

**Oracle Maximum**  
Availability Architecture

An Oracle White Paper  
June 11, 2014

**Oracle® Communications Billing and  
Revenue Management MAA on Oracle  
Engineered Systems**

---

1	Executive Overview .....	3
2	Audience.....	4
3	Introduction.....	5
4	Architecture Overview.....	6
4.1	Technical Architecture.....	7
4.2	Components Used .....	8
4.3	Topology Highlights.....	9
4.4	Site-Specific Deployment Topology.....	10
5	Implementation and Configuration Details.....	11
5.1	Prerequisites .....	11
5.2	Flowchart .....	12
5.3	System Resources Used.....	13
5.4	Detailed Procedures.....	16
6	Monitoring BRM Using Oracle Enterprise Manager .....	20
7	Disaster Recovery Operations Using Oracle Site Guard .....	24
7.1	Disaster Recovery Operations .....	27
7.2	Disaster Recovery Testing .....	27
8	Tests and Results .....	30
9	Benefits from Oracle Engineered Systems for BRM.....	35
10	Summary of BRM MAA Best Practices .....	36
10.1	Best Practices for BRM Database High Availability .....	36
10.2	Best Practices for BRM Application High Availability .....	36
10.3	Best Practices for Disaster Readiness and Recovery.....	37
11	Conclusion.....	38
A	Appendix.....	39
A.1	Creating the ZFS project and shares for BRM MAA .....	39
A.2	Set Up ZFS Remote Replication for BRM Apps .....	41
A.3	Set Up F5 Load Balancer for BRM CMMP Traffic .....	42
A.4	BRM Configuration Files and Scripts Used in BRM MAA ..	47
A.5	Oracle DBFS mounts for BRM Batch Pipeline.....	50
A.6	Disaster Recovery Host Aliasing .....	50
A.7	Oracle Database MAA .....	51
A.8	Benefits of Oracle Site Guard.....	53
B	References .....	56

## 1 Executive Overview

Oracle Maximum Availability Architecture (MAA) is the Oracle best-practices blueprint based on proven Oracle high-availability technologies and recommendations. The goal of MAA is to achieve the optimal high-availability architecture at the lowest cost and complexity. Papers are published at the [MAA home page](#) of the Oracle Technology Network (OTN) website.

Oracle Communications Billing and Revenue Management (BRM) Maximum Availability Architecture is a best-practices blueprint for achieving the optimal BRM high-availability deployment using Oracle high-availability technologies and recommendations. Oracle Communications BRM MAA was implemented on Oracle Exalogic and Oracle Exadata machines and tested to validate best practices and to measure downtime in various outage scenarios. This technical paper provides the following information:

- Oracle Communications BRM MAA architecture along with installation, configuration, and operational best practices
- Monitoring of Oracle Communications BRM MAA using Oracle Enterprise Manager Cloud Control
- Automation of Disaster Recovery operations for BRM using Oracle Site Guard, a component of Oracle Enterprise Manager Cloud Control
- Oracle MAA tests to validate best practices and to measure downtime in various outage scenarios
- Benefits of Oracle Communications BRM MAA on Oracle Exalogic, Oracle Exadata, and F5 Networks BIG-IP LTM Application Delivery Controller

MAA tests demonstrated minimal user impact during typical failure scenarios. In the event of a total site switchover or failure, the disaster recovery site in this MAA exercise could be brought online in as little as 20 minutes, with no data loaded for processing. The total time taken for the new primary site to be fully operational would ultimately depend on the size of the data required to be loaded for processing by the pipeline manager.

## 2 Audience

This document is intended for Oracle Communications Billing and Revenue Management architects and administrators. The reader should be familiar with Oracle Billing and Revenue Management applications, Oracle Exalogic and Oracle Exadata, ZFS storage replication techniques, F5 Networks BIG-IP LTM and GTM, Oracle Enterprise Manager, and Oracle Site Guard. For additional details, refer to the documents listed in the [“References”](#) section.

### 3 Introduction

[Oracle Communications BRM](#) provides end-to-end revenue management solutions for the global communications, media, cloud, and digital services markets, supporting key business processes that incorporate generation, capture, collection, and assurance of revenue. Oracle Communications BRM MAA (version 7.5) was used for the MAA solution described in this paper, and the core functionalities tested were Billing (using Real-Time Pipeline), and Rating (using Batch Pipeline).

This technical paper is organized into the following sections:

- **Architecture Overview** – High-level description of the architecture and key technology components used in the MAA exercise
- **Implementation and Configuration Details** – Flowcharts and procedures to describe the architecture and system resources used in the MAA exercise
- **Monitoring BRM Components Using Oracle EM Enterprise Manager Cloud Control** – Steps to configure the Application Management Pack for Oracle Communications to monitor all BRM targets
- **Disaster Recovery Operations Using Oracle Site Guard** – Configuration of Oracle Site Guard to orchestrate switchover and failover between two BRM sites
- **Tests and Results** – Outline of various tests and results from the MAA exercise
- **Benefits from Oracle Engineered Systems for BRM** – Highlights of the advantages for using Oracle Communications BRM MAA on Exalogic and Exadata
- **Summary of BRM MAA Best Practices** – Checklist of the best practices used
- **Appendix** – Various scripts, screenshots, F5 BIG-IP example configuration and iRule details and detailed explanations of best practices
- **References** – Summary of the external documents referenced by this paper

This paper does not describe any backup or recovery procedures for the Exalogic compute nodes or vServers, network switches, or storage appliances used in the Oracle Communications BRM MAA setup. For such information, refer to Backup and Recovery Best Practices available in the [Oracle Exalogic Elastic Cloud Backup and Recovery Guide](#).

## 4 Architecture Overview

Oracle Communications BRM MAA is a high-availability architecture, providing local high availability layered on top of the MAA best practices for Oracle Database and the high-availability features of BRM applications. The overall MAA architecture includes a secondary site to provide business continuity in the event of a planned shutdown or unplanned failure of the primary site.

Each site in the topology is built using Oracle Engineered Systems (Oracle Exalogic and Oracle Exadata). F5 Networks Local Traffic Manager (LTM) and Global Traffic Manager (GTM) are used for local load balancing and global traffic-routing requirements in the MAA topology. For this paper, Oracle Enterprise Manager Cloud Control 12c is deployed on system resources independent from both of the BRM sites described in the BRM MAA topology.




Figure 1: Oracle Communications BRM MAA with Exadata and Exalogic Disaster Recovery Topology

### 4.1 Technical Architecture

The following figure lays out the technical architecture of the setup used in the BRM MAA exercise.

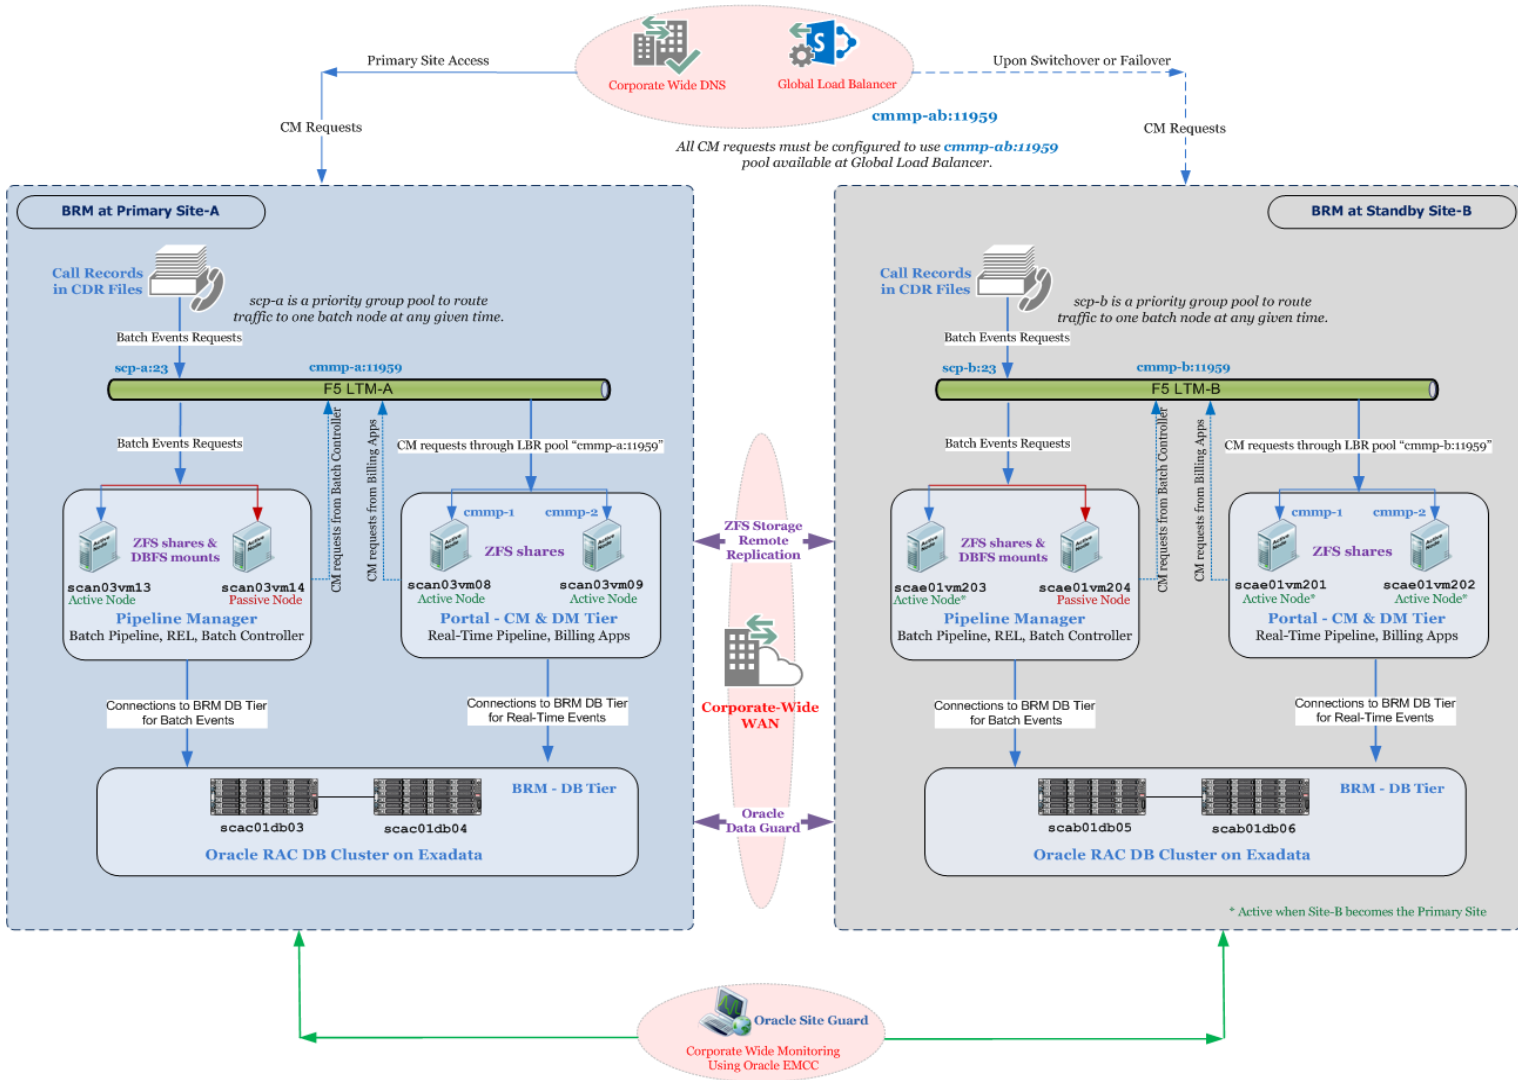


Figure 2: Oracle Communications BRM MAA – Technical Architecture

## 4.2 Components Used

The following hardware and software components were used and configured in the overall BRM MAA topology.

### 4.2.1 Hardware

- Oracle Exadata Database Machine X3-2 – Four RAC nodes, two at each site
- Oracle Exalogic X3-2 – Eight Echo-based vServers from Exalogic VDC, four at each site;
- Oracle ZFS Storage Appliance 7320 – Two sets of storage appliances, included in the Exalogic rack, configured in active-passive mode at each site
- Two sets of F5 Network Load Balancers (LTM and GTM)
- Two nodes for Oracle Enterprise Manager Cloud Control deployment

### 4.2.2 Software and Applications

- Oracle Database 11g Release 2 (11.2.0.3)
- Oracle Communications Billing and Revenue Management (BRM) version 7.5 with the following components
  - Portal Base, which includes Connection Manager (**cm**), Connection Manager Master Process (**cmmp**), Data Manager for Oracle (**dm\_oracle**), Batch Controller, Billing, and Invoice applications.
  - BRM Pipeline, which includes Real-Time Pipeline and Batch Pipeline
  - Oracle BRM Account Sync Tool with EAI Java Server (**eai\_js**)
  - Rated Event Loader (REL)

*Note: While for the MAA exercise, BRM 7.5 was used, the MAA guidelines mentioned in this paper may be used for previous versions of BRM with similar CM and DM architecture as of BRM 7.5.*

- Oracle Enterprise Manager Cloud Control version 12.1.3 with Fusion Middleware Plug-in (12.1.0.5) for Oracle Site Guard
- Application Management Pack for Oracle Communications version 12.1.0.1

*All of the nodes used in the exercise are based on Oracle Enterprise Linux 5*



### 4.3 Topology Highlights

- The overall MAA topology comprises two symmetric sites named **BRM-Site-A** and **BRM-Site-B**. The current production site (BRM-Site-A) is running and active, while the second site (BRM-Site-B) is serving as a standby site and is in passive mode.
- Hosts on each site are Exalogic **vServers** with mount points defined for accessing the shared storage system for each site, over the **IPoIB-vshared-storage** network.
- On both sites the BRM components are deployed on the site's shared storage system. This involves creating all of the operating system user (**pin**) home directories, and configuration data for BRM components, in volumes on the production site's shared storage. "Oracle Database File System (DBFS)" mounts are used for certain BRM batch pipeline directories.
- Oracle ZFS storage replication (**continuous**) is used to copy the configuration data project and shares from the production site's shared storage to the standby site's shared storage.
- **Oracle Data Guard** is used to replicate the BRM database.
- BRM application requests are initially routed to the production site through the F5 GTM and LTM. All BRM CM requests are done through the virtual server set at the F5 GTM, which in turn has the virtual servers created on F5 LTM at each site.
- Oracle Enterprise Manager Cloud Control is used to monitor the availability of the BRM applications and database.
- In the event of any failure or planned outage of the production site, Oracle Site Guard is used to initiate a switchover or failover operation to enable the standby site to assume the production role in the topology.
- F5 LTM is monitoring the health status of the BRM servers in each site. F5 GTM is monitoring the status of all the LTM VIPs, in both Sites. During a failure of the Active site, LTM and GTM monitors detect the failure very quickly. When the Standby LTM sees that the application switchover is successful, it notifies GTM that the Standby site VIP is UP and new BRM application requests are automatically rerouted to the Standby site, with no user or DNS configuration changes needed. At this point, the Standby site has assumed the production role.
- Once BRM-Site-B assumes the production role, ZFS storage replication gets set up in the reverse direction, that is, from storage appliance at BRM-Site-B to storage appliance at BRM-Site-A; Oracle Data Guard also reverses the database role when the BRM database at BRM-Site-A obtains the primary role and the database at BRM-Site-B assumes the role of the new physical standby.

#### 4.4 Site-Specific Deployment Topology

Depending upon the production role assumed by a particular site, the following topology is applicable to that site; Figure 3 shows an example of BRM-Site-A operating as the primary production site.

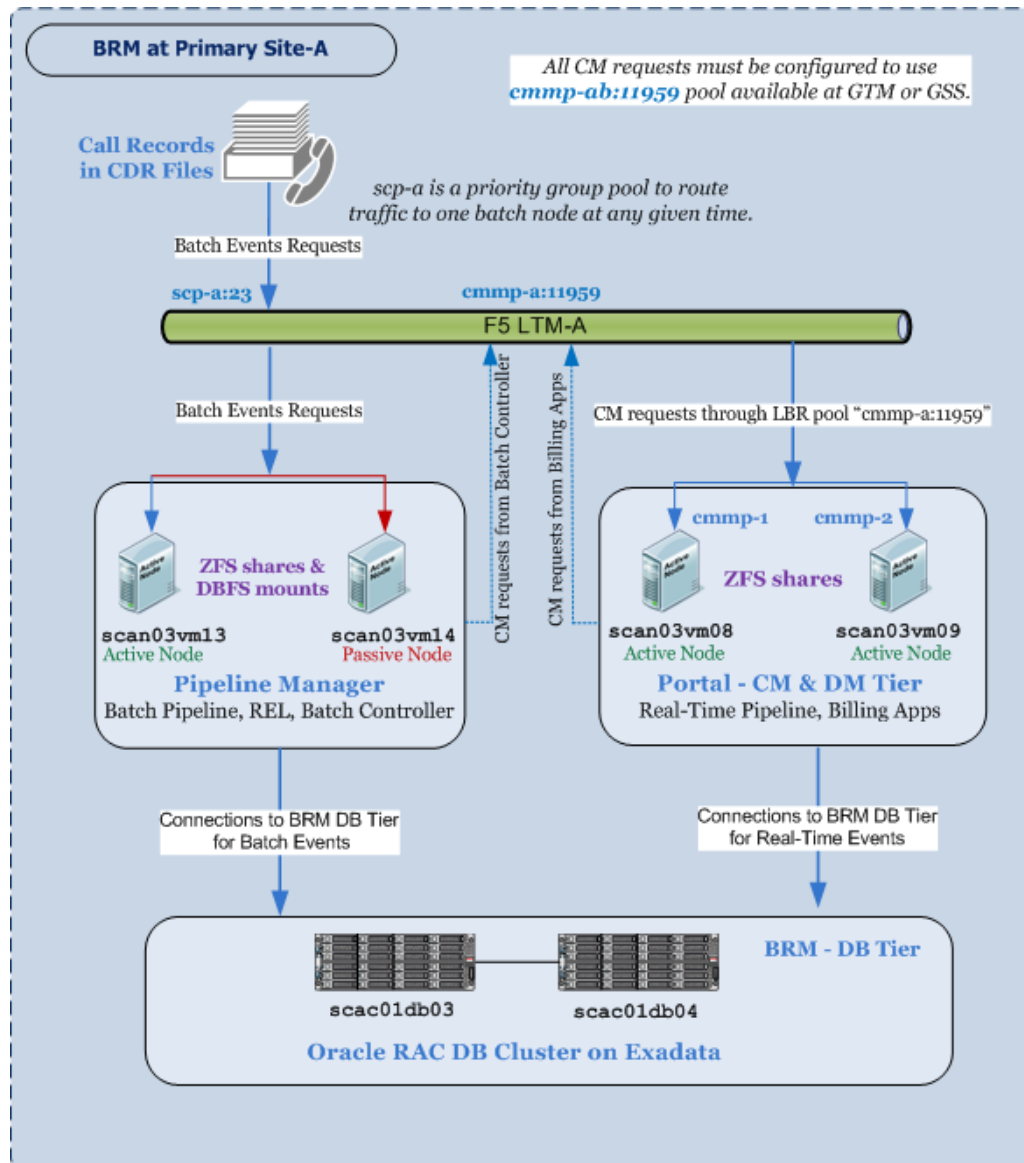


Figure 3: Oracle Communications BRM – Technical Architecture at a Single Site

## 5 Implementation and Configuration Details

As highlighted in the architecture description, the MAA exercise was carried out on Oracle Exadata and Oracle Exalogic. Installation and configuration of Oracle Exadata and Oracle Exalogic is not covered in this white paper. The steps to create an Exalogic virtual Data Center (vDC) and virtual servers, including networks attached to the vServers, are in the [Oracle Exalogic Elastic Cloud Administrator's Guide](#).

### 5.1 Prerequisites

- Create the vServers at each site. For the MAA exercise, Echo-Based vServers were created with Oracle Enterprise Linux 5 as their operating system.
- Create the ZFS project and shares at the primary site ZFS storage appliance. For an example, refer to the Appendix.
- Create an Oracle RAC database at BRM-Site-A, and create a physical standby database at BRM-Site-B; configure Oracle Data Guard between the primary and standby BRM databases. Refer to the Oracle Database MAA section in the Appendix for details.
- Ensure that the load balancers (F5 LTM and GTM) are accessible from each site and that the required virtual server names are registered with the corporate DNS.
- Install Oracle Enterprise Manager Cloud Control on systems independent of BRM Sites system resources.

### 5.2 Flowchart

The following flowchart describes the sequence of major steps for BRM MAA setup.

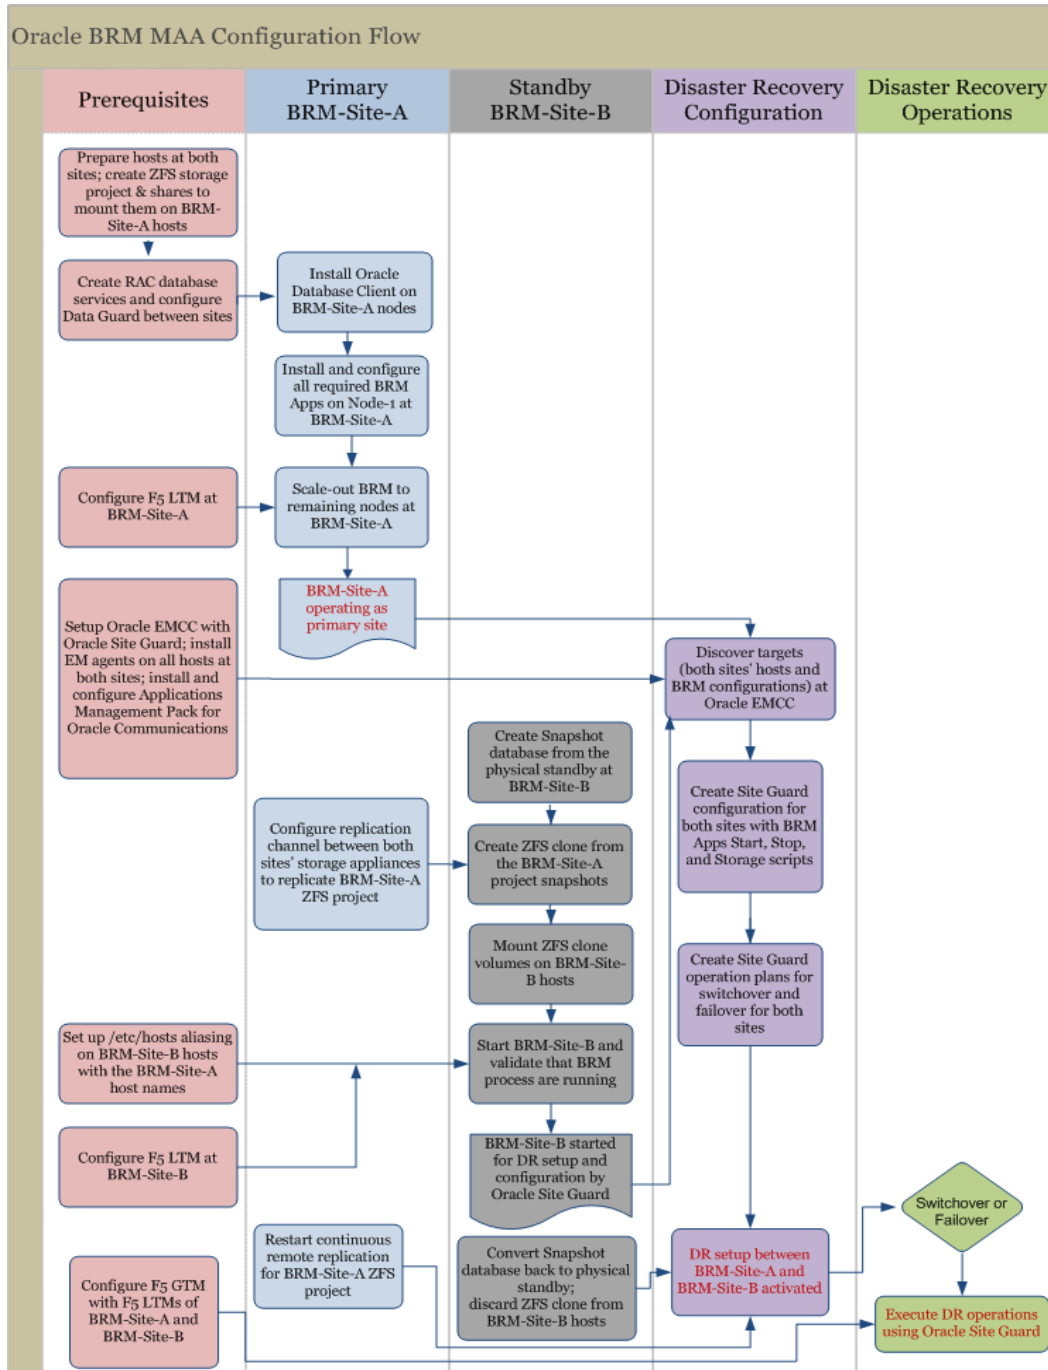


Figure 4: Oracle Communicaitons BRM MAA Configuration Flowchart

### 5.3 System Resources Used

The following systems were used to test the deployment in this MAA paper.

#### Primary Site [BRM-Site-A]

Host Name	EoIB IP	IPoIB IP	Purpose
scan03vm0008	10.133.235.20	192.158.1.40	Node 1 for Portal Base and RTP
scan03vm0009	10.133.235.21	192.158.1.41	Node 2 for Portal Base and RTP
scan03vm0013	10.133.235.22	192.158.1.42	Node 1 for Batch Pipeline, Batch Controller, and REL
scan03vm0014	10.133.235.23	192.158.1.43	Node 2 for Batch Pipeline, Batch Controller, and REL

**Table 1: BRM Apps Hosts at BRM-Site-A**

Host Name	Purpose
scab01db05	RAC Node 1
scab01db06	RAC Node 2
scab0105-vip, scab0106-vip	EoIB VIPs (used in the TNS connect string for BRM Apps)
scac01-scan3	SCAN Name ( <i>not</i> used for BRM Apps DB Connection)

**Table 2: BRM Database Hosts at BRM-Site-A**

Host Name	Net0(IGB0) IP	IPoIB-vserver-shared-storage IP *	Comments
scan03sn01	10.133.41.80	172.47.1.1	Active storage head
scan03sn02	10.133.41.81		Passive storage head
scan03-fe	10.133.10.190		Virtual IP for storage appliances over 1 GbE for remote replication

\*Used for mounting the ZFS shares on the vServers

**Table 3: ZFS Storage Appliances Details at BRM-Site-A**

## Standby Site [BRM-Site-B]

Host Name	EoIB IP	IPoIB IP	Purpose
scae01vm201	10.143.245.1	192.168.1.1	Node 1 for Portal Base and RTP
scae01vm202	10.143.245.2	192.168.1.2	Node 2 for Portal Base and RTP
scae01vm203	10.143.245.3	192.168.1.3	Node 1 for Batch Pipeline, Batch Controller, and REL
scae01vm204	10.143.245.4	192.168.1.4	Node 2 for Batch Pipeline, Batch Controller, and REL

Table 4: BRM Apps Hosts at BRM-Site-B

Host Name	Purpose
scab01db05	RAC Node 1
scab01db06	RAC Node 2
scab0105-vip, scab0106-vip	EoIB VIPs (used in the TNS connect string for BRM Apps)
scac01-scan3	SCAN name ( <i>not</i> used for BRM Apps database connection)

Table 5: BRM Database Hosts at BRM-Site-B

Host Name	Net0(IGB0) IP	IPoIB-vserver- shared-storage IP *	Comments
scae01sn01	10.143.21.60	172.27.2.1	Active storage head
scae01sn02	10.143.21.61		Passive storage head
scae01-fe	10.143.11.170		Virtual IP for storage appliances over 1 GbE for remote replication

\*Used for mounting the ZFS shares on the vServers

Table 6: ZFS Storage Appliances Details at BRM-Site-B

## ZFS Storage Project and Shares

The following table lists the ZFS project and shares available at each site when it assumes the primary role. During the storage replication, the project and shares are available as the **Replica** project and shares at the standby storage appliance. Separate shares for Log files may be created and replicated continuously and any configuration related data may be on shares with scheduled replication.

Project	Shares	Mount Points	Mounted on BRM-Site-A Hosts as Primary	Mounted on BRM-Site-B Hosts as Primary
BRMDR	/export/brm_maa/brm1_rt1	/brm1	scan03vm0008	scae01vm201
	/export/brm_maa/brm1_rt2	/brm1	scan03vm0009	scae01vm202
	/export/brm_maa/brm1_ba1	/brm1	scan03vm0013	scae01vm203
	/export/brm_maa/brm1_ba2	/brm1	scan03vm0014	scae01vm204
	/export/brm_maa/stage_all	/stage	All hosts	All hosts

Table 7: ZFS Storage Project and Shares

## BRM Database Tier Details

The following database details are used in the BRM MAA exercise.

Database Name	Database Role	Purpose
BRMDBP	Primary	Production database ( BRM-Site-A)
BRMDBS	Physical Standby	Standby database ( BRM-Site-B)
BRM DB Service and DB user		
brmdb.us.oracle.com		Database service used by BRM Apps ( FAN Enabled)
pin1		BRM database user

Table 8: BRM Database Tier Details

## F5 Load-Balancer Details

The following table lists the load-balancer details used for Connection Manager (CM) traffic.

Virtual Server Name	Port	Purpose
cmmp-a	11959	On F5 LTM at BRM-Site-A
cmmp-b	11959	On F5 LTM at BRM-Site-B
cmmp-ab	11959	On F5 GTM monitoring both sites
All CM requests from any BRM component uses <b>cmmp-ab:15959</b> in the configuration files.		
Similar to <b>cmmp</b> Virtual Servers, configure <b>scp-ab:23</b> for secure transfer of CDR files to the Batch Pipeline Servers through the load balancer.		

**Table 9: F5 Networks Load-Balancer Details**

### 5.4 Detailed Procedures

As outlined in the flowchart, the procedures followed for the BRM MAA setup are primary site setup, load-balancer configuration, and standby site setup and instantiation.

#### 5.4.1 Primary Site (BRM-Site-A) Setup

1. Ensure that the prerequisites described in Section 4.1 are completed.
2. Mount the shared-storage volumes to their respective nodes as described in Table 7 in Section 4.3.
3. Create the database user **pin1** in the BRM database. For the database user creation script, refer to the Appendix.
4. Install an Oracle Database client on shared storage accessible to all nodes to establish *ORACLE\_HOME*. The database client uses *ORACLE\_HOME* for BRM configuration in the subsequent steps and to start and stop BRM Applications. The mount point used for *ORACLE\_HOME* in the MAA exercise is **/stage**. Note: For Batch Controller, 32 bit Oracle Database Client is used; for Data Manager, 64 bit Oracle Database Client is used.



5. Install and configure all the required BRM components (including the third-party software) from the BRM binaries at Node 1 of BRM-Site-A (scan03vm0008). Before running the `pin_setup` configuration at Node 1 (scan03vm0008), ensure that the **pin\_step.values** file has the required values, including the BRM database details prepared in the prerequisites steps. With the **pin\_setup -all** execution step at Node 1, the BRM schema in BRM database and the BRM configuration files on the Node 1 Apps tier get updated with the required configuration. The BRM components used in the MAA exercise are listed in [section 3.2.2](#). Also refer to the [Oracle Communications Billing and Revenue Management Installation Guide](#) for details about the installation and configuration of BRM. The mount point used here for the BRM Apps installation in the MAA exercise is `/brm1`.
6. Verify the installation by validating the contents of the `PIN_HOME` directory and the values in all the configuration files, such as **pin.conf** and **Infranet.properties** under `$PIN_HOME/sys/*/` and `$PIN_HOME/apps/*/`. This confirms that BRM Apps Node 1 is configured and the components are ready to be started.
7. For scaling out the BRM Apps to multiple nodes for a high-availability configuration, either do a fresh installation (without the 'pin\_setup' configuration) on separate nodes or copy the entire installation directory from Node 1 to the multiple nodes at the respective mount points on the newer nodes. In the MAA exercise, for the BRM Apps directories, the same mount point, `/brm1`, is used. If the copy mode is adopted, replace the host name in each configuration file under the BRM installation directory with the respective node. For example, for Node 2 in the MAA exercise, replace the host name with **scan03vm0009**. Ensure that the host name is replaced in all the configuration files, such as **pin.conf** and **Infranet.properties**, under `$PIN_HOME/sys/*/` and `$PIN_HOME/apps/*/`.
8. Follow the previous step for all of the remaining nodes in the topology. With this step installation and configuration of the BRM Apps on all of the nodes in the architecture are considered complete. Although each node consists of all BRM components, each component described in the topology is started on its respective node. At this step, all of the applications at the primary site (BRM-Site-A) can be started to validate the installation and configuration; however, these applications on all nodes need to be restarted after the F5 GTM configuration

and subsequent update of all the BRM configuration files with the GTM virtual server address.

#### 5.4.2 Load-Balancer Configuration (F5 LTM and GTM)

For load balancing the Connection Manager (CM) traffic at each site, configure the F5 Networks LTM for BRM-Site-A and BRM-Site-B for the designated Connection Manager Master Process (CMMP) nodes according to the details in Table 3. Also refer to the detailed steps in [Appendix A. Section A.3](#) for the LBR configuration for CM traffic.

CMMP provides built-in load-balancing features for the BRM Apps, but it can become a single point of failure if the node where it runs crashes. Extending the load-balancing functionality for CMMP to an external load balancer at each site enhances the high availability of the CM tier for that BRM site. Refer to “Adminstrating a High-Availability BRM System” in the *Oracle Communications Billing and Revenue Management System Administrator's Guide* for a detailed explanation of BRM HA configuration. Any cluster-ware solution may be used for configuring high-availability of Batch Pipeline Servers; however this was not configured in the MAA exercise.

1. Configure virtual server '**cmmp-a:11959**' on LTM at BRM-Site-A for routing CM traffic between the nodes `scan03vm0008` and `scan03vm0009`.
2. Although the standby site is not initiated at this stage, the virtual server '**cmmp-b:11959**' can be configured on LTM at BRM-Site-B for routing CM traffic between the nodes `scae01vm201` and `scae01vm202`.
3. Configure the GTM with virtual server '**cmmp-ab: 11959**' to reroute CM traffic to the BRM site that assumes the primary role. Configuration of F5 GTM is not covered in this white paper.
4. Once **cmmp-ab: 11959** is configured at the GTM, all the BRM configuration files that connect to CM at primary site (BRM-Site-A) are updated with the virtual server address '**cmmp-ab: 11959**'. Refer to [Section A.4.3](#) and [Section A.4.4](#) in Appendix A for sample configuration files.

### 5.4.3 Standby Site (BRM-Site-B) Setup and Instantiation

The standby site can be deployed and partially validated without affecting operations at the primary site. No installation or configuration for BRM components is required at the standby site. When the primary-site storage is replicated to the standby-site storage, the BRM product binaries and configuration installed and configured on the primary site will be replicated at the standby site. However, software installation and configuration is required for the database on the Oracle Exadata Database Machine at the standby site. As mentioned in the prerequisites section, validate that the Data Guard setup is activated between the primary site and standby site for the BRM database. Also ensure that the user (**pin**) home directories are available on each standby node. The following steps are performed to set the standby site (BRM-Site-B).

1. On the standby site, ensure that aliased host names are created that are the same as the physical host and virtual server names used for the peer hosts at the production site. All of the host names of the primary site's applications tier, including all of the virtual server names, are aliased and resolved to the standby site IPs. This can be done at the `/etc/hosts/` directory of the standby site server or by maintaining the names in the standby site DNS if it is separate from the primary site DNS.
2. Perform a manual ZFS replication of the BRM project to create the initial ZFS storage snapshot of the product binaries and configuration on the standby site. The manual update, which can be done through the GUI or the CLI, must be initiated from the ZFS storage appliance on the primary site. On the standby site, validate that the packages have been received by viewing the Projects Replica list. This will list all of the packages that are being received or have been received from the primary site. These newly created clones will be used as local projects to instantiate the standby site while leaving the ongoing replication of the ZFS projects intact. Refer to [Appendix A, Section A.1](#) for an example of creating a ZFS clone from a project snapshot.
3. On the standby site database tier, convert the physical standby database to a snapshot standby database, and bring the standby database online using the following SQL commands:

- a. `SQL> ALTER DATABASE CONVERT TO SNAPSHOT STANDBY;`
- b. `SQL> ALTER DATABASE OPEN;`

4. On the standby site's BRM Apps nodes, mount the shares from the cloned ZFS projects. Use mount points and attributes identical to the ones that were used at the primary site.
5. Start and validate the BRM-Site-B components. At this stage all of the BRM components at BRM-Site-A and BRM-Site-B are started independently. This is not the final state of the DR setup because both sites are not in sync. After the Oracle Enterprise Manager Cloud Control targets and Oracle Site Guard configuration are done, both sites will be brought back in sync with BRM-Site-A, assuming it has the primary role.

## 6 Monitoring BRM Using Oracle Enterprise Manager

Oracle Enterprise Manager Cloud Control is an Oracle integrated enterprise information technology (IT) management platform. It includes the management infrastructure for using the Oracle Site Guard plug-in, which provides disaster recovery services. Because Cloud Control is critical for managing disaster recovery operations, Oracle strongly recommends the following practices for the Oracle Enterprise Manager Cloud Control deployment:

- Ensure that Cloud Control is not deployed on the Exalogic machines at either the primary or standby site. Doing so would make Cloud Control vulnerable to the outages that it is intended to detect and protect against.
- Implement high availability and disaster recovery plans for protecting the Cloud Control deployment.

Oracle Application Management Pack for Oracle Communications consists of an Oracle Enterprise Manager Cloud Control plug-in, which provides application lifecycle services and management capabilities for Oracle Communications applications, including the BRM application suite. The management pack leverages a single solution that is based on Oracle Enterprise Manager Grid Control and extends it by providing a single view and console to manage the BRM application suite. The Application Management Pack for BRM runs on top of the Oracle Enterprise Manager framework and extends Grid Control to add support for BRM system administration.

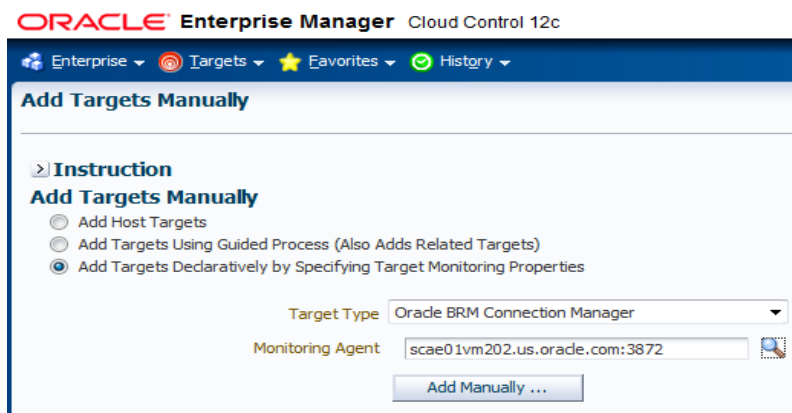
In the BRM MAA exercise, Oracle Application Management Pack for Oracle Communications (version 12.1.0.1) is used. For the BRM components to be monitored from Cloud Control, all of those components need to be managed by the **pin\_ctl** command.

Use the following process flow to implement and use the Application Management Pack for Oracle Communications:

1. Install Oracle Enterprise Manager Cloud Control and the Application Management Pack for Oracle Communications plug-in.
2. Configure Oracle Communications application targets, setting preferred credentials and deploying the Oracle Management Agent to each target.
3. Deploy the Application Management Pack for Oracle Communications agent to each Cloud Control server and target host.
4. Manage Oracle Communications applications with Cloud Control. Managing includes installing, configuring, updating, and monitoring.

For details, refer to the [Oracle Application Management Pack for Oracle Communications System Administrator's Guide](#). An example of creating a target for the BRM target (Connection Manager) in Cloud Control follows.

- a) Add the Oracle Communications targets in Cloud Control by the declarative method of adding targets manually. In the following example, "**Oracle BRM Connection Manager**" is selected in the **Target Type** field.



- b) Enter the BRM component properties and values, and submit them to add the target in Cloud Control, as the following screenshot shows. All of the values are in the `pin.conf` file in the `BRM_HOME` directory.

The screenshot shows the 'Add Oracle BRM Connection Manager' form in Oracle Enterprise Manager Cloud Control 12c. The form is titled 'Add Oracle BRM Connection Manager' and includes the following fields and values:

- Target Name:** CM2-1-Site-B
- Target Type:** Oracle BRM Connection Manager
- Agent:** https://scae01vm202.us.oracle.com:3872/emd/main/
- BRM Credentials:**
  - Credential type:** BRMCredentials
  - User Name:** pin
  - Password:** [Redacted]
  - Confirm Password:** [Redacted]
- Properties:**
  - BRM Component Type:** cm
  - BRM Home Path:** /export/pin/7.5
  - BRM Stage Location:** [Empty]
  - BRM System Type (Prepaid or Postpaid):** Prepaid
  - Classes To Be Partitioned:** [Empty]
  - Create Database Tables (YES/NO):** [Empty]
  - DB Alias:** [Empty]
  - DB Password:** [Empty]
  - DB User:** [Empty]

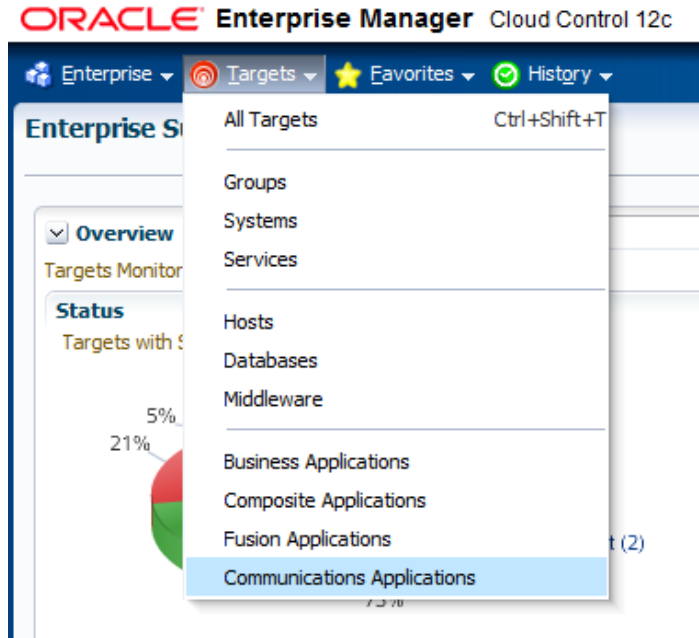
- c) Once the target is added, the BRM component can be monitored, started, or stopped from Cloud Control, as the following screenshot shows.

The screenshot shows the 'Summary' page for the 'CM2-1-Site-B' target in Oracle Enterprise Manager Cloud Control 12c. The page displays the following summary information:

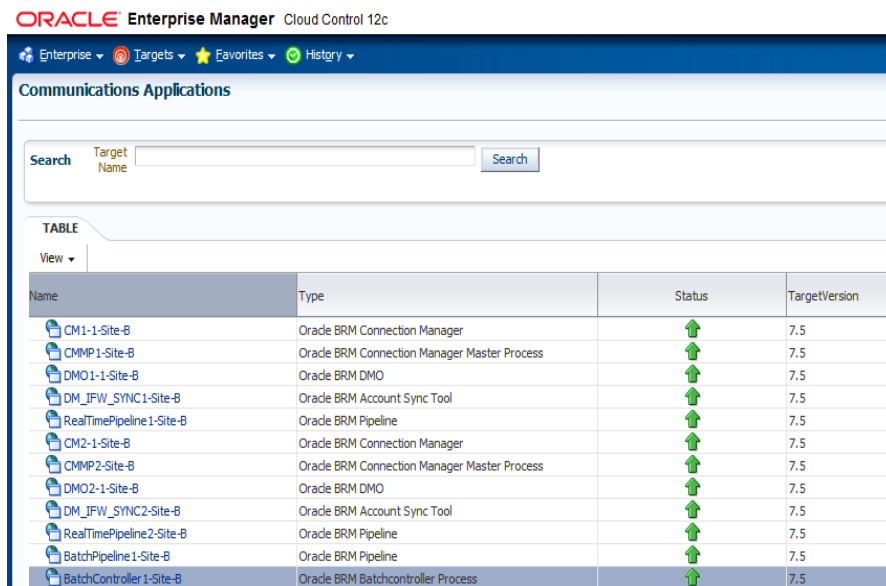
- Instance Name:** CM2-1-Site-B
- BRM Instance Type:** Batch Process
- BRM Home:** /export/pin/7.5
- Component Type:** Oracle BRM Connection Manager
- BRM Log Directory:** /export/pin/7.5/var/cm/
- Port:** 11960
- Version:** 7.5
- Patch Set Level:** [Empty]

At the bottom of the summary, there are two buttons: 'Stop' and 'ReStart'.

- d) To view all Oracle Communications targets, choose **Communications Applications** from the **Targets** menu in Cloud Control.



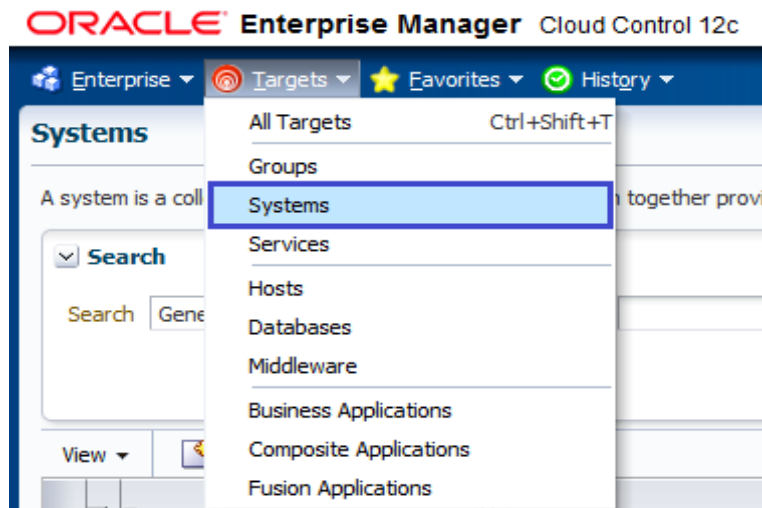
- e) Ensure that all of the BRM component targets of BRM-Site-A and BRM-Site-B are discovered and monitored, as shown in the following screenshot. These targets will be added to the Generic Systems for each site during the Oracle Site Guard configuration, as the next section describes.



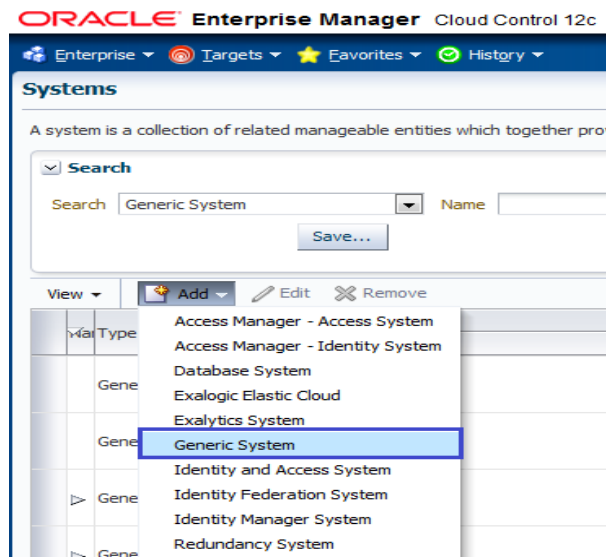
## 7 Disaster Recovery Operations Using Oracle Site Guard

Oracle Site Guard is used to automate the disaster recovery operations of BRM MAA. Each site is represented as a '**Systems**' type of target in Oracle Enterprise Manager Cloud Control. In this exercise the sites are named BRM-Site-A and BRM-Site-B. The following screenshots show the generic system-creation step.

- a) Log in to Cloud Control. Choose **Systems** from the **Targets** menu.



- b) When the System page opens, choose **Generic System** from the **Add** menu.






- c) Ensure that in the **Members** section, the correct targets of BRM components are added for the intended site.

Name	Type	Status	Key Members
CM1-1-Site-B	Oracle BRM Connection Manage	↑	✓
CM2-1-Site-B	Oracle BRM Connection Manage	↑	✓
CMMP1-Site-B	Oracle BRM Connection Manage	↑	✓
CMMP2-Site-B	Oracle BRM Connection Manage	↑	✓
dbmsg.us.oracle.com	Cluster Database	↑	
dbmsg.us.oracle.com_dbmsg1	Database Instance	↑	
dbms.us.oracle.com_dbms2	Database Instance	↑	

- d) 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.

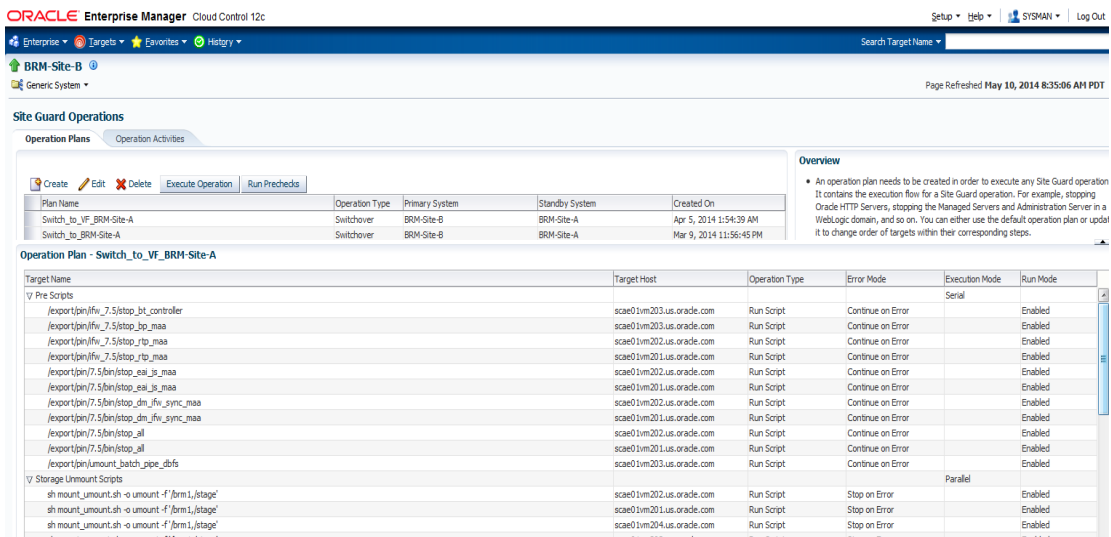
Name	Privilege Propagation	Type	Status	Members
BRM-Site-A	✓	Generic System	↑	Host(6), Listener(5), Oracle BRM Pipeline(4), ... More
BRM-Site-B	✓	Generic System	↓	Host(6), Listener(5), Oracle BRM Pipeline(4), ... More

- e) Create and associate credentials for both sites. Create and associate Pre and Post scripts and storage scripts for the sites. Using Cloud Control, 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 [section A.4.6](#) in Appendix A.



Refer to "[Configuring Oracle Site Guard Operations for Disaster Recovery](#)" in the *Oracle Enterprise Manager Lifecycle Management Administrator's Guide* for instructions about 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.](#)"

A screenshot of an Operation Plan in Oracle Site Guard follows.



- f) After completing the Oracle Site Guard configuration for both sites for BRM deployment, stop the BRM-Site-B site completely and discard the ZFS clone volume **BRM-DR-Clone**. At this stage, ensure that the storage replication for the project BRM\_MAA 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 BRM-Site-A and BRM-Site-B sites is now complete and activated.

## 7.1 Disaster Recovery Operations

The following operations were validated in the MAA exercise.

SITE GUARD OPERATION	DESCRIPTION
Switchover-to-BRM-Site-B	Switch over from primary to standby site
Switchback-to-BRM-Site-A	Switch back to primary site from standby site
Failover-to-BRM-Site-B	Fail over from primary to standby site
Fallback-to-BRM-Site-A	Fail over 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, as follows:

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

## 7.2 Disaster Recovery Testing

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

1. From the **Targets** menu, choose **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, choose **Site Guard** and then **Operations**. The Site Guard Operations page is displayed.
4. Click the **Operation Activities** tab. A table listing all executions of the submitted operation plan is displayed.

### Site Guard Operations

Operation Plans    **Operation Activities**

This table shows list of all submitted operation plan executions. You can see details of each of these activities by clicking on the activity

Activity Name	Plan Name	Primary System	Standby System	Operation Type	Status
SwitchoverSite 1394439131468	Switch_to_BRM-Site-B	BRM-Site-A	BRM-Site-B	Switchover	Succeeded

To monitor detailed procedure steps for an operation-plan activity and their status, click the activity name.

**ORACLE Enterprise Manager** Cloud Control 12c

Enterprise ▾ Targets ▾ Favorites ▾ History ▾

### Provisioning

**Procedure Activity: SwitchoverSite 1394439131468**

▶ Elapsed Time: 11 minutes, 32 seconds

**Procedure Steps**

View ▾    Show All Steps ▾

Select	Name	Type	Status
<input type="checkbox"/>	Run PreChecks	Procedure St	
<input type="checkbox"/>	▶ Run Primary PreScripts	Procedure St	
<input type="checkbox"/>	▶ Stop Primary Site	Procedure St	
<input type="checkbox"/>	▶ Run Primary PostScripts	Procedure St	
<input type="checkbox"/>	▶ Unmount Primary Filesystems	Procedure St	
<input type="checkbox"/>	▶ Switchover Storage	Procedure St	
<input type="checkbox"/>	▶ Mount Standby Filesystems	Procedure St	
<input type="checkbox"/>	▶ Switchover Database	Procedure St	
<input type="checkbox"/>	▶ Run Standby PreScripts	Procedure St	
<input type="checkbox"/>	▶ Start Standby Site	Procedure St	
<input type="checkbox"/>	▶ Run Standby PostScripts	Procedure St	
<input type="checkbox"/>	Update SiteGuard Schema	Computation	

After a successful switchover, BRM-Site-B assumes the role of new primary site.

**ORACLE Enterprise Manager** Cloud Control 12c

Enterprise Targets Favorites History

### Systems

A system is a collection of related manageable entities which together provide one or more business functions. Members of any system can have well-defined relationships amongst

Search Generic System Name Advanced Search Save...

View Add Edit Remove

Name	Privilege Propagation	Type	Status	Members
BRM-Site-A	✓	Generic System	↓	Host(6), Listener(5), Oracle BRM Pipeline(4), ... More
BRM-Site-B	✓	Generic System	↑	Host(6), Listener(5), Oracle BRM Pipeline(4), ... More

While this section of the white paper described the disaster recovery operation tests using Oracle Site Guard, the next section describes the functional tests performed during the local HA and DR tests.

## 8 Tests and Results

While the focus of this paper is to demonstrate the automated disaster recovery tests using Oracle Site Guard, as described in the previous section, various local HA tests were conducted at each BRM site according to the MAA tests guidelines. The following table lists the tests carried out during the MAA exercise. The BRM Oracle Database was loaded with 1 Million subscribers prior to the tests.

Test ID	BRM Apps Module	Test Scenario Description
HA-1	Billing Apps, Real-Time Pipeline	Failover new Billing Apps traffic to available CMMP Nodes and Real-Time Pipeline servers via a Local Load-Balancer VIP
HA-2	Billing Apps, Real-Time Pipeline	Load balance Billing Apps traffic to CMMP Nodes and Real-Time Pipeline servers via a Local Load-Balancer VIP
HA-3	Billing Apps, Real-Time Pipeline	Load balance database traffic routing from DM-Oracle to Oracle RAC-enabled nodes for the Billing Apps processes
HA-4	Billing Apps, Real-Time Pipeline	Failover database traffic routing from DM-Oracle to Oracle RAC-enabled nodes for the Billing Apps processes
HA-5	Batch Pipeline	Route CDR files in the Batch Pipeline directory on Shared Storage with a Virtual IP for the shared-storage access
HA-6	Batch Pipeline	Route CDR files in the Batch Pipeline directory on Shared Storage with a Virtual IP for the shared-storage access -- Failover the Batch Pipeline Node
HA-7	Batch Pipeline	Route CDR files in the Batch Pipeline directory on the Shared Storage Directory with a Virtual IP for the shared-storage access -- Failover the Oracle RAC database instance

Test ID	Steps/Procedure	Results
HA-1	<p>Process billing applications in just one cycle for 100 K subscribers using the following steps:</p> <ul style="list-style-type: none"> <li>- Update 100 K subscriber accounts to be billable from SQL prompt access to the BRM database:  <pre>SQL&gt; update billinfo_t set ACTG_NEXT_T=NEXT_BILL_T where account_obj_id0 in (select distinct account_obj_id0 from service_t where login like '004417104%'); commit;</pre> </li> <li>- Execute the Billing Apps Command:  <pre>date;pin_bill_accts -verbose -active;date</pre> </li> <li>- Check that the pin1. EVENT_BILLING_CYCLE_FOLD_T table is updated.</li> </ul>	<p>1. Check that the pin1. EVENT_BILLING_CYCLE_FOLD_T table is getting updated due to the billing apps processing failed over to the 2nd node.</p>
HA-2	<p>Process billing applications four times using the following steps:</p> <ul style="list-style-type: none"> <li>- Update 100 K subscribers accounts to be billable from SQL prompt access to BRM database:  <pre>SQL&gt; update billinfo_t set ACTG_NEXT_T=NEXT_BILL_T where account_obj_id0 in (select distinct account_obj_id0 from service_t where login like '004417104%'); commit;</pre> </li> <li>- Execute the Billing Apps Command:  <pre>date;pin_bill_accts -verbose -active;date</pre> </li> <li>- Check the CMMP logs on each node.</li> <li>- Check that the pin1. EVENT_BILLING_CYCLE_FOLD_T table is being updated.</li> </ul>	<p>1. Verify that billing apps processing traffic to CMMP nodes are equally distributed, that is, 2 sets are processed by each CMMP node. Tail the CMMP node logs.</p>
HA-3	<p>Process billing applications four times using the following steps:</p> <ul style="list-style-type: none"> <li>- Update 100K subscriber accounts to be billable from SQL prompt access to the BRM database:  <pre>SQL&gt; update billinfo_t set ACTG_NEXT_T=NEXT_BILL_T where account_obj_id0 in (select distinct account_obj_id0 from service_t where login like '004417104%'); commit;</pre> </li> <li>- Execute the Billing Apps Command:  <pre>date;pin_bill_accts -verbose -active;date</pre> </li> <li>- Check the DM-Oracle logs on each node.</li> </ul>	<p>1. Verify that DM-Oracle traffic to Oracle RAC nodes is equally distributed. Check the SQL sessions created on each Oracle RAC node for the 4 billing processes.</p>

	<ul style="list-style-type: none"> <li>- Check the SQL session created on each Oracle RAC node for the preceding 4 billing processes.</li> </ul>	
HA-4	<p>Process billing applications once using the following steps:</p> <ul style="list-style-type: none"> <li>- Update 100K subscriber accounts to be billable from SQL prompt access to the BRM database:  <pre>SQL&gt; update billinfo_t set ACTG_NEXT_T=NEXT_BILL_T where account_obj_id0 in (select distinct account_obj_id0 from service_t where login like '004417104%'); commit;</pre> </li> <li>- Execute the Billing Apps Command:  <pre>date;pin_bill_accts -verbose -active;date</pre> </li> <li>- Check the DM-Oracle logs on each node and find the Oracle RAC database instance to which the preceding billing process is connected.</li> <li>- Stop the preceding found instance of the Oracle RAC database.</li> <li>- Check the DM-Oracle logs again to verify that the processing continues with the 2nd Oracle RAC database instance.</li> <li>- Check that the pin1.  EVENT_BILLING_CYCLE_FOLD_T table is updated for all 100 K subscribers</li> </ul>	<ol style="list-style-type: none"> <li>1. Verify that billing process continues and the pin1.  EVENT_BILLING_CYCLE_FOLD_T table is updated for all 100 K subscribers.</li> </ol>
HA-5	<p>Transfer files with CDR into the IN directory mounted to the Batch Pipeline servers:</p> <ul style="list-style-type: none"> <li>- IN: /stage/pipe1/in</li> <li>- View pipeline logs to check if the CDR files are processed:  \$IFW_HOME/log/pipeline/pipeline_PRE_PROCESS.log0 and pipeline_ALL_RATE.log0</li> <li>- Check the count of pin1.event_t table in the BRM database for successfully processed CDRs.</li> </ul>	<ol style="list-style-type: none"> <li>1. Verify that view pipeline logs do not have any errors for processing the CDR files:  \$IFW_HOME/log/pipeline/pipeline_PRE_PROCESS.log0 and pipeline_ALL_RATE.log0</li> <li>2. Check that the count of the pin1.event_t table has 100 K successfully processed CDRs.</li> </ol>
HA-6	<p>Transfer files with CDR into the IN directory mounted to the Batch Pipeline servers:</p> <ul style="list-style-type: none"> <li>- IN: /stage/pipe1/in</li> <li>- View pipeline logs on the Active (1st) Batch Server to verify that the CDR files are processed:  \$IFW_HOME/log/pipeline/pipeline_PRE_PROCESS.log0 and pipeline_ALL_RATE.log0</li> <li>- Check that the count of the pin1.event_t table in the BRM database is being updated.</li> </ul>	<ol style="list-style-type: none"> <li>1. Verify that view pipeline logs on the 1st Batch Server are processing the CDR files:  \$IFW_HOME/log/pipeline/pipeline_PRE_PROCESS.log0 and pipeline_ALL_RATE.log0</li> <li>2. Check that the count of the pin1.event_t table</li> </ol>



	<ul style="list-style-type: none"> <li>- Stop the 1st Batch Pipeline server processes.</li> <li>- Start the standby (2nd) Batch Pipeline server processes.</li> <li>- Check the 2nd server's Pipeline logs.</li> <li>- Check that the count of the <a href="#">pin1.event_t</a> table in the BRM database is being updated again to finish the whole 100 K CDRs processing.</li> </ul>	<p>stopped increasing after the 1st Batch Server process was stopped.</p> <p>3. After the 2nd Batch Server processes are started, the pipeline logs on the 2nd batch server should resume the batch processing.</p> <p>4. The count of the <a href="#">pin1.event_t</a> table has 100 K successfully processed CDRs.</p>
HA-7	<p>Transfer files with CDR into the IN directory mounted to the Batch Pipeline servers:</p> <ul style="list-style-type: none"> <li>- IN: <a href="#">/stage/pipe1/in</a></li> <li>- View pipeline logs on the Active (1st) Batch Server to verify that the CDR files are processed: <a href="#">\$IFW_HOME/log/pipeline/pipeline_PRE_PROCESS.log0</a> and <a href="#">pipeline_ALL_RATE.log0</a></li> <li>- Check that the count of the <a href="#">pin1.event_t</a> table in the BRM database is being updated.</li> <li>- Check for the Oracle RAC instance to which the preceding batch pipeline session is connected.</li> <li>- Stop the preceding identified Oracle RAC database instance.</li> <li>- Check the Pipeline logs again.</li> <li>- Check that the count of the <a href="#">pin1.event_t</a> table in the BRM database is being updated again to finish the whole 100 K CDRs processing.</li> </ul>	<ol style="list-style-type: none"> <li>1. Verify that view pipeline logs do not have any errors for processing the CDR files: <a href="#">\$IFW_HOME/log/pipeline/pipeline_PRE_PROCESS.log0</a> and <a href="#">pipeline_ALL_RATE.log0</a></li> <li>2. Check that the count of the <a href="#">pin1.event_t</a> table has 100 K successfully processed CDRs</li> </ol>

The following table lists the functional tests done under heavy load during the disaster recovery operations performed by Oracle Site Guard.

Test ID	BRM Apps Module	Test Scenario Description and Results
DR-1	Billing Apps, Real-Time Pipeline	Switchover BRM-Site-A to BRM-Site-A, while a large Billing process is running and processing 100 K + subscribers. After the failover, resume the same Billing process by starting a new thread at BRM-Site-B; The new Billing process at BRM-Site-B continues to process from the thread where it was stopped at BRM-

		Site-A. Check the count of rows in the <a href="#">'EVENT_BILLING_CYCLE_FOLD_T'</a> table, which keeps increasing as it resumed the billing process.
DR-2	Batch Pipeline	Switch over BRM-Site-A to BRM-Site-A while a large batch pipeline is running and processing 100 K EDRs. After the failover, the batch pipeline continues to process the remaining EDRs available in the batch pipeline directory. (This directory was replicated via storage replication from BRM-Site-A to BRM-Site-B.) Check the count of rows in the <a href="#">'EVENT_T'</a> table, which kept increasing as it resumed the batch pipeline processing.

## 9 Benefits from Oracle Engineered Systems for BRM

The Oracle engineered systems bring extreme performance to every layer of the technology stack. These systems are pre-integrated to reduce the cost and complexity of IT infrastructures while increasing productivity and performance. These systems provide optimization at every layer of the stack to simplify data center operations, drive down costs, and accelerate business innovation. These benefits also apply to the Oracle Communications Billing and Revenue Management applications.

1. Scalability and availability of CM processing in a multitier deployment - The Data Processing tier (CM tier), along with the Data Translation tier (DM-Oracle) that handles most of the BRM processing, can be scaled out for high availability and load balancing over the InfiniBand fabric by using Oracle Traffic Director (OTD). In the BRM MAA exercise an F5 load balancer was used instead of OTD. Both Real-Time Pipeline and Batch Pipeline can take advantage of these scalable and HA configurations of the CM tier within the rack itself.
2. BRM Applications to BRM database communications over InfiniBand -- BRM processing with faster database access over Infiniband using a fully Engineered System based platform always adds great value to the overall solution.
3. Simplified management for production and disaster recovery operations -- Oracle Exalogic (with high-speed, fully integrated ZFS appliance), Oracle Exadata Database Machine, and Oracle Enterprise Manager Cloud Control (with Site Guard) together prove to be a well-integrated and tested technology stack for both regular production and disaster-recovery operations.
4. Advanced virtualization (server, storage, and network) using Oracle Exalogic.

## 10 Summary of BRM MAA Best Practices

A summary of the best practices that have been presented in this paper follows, providing a checklist for a BRM MAA implementation.

### 10.1 Best Practices for BRM Database High Availability

The following BRM database best practices should be applied to the primary and secondary sites to achieve highest availability:

- Deploy BRM on an Oracle RAC database for the highest availability and scalability.
- Use Automatic Storage Management to simplify the provisioning and management of database storage.
- Enable Oracle Flashback Database to provide the ability to “rewind” the database in the event of user errors.
- Use Oracle Recovery Manager regularly to back up the BRM database.
- Always use HugePages for BRM databases on Linux. Monitor memory usage, and adjust the workload and parameters accordingly.
- Configure database Dead Connection Detection to actively remove dead connections in the event of a BRM Server node failure.
- For Exadata deployments, configure cluster misscount to 30 seconds to reduce downtime in the event of a database node failure.

### 10.2 Best Practices for BRM Application High Availability

These are the BRM application best practices that should be applied to the primary and secondary sites to achieve highest availability:

- Deploy multiple BRM servers and deploy all critical BRM components in a load-balanced, distributed service, or clustered configuration, so that work can continue in the event of a BRM server node failure.
- Deploy a load balancer in a redundant configuration, and load balance the server load using our recommended logic.

- Deploy the BRM file system on a fault-tolerant filer. Take regular backups of the BRM servers and BRM file system.
- Connect to the database through a FAN-enabled service.

### 10.3 Best Practices for Disaster Readiness and Recovery

These are the best practices for deploying a secondary site and recovery procedures in readiness for a site outage:

- Deploy a second, geographically separated site that can run the BRM workload in the event the primary site is down.
- Use Data Guard to replicate all database changes to a standby database located on the secondary site.
- Take advantage of Oracle Active Data Guard to offload read-only queries to the standby database.
- Enable Oracle Flashback Database so that the old primary database can be quickly reinstated as a standby in the event of site failover.
- Replicate the BRM file system to the secondary site. Develop procedures for how to reverse the direction of replication in the event of failover or switchover and how to clone the replica for site testing.
- Export the BRM file system primary, standby replica, and clones, with different names to avoid mounting the incorrect ones.
- Create different role-based database services for the BRM database in primary, standby, and snapshot standby modes.
- Use Oracle Data Guard Broker to simplify Data Guard administration.
- Use the snapshot standby to provide an updatable replica of the primary database for temporary site testing.
- Use a Global Load Balancer like F5 GTM to manage DNS traffic for BRM Application Requests

## 11 Conclusion

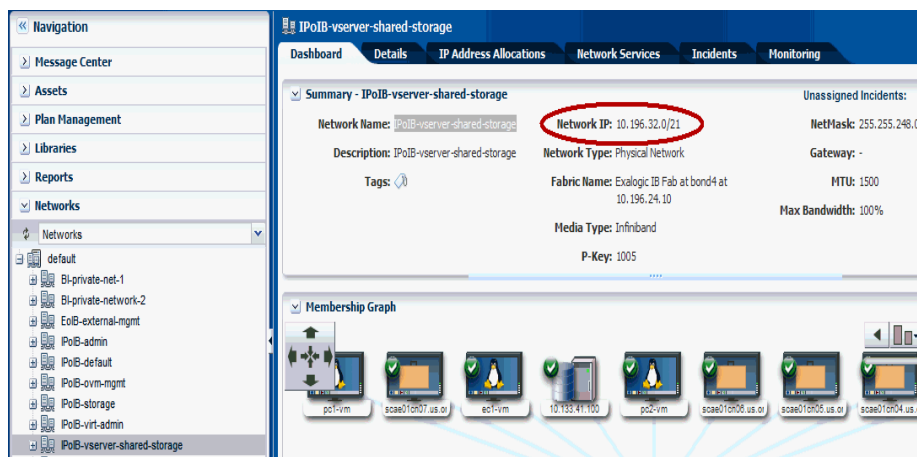
As validated in this MAA exercise, Oracle Communications BRM provides a best-practices blueprint for achieving the optimal BRM high-availability deployment using Oracle high-availability technologies and recommendations. The topology described in this white paper showcase a typical BRM deployment in a disaster-recovery configuration, which can be applicable to different implementation scenarios of BRM.

## A Appendix

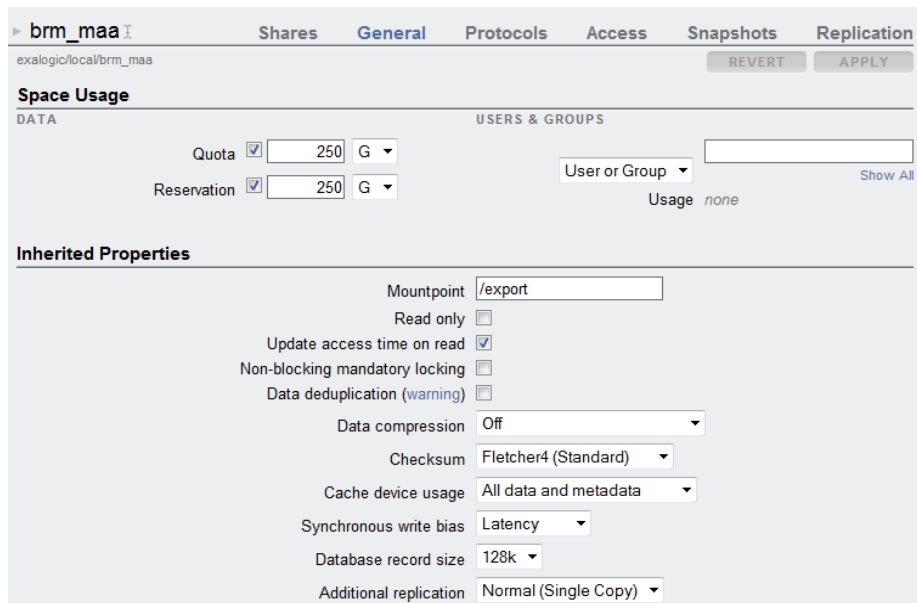
This appendix provides various scripts and screenshots along with the explanation of various best practices adopted in this white paper.

### A.1 Creating the ZFS project and shares for BRM MAA

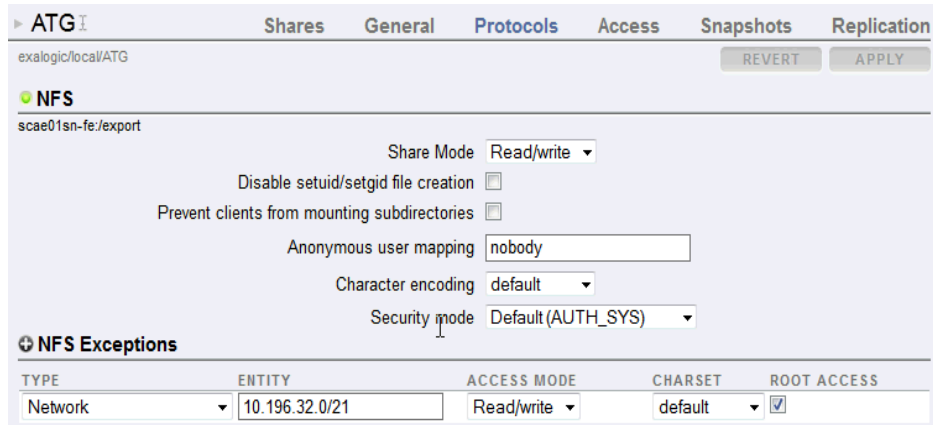
- a. Check the Enterprise Manager Ops Center (EMOC) Networks, specifically at the IPoIB-vserver-shared-storage Network IP, 10.196.32.0/21, to get the NFS Exception value for the ZFS Project.



- b. Create a ZFS project named 'brm\_maa' with the Quota option.



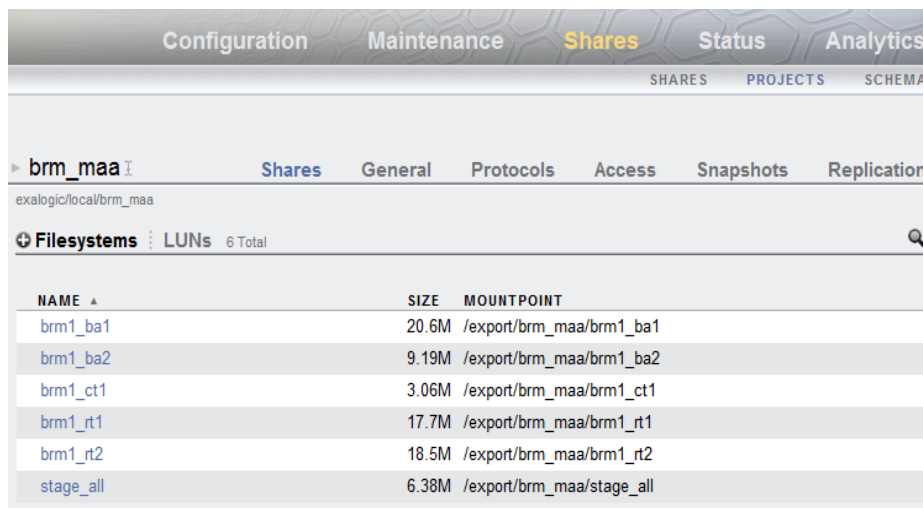
- c. Under the **Protocols** tab and **NFS Exceptions**, add the base host/subnet for the IPoIB-vserver-shared-storage network, as identified in Step a, which is **10.196.32.0/21**.



The screenshot shows the ATG NFS configuration interface. The 'NFS Exceptions' section is expanded, showing a table with the following data:

TYPE	ENTITY	ACCESS MODE	CHARSET	ROOT ACCESS
Network	10.196.32.0/21	Read/write	default	<input checked="" type="checkbox"/>

- d. Specify the Project Root Directory Access using the **pin** (1001) user and **pin** (1001) group. If NIS is in use then the numbers will get the actual user name and group name substituted.
- e. Create the ZFS shares under the **brm\_maa** project. Each share will inherit the project characteristics.



The screenshot shows the Shares configuration interface for the **brm\_maa** project. The 'Shares' tab is selected, and a table lists the following ZFS shares:

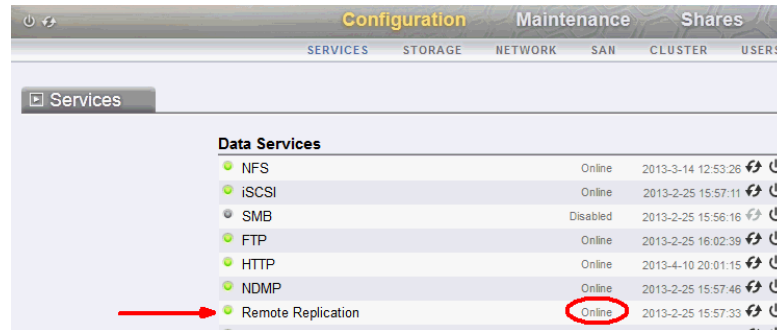
NAME	SIZE	MOUNTPOINT
brm1_ba1	20.6M	/export/brm_maa/brm1_ba1
brm1_ba2	9.19M	/export/brm_maa/brm1_ba2
brm1_ct1	3.06M	/export/brm_maa/brm1_ct1
brm1_rt1	17.7M	/export/brm_maa/brm1_rt1
brm1_rt2	18.5M	/export/brm_maa/brm1_rt2
stage_all	6.38M	/export/brm_maa/stage_all

Also refer to “[Setting up Access to the ZFS Storage Appliance for a vServer](#)” section in *Oracle Exalogic Elastic Cloud Administrator's Guide*.

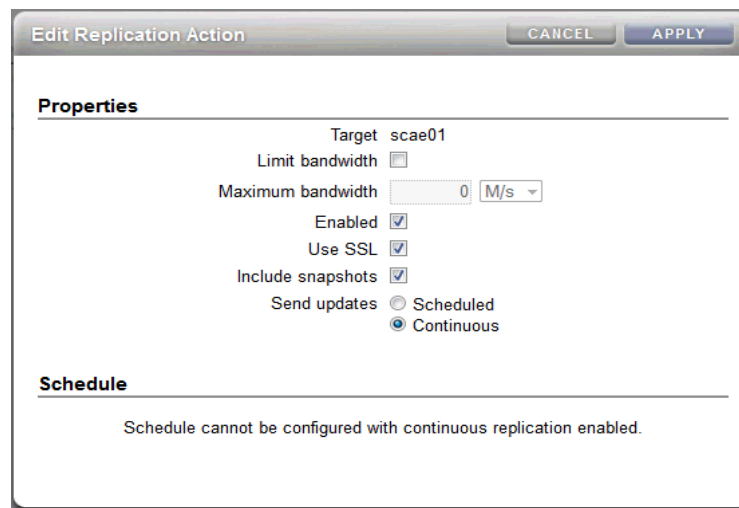


## A.2 Set Up ZFS Remote Replication for BRM Apps

- a. From the Configuration menu, choose **Services**, and verify that the status of the Remote Replication service is **Online**.



- b. Add the replication target if necessary. To do so click **Remote Replication** in the preceding screen. Ensure only one unique replication target is created in the Remote Replication page.
- c. Select the virtual host that floats to the active ZFS clustered head. Turn on replication for the ZFS Project **brm\_maa**.



- d. Verify that the shares are being replicated.
- e. Log in to the target ZFS Storage Appliance and look at the replicas.

### A.3 Set Up F5 Load Balancer for BRM CMMP Traffic

Obtain an IP address for the virtual servers as listed in Table 9: *F5 Networks Load-Balancer Details*. The following is an example of creating a virtual server at F5 LTM at BRM-Site-A.

- a. Create a TCP\_HALF open monitor named '**brm\_tcp\_half\_open**'.

The screenshot shows the configuration page for a monitor named 'brm\_tcp\_half\_open'. The breadcrumb path is 'Local Traffic >> Monitors >> brm\_tcp\_half\_open'. There are tabs for 'Properties' (selected) and 'Instances'. Below the tabs is a 'General Properties' table:


Name	brm_tcp_half_open
Partition / Path	Common
Description	BRM TCP Half Open Monitor
Type	TCP Half Open
Parent Monitor	tcp_half_open

Below the table is a 'Configuration' section with a dropdown menu set to 'Basic'. It contains three rows of configuration options:

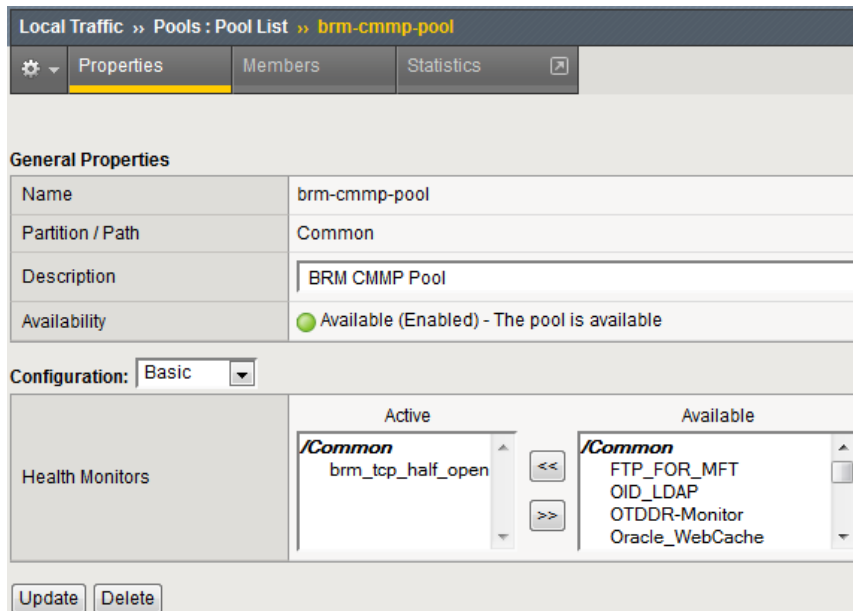
Interval	Specify... [v]	5	seconds
Timeout	Specify... [v]	16	seconds
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No		

The TCP Half-Open monitor is used to check if the BRM process is listening on the TCP port, but does not complete the TCP 3-way handshake. The TCP Half-Open Syn + Syn-Ack monitor is used to prevent error messages in the BRM application logs. Using other monitors will create these "application connection error" messages in the logs, because the BRM process expects data after every TCP handshake. The LTM monitor does not send this data, so BRM will log this (incorrectly) as an error.

- b. Add the BRM Apps nodes to the Nodes section of the Local Traffic manager, and associate the monitor '**brm\_tcp\_half\_open**' with the CMMP nodes.



- c. Create a pool named '**brm-cmmp-pool**', and associate the two CMMP nodes and the pool with the **brm\_tcp\_half\_open** monitor.



Local Traffic » Pools : Pool List » brm-cmmp-pool

Properties Members Statistics

**Load Balancing**

Load Balancing Method: Round Robin

Priority Group Activation: Disabled

Update

**Current Members**

<input type="checkbox"/>	Status	Member
<input type="checkbox"/>	<span style="color: green;">●</span>	scae01vm201.us.oracle.com:11959
<input type="checkbox"/>	<span style="color: green;">●</span>	scae01vm202:11959

d. Create an iRule for BRM as follows.

Main Help About Local Traffic » iRules : iRule List » New iRule...

Statistics iApp Local Traffic

- Network Map
- Virtual Servers
- Policies
- Profiles
- iRules**
- Pools
- Nodes
- Monitors
- Traffic Class
- Address Translation
- DNS Express Zones
- DNS Caches

Acceleration Device Management Network

**Properties**

Name: brm-iRule

**Definition**

```

### brm-iRule ###
when CLIENT_ACCEPTED {
  # The current_pool is derived from the virtual server's "Default Pool"
  set current_pool [LB::server pool]
  # Count of rerelect attempts
  set rerelects 0
}

when LB_FAILED {
  # Rerelect another pool member if the connection fails.
  # Limit rerelects to the number of active members.
  # To be active, a member must be enabled with status available or unknown.
  if { $rerelects < [active_members $current_pool] } {
    incr rerelects
    log local0.info "Client [IP::client_addr]:[TCP::client_port]: Selected
pool member [LB::server addr]:[LB::server port] not responding. Rerelect
attempt $rerelects."
    LB::mode rr
    LB::rerelect pool $current_pool
  } else {
    log local0.error "Client [IP::client_addr]:[TCP::client_port]: Failed
to connect after $rerelects rerelects."
  }
}
    
```

Extend Text Area  
 Wrap Text

The following code is used in this iRule:

```
## brm-iRule ###
when CLIENT_ACCEPTED {
    # The current pool is derived from the virtual server's "Default Pool"
    set current_pool [LB::server pool]
    # Count of reselect attempts
    set reselects 0
}

when LB_FAILED {
    # Reselect another pool member if the connection fails.
    # Limit reselects to the number of active members.
    # To be active, a member must be enabled with status available or
    unknown.
    if { $reselects < [active_members $current_pool] } {
        incr reselects
        log local0.info "Client [IP::client_addr]:[TCP::client_port]: Selected
pool member [LB::server addr]:[LB::server port] not responding. Reselect
attempt $reselects."
        LB::mode rr
        LB::reselect pool $current_pool
    } else {
        log local0.error "Client [IP::client_addr]:[TCP::client_port]: Failed
to connect after $reselects reselects."
    }
}

###
```

- e. Create a virtual server named '**cmmp-a**' for CMMP at BRM-Site-A. Associate **brm-iRule** and the **brm-cmmp-pool** with the virtual server.

Local Traffic >> Virtual Servers : Virtual Server List >> New Virtual Server...

**General Properties**

Name	cmmp-a
Description	Virtual Server for CMMP at BRM-Site-A
Type	Standard
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: cmmp-a.us.oracle.com
Service Port	11959 Other:
State	Enabled

Configuration: Basic

Protocol	TCP
SNAT Pool	Auto Map

**Resources**

iRules	Enabled /Common brm-iRule
Default Pool	+ brm-cmmp-pool

Ava  
beehive\_HTT  
ccuc\_httpoht  
fc\_debug  
fc\_replace\_p  
gc11ui\_http\_

Key points to note here for the BRM Virtual Server configuration are to associate the pool with the **tcp\_half\_open** monitor, select **AutoMap** for SNAT pool, and have the associated iRule according to the code described earlier. Although a default TCP Profile ( with TCP Idle Timeout of 300 seconds) was used in the MAA exercise, for long lived connections, a new TCP profile ( with TCP Idle Timeout of 1800 or 3600 seconds) can be used for both client and server profile of the virtual server.

Similar to the **cmmp-a** virtual server, create a similar virtual server named **cmmp-b** at the LTM of BRM-Site-B. Once the two local virtual servers for cmmp are created, F5-GTM needs to be configured with the virtual server **cmmp-ab** for Disaster Recovery configuration.

## A.4 BRM Configuration Files and Scripts Used in BRM MAA

### A.4.1 BRM DB User (pin) creation script

```
## Create BRM DB User Script###
#!/bin/sh
for i in 1
do
(
USERID="PIN${i}"; export USERID
PASSWD="pin${i}"; export PASSWD
DBNAME="brmdbp"; export DBNAME
SYSPASSWD="welcome1"; export SYSPASSWD
sqlplus sys/${SYSPASSWD}@${DBNAME} as sysdba <<EOF
connect ${USERID}/${PASSWD}@${DBNAME
exec DBMS_AQADM.DROP_QUEUE_TABLE('ACK_QUEUE_T',true);
exec DBMS_AQADM.DROP_QUEUE_TABLE('IFW_SYNC',true);
exec DBMS_AQADM.DROP_QUEUE_TABLE('IFW_SYNC_ROUTER',true);
connect sys/${SYSPASSWD}@${DBNAME} as sysdba
drop user ${USERID} cascade;
create user ${USERID} identified by ${PASSWD} default tablespace
${USERID}00 temporary tablespace pinltemp;
grant resource,connect,create session,create procedure,dbfs_role to
${USERID};
grant dba to ${USERID};
grant create table, create view, create snapshot, create any snapshot,
create trigger to ${USERID};
grant execute on dbms_aq to ${USERID};
grant execute on dbms_aqadm to ${USERID};
grant execute on dbms_lock to ${USERID};
grant select on sys.gv_\$aq to ${USERID};
call dbms_java.grant_permission(
'${USERID}','SYS:java.net.SocketPermission','*', 'connect,resolve' );
call dbms_java.grant_permission( '${USERID}','SYS:java.io.FilePermission',
'<<ALL FILES>>', 'read,write');
commit;
connect pin1/pin1@${DBNAME
grant all privileges to ${USERID};
exit
EOF
)
Done
```

### A.4.2 TNS Name used for BRM DB connection

```
## TNS Names used in BRM MAA Exercise ###

RACHA =
(DESCRIPTION =
(ADDRESS = (PROTOCOL = TCP)(HOST = scac0103-vip)(PORT = 1521))
(ADDRESS = (PROTOCOL = TCP)(HOST = scac0104-vip)(PORT = 1521))
(LOAD_BALANCE = YES)
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = brmdb.us.oracle.com)
(FAILOVER_MODE =
(TYPE = SELECT)
(METHOD = BASIC)
(RETRIES = 180)
(DELAY = 5)
)
)
)
```

#### A.4.3 pin\_ctl.conf in the Billing Apps with CMMP LBR Address

Locate the `pin.conf` configuration file for the Billing Apps at `$PIN_HOME/apps/pin_bill1d`. In the BRM MAA exercise, Billing Apps nodes at BRM-Site-A are `scan03vm0008` and `scan03vm0009`. Update the file to include the F5 GTM virtual server as follows:

```
#- nap cm_ptr ip scan03vm0008-eoib2 11960
#- nap cm_ptr ip scan03vm0008-eoib2 11959
- nap cm_ptr ip cmmp-ab.us.oracle.com 11959
```

#### A.4.4 Infranet.properties File of Batch Controller with CMMP LBR Address

Locate the `Infranet.properties` configuration file for Batch Controller at `$PIN_HOME/apps/batch_controller`. In the BRM MAA exercise, Batch Controller nodes are `scan03vm0013` and `scan03vm0014`. Update the file to include the F5 GTM virtual server as follows:

```
# Infranet CM connection configuration

infranet.connection
pcp://root.0.0.0.1:&aes|<encrypted DB user Password
>@cmmp-ab.us.oracle.com.:11959/service/pcm_client
```

#### A.4.5 dm\_oracle svcname configuration

For `dm_oracle` to process the FAN notifications, the FAN enabled database service name needs to be configured using `svcname` setting. If the correct service name is not configured in the `pin.conf`, `dm_oracle` will not handle the received FAN notifications.

The setting that needed to be configured in the MAA exercise was:

```
- dm sm_svcname brmdb.us.oracle.com
```



#### A.4.6 Start and Stop Scripts for BRM Components

Based on the BRM components used prepare the start and stop scripts. Test these scripts prior to using them in the Oracle Site Guard configuration. Follow the sequence of starting of the BRM components as follows:

1. Start the BRM Database.
2. Start Data Manager for Oracle components on all nodes (dm\_oracle).
3. Start Connection Manager (CM) on all nodes (cm).
4. Start the Connection Manager Master Process (cmmp) on all nodes.
5. Start any remaining BRM components, including Real-Time Pipeline and Batch Pipeline.

To stop the BRM components follow the reverse order of the starting sequence.

**Site Guard Configuration**

General Credentials Pre/Post Scripts Storage Scripts

Pre and Post scripts are custom scripts associated with a Site. A script can be associated with more than one host target in the site. They are executed as part of the work flow - Pre-Script are executed as the first step and Post script are executed as the last step.

- For example, script.sh -param1 value1 -param2 value2

Switchover and Failover operation types will be shown when Site Guard configuration has primary and one or more standby sites.

+ Add Edit Delete

Script Path	Script Type	Operation	Role	Target Hosts	Run On	Cred Type
/export/pin/umount_batch_pipe_dbfs	Pre-Script	Switchover	Primary	scan03vm0013-eoib2.us.oracle.com	All Hosts	Privile
/export/pin/7.5/bin/stop_all	Pre-Script	Switchover	Primary	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/7.5/bin/stop_dm_fw_sync_maa	Pre-Script	Switchover	Primary	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/fw_7.5/stop_rtp_maa	Pre-Script	Switchover	Primary	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/fw_7.5/stop_bp_maa	Pre-Script	Switchover	Primary	scan03vm0013-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/fw_7.5/stop_bt_controller	Pre-Script	Switchover	Primary	scan03vm0013-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/7.5/bin/stop_eai_js_maa	Pre-Script	Switchover	Primary	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/7.5/bin/start_db_queues	Post-Script	Switchover	Standby	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	Any Host	Norm
/export/pin/7.5/bin/start_all	Post-Script	Switchover	Standby	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/7.5/bin/start_dm_fw_sync_maa	Post-Script	Switchover	Standby	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/fw_7.5/start_rtp_maa	Post-Script	Switchover	Standby	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/fw_7.5/start_bp_maa	Post-Script	Switchover	Standby	scan03vm0013-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/fw_7.5/start_bt_controller	Post-Script	Switchover	Standby	scan03vm0013-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/7.5/bin/start_eai_js_maa	Post-Script	Switchover	Standby	scan03vm0008-eoib2.us.oracle.com, scan03vm0009-eoib2.us.oracle.com	All Hosts	Norm
/export/pin/mount_batch_pipe_dbfs	Post-Script	Switchover	Standby	scan03vm0013-eoib2.us.oracle.com	All Hosts	Norm

#### A.4.7 ZFS Storage Scripts

All of the ZFS storage scripts in this deployment come bundled with Oracle Site Guard. The following scripts are used in the MAA exercise.

Storage Role-Reversal Scripts for Switchover and Failover.

```
sh zfs_storage_role_reversal.sh -t scan03-fe -h scae01-fe -j brm_maa -p exalogic -q exalogic -o Switchover -c Y -f Y -e 1800 -l N
```

```
sh zfs_storage_role_reversal.sh -t scan03-fe -h scae01-fe -j brm_maa -p exalogic -q exalogic -o Failover -c Y -f Y -e 1800 -l N
```

## Mounting and Unmounting of Shares

```
sh mount_umount.sh -o mount -f '/brm1,/stage'  
sh mount_umount.sh -o umount -f '/brm1,/stage'
```

### A.5 Oracle DBFS mounts for BRM Batch Pipeline

Refer to the [Oracle DBFS documentation](#) for information about setting up and mounting the DBFS shares for the BRM Pipeline.

Mount DBFS shares using the **pin** OS user. Oracle Wallet is used to avoid user name and password prompts.

```
$ORACLE_HOME/bin/dbfs_client -o wallet /@racha1 /u01/dbfs
```

Unmount DBFS shares as the **root** user.

```
# fusermount -u /u01/dbfs
```

### A.6 Disaster Recovery Host Aliasing

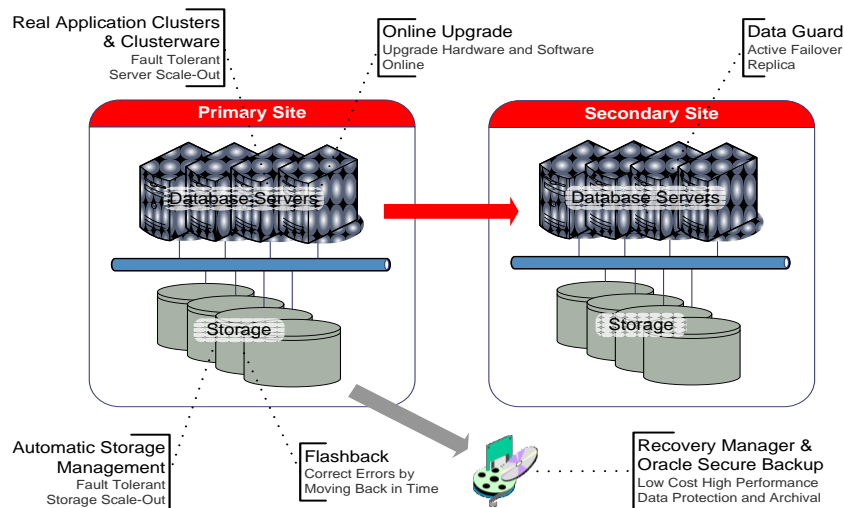
In a disaster recovery topology, the production site host names must be resolvable to the IP addresses of the corresponding peer systems at the standby site. Therefore, it is important to plan the host names for the production site and standby site. After a role transition (failover or switchover) from a primary site to a standby site, the alias host names for the application tier hosts on the standby site become active. You do not need to reconfigure host names for the hosts on the standby site because you set up aliases on the standby site. Also see “[Network Considerations](#)” and “[Planning Host Names](#)” in the *Oracle® Fusion Middleware Disaster Recovery Guide*.

In the BRM MAA exercise a single (global) DNS server was used so the disaster recovery site **/etc/hosts** files had to be updated with host aliases.

## A.7 Oracle Database MAA

To achieve maximum BRM application availability, Oracle recommends deploying BRM applications on an Oracle Database MAA foundation, as depicted in the next figure, which includes the database technologies described in the following text.

- [Oracle Real Application Clusters](#)
- [Oracle Data Guard and Online Upgrade](#)
- [Oracle Flashback Database](#)



### A.7.1 Oracle Real Application Clusters

Oracle Real Application Clusters (Oracle RAC) allows the Oracle Database to run any packaged or custom application unchanged across a set of clustered nodes. This capability provides the highest levels of availability and the most flexible scalability. If a clustered node fails the Oracle Database instance will continue running on the surviving nodes. When more processing power is needed another node can be added without interrupting user access to data. For more information see the [Oracle Real Application Clusters Administration and Deployment Guide](#).

### A.7.2 Oracle Data Guard and Online Upgrade

Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle

Database instances to survive failures, disasters, user errors, and data corruption. Data Guard maintains these standby databases as transaction-based, consistent copies of the production database. If the production database becomes unavailable due to a planned or unplanned outage, Data Guard can switch any standby database to the production role, thus greatly reducing the application downtime caused by the outage. Data Guard can be used with traditional backup, restore, and clustering solutions to provide a high level of data protection and data availability. See also [Oracle Data Guard Concepts and Administration](#).

BRM supports both physical and logical standby databases. A physical standby database provides a physically identical copy of the primary database, with on-disk database structures that are identical to the primary database on a block-for-block basis. A physical standby database is kept synchronized with the primary database through Redo Apply, which recovers the redo data received from the primary database and applies the redo to the physical standby database.

With a single command, a physical standby database can be converted into a Snapshot Standby and become an independent database open for read-write access, ideal for QA and other testing. The Snapshot Standby continues to receive and archive redo data from the primary database while it is open for read-write access, thus protecting primary data at all times. When testing is complete, a single command will convert the snapshot back into a standby database and automatically resynchronize it with the primary.

It is possible to deploy a local standby database at the primary site as well as a remote standby at a secondary site. A local standby offers the advantage that a failover can be performed while the BRM servers continue running, and it can be done almost transparently to the end users. This standby also offers the ability to perform an online database upgrade without the need to switch to another site. Oracle recommends that both a local and remote standby be deployed for maximum availability.

### A.7.3 Oracle Flashback Database

Oracle Flashback Database provides a more efficient alternative to database point-in-time recovery. With Oracle Flashback Database, current data files can be reverted to their contents at a past time. The result is much like restoring data from data-file backups and executing point-in-time database recovery. However, Flashback Database skips the data file restoration and most of the applications of redo data.

Enabling Oracle Flashback Database provides the following benefits:

- Eliminates the time to restore a backup when fixing a human error that has a database-wide impact.
- Because human errors can be quickly undone, it allows standby databases to use real-time apply to synchronize with the primary database.
- Allows quick standby database reinstatement after a database failover.

#### A.7.4 Configure Dead Connection Detection

When a BRM server node fails suddenly there may not be time for the operating system to reset the TCP connections, and as a result the connections on the database server will remain open. To clean up the “dead” connections, Oracle recommends that Dead Connection Detection be configured. See [MOS ID 151972.1 “\(Dead Connection Detection \(DCD\) Explained”](#) for details.

Making these configuration changes may have an adverse effect on network utilization, so all changes should be tested and monitored carefully.

#### A.7.5 Reduce Timeout on Oracle RAC Node Failure (Exadata Only)

On Exadata it is possible to failover more quickly in the event of an Oracle RAC node failure by reducing the **misscount** parameter. This parameter defines how long to wait after a node becomes unresponsive before evicting the node from the cluster. The parameter should not be set to less than **'30'** (30 seconds). To update the CSS **misscount** setting, log in as the **root** user on one of the database servers, and run this command:

```
$GRID_HOME/bin/crsctl set css misscount 30
```

## A.8 Benefits of Oracle Site Guard

The following are key differentiators that Oracle Site Guard offers as compared to any home-grown scripting solution.

1. **Centralized logging** - Logs of all the remote operations are captured and persisted in the Oracle Enterprise Manager repository. The logs are properly categorized in different buckets, and it is very handy for troubleshooting issues.

2. **Dynamic binding of credentials** - Site Guard directly fetches the required credentials at run time from the Oracle Enterprise Manager credential framework. Any change to the

underlying credentials doesn't affect the DR workflow. These credentials are highly secured. Credentials are very difficult to manage and create a huge security violation with any scripting solution.

3. **Restartable options** - Site Guard provides an option to stop, suspend, or resume the DR operation at any point during its execution.

4. **Prechecks** - Site Guard runs extensive checks before performing any DR operation. The checks assure that the end-to-end DR operation would complete successfully and provide a lot of confidence to the DR administrator for carrying out the DR operation. More details about prechecks is captured in the following text.

5. **Error management and troubleshooting** - Site Guard provides options to set error modes for every step in the end-to-end DR operation. The error mode can be either "CONTINUE ON ERROR" or "STOP ON ERROR". This provides great flexibility and resilience to failures at the same time it provide more robustness.

6. **Retry/Timeout mechanisms** - Site Guard provides options to retry any failed step. If the issue is known to the user, the errors can be ignored as well. Site Guard also provide a fine-grained timeout mechanism for all remote operations.

7. **Rerun safe** - The DR operation is rerun safely, and Site Guard internally generates no-ops for all steps that were previously successful.

8. **HA support for critical operations** - If a remote operation or action is critical, Site Guard provides options to execute them from any host. For example, a script can be configured to get executed from Host A (and also from Host B, Host C, and so on). If Host A is down or not reachable, the script gets automatically executed from Host B or Host C, depending on the availability. This is important for running mission-critical operations like Oracle RAC database or storage role reversal. This gets really tricky with home-grown scripts when it comes to monitoring and manageability of remote host and database instances.

9. **Periodic health checks and reporting** - The end-to-end DR operation can be validated periodically at scheduled intervals, and the user gets notified if Site Guard detects any errors in the underlying DR deployment. This provides lot of confidence and peace of mind to the site administrator and assures that the deployment is ready to get recovered from any kind of failures and data-center outages. The checks also assert replication (database and ZFS storage) lags against preconfigured thresholds (SLA).

10. **Option to disable any step** - Any step can be disabled and get excluded from execution at run time.

11. **Parallel execution** - Most of the remote operations are executed in parallel (where possible). The execution can be made **Serial** as well.

12. **Extensibility** - Site Guard is highly extensible and exposes call-out mechanisms to execute any custom scripts as part of the DR workflow. These callouts are very flexible and provide support to be executed virtually anywhere during the DR operation.

13. **Auto discovery of topology** - Site Guard discovers the topology and its properties from Oracle Enterprise Manager, and the DR workflow can be automatically generated. This removes the complexity and burden from the users to determine "where to run what" and feeding in all the necessary information like Oracle homes, SIDs, and other properties required for the remote operation. This gets super complicated and impossible to manage with home-grown scripts.

14. **Handling topology changes like Scale up / Scale out / Scale down** - Site Guard can automatically handle any changes (scale up, scale down, or scale out) to the underlying topology.

15. **Reduce application down time** - With one-click Switchover/Failover, Site Guard helps to drastically reduce the overall application down time. Site Guard enables businesses to meet aggressive RTOs.

16. **Out-of-the- box support for ZFS storage** - Site Guard provides complete out-of-the-box support to perform Switchover and Failover of ZFS volumes.

17. **Reduce recovery time** -Site Guard helps reduce human error.

18. **Reduced skill set required to perform DR operations** - No application, replication, or infrastructure experts are needed onsite when disaster happens.

## B References

1. [Oracle Maximum Availability Architecture Web site](#)
2. [MAA Best Practices for Oracle Exadata Database Machine](#)
3. [Oracle Exalogic Elastic Cloud Administrator's Guide](#)
4. [Oracle Fusion Middleware Disaster Recovery using Oracle's Sun ZFS Storage Appliance](#)
5. [Disaster Recovery for Oracle Exalogic Elastic Cloud with Oracle Exadata Database Machine](#)
6. [Oracle Communications Billing and Revenue Documentation](#)
7. [Automating Disaster Recovery using Oracle Site Guard for Oracle Exalogic](#)
8. [F5 Networks BIG-IP Product Modules](#)
9. [F5 Networks BIG-IP Global Traffic Manager](#)
10. [Oracle Landing Page at F5 Networks](#)





Oracle Communications BRM MAA on Oracle  
Engineered Systems  
June 2014

*Authors:* Lingaraj Nayak, Jessica Mao  
*Contributors:* Shari Yamaguchi, Mahesh Desai,  
Latha Krishnaswamy, Jitendra Yadav, Praveen  
Sampath, Chris Akker (F5 Networks)

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

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



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

Copyright © 2014, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

**Hardware and Software, Engineered to Work Together**