovery Point Objective (RPO)	Maximum age of the data you want to be able 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 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.

The following table provides terminology for the Sun ZFS Storage Appliance.

Term	**Description**
Source	The site being replicated from, usually the primary site.
Target	The site being replicated to, usually the standby site. A target can receive one or more packages from one or more Sun ZFS Storage 7320 Appliances. In the Oracle Fusion Middleware infrastructure, the target site is the standby site.
Replica or package	The replicated copy of the project at the target site. It cannot be accessed directly. To be accessed, the replica has to be cloned, and the clone can be accessed for read and write operations
Snapshot	Point-in-time read-only copy of a 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	A process to access the replica at the target. A new project is created. All the shares, snapshots, clones, and so on are accessible under the cloned project.
Role reversal	The direction of the replication is reversed from source → target to target → source for a package.

References

1. Oracle Maximum Availability Architecture website
   http://www.oracle.com/technetwork/database/features/availability/maa-090890.html
2. *Oracle Fusion Middleware Disaster Recovery Guide*

   http://download.oracle.com/docs/cd/E14571_01/doc.1111/e15250/toc.htm
3. *Oracle Exalogic Enterprise Deployment Guide*

   http://download.oracle.com/docs/cd/E18476_01/doc.220/e18479/toc.htm
4. Oracle Exalogic Documentation Library

   http://download.oracle.com/docs/cd/E18476_01/index.htm
5. Oracle Traffic Director Documentation Library

   http://docs.oracle.com/cd/E23389_01/index.htm
6. *Oracle Fusion Middleware High Availability Guide*

   http://docs.oracle.com/cd/E14571_01/core.1111/e10106/toc.htm
7. Oracle Enterprise Manager Cloud Control Library

   http://docs.oracle.com/cd/E24628_01/index.htm
8. Using Oracle Site Guard (*Oracle Enterprise Manager Lifecycle Management Administrator's Guide*)

   http://docs.oracle.com/cd/E24628_01/em.121/e27046/site_guard.htm
9. Automating Disaster Recovery using Oracle Site Guard for Oracle Exalogic

   http://www.oracle.com/technetwork/database/availability/maa-site-guard-exalogic-exadata-1978799.pdf
10. Oracle Exalogic Backup and Recovery Best Practices

    http://www.oracle.com/technetwork/database/features/availability/maa-exalogic-br-1529241.pdf

Disaster Recovery Solution for
Oracle Traffic Director
September 2013

Author: Lingaraj Nayak
Contributors: Pradeep Bhat, Sriram Natarajan, Bonnie Vaughan

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

0109