

EMC IT's "ON-RAMP" TO THE JOURNEY TO THE PRIVATE CLOUD

Replatform to an Open, Scalable Infrastructure

Abstract

This white paper illustrates EMC IT's system replatform from a legacy infrastructure (Sun SPARC and Solaris) to an open platform (Intel and Linux), Cisco Unified Computing System (UCS), and EMC IT Oracle 11i CRM RAC approach. The migration resulted in up to 20 times performance gains and \$5 million to \$7 million in savings.

February 2011

Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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 h8170

Table of Contents

Executive summary	4
Audience.....	4
EMC IT Oracle E-Business Suite challenges	5
EMC IT's replatform approach	6
Current architecture.....	6
EMC IT's challenges.....	6
Replatform business drivers	7
Replatform highlights	8
Server platform migration	8
High-level migration milestones	8
Migration drivers	8
How was it done: A phased approach.....	9
Replatform production architecture	9
UCS enablers.....	9
Design principles for the EMC IT UCS domain	10
Production and performance test RAC database servers	10
UCS delivered benefits	13
EMC IT production storage.....	13
EMC IT Oracle migration.....	15
EMC IT replatform migration approaches	15
Detailed steps	17
Oracle-certified method (April 2010)	21
Replatform lessons learned	22
Oracle lessons learned	23
EMC IT Oracle best practices.....	24
Other realized EMC benefits	25
Replatform improvements	27
Business performance gains	28
Business cost savings benefits	31
Conclusion	31
References	32

Executive summary

This white paper shows how EMC IT migrated its Oracle 11i CRM E-Business Suite database servers from a Sun E25K/Solaris infrastructure to Cisco's Unified Computing System (UCS) platform. In addition it highlights the methodology developed by IT while migrating the database from RISC/SPARC to the target UCS (Intel and Linux) platform.

The migration creates a “bridge” for EMC to move from a legacy platform to an open platform that will enable and accelerate the following:

- The ability to get to an open and scalable computing platform UCS (x86). This gives EMC IT a platform to commence its [Journey to the Private Cloud](#).
- A working, documented EMC IT procedure to migrate EMC's 8 TB Oracle CRM E-Business Suite production database to an open platform (from Solaris/SPARC to Cisco's UCS/Linux/x86).
- CAPEX/OPEX savings of \$5 million to \$7 million in the following areas:
 - Data center environmental costs such as power, cooling, and floor space
 - Software licenses
 - Software and hardware maintenance
 - Support for the environment

The new Cisco UCS production infrastructure is up to 20 times better performing than the legacy Solaris platform, reducing end-user response times and batch runtimes by over 60 percent.

Audience

This white paper is intended for CIOs, system architects, Oracle architects, storage architects, Oracle Applications DBAs, and server and network administrators.

EMC IT Oracle E-Business Suite 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 core revenue-generating functions (\$17 billion in revenue for 2010) and are the sources for its Business Intelligence/Data Warehousing infrastructure. The following is a snapshot of the infrastructure and the EMC user community who utilize it:

- An ERP solution, supporting 20,000 employees with 2,000 concurrent users.
- A CRM solution, supporting 40,000 named users worldwide with peaks of over 4,000 concurrent users. This implementation is one of the top four Oracle Applications transactional systems in the world.

- Very high usage penetration within the enterprise
 - 40,000+ Named Users
 - 4,000+ Concurrent Users at peak periods
- 80+ Application Tiers –VMware Virtual Machines (VM)/Linux
- Four-node RAC Architecture - 128 cores
- Oracle Database 10g R2, Red Hat AS 5
- Extremely dense module usage
 - Largest users of Service, Install Base and Contracts, Oracle Sales, Configurator
 - One of the largest users of iStore, Projects, TCA, Order Management, etc.
- Application Continuity Tool – keeping 11i available during maintenance

One of the largest deployments of Oracle E-Business Suites in the world.

Business Statistics	
Quotes	2,263,231
Orders	674,458
Service Requests	9,210,593
Service Contracts	1,391,325
Projects	110,512
IB instances	59,937,589
Material Transactions	12,305,357
Parties	9,429,011
Relationships	10,356,648
Time Cards	7,336,539

System Statistics	
DB Size	8 TB
Number of Rows	8.8 Billion
IOPS	79K
Interconnect Traffic	28,000-40,000 Blocks/sec
Daily Volume	
Database Transactions	1824/Sec - 520 Million/Day
Archive Log	1.8 TB
Conc. Jobs	91K
Workflow Events	250K
Workflow Roles	20,193,476



Figure 1. EMC's current E-Business Suite statistics and volumes

EMC IT's approach was to perform a replatform from Sun Solaris/SPARC to Cisco's UCS platform. The UCS platform is an open expandable infrastructure, enabled by EMC technologies and best practices with the following components:

- Open – Linux/x86 platform (UCS) to replace legacy Solaris and legacy SPARC hardware

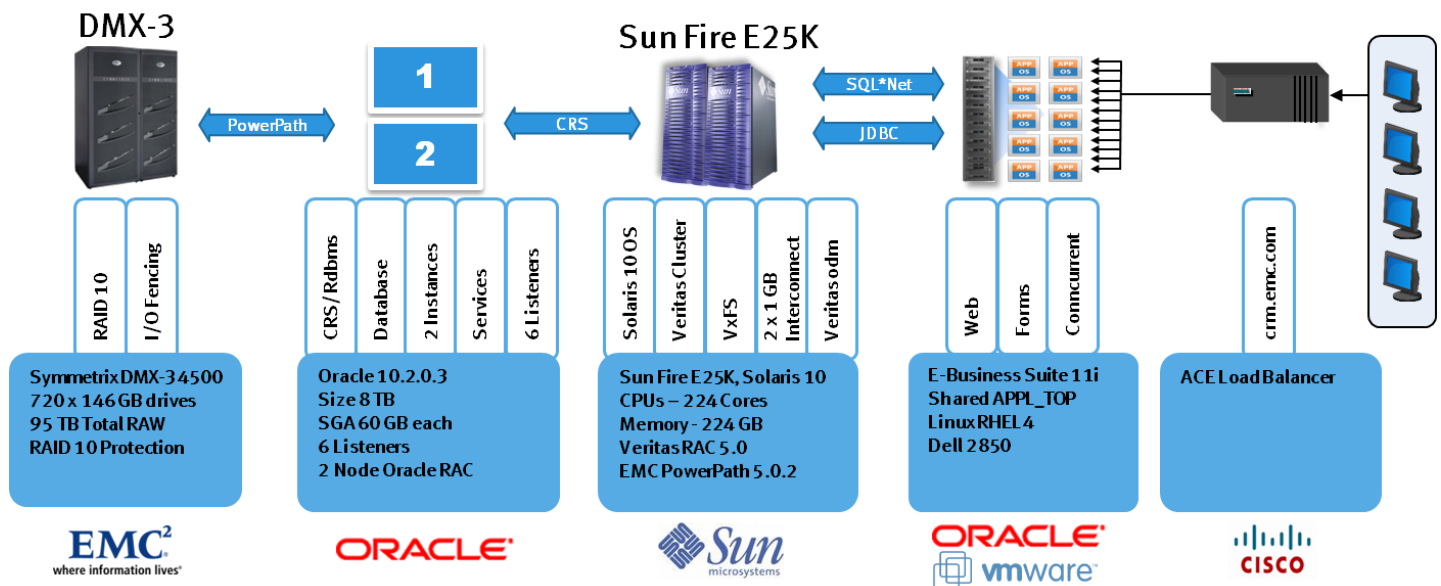
- Scalable – An Oracle Real Application Clusters (RAC) physical deployment to create a database grid that has the ability to horizontally scale out (by adding RAC nodes)
- Platform – Cisco’s UCS platform

EMC IT’s replatform approach

The following section identifies the EMC legacy infrastructure of the EMC IT CRM E-Business Suite database server deployment.

Current architecture

Figure 2 shows the legacy components that EMC IT deployed for its Oracle CRM (11.5.10 CU2) production infrastructure:



System Statistics	
Database Size	8 TB
Number of Rows	8.8 Billion
IOPS	61K
Interconnect Traffic	3000-4000 Blocks/second
Interconnect Latency	12-18 Milli-Seconds
CPU Utilization	80+%

Daily Volume	
Database Transactions	1098/Sec- 620 Million/Day
Archive Logs	1.8 TB
Concurrent Jobs	91K
Workflow Events	250K
Workflow Roles	20,193,496

Figure 2. EMC’s legacy E-Business Suite infrastructure and statistics

EMC IT’s challenges

Challenges of the legacy infrastructure include the following:

- Need to migrate away from proprietary hardware and UNIX to open hardware and Linux

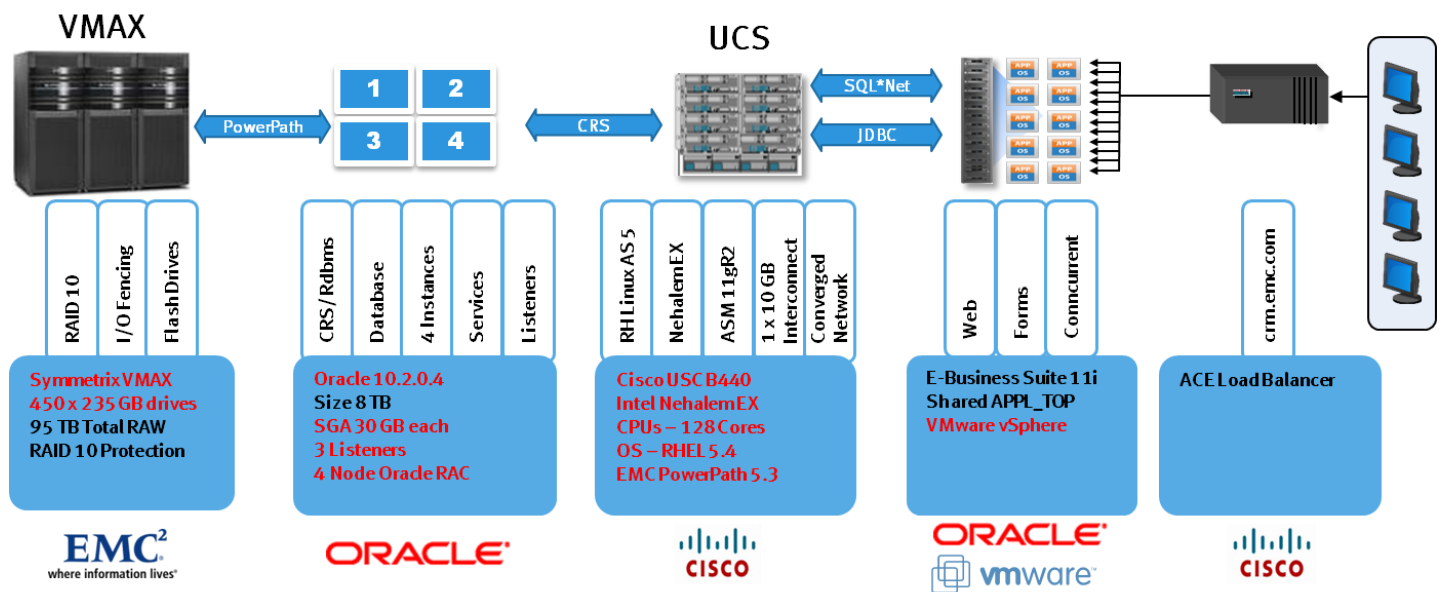
- Aging infrastructure
- Legacy hardware has not kept pace and has poor transactional performance
- Tied to a single vendor
- Legacy platform is “end of life”

Replatform business drivers

The following high-level replatform drivers illustrate the need to migrate EMC from a legacy deployment platform to an open, performant UCS platform.

- A variety of Sun SPARC server technologies deployed in EMC data centers (880, 25K)
- Currently deployed Sun platform was end of life
- High operational cost to support a Sun SPARC platform
- Business service-level objectives could not be met on the current Sun SPARC infrastructure

Figure 3 is a high-level illustration of the EMC UCS replatform infrastructure architecture deployment. (The red text highlights new components.)



System Statistics	
Database Size	8 TB
Number of Rows	8.8 Billion
IOPS	79K
Interconnect Traffic	28,000-40,000 Blocks/second
Interconnect Latency	.5-2Milli-Seconds
CPU Utilization	10%

Daily Volume	
Database Transactions	1824/Sec - 520 Million/Day
Archive Logs	1.8 TB
Concurrent Jobs	91K
Workflow Events	250K
Workflow Roles	20,193,496

Figure 3. EMC’s high-level replatform production deployment infrastructure and statistics

Replatform highlights

Enabling an open, expandable E-Business Suite database infrastructure to support EMC, 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.

The following are the highlights of the production replatform.

Server platform migration

These sections detail the [high-level migration milestones](#), [migration drivers](#), and [“How it was done” approach](#) to migrate to an open and scalable platform. They highlight the vendor discovery and phased migration of incremental Tier 1 applications projects to gain experience and knowledge on open software (Linux), open infrastructure (x86), and successively more complex Tier 1 migrations.

High-level migration milestones

The following are the high-level migration milestones:

- **October 2009 – January 2010:** Proof-of-concept (POC) equipment procured and configured
- **February – March 2010:** Initial load tests run on the POC environment
- **April – May 2010:** POC gear and initial procurement used to build a dev and test environment
- **Mid-July 2010:** Test/performance/production equipment arrives
- **July – September 2010:** Environment builds
- **August – October 2010:** Functional/regression testing
- **September – October 2010:** Performance testing
- **October 16, 2010:** Production system is live

Migrating a \$17 billion business’ legacy IT platform to an open hardware and software platform requires planning, a knowledgeable IT staff, and partners. EMC (Symmetrix VMAX™, TimeFinder®, and SRDF®) and partner (Cisco’s UCS) technologies provided solutions to complete this migration in less than 12 months.

Migration drivers

Movement from its legacy platform was precipitated by EMC IT’s realization that to become both a business- and technology-*agile* organization it would need to move to an open hardware and software platform (x86 and Linux).

A roadmap discovery of all vendors was conducted, to show that their platforms, both operating system and hardware, were ready for the current EMC 11.5.10 workload as well as EMC’s future workload. This discovery phase of the vendors’ roadmaps was the initial step in a phased project approach by EMC IT.

How was it done: A phased approach

EMC first got its knowledge and experience of migrating from legacy infrastructures to open infrastructures with three EMC IT Tier 1 mission-critical applications of increasing size and scope:

- **Virtualization of the EMC CRM 11.5.10 CU2 Middle Tier infrastructure** – EMC/Oracle 11.5.10 CU2 Middle Tier Production Applications servers went from 71 physical Dell servers to 80 Middle Tier Productions Applications servers’ virtual machines on eight ESX® servers.

Note: Today, all EMC/Oracle 11.5.10 CU2 Middle Tier Applications (originally 210 physical servers) now execute as VMware virtual machines (VMs) on 20 multitenancy ESX servers (production/development/test/patch/performance), which are Dell R900 class servers.

- **Global R12 implementation with Asia** – This project was a global project for Asia that virtualized the R12 Oracle Applications stack with VMware on a Linux/x86 infrastructure.
- **Oracle BI Grid** – This project 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) that supports EMC’s virtualized BI toolset reporting architecture (VMware, OBIEE, and BIP).

With those three models completed, EMC IT understood critical issues of the migrations, including the people, processes, and technology required when migrating to an open, scalable platform. Those three successes gave EMC IT the confidence that led to the replatform of the Tier 1 mission-critical Oracle CRM database infrastructure (11.5.10 CU2).

Replatform production architecture

The following section discusses:

- [Cisco UCS enablers to EMC IT](#)
- [EMC IT UCS design principles](#)
- [EMC IT production UCS deployment architecture](#)
- [UCS delivered benefits](#)
- [EMC IT production storage](#)

UCS enablers

Cisco’s UCS is a strategic platform for EMC IT. At a high level it provides the following for the EMC vision of [Journey to the Private Cloud](#) computing:

UCS “unifies”

- It uses industry-standard x86.

- You only need to wire once for SAN, NAS, and iSCSI.
- Virtualization adds control, scale, and performance.
- UCS profiles enable “swappable” blades.

Embedded management

- Increase scalability without added complexity
- Utilize dynamic resource provisioning
- Integrate with a broad partner ecosystem

Energy-efficient

- Use fewer servers, switches, adapters, and cables via “unified fabric”
- Lower power and cooling requirements
- Increase compute efficiency by removing I/O and memory bottlenecks

Design principles for the EMC IT UCS domain

A standard UCS configuration was used with the following design attributes:

- **Little or no customizations of the UCS.** It is a standard EMC IT UCS building block for data center and disaster recovery enabled by Cisco’s UCS policy, templates, and profiles
- **High availability of the UCS infrastructure** with a no-single-point-of-failure design within the data center (see Figure 4)
- **Maximum I/O throughput** of 80 Gb/s per chassis for both the fabric interconnect and external fabric:
 - Maximizing FEX to FI cabling and expanding FC capabilities with a six-port expansion module
 - Ability for Oversubscription > 80 Gb/s is enabled by the USC FEX and fabric interconnect mentioned above
- **Foundation for a cloud architecture** with the ability to add or remove UCS components such as:
 - Compute
 - I/O fabric (external and interconnect)

Production and performance test RAC database servers

Six Cisco UCS B440 blades were deployed and configured using Oracle RAC to protect against physical node failure. Four Cisco UCS B440 blades are used today.

The cluster is designed to function with more than or less than six blades but in the event that an extended blade outage exists or short-term compute power is required,

EMC IT can attach the failed server profile to another spare blade from the UCS B200 or B440 pool.

The six-server chassis is spread across three cabinets, each having redundant power feeds. Each B440 blade contains dual CNAs, each presenting multiple NIC and FC devices.

The public and interconnect NICs will be further protected by using N+1 interfaces on each server and will be configured with Linux bonding.

Access to the storage devices is protected by using N+1 FC interfaces on each server and those will be configured with EMC PowerPath®. This will protect from a CNA card or port failure. Each server will boot from mirrored SAN devices.

The following is the high-level EMC IT production “UCS Domain” production configuration.




Figure 4. EMC’s replatform production UCS deployment infrastructure

The following tables show the detailed components of the deployed Cisco/UCS infrastructure:

Table 1. UCS Domain replatform components

Vendor	Quantity	Item	Description
Cisco	6	UCS 5108 blade server chassis	Fully configured with UCS 2104XP fabric extenders
Cisco	28	UCS M81KR	“Palo” CNA, 2-port 10 Gb
Cisco	4	UCS M71KR-E	“Menlo” CNA, 2-port 10 Gb
Cisco	2	UCS 6140XP 40-port fabric interconnect	Fully configured with fans and power supplies, and all ports licensed
Cisco	14	UCS B200 M1 blade server	Half-width, 2-socket/4-core x5680 2.93 GHz, 130W, 8 MB cache, 96 GB of memory using 18 x 8 GB DIMMs
Cisco	12	UCS B440 blade server	Full-width, 4-socket/8-core x7560 2.26 GHz, 130W, 24 MB cache, 128 GB of memory using 32 x 4 GB DIMMs

Table 2. Network switches components

Vendor	Quantity	Item	Description
Cisco	2	7108	LAN switches, to be shared with other domains through separate VLANs
Cisco	2	Nexus 5020	Aggregation-level switches. 6100 UCS switches will connect to these for backup SAN traffic and for NAS.

Table 3. Storage switches components

Vendor	Quantity	Item	Description
Cisco	2	MDS 9513	SAN switches, to be shared with other domains through separate VSANs

Table 4. EMC storage components

Vendor	Quantity	Item	Description
EMC	1	VMAX	4 Engines - Production
EMC	1	VMAX	4 Engines – Dev, Test, and Performance

UCS delivered benefits

The following are the high-level benefits for EMC to deploy on the Cisco UCS platform:

- Significantly lowers hardware, heating, cooling, service, and support costs (see the [Business cost saving benefits](#) section)
- Better performance, especially transactional performance improvements
- Intel Xeon has “changed the equation”
- No longer tied to a proprietary vendor
- A virtualization enabler, allowing the move to virtualization and [EMC's Journey to the Private Cloud](#)
- Allows greater agility and flexibility
- Simplifies the data center
- High availability is no longer cost-prohibitive
- Flexibility to grow as business and IT requirements grow

EMC IT production storage

The following is a high-level description of the configuration and connectivity of the Symmetrix VMAX storage array used for production:

- The configuration included 450 devices (Enterprise Flash Drives, Fibre Channel, SATA).
- The disk format is RAID 1 and RAID 5 with concatenated metavolumes.
- SAN connectivity is a 8 Gb infrastructure and 4 Gb Fibre Channel to the SAN.

Table 5 illustrates the file system layout on the Symmetrix VMAX array connected to the production Cisco UCS platform (**only two nodes shown**).

Table 5. Symmetrix VMAX storage layout

Mount Point (# of mount points)	Meta Size	Meta/FS	# of FS	GB	# m/c	m/c name	Notes
DB Tier							
crs/asm	35	1	1	70	2	node1, node2	ext3
oracle	35	1	1	70	2	node1, node2	ext3
oraconfig	2	1	1	2	1	node1, node2	Cluster disk shared (OCFS)
oradata/dataXX	210	1	46	9660	1	node1, node2	ASM disks shared
oradata/redoXX	10	1	8	80	1	node1, node2	ASM disks shared
arch	210	1	1	210	1	node1, node2	ASM disks shared
temp	100	1	1	100	1	node1, node2	ASM disks shared
imap	10	1	1	10	1	node1, node2	ext3, floats non- cluster FS
ctrlm	8	1	1	8	1	node1, node2	ext3, floats non- cluster FS
splex	70	1	1	70	1	node1, node2	ext3, floats non- cluster FS
OCR/Vote disks	1	1	1	41	4	node1, node2	ASM disks shared
Total (DB tier)				10284			

EMC IT Oracle migration

The following section details the method used by EMC IT to migrate its legacy Oracle Applications CRM Oracle database (8 TB) to an open, scalable platform.

The following will be discussed:

- [High-level EMC IT migration approaches](#)
- [Detailed steps on the EMC IT method](#)
- [Oracle-certified method \(April 2010\)](#)

EMC IT replatform migration approaches

Oracle provides many methods to move from legacy architecture (Big Endian) to an open x86 architecture (Little Endian) that are beyond the scope of this discussion. The following are the two methods that EMC IT utilized:

- Import/Export
- Transportable Tablespaces

Import/Export

EMC IT used this method in the Oracle BI Grid replatform. The white paper [EMC IT's Replatform Proof of Concept to an Open, Scalable Platform - Applied Technology](#) explains how import/export was used to migrate to an open (x86) platform.

Transportable Tablespaces (TTS)

The Transportable Tablespace feature allows users to move a non-system tablespace across Oracle databases. It's an efficient way to move bulk data between databases. If the source platform and the target platform are of different endianness, then an additional conversion step must be done on either the source or target platform to convert the tablespace being transported to the target format. If they are of the same endianness, then no conversion is necessary and tablespaces can be transported as if they were on the same platform.

Moving data using transportable tablespaces can be much faster than performing either an export/import or unload/load of the same data. This is because transporting a tablespace only requires the copying of datafiles from the source to the destination and then integrating the tablespace structural information.

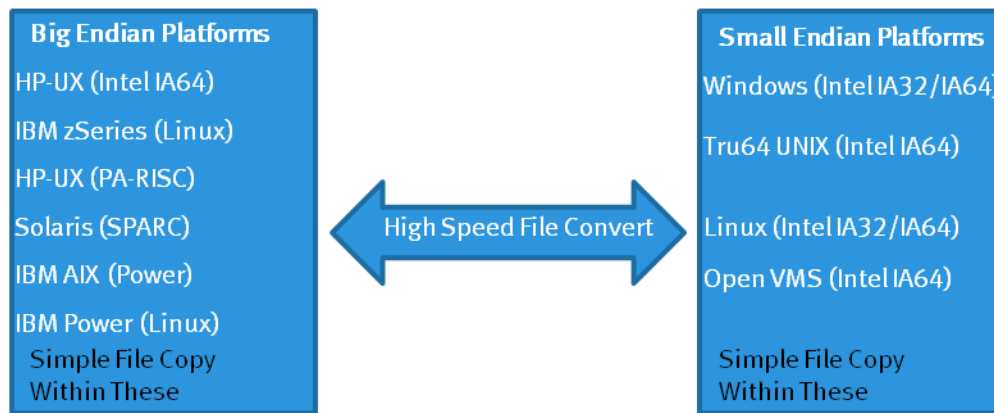


Figure 5. Platform migrations (Sun SPARC) to a Linux platform (x86/UCS)

EMC developed a method to migrate its CRM 11i Oracle Applications E-Business Suite from legacy UNIX to Linux via the RMAN datafile Convert. Figure 6 illustrates the high-level steps in the migration method.

Sun Fire E25K




Figure 6. EMC’s IT Oracle Applications E-Business Suite database migration method

The following are the high-level steps:

1. NFS mount the datafile LUNs to the target UCS PRODUCTION. Converting directly from the NFS mounts saves the time and space that would be required to copy them to the target database server.
2. Put the database in “read-only” mode.
3. Create a Transportable Tablespace set.
4. Export the Transportable Tablespace metadata.
5. Export the E-Business APPS user – norows to get all of the database code, synonyms, triggers, and sequences.
6. Export the SYSTEM schema to get all of the temporary tables.
7. Use the RMAN Convert datafile directly into ASM.
8. Import the Transportable Tablespace metadata.
9. Import the E-Business APPS user – ignore='Y'.

10. Import the SYSTEM Schema – ignore='Y'.
11. Rebuild Queue Tables.
12. Rebuild all Spatial Indexes.

Detailed steps

The following are the detailed steps laid out (or executed) by EMC IT to migrate its Oracle Applications CRM instance from legacy architecture (SPARC/Big Endian) to an open architecture (Linux/x86/Little Endian).

Using Oracle's Metalink note 454574.1 as a baseline approach. EMC IT enhanced it by noting issues they discovered and also strategies for accelerating certain steps (**Bold Text**):

Step 1) EMC IT chose to use NFS mounting of the source LUNs. This allowed EMC IT to migrate the datafiles without having to first stage them onto temporary storage in the new cluster. Copying 8 TB of database files can take over 12 hours to complete. **This reduced or removed hours from the migration process.** A more efficient method would have been to mount the actual LUNs directly to the new cluster. However, since the source filesystem was Veritas, EMC IT chose not to do this, due to the cost of the license. This cost some valuable time, as the NFS protocol was a bottleneck.

Step 2) Mark the tablespaces within the transportable set read only. For generating the script to mark the tablespaces read only, we used the following:

```
select 'alter tablespace '||tablespace_name||' read only;' from dba_tablespaces
where tablespace_name not in
('TEMP', 'SYSTEM', 'SYSAUX', 'CTXD', 'UNDOTBS1A', 'UNDOTBS2A') and contents <>
'UNDO';
```

Step 3) Create a Transportable Tablespace

- a. Drop unused columns; if missed there will be loss of data after import.


```
select * from sys.dba_unused_col_tabs;
alter table <table name> drop unused columns;
```
- b. Create a Transportable Tablespace set for EBS (vanilla and custom) schemas and make sure violations are addressed. RDBMS schemas are disregarded in the creation of the transportable set.


```
declare
tbls varchar2(4000);
begin
tbls:='<list of tablespaces separated by commas>'
dbms_tts.transport_set_check( tbls,true);
end;
select count(*) from transport_set_violations;
```

- Step 4) Export the Transportable Tablespace metadata.
- a. Export stats for all the schemas except sys, system, and sysaux from the source. **This reduces the downtime during import. (You will not have to regenerate on the target a minimum of 20-plus hours in the EMC IT Oracle Applications CRM environment.)**
 - b. Apply expdp patches in the source (6354875, 5249074, and 5120780). These patches are required not only to improve the performance of expdp processes but also for fixing a bug that gets exposed during the impdp process. **Confirm patch 6730429 is applied on the target before starting the impdp process there. Missing this patch can cause loss of the database.**
 - c. Export the Transportable Tablespaces set - tablespaces from which metadata will be unloaded:


```

      directory=dmpdir
      dumpfile=xttsmeta.dmp
      filesize=5g
      logfile=expxtts.log
      transport_tablespaces=
      ABMD,
      ABMX,
      ...
      
```

- Step 5) Export the E-Business APPS user.
- a. Export miscellaneous data (packages sequences, procedures):


```

      directory=dmpdir
      dumpfile=xttsmsc.dmp
      full=y
      filesize=5g
      logfile=expmsc.log
      exclude=TABLE,INDEX,CONSTRAINT,COMMENT,RLS_POLICY,
      MATERIALIZED_VIEW,MATERIALIZED_VIEW_LOG
      
```
 - b. Export the RLS policy for the full database. This is required when setting up security policies for EBS. **The expdp-impdp steps of the Oracle Metalink documented process do only the RLS policies of the Transportable Tablespaces. This caused functionality of the Territory administration to be broken.**

```

      directory=dmpdir
      dumpfile=rls.dmp
      full=y
      logfile=rls.log
      INCLUDE=RLS_POLICY
      
```
 - c. Get the AQ information from dba_queues to restart them later on the target system.

- Step 6) Export the SYSTEM schema
- a. Export SYSTEM tablespace schemas:

```
directory=dmpdir
dumpfile=xttssys.dmp
filesize=5g
logfile=expsys.log
SCHEMAS=
DBSNMP,
ORDSYS,
ORDPLUGINS,
OUTLN,
MDSYS,
EIM_CONTINGENCY_RPT,
OAI,
SYS,
SYSTEM
....
```

b. Export global temporary tables:

```
file=<Path>/gtab.dmp
filesize=5g
log=<Path>/expgtab.log
tables=(
'BIC.BIC_TEMP_MEASURES',
'BIC.BIC_TEMP_RETENTION',
....
```

Step 7)

EMC IT used RMAN to convert each datafile. This was necessary due to the conversion occurring on the target servers. It also allowed us to utilize the increased speed of the target hosts and better parallelism. The datafiles were split into multiple threads to run on each target host. **EMC IT utilized RMAN parallelism, rather than channels, for the best throughput** (EMC IT use of six hosts (multiple threads per hosts) versus a single host and multiple channels). An example command is provided below:

```
rman target / <<EOF> logfile_name.log
run
{
convert datafile
'/convert/data39/cnxxx_26.dbf',
'/convert/data39/cnxxx_27.dbf',
'/convert/data29/csidxx_26.dbf',
'/convert/data31/csidxx_27.dbf',
'/convert/data31/csidxx_28.dbf',
'/convert/data33/csidxx_29.dbf',
'/convert/data15/emcops05.dbf',
'/convert/data33/csidxx_30.dbf',
'/convert/data04/staged05.dbf',
```

```

'/convert/data13/cscx03.dbf',
'/convert/data16/wshd03.dbf'
from platform 'Solaris[tm] OE (64-bit)'
PARALLELISM=16
db_file_name_convert='/convert/data01','+CRMDBP1_DATA','/convert/d
ata02','+CRMDBP1_DATA','/convert/data03','+CRMDBP1_DATA','/conver
t/data04','+CRMDBP1_DATA','/convert/data05','+CRMDBP1_DATA','/c
onvert/data06','+CRMDBP1_DATA','/convert/data07','+CRMDBP1_DATA','/c
onvert/data08','+CRMDBP1_DATA','/convert/data09','+CRMDBP1_DATA',
'/convert/data10','+CRMDBP1_DATA','/convert/data11','+CRMDBP1_DAT
A','/convert/data12','+CRMDBP1_DATA','/convert/data13','+CRMDBP1_
DATA','/convert/data14','+CRMDBP1_DATA','/convert/data15','+CRMDB
P1_DATA','/convert/data16';
}
quit
EOF

```

- Step 8) Import the Transportable Tablespace metadata.
- a. Import miscellaneous data. Import all objects in the system tablespace except procedural objects:


```

directory=dmpdir
dumpfile=xttsmsc.dmp
full=y
logfile=impmsc.log
EXCLUDE=PROCOBJ
remap_tablespace=
ABMD:SYSTEM,
....

```
 - b. Plug in the datafiles (Transportable Tablespaces set) to the target database. **Make sure to exclude table statistics and index statistics; this saved more than 20 hours.**

```

directory=dmpdir
dumpfile=xttsmeta.dmp
logfile=impxtts.log
exclude=TYPE
exclude=TABLE_STATISTICS
exclude=INDEX_STATISTICS
transport_datafiles=
'<datafile1 path>/<datafile1 name>',
'<datafile2 path>/<datafile2 name>',
....

```
 - c. Make the Transportable Tablespaces read-write on the target:

```
select 'alter tablespace '||tablespace_name||' read write;' from
dba_tablespaces
where tablespace_name not in
('TEMP','SYSTEM','SYSAUX','CTXD','UNDOTBS1A','UNDOTBS2A') and
contents <> 'UNDO';
```

- Step 9) Import the E-Business APPS user.
- Update the default tablespace and quota on tablespaces of all users and match them with the source.
 - Import the procedural objects:
directory=dmpdir
dumpfile=xttsmsc.dmp
full=y
logfile=impmsc-include.log
INCLUDE=PROCOBJ
 - Import the RLS policies.
 - Restart Queue Tables and rebuild all Spatial Indexes.
 - Import the schema statistics taken on the source.
- Step 10) Import the system schema.
- Import the system schema to the target:
directory=dmpdir
dumpfile=xttssys.dmp
full=y
logfile=impsys.log
 - Import the global tables:
file=<Path>/gtab.dmp
filesize=5g
full=y
log=<Path>/impgtab.log
- Step 11) Rebuild Queue Tables.
- Step 12) Rebuild all Spatial Indexes.

Oracle-certified method (April 2010)

EMC IT joined the early adopters program in the first quarter of 2010 for using Transportable Tablespaces to migrate cross-platform. In addition, Oracle Metalink note 454574.1 was released in the second quarter to serve as an officially supported method. This method only supports Oracle E-Business Suite 11.5.10.2 on Oracle 10.2.0.4. This method, as well as EMC IT's, will work on other versions from 10gR1 onward.

The Metalink note fails to disclose the problems with using datapump export/import. EMC previously utilized the non-datapump export/import due to these problems. Many patches for expdp and impdp provided some benefit (as detailed in [Detailed steps](#)).

Replatform lessons learned

The following are some lessons learned during the replatform to an open, scalable platform from a legacy platform:

Platform (UCS) changes

- Physical disks were removed from all blades and the blades were configured to boot from SAN. This required setup of the kickstart configuration to build the boot drives using LVM since we would be using PowerPath across all devices.
- Accelerated provisioning:
 - Service profile templates were used to create service profile instances.
 - Firmware upgrades for blades were done by using UCS policies and re-binding service profiles to templates.
- Firmware issues were experienced using initial releases, and were resolved with the current update.

Operating system changes

- Red Hat 5.3 contained a bug that could lead to silent filesystem corruption. We upgraded all nodes to Red Hat AS 5.4.
- Shut off kudzu to prevent any MAC changes from altering device paths.
- Reconfigured all servers to send kdump core files to a central server using scp.
- Needed the latest enic/fnic drivers (early code release issues).

Infrastructure changes

- The new Nexus 7000 switches no longer supported the automatic forwarding of NIS broadcast packets so we had to hard-code NIS servers as opposed to using a broadcast configuration. We will change over to LDAP authentication.
- We used a Compute by Purpose design and set up blade pools by purpose (prod, dev, B200, B440, and so on).

Operational changes

- Changes were made for ease of management via a single Restricted Access Login.
 - Set up a single Cisco TAC account to funnel all calls and provide a central contact point.
 - Needed to create a “KVM only” (Console) UCS access account for command center staff.

- Could not set up Active Directory (AD) integration, due to a requirement to the modified AD schema (fixed in UCS firmware 1.4).

Oracle Applications deployment changes

- A separate proxy backup host was installed outside of UCS.

Oracle lessons learned

The following section contains some of the highlights of the lessons learned regarding Oracle from this migration.

Memory

Migrating from Sun Solaris to Red Hat Linux caused EMC IT to need to implement the Linux kernel feature, HugePages. HugePages is a feature integrated into the Linux kernel 2.6. It is a method to have a larger page size that is useful for working with very large memory. The entire SGA must fit within the reserved HugePages or the SGA will not utilize the configured HugePages and that memory will be wasted. (The [References](#) section has links to more details.)

Performance

EMC IT discovered that it had not accounted for the extremely large bursts of data changes that this new platform (Cisco UCS) was capable of executing. Batch jobs that generated over 100 GB of redo would spread those changes over many hours on the legacy platform (Sun). With the new capabilities of the replatform (UCS) these jobs now complete in a fraction of the time. The old platform could generate only a maximum of 1 GB per minute of redo changes. Since it takes approximately one minute to perform a checkpoint in EMC’s CRM database (8 TB) and the old system would generate 1 GB of redo per minute, this gave EMC a minimum redo size requirement of 2 GB. The new system (UCS) is capable of significantly more throughput, and the EMC CRM workload now generates **8 GB** of redo per minute. **This did not allow the redo logs to be archived, as the checkpoint would not complete and ultimately led to a hung archiver process.**

The same issue also led EMC to undersize the archive area. Due to the significantly larger bursts of data changes, the archive area, 630 GB, was filling up in hours and RMAN was barely able to keep up with it. The archive area was resized to 2 TB, in order to allow for approximately one day’s worth of archive logs.

The following table highlights the lessons learned from this migration.

Table 6. Oracle lessons learned

Challenge	Resolution
Memory	Increased HugePages and reduced SGA (to fit within huge memory) on database servers
Service readiness	New training and knowledge transition sessions

Operational	TNS and JDBC connection updates for external feeds Firewall rule changes Clone scripts
Performance	Increased REDO size from 2 GB to 8 GB to reduce contention Applied the Database Rollup Patch set Expanded the archive area to 2 TB

EMC IT Oracle best practices

The following highlights the best practices used in this migration.

Platform (UCS)

- EMC IT common design pattern and templates were applied for the UCS domain. The templates made the platform disaster recovery-ready and provided elastic computing.
- A maximum I/O throughput of 80 Gb/s was established for the deployed UCS infrastructure.
- The data center was made HA-ready with no single point of failure.

Operating system

- A tweak was made to RAC interconnect networking. Jumbo frames was applied to Red Hat configurations.
- We specified an override for maximum size that can be locked in memory (ulimit -l) in /etc/init.d/ohasd for CRS startup.
- As per best practices, we changed /etc/sysconfig/o2cb.
 - No Oracle spfiles were used in init.ora.
 - This benefits a single copy of init.ora and the ability to document changes to init.ora.
 - This was used for Oracle Maintenance scripts.
- Modifications were made to /etc/security/limits.conf for memory locking, total open files, and processes.

Oracle Applications deployment

- Huge Pages was configured in Linux to match the SGA size.
- REDO logs were sized to allow for checkpoints to complete prior to switching to the next online redo log.
- Applications Partitioning was applied through Oracle Services (batch versus online users).
- EMC's Enterprise Flash Drives (EFDs) were utilized for hot objects (tables/indexes). EFDs perform exceedingly well for Random Reads (Random

Read Access/Data – Indexes, for example). This is particularly true in RAC and can significantly reduce interconnect latency times for extremely hot objects.

EMC utilized a tool, DB*CLASSIFY, from our partner Zettapoint both to confirm EMC IT knowledge and to automate the analysis of the “hot objects” promotion to EFDs with Zettapoint’s DB*CLASSIFY. EMC IT has tested this tool to compare to its own recommendations. EMC IT utilized the AWR reports, application knowledge, and wait events in a long and painstaking process and compared the results to the output of Zettapoint ‘s DB*CLASSIFY. The results were almost identical; however DB*CLASSIFY had several objects that EMC IT had missed.

- IPC modifications were made to /etc/sysctl.conf based on Oracle recommendations and IT best practices.

Infrastructure changes

- An interconnect was set up to use Fabric B as the default primary in all cases to avoid routing through the 7000 switch.

Operations changes

- New training and a knowledge transition session were provided to the EMC IT teams

Other realized EMC benefits

Migrating to Symmetrix VMAX allowed EMC IT to better utilize the EMC replication technologies TimeFinder and SRDF to support EMC IT’s Development, Test and Performance environments, both in creating a refresh of Oracle instances and backups of those instances in seconds (see [References](#)).

TimeFinder

TimeFinder is a family of EMC replication products that operates in a single Symmetrix array and nondisruptively creates and manages point-in-time copies of data volumes. TimeFinder runs in Symmetrix Enginuity™ but is controlled by TimeFinder software running on an attached host. It can be administered by the user through Solutions Enabler Command Line Interface, Symmetrix Management Console (SMC), EMC Ionix™ ControlCenter (ECC), or Mainframe Enablers. TimeFinder includes the following sets of products:

TimeFinder/Clone

TimeFinder/Clone provides single or multiple point-in-time copies of full volumes or individual datasets. Cloned data is available to a host immediately upon activation, even if the copy process has not completed.

TimeFinder/Snap

TimeFinder/Snap provides pointer-only-based replicas simultaneously on multiple target devices from a single source device. With a space-saving TimeFinder/Snap only

changed data is written to a pool of save devices. Data reconstruction is from the source device and the pointers into the change tracking save pool.

TimeFinder Consistency Groups

TimeFinder Consistency Groups is a flow control mechanism that is employed to create copies of source volumes at an instantaneous point in time. TF/CG utilizes an Enginuity feature call Enginuity Consistency Assist (ECA) to momentarily halt writes to source volumes during the creation of the copy.

SRDF

The Symmetrix Remote Data Facility (SRDF) family of remote mirroring software offers various levels of Symmetrix-based business continuity and disaster recovery solutions. The SRDF product family offers the capability to maintain multiple, host-independent, mirrored copies of data. Symmetrix systems participating in SRDF can be in the same room, in different buildings within the same campus, or hundreds to thousands of kilometers apart.

EMC IT is now able to offline backups of EMC IT's two Dev and two Test environments through its proxy backup server. EMC IT utilizes EMC Snap replication technologies (snaps) to take a consistent backup of the 8 TB Dev and Test databases, without impacting those databases and with very little additional disk space.

EMC IT utilized both local (TimeFinder) and remote replication (SRDF) to enable and accelerate the following:

- EMC IT's Oracle Applications Instance Lifecycle – Replication of Oracle instances utilized EMC's TimeFinder/Snap and TimeFinder/Clone to create “gold copy” instances and utilized other mount hosts to minimize impact to the production environment
- Use of SRDF to create EMC IT's Oracle Applications Disaster Recovery (DR) environment
- Backups

The following high-level diagram illustrates the use of TimeFinder for clones, snaps, and SRDF to accelerate EMC IT lifecycle deployment of Oracle instances for one of the largest Oracle Applications deployments in the world.

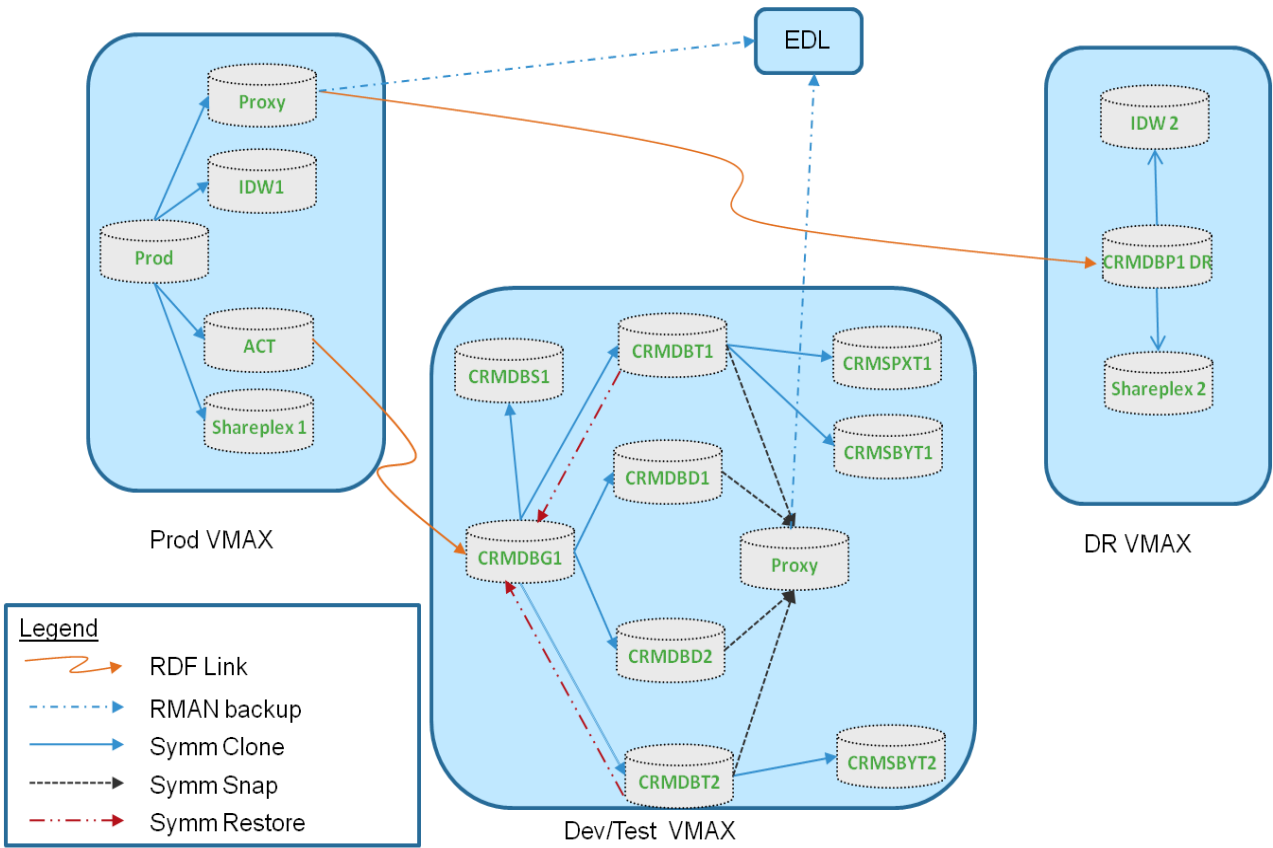
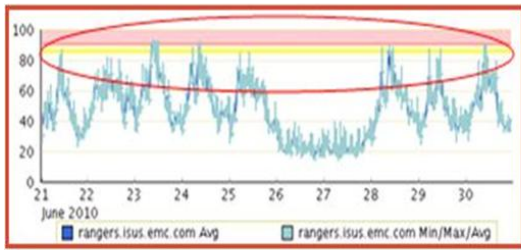


Figure 7. TimeFinder/SRDF for Oracle Applications lifecycle management instances

Replatform improvements

The following diagrams illustrate the replatform improvements gained by moving from a legacy (Sun SPARC) platform to an open, scalable (UCS/Linux) platform:

CPU Utilization



Re-Platform

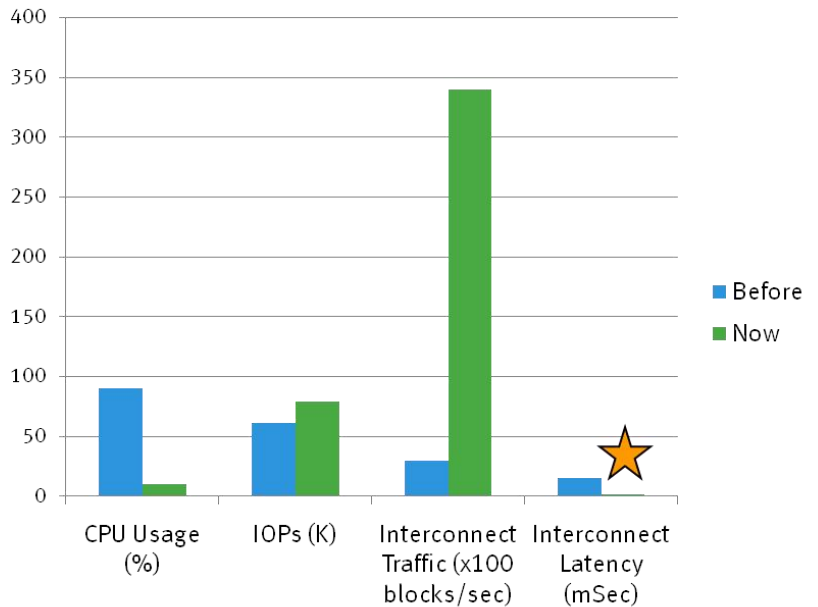
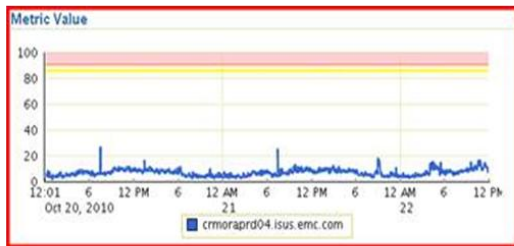


Figure 8. UCS replatform system improvements (CPU/IOPS/RAC interconnect)

The diagrams illustrate the following:

- The legacy platform had daily CPU spikes from 40 percent to over 90 percent utilization.
- The UCS platform does not have the extreme CPU spikes (5 percent to 20 percent) that the legacy platform does.
- CPU usage and RAC interconnect latency were lower.
- IOPS and RAC interconnect throughput increased.

Business performance gains

The following section highlights results from the performance test that show the improvements on the deployed UCS domain infrastructure and a Symmetrix VMAX.

Table 7 highlights important activities that are critical to EMC's business and the business transaction times with the legacy system (the column Base Avg Time) compared to times with the UCS replatform server.

Also in the table:

- CSI stands for Customer Service actions.
- CXP stands for Channel Express, which means external orders for EMC.
- DXP stands for Direct Express, which means internal orders generated by EMC Sales teams.

Table 7. Highlighted business transactions times - legacy versus replatform server (UCS)

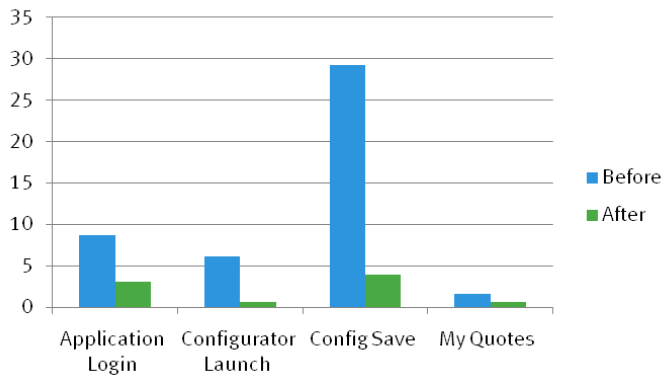
Script	Description	Time (sec)		Improvement
		Legacy	Nehalem	
CSI - Browse	Refresh Customer Service Group	2.0	1.0	2.0X
	Refresh Open Ticket Queue	0.6	0.5	1.2X
CSI - Contracts	Renewal Price Quote	11.5	4.5	2.6X
	Re-price Contract	1.9	0.8	2.4X
CSI – Field Service	Load Universal Work Queue	21.0	9.6	2.2X
	Save Service Request	2.8	1.0	2.8X
CXP – Channel Express	Login	2.3	0.4	5.8X
	Place Order	27.8	1.2	23.6X
	Save Configuration	27.4	3.5	7.8X
	Launch Configurator	8.5	3.0	2.8X
DXP – Direct Express	Add a Product	5.6	1.2	4.7X
	Price a Quote	196.3	18.2	10.8X
	Approve Quote	33.2	5.4	6.1X
TSE – Cost Budgets	Create Cost Budget	4.3	1.0	4.3X
TSE – Timecards	Save Timecard	1.3	0.4	3.3X

Improvements with the new system were an “order of magnitude” over the legacy infrastructure:

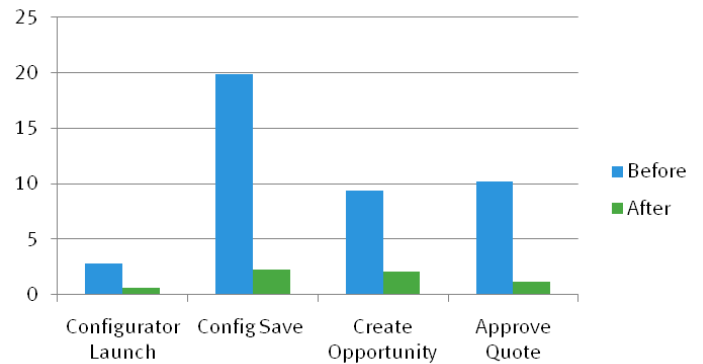
- For Customer Service (CSI) actions, the “Renewal Price Quote” transactions improved by over 200 percent from 11.5 to 4.5 seconds.
- For Channel Express (CXP) actions, the “Place Order” transactions improved by over 2,300 percent from 27.8 to 1.2 seconds.
- For Direct Express (DXP) actions, the “Price a Quote” transactions improved by over 1,000 percent from 196.3 to 18.2 seconds. The “Save Configuration” transactions improved by 780 percent from 27.4 to 3.9 seconds.

The following figures, from an EMC business view, illustrate the improvements gained in EMC’s Oracle CRM 11i (11.5.10 CU2) transactional and batch business processes.

CXP Transaction Times (sec)




DXP Transaction Times (sec)

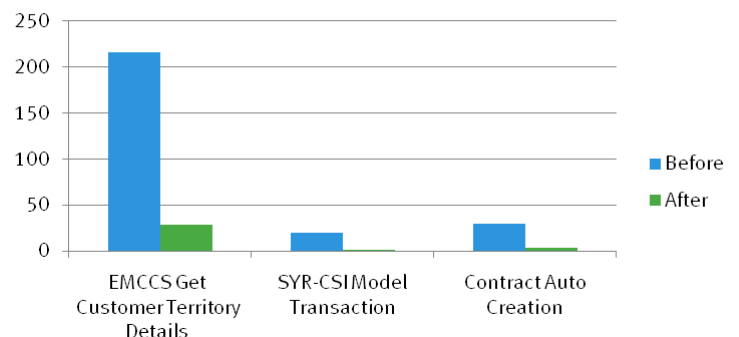


★ 50%-90% reduction in times for online transactions (i.e. 2-10 times faster)
Figure 9. Replatform Oracle 11i online improvements

Sales Job Timings (sec)



CS Job Timings (sec)



★ 85% - 95% reduction in transaction times for the above jobs (i.e. up to 20 times faster)

ACT Stats

Replay duration will be about 4 times faster

- Total ACT Transactions : 392,806
- Used to require 1 hour of replay for every 3 hours of downtime
- Now requires 1 hour of replay for every 12 hours of downtime

Reduces Future 11i Maintenance windows by 20%

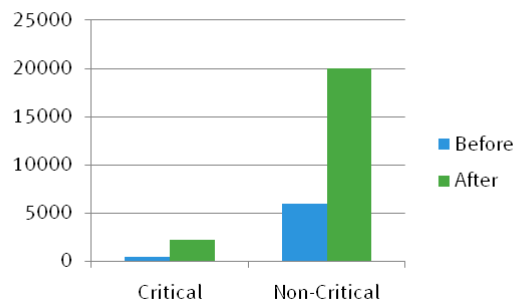
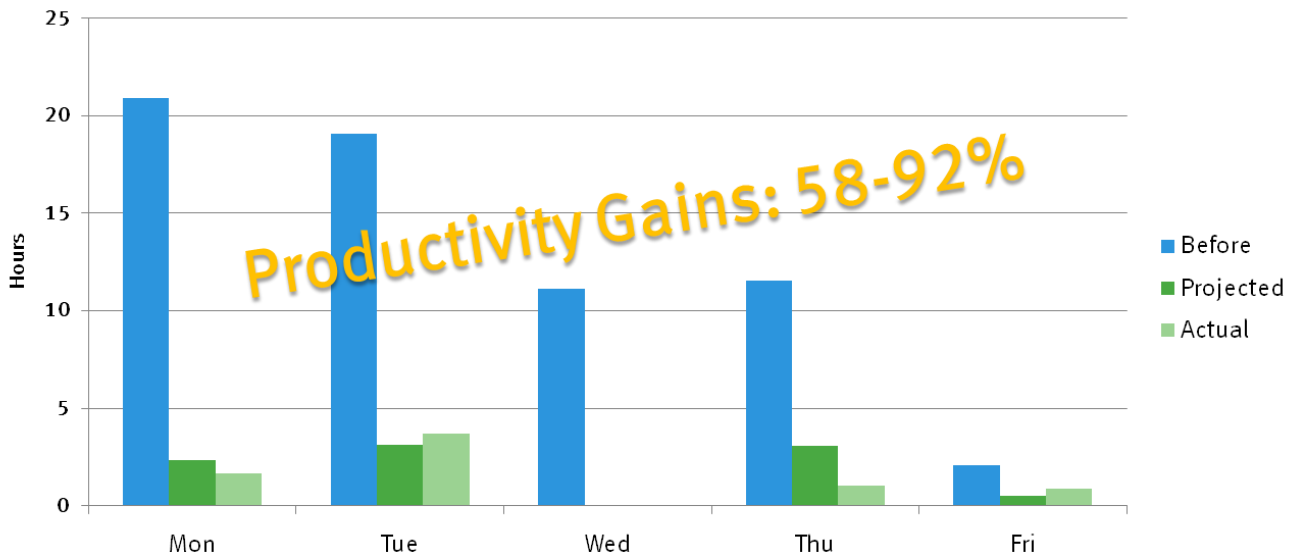


Figure 10. Replatform Oracle 11i batch improvements



Transactions Monitored

- 1) INT-CXP-Application Login
- 2) INT-CXP-Configurator Launch
- 3) INT-CXP-Config Save
- 4) INT-CXP-My Quotes
- 5) INT-DXP-Configurator Launch
- 6) INT-DXP-Config Save

Figure 11. Replatform Oracle 11i U.S. end-user experience improvements

Business cost savings benefits

The following is a summary of the cost savings realized in this migration from legacy to an open, scalable platform:

Operational and capital expense - approximately \$5 million to \$7 million dollars

- Legacy hardware replacement
- Software and hardware maintenance
- Third-party software licenses
- Environmental savings in power, cooling, and space
- Support for the environment

Conclusion

The replatform of EMC IT's production Oracle 11i CRM RAC database infrastructure to Cisco's UCS has accomplished the following for EMC IT:

- The UCS platform gives EMC IT the ability to get to an open and scalable computing platform that will provide a strategic architecture for EMC IT and create the following:
 - Incremental growth as needed with the UCS platform and Oracle RAC
 - An “on-ramp” to enable acceleration to EMC’s Journey to the Private Cloud computing. This enabled EMC IT to use UCS as a foundation platform for “IT as a Service.”
- EMC business performance gains from 1x to greater than 20x faster times for both CRM and batch transactions.
- A working, documented EMC IT procedure to migrate its CRM Oracle Applications 8 TB production database to another platform (from Solaris/SPARC to Linux/UCS/x86).
- \$5 million to \$7 million in CAPEX/OPEX savings via the following:
 - Data center environmentals (power, cooling, and space)
 - Software licenses
 - Maintenance (software/hardware)
 - Support for the environment

References

The following URLs direct you to more detailed information on the following subjects:

Cisco Unified Computing System (UCS) page on the Cisco website

<http://www.cisco.com/en/US/netsol/ns944/index.html>

Cisco UCS, Oracle, and EMC

Oracle Design Guide *Deploying Oracle Real Application Clusters on Cisco UCS*

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns1071/ucs_oracle.html

EMC IT Oracle BI grid migration

EMC white paper *EMC IT’s Migration to the Open, Expandable Oracle BI Grid*

<http://www.emc.com/collateral/hardware/white-papers/h7063-emc-it-migration-oracle-bi-grid-wp.pdf>

HugePages

“Administering Oracle Database on Linux”

http://download.oracle.com/docs/cd/B19306_01/server.102/b15658/appc_linux.htm

“Performance Tuning: HugePages in Linux”

<http://www.pythian.com/news/1326/performance-tuning-hugepages-in-linux/#more-1326>

“Configuring Linux Hugepages for Oracle Database Is Just Too Difficult! Isn’t It? Part I”

<http://kevinclosson.wordpress.com/2010/09/28/configuring-linux-hugepages-for-oracle-database-is-just-too-difficult-part-i/>

“Configuring Linux Hugepages for Oracle Database Is Just Too Difficult! Isn’t It? Part II”

<http://kevinclosson.wordpress.com/2010/10/21/configuring-linux-hugepages-for-oracle-database-is-just-too-difficult-isn%E2%80%99t-it-part-%E2%80%93ii/>

“Configuring Linux Hugepages for Oracle Database Is Just Too Difficult! Isn’t It? Part III”

<http://kevinclosson.wordpress.com/2010/11/18/configuring-linux-hugepages-for-oracle-database-is-just-too-difficult-isn%E2%80%99t-it-part-%E2%80%93iii-do-you-really-want-to-configure-the-absolute-minimum-hugepages/>

EMC technologies best practices with Oracle

Oracle Databases on EMC Symmetrix Storage Systems TechBook

<http://www.emc.com/collateral/hardware/solution-overview/h2603-oracle-db-emc-symmetrix-stor-sys-wp-ldv.pdf>

EMC white paper *EMC Symmetrix VMAX Using EMC SRDF/TimeFinder and Oracle Database 10g/11g*

<http://www.emc.com/collateral/hardware/white-papers/h6210-symmetrix-vmax-srdf-timefinder-oracle-database-wp.pdf>

Acknowledgments

The author would like to thank EMC IT’s replatform teams for assistance in the creation of this white paper.