

EMC IT's Migration to the Open, Expandable Oracle BI Grid

Applied Technology

Abstract

This white paper illustrates EMC IT's migration from a legacy BI/DW infrastructure (Fujitsu and Solaris) to an open platform (x86 and Linux), expandable infrastructure (consolidated physical Oracle RAC and Grid) supporting EMC's virtualized BI toolset reporting architecture (VMware, OBIEE, and BIP).

April 2011

Copyright © 2010, 2011 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

VMware and ESX are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h7063.1

EMC IT's Migration to the
Open, Expandable Oracle BI Grid
Applied Technology

Table of Contents

Executive summary	4
Introduction	4
Audience	4
EMC IT BI/DW challenges.....	4
EMC IT's phased migration approach.....	6
Legacy architecture.....	6
Oracle BI Grid data architecture	8
EMC IT's Oracle BI Grid infrastructure architectures	9
BI/DW Grid enabling technology overview	10
Recommended disk configuration	11
Today's method	11
Oracle SQL scripts – Moving the tablespace and database objects	14
Future method.....	15
EMC deployment practices enabling the EMC IT Oracle BI Grid	16
Gold Copy.....	17
Test and development.....	17
Disaster recovery and performance	18
Oracle BI Grid architecture migration highlights	20
Virtualization of the BI toolset infrastructure (2008).....	20
EMC's legacy BI infrastructure	20
EMC's virtualized BI infrastructure	21
EMC's virtualized BI infrastructure (OBIEE and BIP) benefits	23
Pre-Oracle Grid proof of concept/discovery (Q4 2008/Q1 2009)	23
Oracle BI Grid design (Q2 2009)	23
BI Grid build (Q3 2009)	23
BI Grid deploy (Q4 2009)	24
Migration-to-deployment infrastructure	24
Migration-to-deployment process.....	24
Business performance gains	25
Conclusion	25
References	26

Executive summary

EMC is a leader in information infrastructure. Over time EMC, which creates critical reporting for EMC's Sales, Manufacturing, Customer Services, Professional Services and Marketing functions, came to a crossroad with its business intelligence (BI) infrastructure. It was an aging, redundant, and inconsistent architecture. It was missing key business service level objectives with respect to reporting and there were many data sources and no single source of record for some BI functions.

EMC IT started a migration from a closed legacy infrastructure both in hardware and software to an open platform (Linux) with an Oracle RAC deployed to an expandable performant-balanced hardware infrastructure. EMC utilized its enabling and accelerating hardware (Symmetrix VMAX™ storage array) and local and remote replication technologies (TimeFinder® and SRDF®) and implemented EMC Oracle best practices to design, develop, and deploy its Oracle BI Grid.

This white paper will show the migration to the Oracle BI Grid, and how EMC hardware and software were an accelerator and enabler both in the migration and deployment to provide EMC with the following BI improvements with a single Global Data Warehouse (10 TB):

- 10x improvement in Day 1 performance
- Improvement by 180 percent in batch job performance
- Datamart update times were cut in half
- Improvement by 200 percent in dashboard rendering and drill-down performance
- 2x to 3x performance improvement in reporting cube build time

In summary, the new Oracle BI Grid infrastructure is two to three times more performant, reduced time on BI functions by 60 percent, and saved three times the storage footprint.

Introduction

This white paper illustrates EMC IT's approach to a migration from a legacy/closed implementation of a BI/DW infrastructure on a Solaris platform to an open (Linux), expandable platform via an Oracle RAC Grid architecture and support for virtualization (VMware) BI infrastructure/components (OBIEE/BIP).

Audience

This white paper is focused on the CIO, system architect, Oracle architect, storage architect, and supporting staff, focusing on Oracle Applications DBAs, server administrators, and network administrators.

EMC IT BI/DW challenges

EMC, like many large enterprises, has deployed enterprise-scale implementations of Oracle's ERP and CRM solutions to enable its business in Manufacturing, Finance, Quoting, Customer Service, Professional Services, Sales, and Marketing.

EMC's two enterprise-scale mission-critical systems support EMC's core revenue-generating functions and are the sources for their BI/DW infrastructure. Here is a snapshot of the EMC user community who utilizes this infrastructure:

- An ERP solution, supporting 20,000 employees with 2,000 concurrent users
- A CRM solution, supporting 36,000 named users worldwide with 3,500 concurrent users. This implementation is one of the top five Oracle Applications transactional systems in the world, depending on the modules that are used.

With time, EMC's business intelligence and data warehouse grew as business drivers and clients increased and developed the following challenges:

- Redundant data
 - Five major data warehouse environments with a high overlap of data
 - Increased risk of data discrepancies due to multiple code bases
- Redundant infrastructure
 - Multiple hardware and DB environments hosting the same data
 - Increased support and run costs (storage, backup, and so on)
- Inconsistent architecture
 - Inconsistent HA and DR solutions
 - Inconsistent backup/recovery solutions
 - Current storage arrays do not support ILM tiering/best practices
 - Multiple versions of the database, operating system, and the underlying infrastructure
 - Multiple technologies to refresh testing and development environments
- Aging infrastructure
 - Unable to scale and meet increasing business demand
 - Majority of infrastructures are at or near end of service life

The EMC IT team was at a crossroad on how to resolve these challenges. As the TDWI white paper *Next Generation Data Warehouse Platforms* illustrates, "Technical deficiencies can make an old Data Warehouse platform ripe for replacement."

There were "Data Driver" challenges in the legacy data warehouse:

- Reduce data redundancy - EMC business owners needed to get to one Common Data Model versus five data models in the legacy data warehouse architecture
- Quantity of data - EMC business owners needed a more frequent refresh and quality of data versus the legacy data warehouse refresh rate (four times a week)
- Availability of data - EMC needed a disaster recovery infrastructure

EMC IT's approach was a phased migration to an open, expandable BI infrastructure enabled by EMC technologies and best practices with the following components:

- **Open** – Linux/x86 servers to replace legacy Solaris and legacy SPARC hardware
- **Expandability** – An Oracle RAC physical deployment to create the BI Grid that has the ability to grow via scaling out (adding RAC nodes)

This Oracle BI Grid is the source for the virtualized (VMware) BI infrastructure of Oracle OBIEE and BIP, which utilized the physical Oracle BI Grid infrastructure (see the "Oracle BI Grid architecture migration highlights" section).

The following EMC components are the Oracle BI Grid accelerators/enablers for the phased deployment and operational effectiveness for EMC's BI Grid production, test and development, and disaster recovery and performance infrastructures:

- EMC Symmetrix VMAX
- EMC Enterprise Flash Drives (EFD)
- EMC Virtual LUN with Auto-provisioning
- EMC TimeFinder/Clone (BCV) and /Snap
- EMC SRDF

-
- EMC Disk Library (EDL)
 - EMC NetWorker®
 - EMC PowerPath®
 - EMC Celerra® (NAS)
 - EMC Ionix™ ControlCenter®
 - EMC Smarts®

EMC IT's phased migration approach

The following section identifies the EMC legacy infrastructure of the Integrated Data Warehouses (IDW).

Legacy architecture

The legacy IDWs provided decision-support data to Finance, Sales, Marketing, Customer Service, and Manufacturing business users.

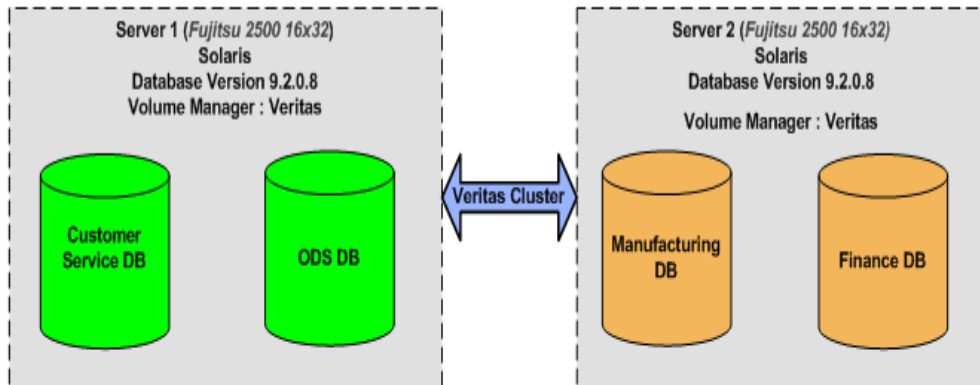
The following four legacy IDW instances plus one Near Real Time instance were supporting the following EMC business environments/functions:

- **Operational Data Store** - IDW extracts data from several main sources like Oracle Applications (Financial, CRM) and writes to a common database, which is termed an Operational Data Store. This data store is a subject-oriented, integrated, frequently augmented store of detailed data in the Enterprise Data Warehouse.
- **Finance** – The financial datamart contains history, details, and summary data of EMC-sensitive financial information that is used by the Finance, Sales, and Marketing groups.
- **Manufacturing** – The manufacturing datamart contains history, detail, and summary data of EMC products and manufacturing areas that is used by Global Manufacturing users.
- **Customer Service** - Customer Service currently contains Change Control Automation data. It contains and tracks the customer's microcode level of each EMC array.
- **Near Real Time** – NRT is a revenue and accounting system. NRT systems support analytical and operational reporting for primarily the Sales and Finance functions.

The following figure and table illustrate these challenges with regard to the following:

- Redundant infrastructure
- Inconsistent architecture
- Aging infrastructure

IDW Production Environment



**NRT Production Environment
Oracle Database version 9.2.0.8, Solaris**




Figure 1. EMC legacy data warehouse infrastructure

The following table shows the EMC legacy components.

Table 1. EMC legacy data warehouse components

Technology	Components
Database	Oracle 9.2.0.8
File management	Veritas vxfs
Character set	WE8ISO8859P1
Operating system	Solaris 5.10
Database server	Fujitsu 2500 with two High Availability Domains
Storage array	EMC Symmetrix DMX™ 200P

The above components reflect what EMC IT was experiencing in the legacy deployment:

- Six aging servers to support BI/DW
- A dated storage array
- No disaster recovery
- Older version of Oracle and the language set was not flexible

Oracle BI Grid data architecture

The following is a high-level diagram that illustrates the benefits for EMC of moving from four separate data warehouse instances and one Near Real Time (Financial) instance to one Global Data Warehouse (GDW)/ Oracle BI Grid:

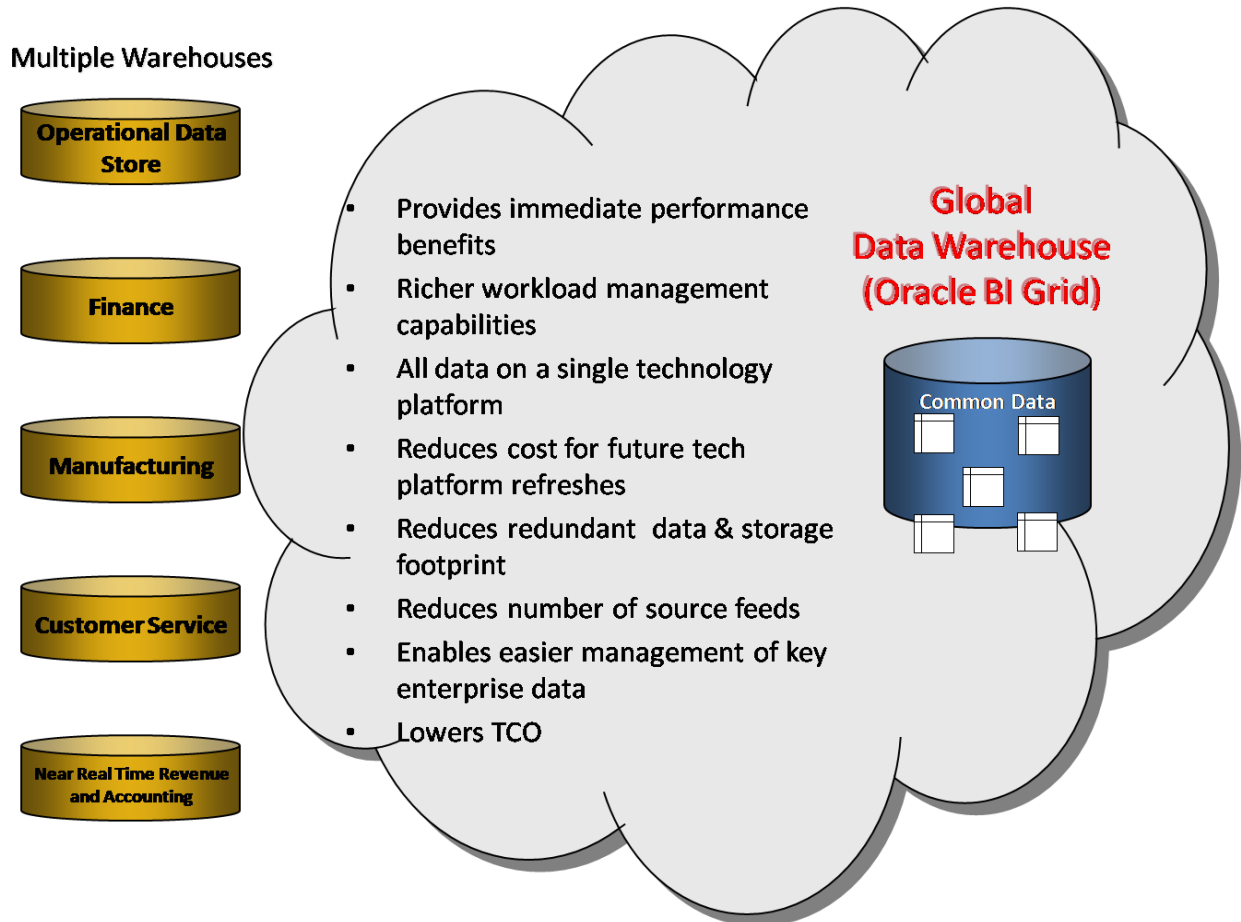


Figure 2. Oracle BI Grid business drivers

As part of the BI Grid integration, four data warehouses each running on separate Solaris-based servers were merged into one data warehouse running on a six-node Linux-based Oracle Real Application Cluster. Additionally, several financial data warehouses that also resided on individual servers were brought into the Oracle BI Grid.

The immediate benefits were the following:

- A common, flexible, easily expandable and high-performing database environment
- Utilization of a tiered ILM storage strategy that used Enterprise Flash Drives, and Fibre Channel and SATA disk drives on Symmetrix VMAX

The foundation was in place to proceed with follow-up efforts to consolidate common and redundant data and to merge customer service and human resource data warehouses.

EMC IT's Oracle BI Grid infrastructure architectures

The following is a high-level illustration of the EMC Oracle BI Grid infrastructure architectures with all four deployments:

- Production
- Disaster recovery
- Performance
- Test and development

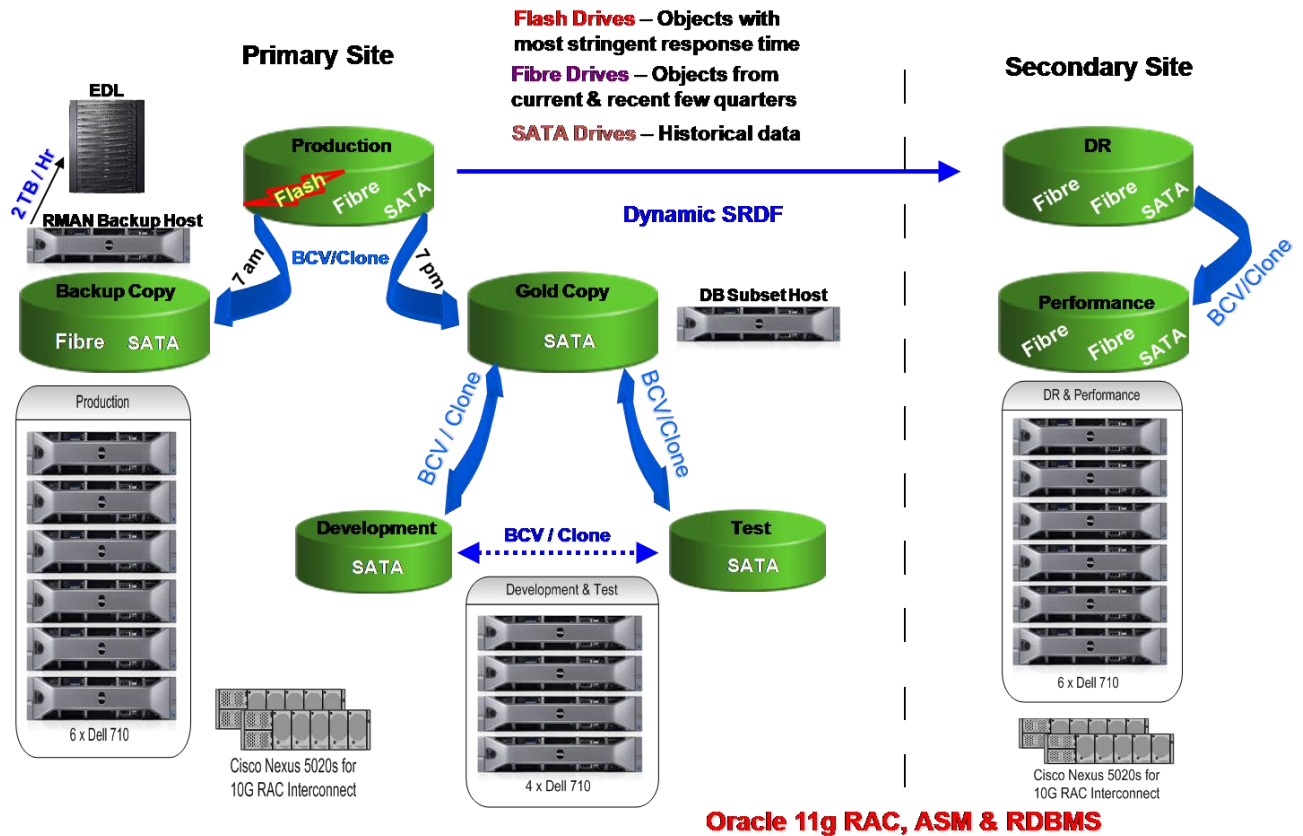


Figure 3. EMC's Oracle BI Grid four deployment infrastructures

Each blue arrow represents the accelerators/enablers of EMC software and hardware that made these architectures efficient and agile for EMC. They will be highlighted in the “BI/DW Grid enabling technology overview” section.

The following table illustrates the technologies used in the construction and deployment of EMC IT's BI Grid:

Table 2. EMC Oracle BI Grid data warehouse components

Technology	Components
Oracle software	Oracle Database 11.1.0.7 <ul style="list-style-type: none"> • RAC • ASM • Partitioning
Character set	AL32UTF8, which enables multiple languages

Operating system	Red Hat Enterprise Linux 5.3
Database servers	Dell servers (Today R710) 6 - Production 6 – Disaster recovery/performance testing 4 – Test and development RAC interconnect 10g Ethernet – Cisco Nexus 5020
Storage array	EMC Symmetrix VMAX with EFDs, FC, and SATA
EMC software	EMC Virtual LUN with Auto-provisioning EMC TimeFinder/Clone and /Snap EMC SRDF EMC Disk Library (EDL) EMC NetWorker EMC PowerPath EMC Celerra (NAS) EMC Ionix ControlCenter EMC Smarts

BI/DW Grid enabling technology overview

The following sections detail the enabling EMC technology components used in this migration to the Oracle BI Grid that accelerated the building, deploying, and ongoing go-live operations.

The diagram illustrates the enabling/accelerating EMC components used in production, test and development, disaster recovery, and performance.

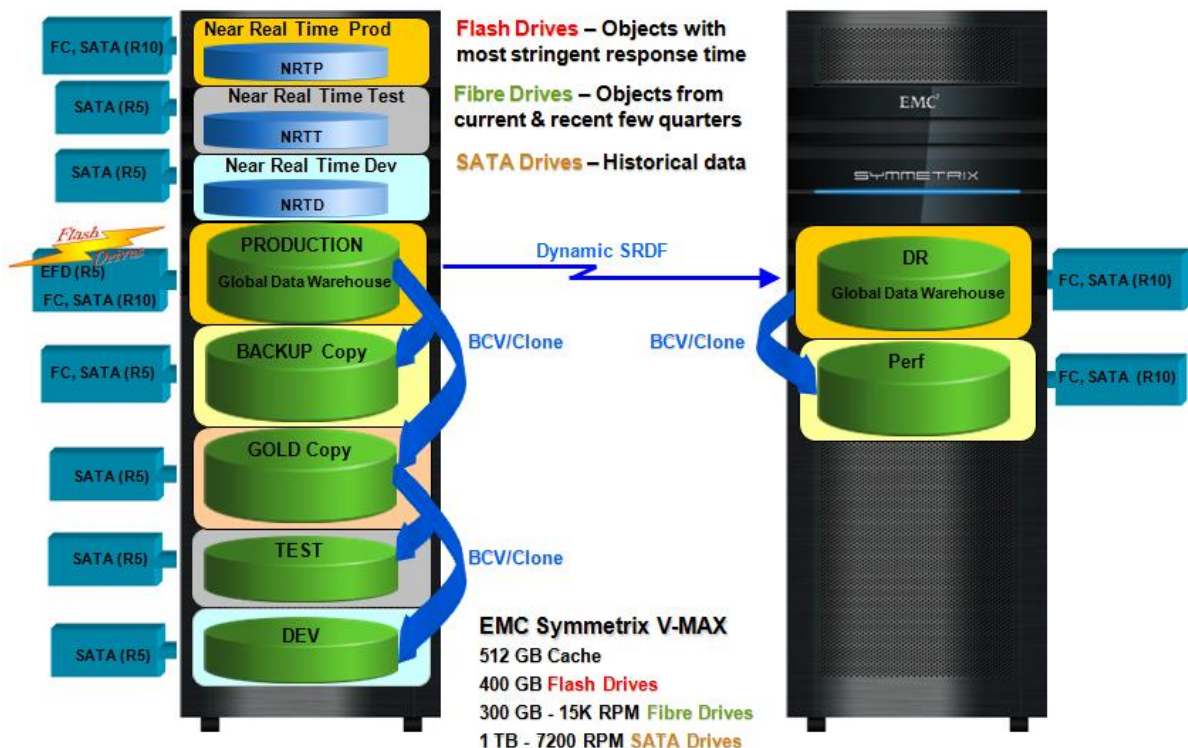


Figure 4. EMC's Oracle BI Grid production enabling technologies

Recommended disk configuration

The figure also illustrates the recommended disk configuration (type and RAID) that creates the needed performance for EMC IT's Oracle BI Grid.

A Symmetrix VMAX was used. Symmetrix VMAX storage arrays accelerate and enhance the ability for the BI Grid to do "storage tiering." Storage tiering is the approach in a multi-typed storage configuration (Enterprise Flash Drive, Fibre Channel, and SATA) to place the right data on the right storage type at the right time.

Today, the Oracle BI Grid is using the following storage types and disk configurations:

- Enterprise Flash Drives with RAID 5 were used for database objects with the most stringent response times
- Fibre Channel drives with RAID 10 were used for objects from current and recent few quarters
- SATA drives with RAID 10 were used for historical data

Today's method

Today, EMC's Oracle BI Grid utilizes ZettaPoint's DBclassify tool and EMC IT's data warehouse operational knowledge to identify which database objects should reside on which tier, and uses Oracle SQL scripts to move the data to the right tier.

Within the Oracle BI Grid, ZettaPoint's DBclassify accelerates and enables the identification of database objects that are candidates to be moved from one storage type/tier to another (that is, from Fibre to EFD).

DBclassify has the ability to take into consideration the actual, detailed data access and usage characteristics of Oracle database workloads, and then classify the data by its usage profile or "temperature." Only data that is classified as mission-critical ("hot") and is frequently accessed by high-value applications and users should be stored on Enterprise Flash Drives (EFD), or Tier 0. Data that is less frequently accessed but still important ("warm") with less stringent service level requirements should be relegated to lower-performing, higher-capacity storage device Fibre Channel, or Tier 1. Dormant data ("cold"), which is typically historical or static, may be moved to SATA archival storage, or Tier 2. The benefits of object-level analysis include minimizing I/O, assuring proper asset utilization, and leveraging new technologies like Flash and platforms such as EMC Symmetrix VMAX.

Using DBclassify to specifically identify which objects should be migrated to Tier 0 is much more efficient than using tablespace or LUN-level measurements. Tablespaces and LUN are more likely to contain many objects, but not all of these objects are I/O bound and, therefore, may not be good candidates for occupying the most expensive storage tier. Moving the entire tablespace to Tier 0 can be wasteful but moving only the objects that cause the I/O bottleneck minimizes the probability of wasting space on a relatively costly Tier 0 device (EFD).

DBclassify is completely transparent to the application so there is no need to change application code. Migrating objects to Tier 0 can be done by creating a dedicated tablespace located on Flash and then using a few Oracle commands to move them to the new tablespace.

The DBclassify "Object Analysis" report provides recommendations regarding the proper object storage tier placement based on I/O characteristics. The "Report Filter" illustration below is used to focus on activity coming from interactive users and give more priority to the objects they are accessing.

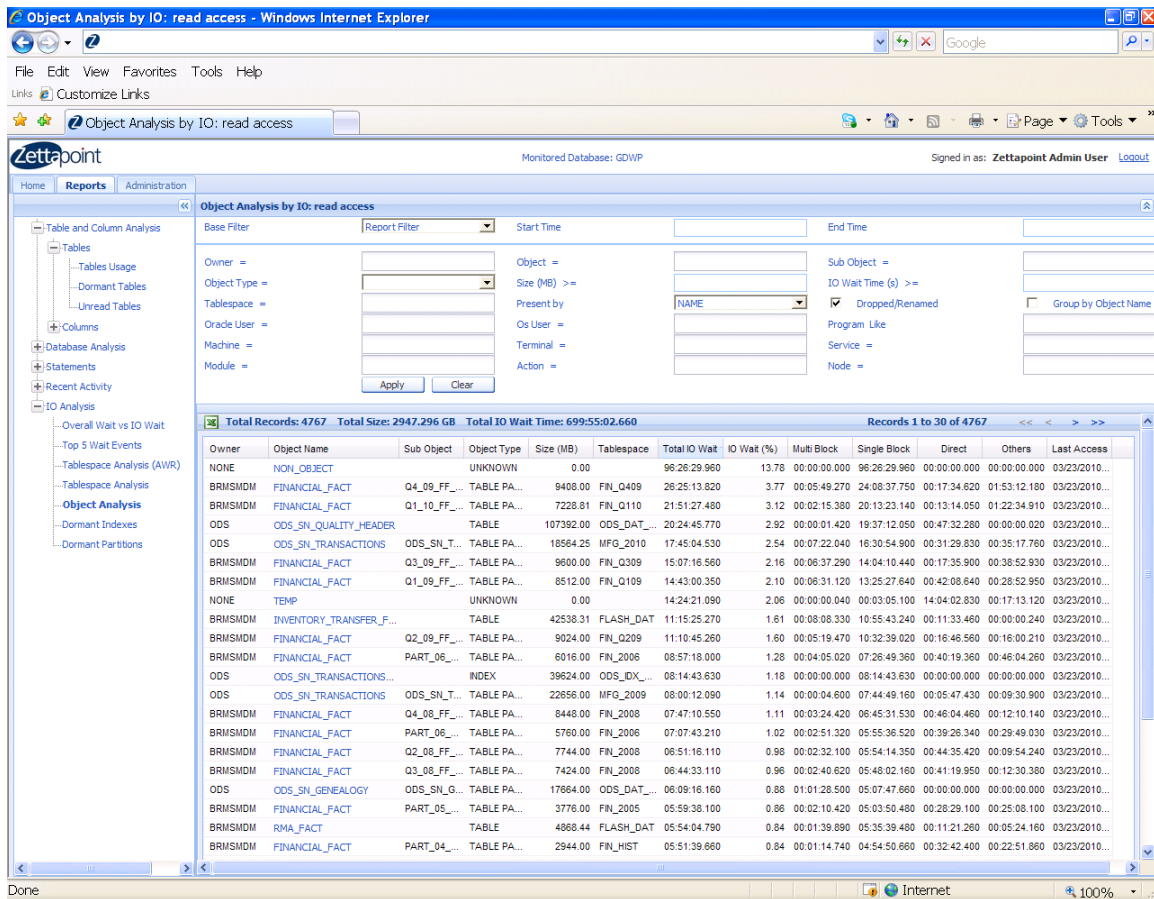


Figure 5. Zettapoint's DBclassify Object Analysis report

Using DBclassify in the Oracle BI Grid deployment, tablespace analysis was done to better understand current tiering versus optimal tiering.

The next diagram illustrates the ability to create a report filter on the current tablespace usage:

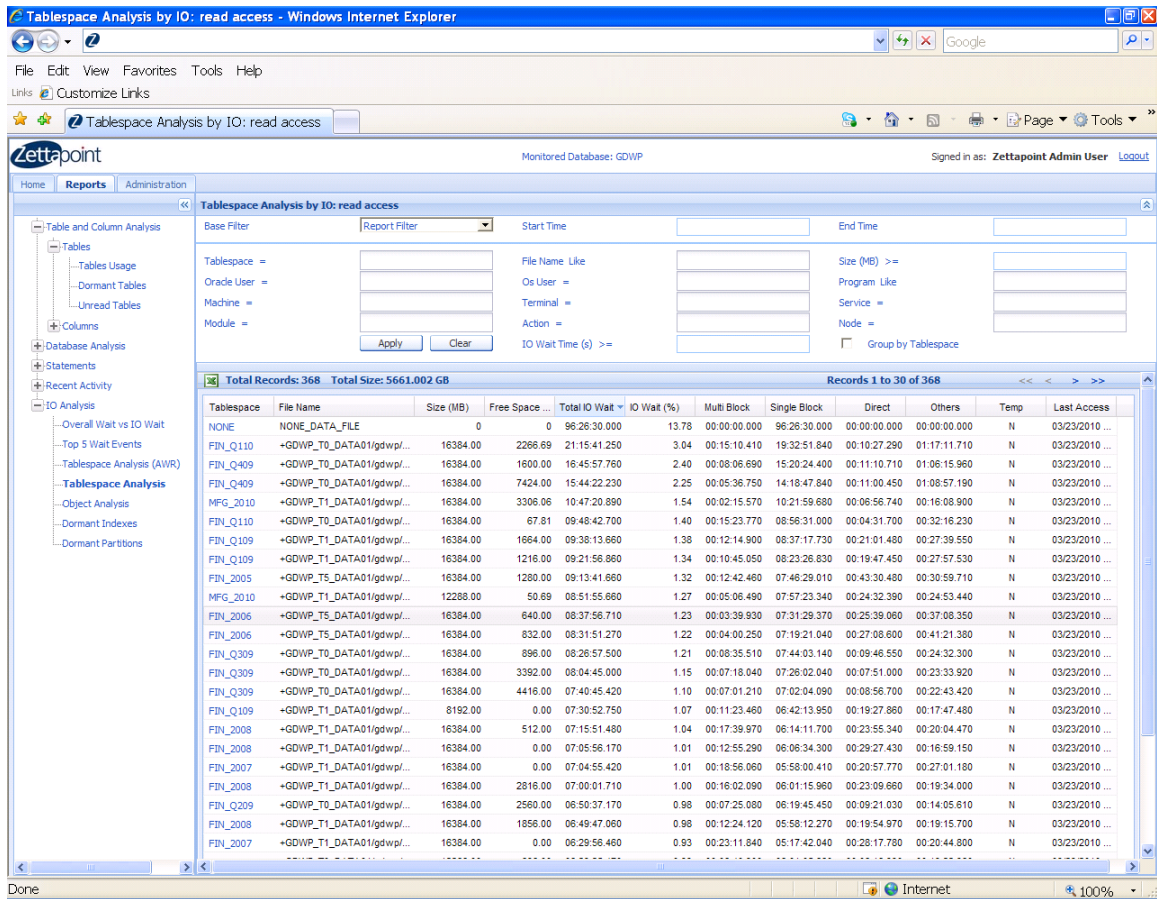


Figure 6. Zettapoint's DBclassify Tablespace Analysis report

As deployed in the Oracle BI Grid, DBclassify enables the following:

- Relates the usage of database objects (tables, columns, indexes, partitions) and the I/O profile to the specific processes and users who utilize them
- Identify who is using the data, how, and in what frequency
- Identify what objects should be placed on what storage tiers and the impact on business
- Identify dormant objects that are not being used at all and can be removed from the ETL process

Oracle SQL scripts – Moving the tablespace and database objects

The following “pseudo code” scripts show the methods to move the following:

- Tablespace
- Database objects

Tablespace

The following is an example of moving a tablespace using the object’s identification by ZettaPoint’s DBclassify.

Moving Tablespace FIN_Q409 to Flash disks from Fibre Channel disks

In this example the FIN_Q409 tablespace contains tables with the financial information of Q4 2009, which is very critical for year-end financial report generation. It was moved to Enterprise Flash Drives (EFDs) to meet the demands of faster response times.

1. Create directories of the source and destination to move datafiles. Here, the source is Fibre Channel disks ASM diskgroup GDWT_T1_DATA01, and the destination is the Flash disks’ diskgroup (GDWT_T0_DATA01).

```
SQL> create or replace directory asm_source as '+GDWT_T1_DATA01/gdwt/datafile';
SQL> create or replace directory asm_destn as '+GDWT_T0_DATA01/gdwt/datafile';
```

2. Bring the tablespace FIN_Q409 offline.

```
SQL> ALTER TABLESPACE FIN_Q409 OFFLINE;
```

3. Move the datafile using the DBMS_FILE_TRANSFER.COPY_FILE procedure.

```
SQL> BEGIN
      DBMS_FILE_TRANSFER.COPY_FILE(
        source_directory_object => 'asm_source',
        source_file_name => 'fin_q409.633.698768089',
        destination_directory_object => 'asm_destn',
        destination_file_name => 'fin_q4091');
    END;
/
```

4. Rename the datafile (its new location on Flash disks).

```
SQL> ALTER DATABASE RENAME FILE
'+GDWT_T1_DATA01/gdwt/datafile/fin_q409.633.698768089' TO
'+GDWT_T0_DATA01/gdwt/datafile/fin_q4091';
```

5. Bring the datafile online.

```
SQL> ALTER DATABASE DATAFILE '+GDWT_T0_DATA01/gdwt/datafile/fin_q4091'
ONLINE;
```

6. Verify the location and status of the datafile.

```
SQL> SELECT name ,status from v$datafile WHERE
name='+GDWT_T0_DATA01/gdwt/datafile/fin_q4091';
```

Note: Tablespace FIN_Q409 had only one datafile. Repeat steps 3, 4, 5, and 6 for more datafiles.

7. Bring Tablespace FIN_Q409 online.

```
SQL> ALTER TABLESPACE FIN_Q409 ONLINE;
```

Repeat steps 1 to 7 if more tablespaces need to be moved.

Database objects

The following is example “pseudo code” to move individual objects:

- Move an individual table from Fibre Channel to Flash as below:

```
alter table <table_name> move tablespace flash_dat;
```
- Move an individual index from Fibre Channel to Flash as below:

```
alter index <index_name> rebuild tablespace flash_indx;
```

Future method

In the future the Oracle BI Grid will utilize the EMC technologies Virtual LUN and Fully Automated Storage Tiering (FAST) with DBclassify.

Virtual LUN migration technology provides users with the ability to move Symmetrix VMAX logical devices between drive types, such as EFDs, Fibre Channel drives, or SATA drives, and at the same time change their RAID protection.

Virtual LUN migration occurs independent of host operating systems or applications, and during the migration the devices remain fully accessible to database transactions. Virtual LUN migration is fully integrated with Symmetrix replication technology and maintains consistency of source/target device relationships in replications such as SRDF, TimeFinder/Clone, TimeFinder/Snap, or Open Replicator.

Virtual LUN migration assists customers to implement an Information Lifecycle Management (ILM) strategy for their databases, such as the move of the entire database, tablespaces, partitions, or ASM diskgroups between storage types. It also allows changes in service levels and performance requirements to application data. For example, often application storage is provisioned before clear performance requirements are known. At a later time, once the requirements are better understood, it is easy to make any adjustment to increase user experience and ROI using the correct storage type.

EMC FAST (introduced in the Engenuity™ 5874 Q4 service release) is Symmetrix software that utilizes intelligent algorithms to continuously analyze device I/O activity and generate plans for migrating devices for the purposes of allocating or re-allocating application data across different performance storage types within a Symmetrix array.

FAST proactively monitors workloads at the Symmetrix device (LUN) level in order to identify “busy” devices that would benefit from being moved to higher-performing drives such as EFD. FAST will also identify less “busy” devices that could be relocated to higher-capacity, more cost-effective storage such as SATA drives without altering performance.

EMC deployment practices enabling the EMC IT Oracle BI Grid

The following were the EMC deployment practices and components used in the production, test and development, and disaster recovery and performance site architectures.

Production

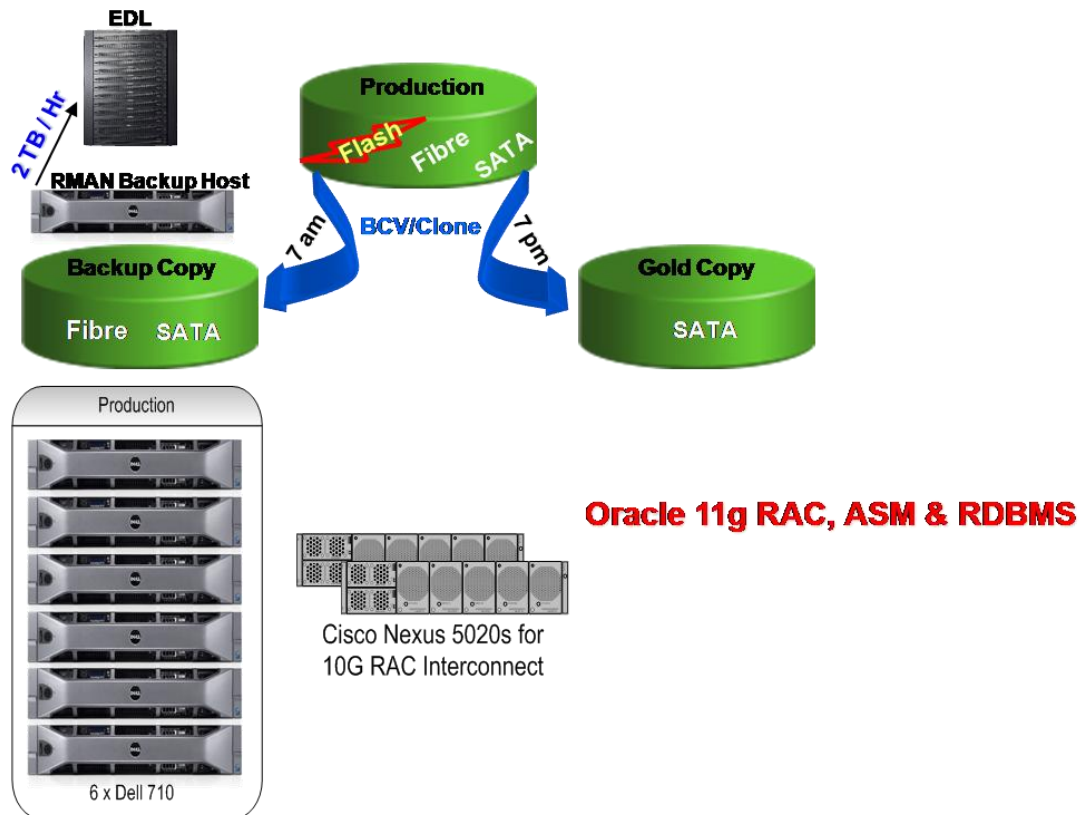


Figure 7. EMC's Oracle BI Grid production deployment model

Production was enabled by the following EMC technologies and best practices.

The diagram illustrates that by deploying a dedicated backup infrastructure the EMC IT Oracle BI Grid is able to meet the SLA defined by the business organizations that it could not in the legacy deployment. The EMC white paper *Reducing Backup Window and Recovery Time with Oracle Database 11g RMAN and EMC TimeFinder/Clone* has more information.

A backup clone copy via EMC TimeFinder technology is created every day starting at 7 a.m. that takes about 17 to 20 minutes for a 10 TB database. This includes unmounting ASM diskgroups on the backup host, putting the production database in hot backup mode, activating the clone (incremental), and then taking the database out of hot backup mode, mounting back the ASM diskgroups, and so on for single-node backup host.

RMAN database backup is executed from the backup node, which acts as a storage node for EMC Disk Library (EDL). The backup server has two dedicated dual-port HBAs that are connected to EDL. This configuration provides 2 TB/hour throughput.

-
- The archive log gets backed up to EDL directly from a production host.
 - Fibre Channel disks and SATA disks are used.

Gold Copy

The use of “Gold Copy architecture” of the GDW instance via EMC TimeFinder technology creates rapid clones, an agile provision process that enables the following environments: production, test, development, DR, and performance.

TimeFinder allows users to nondisruptively create and manage point-in-time copies of data (local replication). This allows operational processes, such as backup, reporting, and application testing, to be performed independently of the source application to maximize service levels, without impacting performance or availability.

TimeFinder/Clone was used in this use case. It creates highly functional, high-performance, pointer-based, full-volume copies of Symmetrix DMX™ volumes that can be used as point-in-time copies for data warehouse refreshes, backups, online restores, and even volume migrations.

There are six purposes for a Gold Copy host (100% SATA disks):

- A production clone is created daily in the evening (7 p.m.). This acts as another recovery point for the production database.
- Test and development databases are refreshed from the Gold Copy using EMC Symmetrix/VMAX clone technology.
- The Gold Copy is used as a surgical restore place for the test and development database.
- Production database maintenance tasks are tested and validated prior to production implementation by DBAs.
- Data masking is implemented prior to test and development refresh (future plan).
- Data subsetting is implemented prior to test and development refresh (future plan).

Test and development

Test and development environments were enabled by the following EMC technologies and deployment practices:

- 100% deployment of SATA disks in this environment
- Development clone/refresh is enabled by TimeFinder

Development/Test

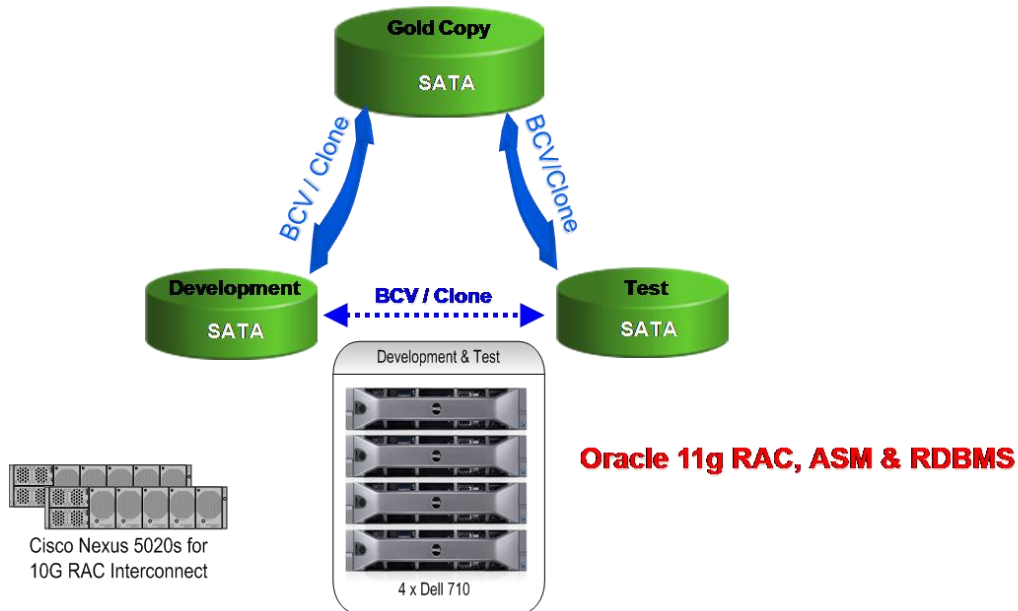


Figure 8. EMC's Oracle BI Grid test/development deployment model

Disaster recovery and performance

Disaster recovery and performance environments were enabled by the following EMC technologies and deployment practices:

- Production database redo and archive log diskgroups are Dynamic SRDF to the DR site.
- Data diskgroups are initially RDF'ed then hourly archive roll forward is used to keep the DR in sync with production. This is accomplished by doing a snap of the archive log diskgroup at DR then presenting the snap to the DR host.
- The performance database is a clone from DR.

Disaster Recovery/Performance

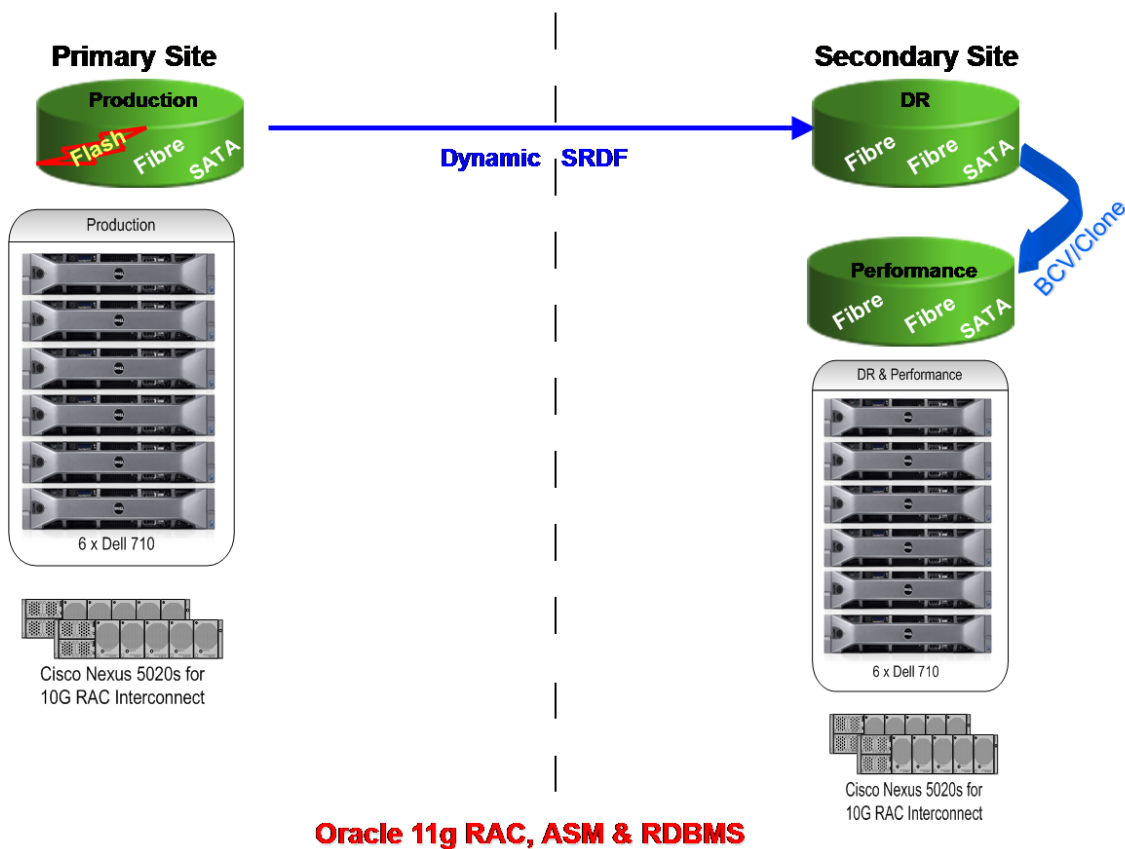


Figure 9. EMC's Oracle BI Grid disaster recovery/performance deployment model

Dynamic SRDF used as a disaster recovery deployment practice

EMC SRDF is remote mirroring software, which can be used to create a remote copy of the customer's entire database and application environment. SRDF can be used to synchronously ship all customer information on a write-by-write basis to the target location (Synchronous SRDF or SRDF/S). It can also be used asynchronously (SRDF/A) to provide Oracle restartable images of the database, utilizing reduced network bandwidth (as compared with SRDF/S solutions). SRDF/A does not have a latency overhead on the database regardless of database size or replication distance.

Oracle BI Grid architecture migration highlights

Enabling an open, expandable BI/DW to support a leading information vendor is done in a phased approach to see how the new technology, new infrastructures (physical/database), and operational process can be improved. This was a yearlong activity from pre-POC to migration to production go-live.

The following are the highlights of the migration to the EMC IT Oracle BI Grid.

Virtualization of the BI toolset infrastructure (2008)

A precursor project before the migration of the Oracle BI Grid was the virtualization of the BI infrastructure.

This section briefly illustrates the previous infrastructures, the post-virtualized BI infrastructure that utilizes the Oracle BI Grid (physical) deployment, and the benefits to moving from a physical BI infrastructure to a virtual BI infrastructure.

EMC's legacy BI infrastructure

The legacy BI infrastructure used two physical environments: using Brio version 6 and Brio version 8.

- Brio version 6 environment. As shown in Figure 10 this deployment was comprised of over 25 servers.

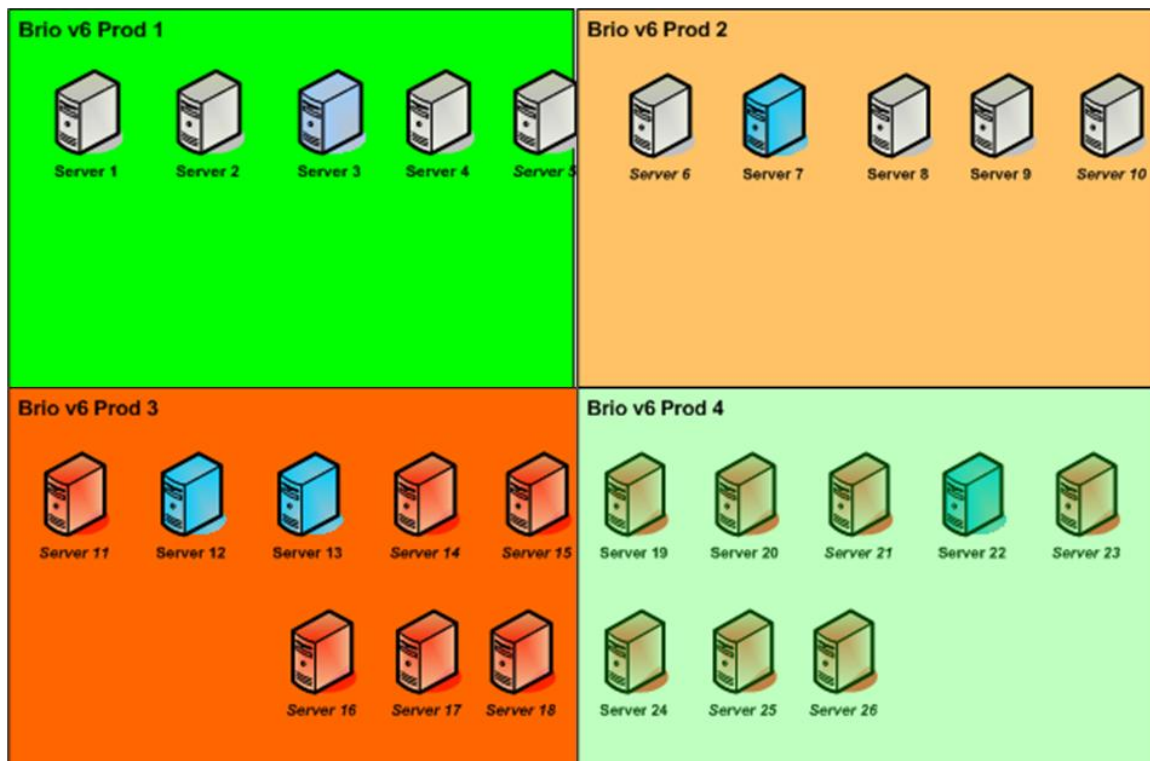


Figure 10. EMC's legacy BI Brio version 6 infrastructure

- Brio version 8 environment. As shown in Figure 11, this deployment was comprised of eight servers.

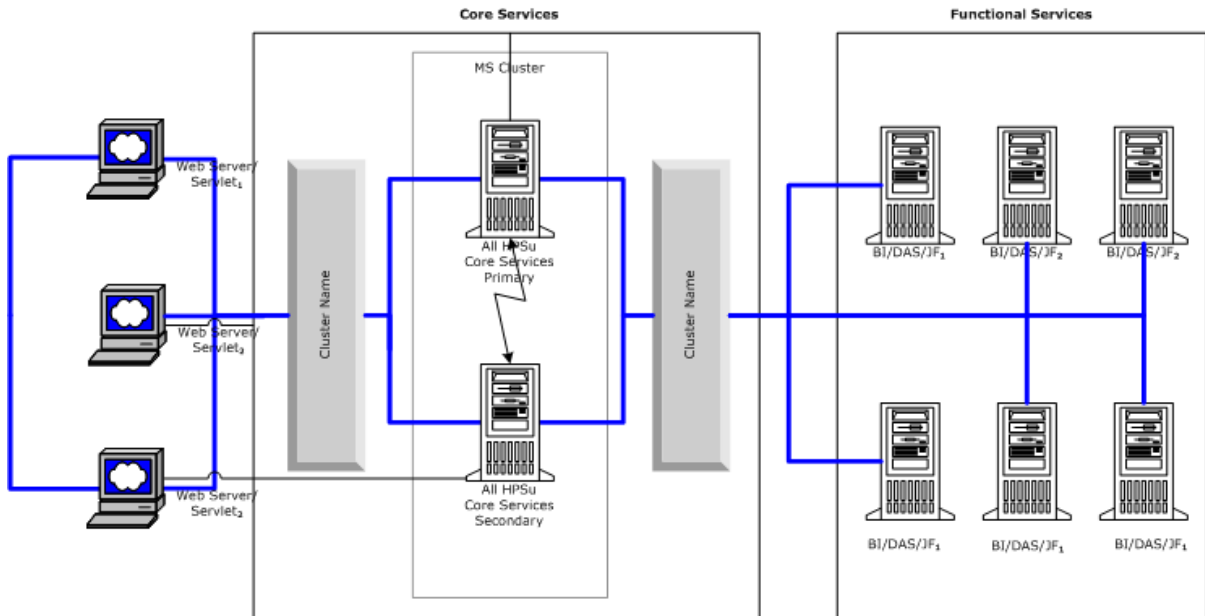


Figure 11. EMC's legacy BI Brio version 8 infrastructure

EMC's virtualized BI infrastructure

The method to size the new virtualized infrastructure first used the standard Oracle OBIEE sizing tool. It recommended 17 physical servers. EMC asked VMware what would the sizing model would be to implement the OBIEE/BIP infrastructure. VMware recommended 12-15 virtual machines (VM) of a size of two vCPUs and 4 GB virtual memory. The POC that was conducted showed a VM of two vCPUs and 4 GB of memory. The production-deployed infrastructure is eight VMs for the Oracle OBIEE infrastructure and four VMs for the Business Intelligence Publisher (BIP). It is deployed on VMware® ESX® Server 3.5. The VM have two vCPUs and 4 GB of virtual memory.

Virtualized (VMware) OBIEE and BIP production infrastructure

Figure 12 shows the virtualized OBIEE and BIP production deployment across two physical data centers. The OBIEE and BI applications use load-balancing appliances. The deployment is also utilizing OBIEE clustering and BIP clustering to provide availability. Additionally the ESX servers host other virtualized applications, therefore OBIEE and BIP are deployed in a multitenant ESX infrastructure.

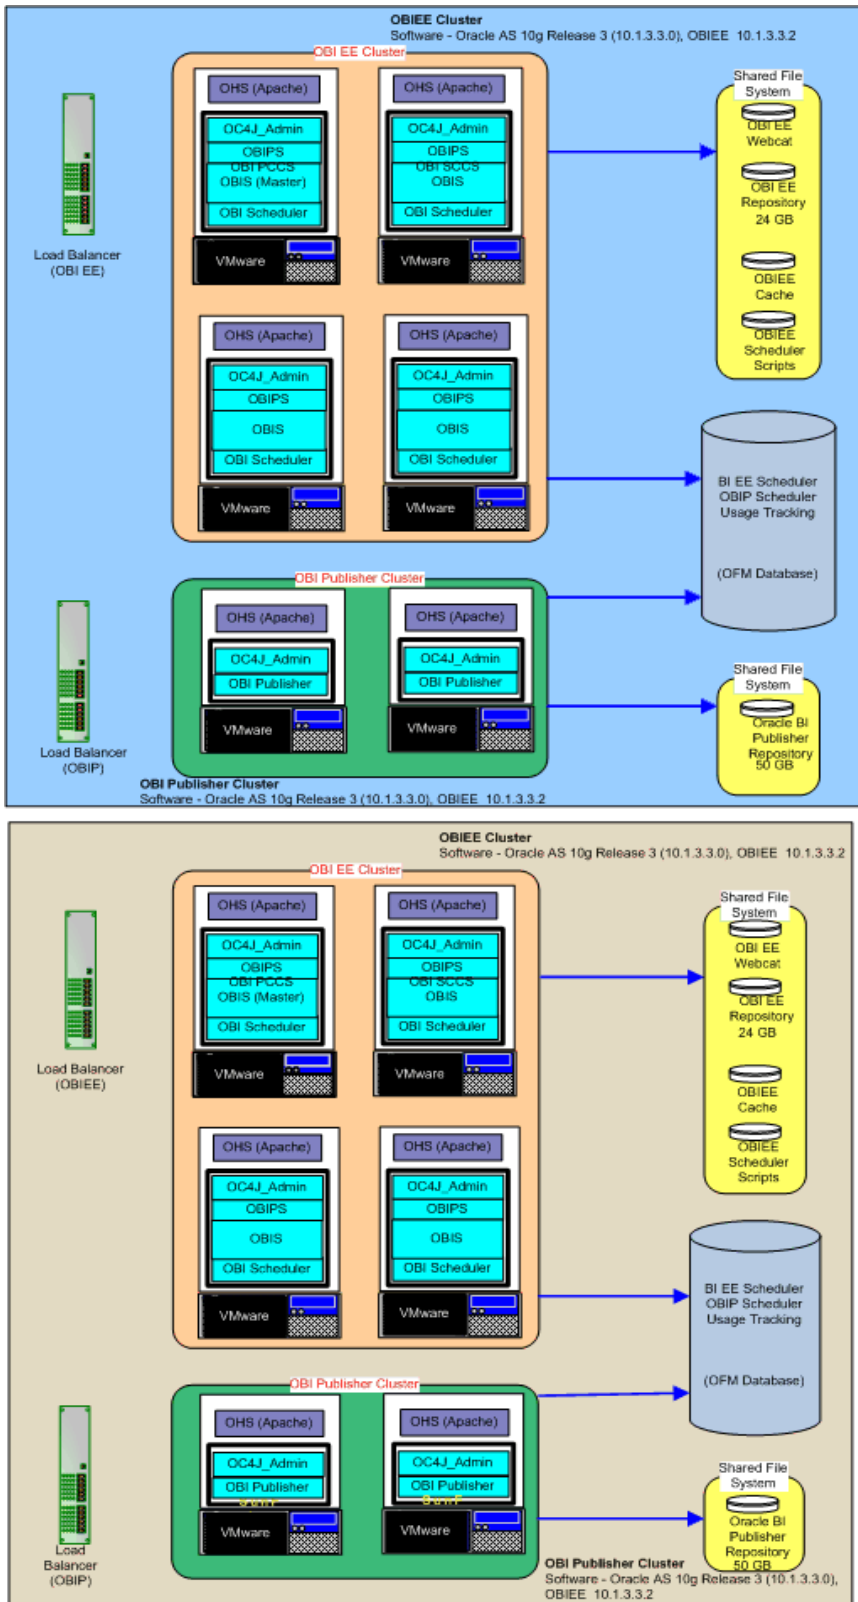


Figure 12. Virtualized (VMware) OBIEE and BIP production infrastructure

EMC's virtualized BI infrastructure (OBIEE and BIP) benefits

The benefits for the virtualized (VMware) BI Infrastructure are the following:

- Server procurement, maintenance, and associated operating system licensing avoidance for eight physical servers
- Thirty percent reduction in FTE support of the virtualized BI infrastructure
- “Cloud-like” provisioning ability enabled by VMware
- No issue with OBIEE or BIP virtualized

Pre-Oracle Grid proof of concept/discovery (Q4 2008/Q1 2009)

Working with legacy technology the EMC IT BI Grid team needed to do a “new technology baseline.” This baseline was the following.

- Migration from a legacy infrastructure to an open one (Sun to Linux). The BI Grid team used Dell white papers to migrate from UNIX to Linux for the migration from Solaris/SPARC to Red Hat/x86 (see the “References” section).
- Migration to new Oracle technologies (11g/RAC/ASM, 10 Gigabit Ethernet, RAC interconnect).
- The POC storage layout used was from *Oracle Data Warehouse Sizing with EMC Symmetrix DMX-4 and Dell R900* (see the “References” section).
- POC migration steps were taken from the EMC white paper *Cross-Platform Oracle Database Migration Using Transportable Tablespaces and EMC Open Replicator for Symmetrix* (see the “References” section).
- The BI Grid team configured Oracle 11g RAC with ASM.

The baseline of knowledge and experience instilled in the team that this was the correct infrastructure to deploy.

Oracle BI Grid design (Q2 2009)

During the planning and design of the BI Grid, a POC took place that gave IT exposure, experience, and insight into technologies and architecture proposed but not yet “EMC IT proven.”

The POC aligned three groups: EMC IT, EMC Engineering, and Oracle. The IT group wanted to understand and evaluate their current customer service data warehouse workload running on an 11g RAC/Linux environment with Oracle's Automatic Storage Management (ASM).

EMC IT's goal was to uncover potential issues or “show-stoppers” and gauge the performance, scalability, and high availability benefits of this new architecture. The EMC Engineering group wanted a real warehouse experience to reinforce and refine best practice recommendations regarding ASM tuning and additionally to evaluate the benefits of using EFDs for key hot objects in a data warehouse. Oracle spent considerable time working closely with EMC to set up its Real Application Testing Tool (RAT) for this project. RAT captures actual workload from a 10g or 11g database and allows replay of that workload on another 11g database.

BI Grid build (Q3 2009)

EMC IT and Engineering worked jointly to migrate, test, evaluate, and present findings and recommendations. These recommendations highlight an Oracle Optimized Warehouse story that is based on the principle to scale out rather than scale up. EMC Flash technology was used to show how a data warehouse can benefit from the use of Flash drives on isolated hot objects in the database. Oracle ASM tuning was introduced to evaluate and further strengthen tuning recommendations already articulated by EMC Engineering. Oracle's RAT 11g feature was used to “capture” workload from the existing 10g customer service data warehouse and “replay” that workload on the 11g RAC POC.

BI Grid deploy (Q4 2009)

Migration-to-deployment infrastructure

The Global Data Warehouse (GDW) was designed using an ILM storage tiering strategy enabled with a Symmetrix VMAX array, with each tier layered according to EMC best practices for a data warehouse. Data was imported into the appropriate tier according to performance characteristics and the significance of the data. Once the data was fully imported, a “SWAT team” validated the data while the data was replicated using SRDF to the production cluster, which was in a different data center.

Migration-to-deployment process

1. Establish and split TimeFinder Business Continuity Volume (BCV) clones from standard volumes from the legacy array (Symmetrix 2000p)
2. Mount BCV volumes on the BCV host and start up the databases
3. Export the data through pipe (|), which compressed, ssh'd, and uncompressed at the target site
4. Import the data into the new database
5. Replicate the data using SRDF from the new DR Oracle BI Grid infrastructure to the new Oracle BI Grid production infrastructure

The following diagram illustrates the process/steps that were used to migrate the legacy data warehouse to the new Oracle BI Grid data warehouse:

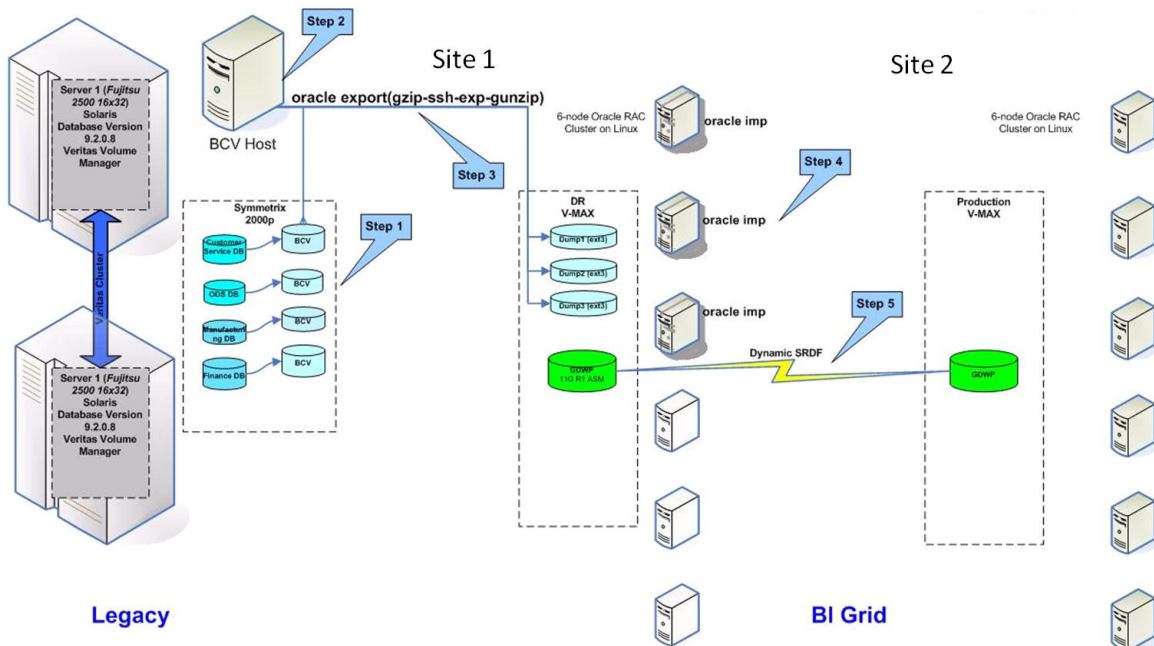


Figure 13. EMC Oracle BI Grid – Migration from legacy to production

Business performance gains

Since the migration to the Oracle BI Grid, tremendous performance gains have been realized across all measurable areas. The following are some high-level business performance improvements:

- Business operations and quarter-end service level objectives (SLO)
 - Before: Frequently missed the 8 a.m. SLO. Often as late as 4 or 5 p.m. on Day 1
 - After: Have not missed the 8 a.m. SLO and typically have beat it by more than three hours
- Oracle BI Grid improvements in batch data loads
 - Reduction in source-to-GDW load times of 50 percent to 75 percent
 - Reductions in cube build times of 75 percent to 95 percent. One cube dropped from three hours to hour minutes.
- For BI reporting, report and dashboard rendering time was reduced up to 78 percent

Conclusion

The migration from a legacy and closed data warehouse architecture to an open and expandable data warehouse architecture was enabled and accelerated by the following.

EMC's IT deployment of the Oracle BI Grid utilizing the "building blocks" (see the "References" section) has created an open and expandable, balanced hardware infrastructure enabled by EMC hardware and software with Oracle's RAC technology, which produced the following:

- Smooth and seamless cutover spanning 110 hours
- 15 billion rows and 10 TB of data were migrated from UNIX to Linux
- First major transaction-intensive system was migrated to an x86-based architecture
- Database architecture ready for multiple languages
- Services configured for each major business function
- First major system on a Symmetrix VMAX architecture with EFDs
- Migration from rule-based to cost-based optimizer
- Database backup offloaded from production at 2 TB per hour throughput

The gains can be summarized in the following diagram:

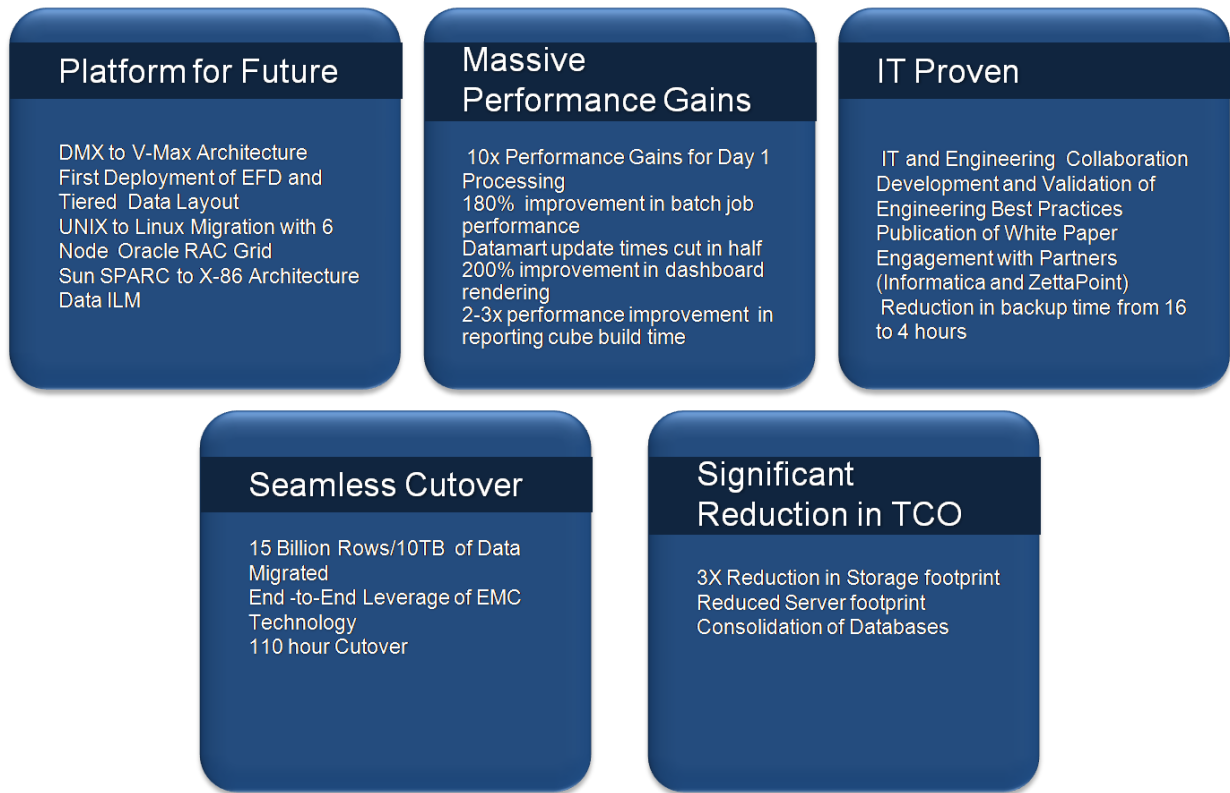


Figure 14. EMC Oracle BI Grid summary

References

The following “building blocks” enabled the construction of the Oracle BI Grid. They include EMC white papers and an EMC Proven™ Solution:

- *Next Generation Data Warehouse Platforms*, TDWI Best Practices Report
<http://www.oracle.com/database/docs/tdwi-nextgen-platforms.pdf>
- “UNIX to Linux Migration,” Dell Enterprise Technology Center TechCenter Wiki
<http://www.delltechcenter.com/page/UNIX+to+Linux+Migration>
- *Cross-Platform Oracle Database Migration Using Transportable Tablespaces and EMC Open Replicator for Symmetrix*, EMC white paper
<http://www.emc.com/collateral/hardware/specification-sheet/h2510-cross-pltfrm-oracle-db-mig-symmetrix-wp-ldv.pdf>
- *EMC Symmetrix VMAX Using EMC SRDF/TimeFinder and Oracle Database 10g/11g*, EMC white paper
<http://www.emc.com/collateral/hardware/white-papers/h6210-symmetrix-v-max-srdf-timefinder-oracle-database-wp.pdf>

-
- *Reducing Backup Window and Recovery Time with Oracle Database 11g RMAN and EMC TimeFinder/Clone*, EMC white paper
<http://www.emc.com/collateral/software/white-papers/h4262-red-backup-wndw-rcvry-time-oracle-11g-rman-timefinder-clone-wp.pdf>
 - *Implementing TimeFinder Multiple Snapshots with Virtual Provisioning for Oracle 11g Databases on EMC Symmetrix VMAX*, EMC white paper
<http://www.emc.com/collateral/software/white-papers/h5635-timefinder-snap-multiple-snapshots-wp.pdf>
 - *Oracle Data Warehouse Sizing with EMC Symmetrix DMX-4 and Dell R900*, EMC white paper
<http://www.emc.com/collateral/hardware/white-papers/h6015-oracle-data-warehouse-sizing-dmx-4-dell-wp.pdf>

Acknowledgments

The author would like to thank EMC IT's BI/DW team and the BI Tools (OBIEE and BIP) team for assisting in the creation of this white paper.