An Oracle White Paper
February 2013

Beijing Mobile
Oracle Exadata Technical Case Study

Overview

BMCC (Beijing Mobile Communications Corporation, or Beijing Mobile) is one of 32 operating subsidiaries of China Mobile Limited. China Mobile provides mobile voice, data, IP telephone and multimedia services to more than 660,000,000 customers covering every province in China plus Hong Kong. As the leading mobile services provider in China, China Mobile boasts the world's largest mobile network and the world's largest mobile customer base. In 2010, the Company was once again selected as one of the "FT Global 500" by Financial Times and one of "The World's 2000 Biggest Public Companies" by Forbes magazine.

Beijing Mobile chose the Oracle Exadata Database Machine (Exadata) in 2011 as a consolidation destination for 18 separate database servers supporting the company's Operational Support Systems (OSS). The project was designed to rationalize the mix of different hardware platforms, operating systems and databases, resulting in less administrative effort and better resource utilization. At the same time, the Exadata performance was expected to overcome previous I/O challenges. Finally, the use of Oracle Active Data Guard between Exadata Database Machines would ensure high availability of the database services, which hadn't existed previously.

Intended Audience

This paper reviews Beijing Mobile's use of Oracle Exadata Database Machine and provides configuration details and benefits specific to the deployment. Readers are assumed to have experience with Oracle Database 11g technologies, familiarity with the Oracle Maximum Availability Architecture framework (MAA), and a general technical understanding of Oracle Exadata Database Machine. When referenced in this paper, in-depth background on these topics will be deferred, as they are covered in other documentation and technical white papers available on the Oracle Technology Network[1]. Where applicable, links to additional technical information are provided in footnotes that accompany each section. Appendix A also provides a list of the acronyms used in this paper.

Introduction

China Mobile[2] operates the world's largest mobile communications network, with over 600,000 mobile base stations in China, and a subscriber base equal to nearly half the population in China. China Mobile's network is operated by 31 subsidiaries in China, including Beijing Mobile, which serves the Beijing metropolitan area of over 20 million residents.

Beijing Mobile has a long history with the Oracle Database. The company has been an Oracle customer since the mid-1990s.

Beijing Mobile's OSS Database Consolidation

Beijing Mobile selected the Oracle Exadata Database Machine X2-2 Half-Rack configuration as a consolidation destination for its 18 separate OSS (Operational Support System) database servers. A Quarter-Rack X2-2 configuration was later chosen as a standby system for HA (High Availability), and became the primary system for one of the Oracle databases. The 18 databases were migrated as

[1] http://www.oracle.com/technetwork/database/exadata/index.html

[2] http://www.chinamobileltd.com/

schemas into 4 Oracle database instances on the Exadata systems. Each database was implemented as a 2-node Real Application Cluster (RAC) database, ensuring availability of each database in the event of a compute node failure. See Figure 1 for the logical database layout.

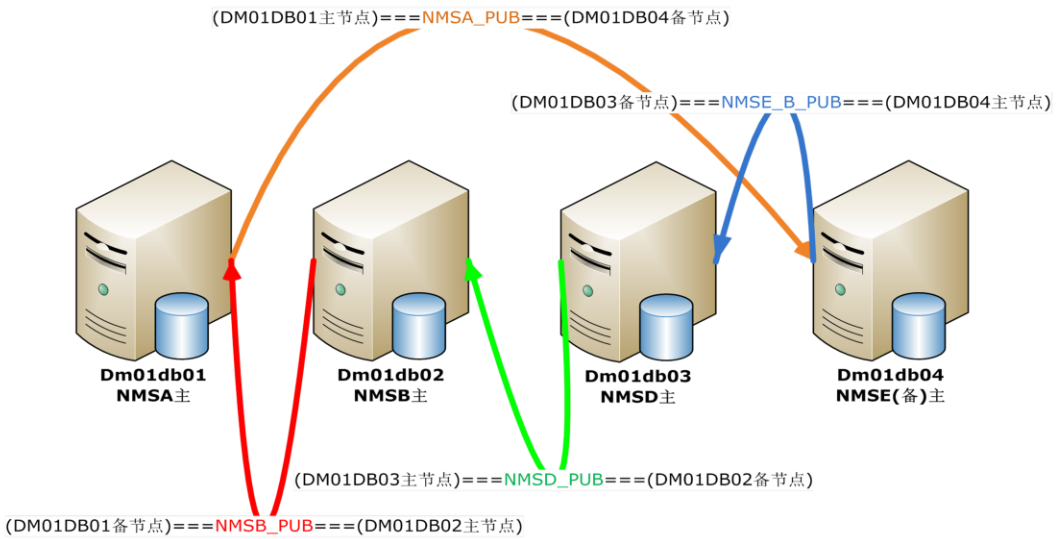

FIGURE 1 – LOGICAL DATABASE ARCHITECTURE

The applications and databases supported by the Exadata consolidation are as follows:

Application Databases Consolidated into NMSA DB[3]

EMOS

The E-Maintenance and Operating System (EMOS) is used by the Beijing Mobile network maintenance department to direct, schedule and manage daily production tasks. This is the primary and most important business system supported by Exadata.

[3] NMS = Network Management System. There are four NMS databases in this implementation; NMSA, NMSB, NMSC and NMSD.

Transmission Network System

Used for the scheduling, management and inspection of the Transmission Network. It includes the following functional modules:

- Topology Management
- Fault Management
- Performance Management
- Configuration Management
- Resource Management
- Resource Scheduling
- Automatic Inspection
- Spare Parts and Instrument Management
- Protection Management
- Report Management
- System Management
- Safety Management

Information Publishing Platform

Used for the configuration and publishing of indicators, monitoring, permissions, SMS and MMS.

Application Databases Consolidated into NMSB DB

Comprehensive Alarm System

A comprehensive monitoring platform used for alarm collection and presentation, covering every domain of the mobile communications network, containing different types of alarms for 2G and 3G network elements, performance alarms, dynamic environment alarms, dials measure alarms, data services alarms, etc. The data is used to analyze and enhance the network quality and the stability of the proprietary mobile network.

Centralized Operation and Maintenance Platform

Provides the following basic functions:

- Adapter connection
- Instruction issuing and analysis
- Intelligent inspection for the main functions
- Bureau data verification and new base station building

GIS System

This application is responsible for the display and analysis of geographic information for all resources, alarms and performance.

Uniform Collection Platform

Used for the statistics of mobile network data, MSC, BSC, base stations & community network, monitoring the network load, network performance analysis & statistics, providing network performance data used by the network optimization and maintenance departments.

Application Databases Consolidated into NMSC DB

OLAP / DW

The Data Warehouse is sourced from the Signaling System, Synthesis Analysis System and in the future, the Unified Collection History System. The data is used by the Beijing Mobile internal network department and China Mobile Group.

Application Databases Consolidated into NMSD DB

Network Customer Service

Used by customer service representatives to resolve the China Mobile customer registration complaints, support information consulting and provide simple statistical query functions.

Comprehensive Analysis

Used for the acquisition of the original voice, SMS, VLR and related information.

Base Station Monitoring

Used to query real-time and historical base station data, including real-time and historical alarms, and for viewing related video.

Consolidation Approach

Consolidation of separate databases into a common database can be challenging, particularly if consistent naming standards and definitions were not followed prior to the consolidation. Name "collisions" are often a big obstacle, as not only will tables and column names need changing, but all the applications that reference the names that have changed. Beijing Mobile was able to map the schemas from the 18 previous databases to the 4 consolidated databases without making any changes to the schema object names. For example, the schema 'user001' on the previous database was mapped to the same schema 'user001' on the new Exadata database.

It was fortunate that no duplicated schema names were found in these original database systems. Beijing Mobile does not have a unified naming standard, so if they find a duplicated schema name(s) in any additional consolidations in the future, they will create a new database or alter the conflicting names to fit into an existing database as a new schema.

Another challenge can be the conversion and loading of the existing data into the new system. This could entail a format conversion if the before and after systems use a different internal data representation. It will also likely require changes to the overall database settings, such as the amount of memory allocated for database buffers and log files, and perhaps configuration settings that relate to the number of users accessing the consolidated database.

Although there were no naming collisions, Beijing Mobile made a small number of changes to their database and applications to accommodate the consolidation onto Exadata, as follows:

1. Removed the limit on the number of failed logins to the database.

2. Set the maximum connection time limit and maximum idle time limit for schemas with short connection times.

3. Modified the db connection string in the application configurations due to the target db change.

4. Replaced the "SID" (Oracle Database Instance Identifier) with "ServiceName" in the client connection.

5. To take advantage of Hybrid Columnar Compression in Exadata, added a "HINT" in the respective SQL statements.

Data Migration Methods

Multiple approaches were used to migrate the 18 databases onto Exadata, depending on the size of the database and the amount of downtime allowed by the business.

Export-Import

This approach is recommended when transporting data between systems with different architectures. Export-Import (and more recently, Datapump export-import) converts the database into a neutral Oracle format that can be imported into the target system. In the case of Datapump export-import, this can be done directly over the network with no intermediate files. Depending on the size of the database, both of these methods were employed.

Oracle GoldenGate

The databases that had to be available with near-zero downtime were migrated using the Oracle GoldenGate product that applies changes in real-time from the source database to the Exadata database. Since GoldenGate is usable on heterogeneous system architectures, it is the most flexible data replication tool available for the Oracle Database.

Application-based Synchronization

Some of the databases contained historical data that didn't need to be migrated to Exadata. In those cases, the applications that updated the database were modified to insert new data into both the current

database and the Exadata version of that database.

Backup and High-Availability Strategy

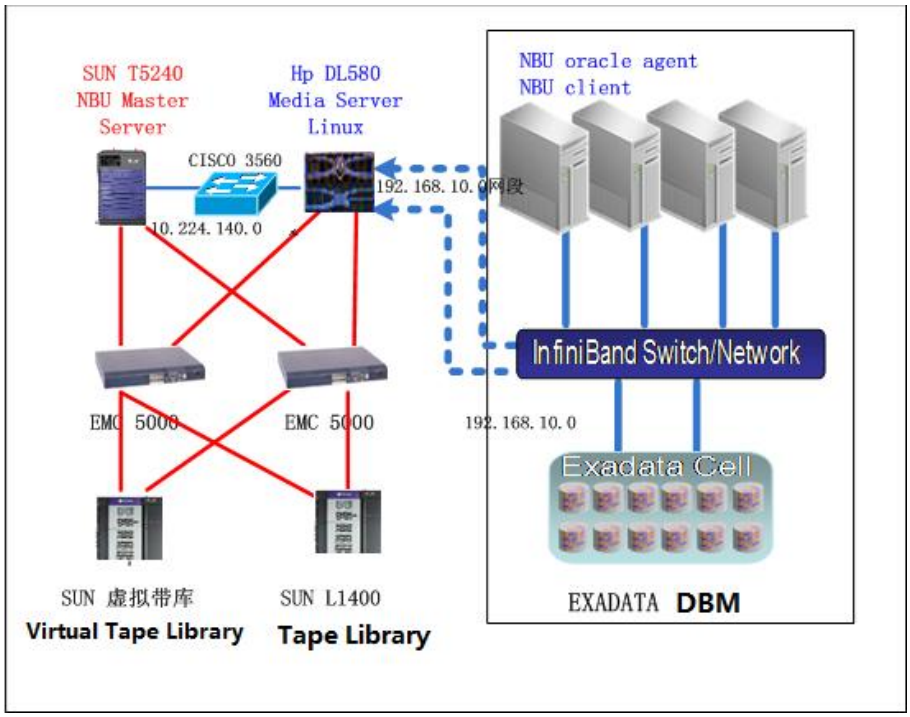

FIGURE 2 – DATABASE BACKUP CONFIGURATION

Of the four NMS databases, only NMSA and NMSB require regular backups, as the other two databases are easily recreated from other data sources. As depicted in Figure 2, NMS databases are backed up directly to tape, whereas other data could use a VTL backup target. NMSA is 1 TB in size and NMSB is 4 TB in size.

A full backup is performed for each of the databases on the weekend, with incremental backups taken daily. All backups are maintained for 2 weeks.

Beijing Mobile maintains an Active Data Guard configuration for failover, in the event of a primary system outage. Thus, backups are maintained only for disaster recovery purposes. Both of the Exadata systems serve as a primary database platform and a standby platform. The Active Data Guard standby databases support a number of read-only applications, including reports, queries and database backups.

The applications used on the ADG standby databases are listed below (see Figure 3):

- EMOS
- Transmission Network System
- Information Portal

- Business Quality Monitoring
- Third Party Operation and Maintenance
- Uniform Collection Platform
- GIS
- Centralized Operation and Maintenance Platform
- Comprehensive Alarm System Report
- Comprehensive Resources System

High-Availability with Active Data Guard

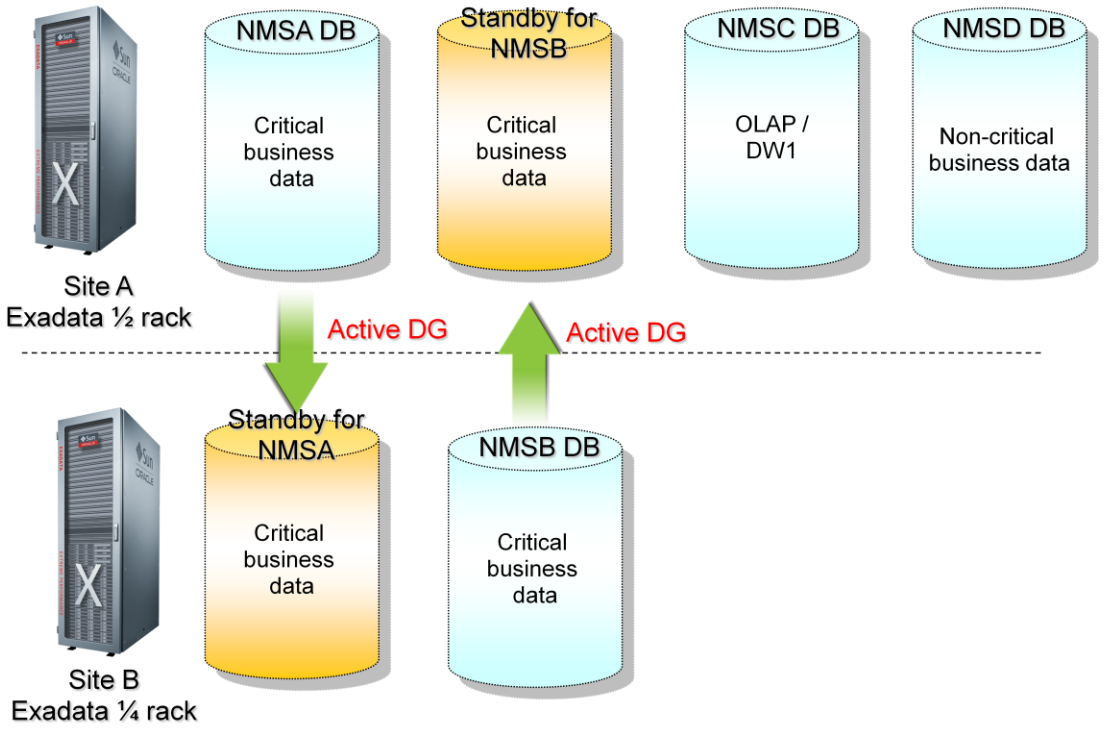

FIGURE 3 – HIGH-AVAILABILITY OVERVIEW

System Monitoring

Beijing Mobile uses Oracle Enterprise Manager 12c to monitor and manage their Exadata environment. The administration team performs three scheduled checkups of Exadata daily. They check for system and database log messages, verify the space usage, check for long-running sessions, and visually inspect the Exadata system status.

The database team generates and analyzes regular database AWR reports. If they discover poor performing SQL statements, they ask the respective team to optimize the SQL statements.

Benefits From Using The Exadata Database Machine

Beijing Mobile's consolidation onto Exadata reduced storage and improved application performance, along with a simpler environment and reduced operating costs.

Prior to Exadata, the 18 databases consumed a total of over 20 terabytes of storage. With Oracle's Hybrid Columnar Compression, most tables that were compressed experienced an average compression of over 8x, and the total storage consumed after migration was roughly 5 terabytes, or one quarter of the storage used by the prior systems.

The reduction in query response times for the applications accessing databases on Exadata averaged between 5x and 8x, making the application users significantly more productive. Similarly, testing of common SQL statements showed improvements of up to 12x shorter execution times.

By consolidating 18 database servers onto two Exadata systems, Beijing Mobile was able to free up approximately 5 full-time administrators, as compared to the prior 18 database servers that required more than 10 administrative staff to manage. The reduction in operating costs (power, cooling, real estate) between 18 separate database servers and storage and two integrated Exadata systems, was also significant. These original 18 database servers and related storage were repurposed for other applications.

Since Beijing Mobile used Exadata to consolidate existing database servers, there was no additional cost for the Oracle Database and related software licenses. Per Oracle's policy, those licenses were transferred to the Exadata system.

Conclusion and Lessons Learned

Beijing Mobile found it easy to consolidate their 18 databases into 4 databases on Exadata, using a variety of options to move the data, depending on speed and uptime needs. Having corporate standards for naming, or a Master Data Management implementation, is recommended for all companies, to prevent significant rework should a consolidation project such as this occur.

Beijing Mobile began this effort with one Exadata Database Machine and added a second Exadata system after going into production, thus enabling a more highly available configuration without any changes to their applications or database structures. The simple addition of Active Data Guard enabled both Exadata systems to perform both primary and standby roles, and enabled much higher availability for both planned maintenance and unplanned outages.

The improvements in performance and storage efficiency were dramatic, and enabled Beijing Mobile to convert 18 previous database servers into two Exadata systems, saving substantial data center and administrator resources. In the future, additional Exadata Database Machines or Exadata Storage Expansion racks can be easily added to extend the Exadata configuration for more compute and/or storage capacity.

Appendix A

- AWR = Automatic Workload Repository
- BMCC = Beijing Mobile Communications Corporation
- BSC = Base Station Controller
- DR = Disaster Recovery
- GIS = Geographical Information System
- HA = High Availability
- HCC = Hybrid Columnar Compression
- MAA = Maximum Availability Architecture
- MMS = Multimedia Messaging Service
- MSC = Mobile Switching Center
- NMS = Network Management System
- OLAP = OnLine Analytical Processing
- OSS = Operational Support System
- RAC = Real Application Clusters
- RMAN = Oracle Recovery Manager
- SLA = Service Level Agreement
- SMS = Short Message Service
- SQL = Structured Query Language
- VLR = Visitor Location Register

BMCC – Oracle Exadata Case Study
February 2013

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

Hardware and Software, Engineered to Work Together