



Cisco UCS Migration Guide Oracle PeopleSoft Migration Solaris to RHEL

February 2012



Contents

1. Executive Summary.....	4
2. Cisco Unified Computing System	4
Main Differentiating Technologies.....	4
2.1 Unified Fabric.....	4
2.2 Embedded Multirole Management.....	4
2.3 Cisco Extended Memory Technology	5
2.4 Cisco Data Center Virtual Machine Fabric Extender Virtualization Support and Virtualization Adapter	5
2.5 Dynamic Provisioning with Service Profiles	5
2.6 Cisco UCS Manager	5
2.7 Cisco UCS Fabric Interconnects.....	7
2.8 Cisco UCS 2100 and 2200 Series Fabric Extenders	8
2.9 Cisco UCS M81KR Virtual Interface Card	8
2.10 Cisco UCS 5100 Series Blade Server Chassis.....	8
2.11 Cisco UCS B200 M2 Blade Servers	9
2.12 Cisco UCS B250 M2 Extended Memory Blade Servers.....	9
3. EMC VNX Storage	9
3.1 EMC VNX Storage Platforms	9
3.2 EMC FAST Cache Technology	10
3.3 EMC FAST Virtual Pools.....	11
4. Background—Why Migrate?	12
5. The Migration: Solaris to Linux	13
5.1 Prerequisites.....	16
5.2 Target	17
5.3 Solaris Kernel Setting:	18
5.4 Linux Kernel Setting:.....	18
6. Preparing Source System	18
7. Creating Target System.....	19
8. Export Source Database Metadata.....	21
9. Perform the Transport.....	21
10. Import Tablespaces into Target Database	22
11. Import Source Database Metadata into Target Database	26
12. Migrating the Web and Application Layer	26
13. Post-Migration Activities.....	26
14. Validation of the Migrated Environment.....	27
Appendix A.....	29
Hardware and Software	29
Appendix B.....	31
Installation of Oracle PeopleSoft on Sun Solaris	31
Installations in Web Server.....	31
Installations in Application Server.....	31
Installations in Database Server	32
Configuration Sequence	32
Appendix C.....	32
Appendix D.....	35
Installation of PeopleSoft on Cisco UCS	35

Appendix E	38
Storage Configuration	38
Multipathing	39
14.1 PowerPath/VE:	39
14.2 Configuring LUNs with Unisphere.....	40
Appendix F	42
Configuring Cisco Nexus Switches	42
Creating Zone and Zoneset	43
Appendix G.....	44
Discover LUNs and Install RHEL	44
References Links	51
References Documents	52

1. Executive Summary

Enterprise Resource Planning (ERP) is a business-critical application that has a wide-ranging, intra-company impact across business units, departments, and functions. Organizations deploying ERP software, such as Oracle PeopleSoft, are under pressure to reduce cost, minimize risk, control change by accelerating deployments, and increase the availability of their application.

Increasing complexity in the data center is pushing companies to release products designed to help organizations standardize, simplify, and automate the management of these systems. Due to the nature and scope of an ERP application, it has organization-wide implications, ranging from tactical cost-cutting targets to long-term strategic initiatives.

Organizations are dramatically downsizing their average deployed server, while radically increasing the total number of servers. In some cases, this includes a migration of the PeopleSoft database or the underlying database operating system. This provides organizations with an even greater opportunity for cost-savings and business-value generation.

Apart from cost-cutting targets, capacity planning helps organizations to better use and optimize memory, boost overall CPU resource utilization, and simplify deployments that take a large amount of time. Overall, the greatest advantage is to be able to predict IT costs more accurately.

2. Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS™) is a set of preintegrated data center components that includes blade servers, adapters, fabric interconnects, and extenders integrated under a common embedded management system. This model results in far fewer system components and much better manageability, operation efficiency, and flexibility than comparable data center platforms.

Main Differentiating Technologies

The main differentiating technologies described here are what make Cisco UCS unique and give it advantages over competing offerings. The technologies presented here are high level, and the discussions do not include the technologies (such as Fibre Channel over Ethernet [FCoE]) that support these high-level elements.

2.1 Unified Fabric

Unified fabric can dramatically reduce the number of network adapters, blade-server switches, cables, and management touch points by passing all network traffic to parent fabric interconnects, where it can be prioritized, processed, and managed centrally. This approach improves performance, agility, and efficiency and dramatically reduces the number of devices that need to be powered, cooled, secured, and managed.

2.2 Embedded Multirole Management

Cisco UCS Manager is a centralized management application that is embedded on the fabric switch. Cisco UCS Manager controls all Cisco UCS elements within a single redundant management domain. These elements include all aspects of system configuration and operation, eliminating the need to use multiple, separate element managers for each system component. Massive reduction in the number of management modules and consoles and in the proliferation of agents resident on all the hardware (which must be separately managed and updated) are important deliverables of Cisco UCS. Cisco UCS Manager, using role-based access and visibility, helps enable cross-function communication efficiency, promoting collaboration between data center roles for increased productivity.

2.3 Cisco Extended Memory Technology

Significantly enhancing the available memory capacity of some Cisco UCS servers, Cisco® Extended Memory Technology helps increase performance for demanding virtualization and large-data-set workloads. Data centers can now deploy very high virtual machine densities on individual servers as well as provide resident memory capacity for data bases that need only two processors but can dramatically benefit from more memory. The high-memory dual in-line memory module (DIMM) slot count also lets users more cost-effectively scale this capacity using smaller, less costly DIMMs.

2.4 Cisco Data Center Virtual Machine Fabric Extender Virtualization Support and Virtualization Adapter

With Cisco Data Center Virtual Machine Fabric Extender (VM-FEX), virtual machines have virtual links that allow them to be managed in the same way as physical links. Virtual links can be centrally configured and managed without the complexity of traditional systems that interpose multiple switching layers in virtualized environments. I/O configurations and network profiles move along with virtual machines, helping increase security and efficiency while reducing complexity. Cisco Data Center VM-FEX helps improve performance and reduce network interface card (NIC) infrastructure.

2.5 Dynamic Provisioning with Service Profiles

Cisco UCS Manager delivers service profiles, which contain abstracted server-state information, creating an environment in which everything unique about a server is stored in the fabric, and the physical server is simply another resource to be assigned. Cisco UCS Manager implements role- and policy-based management focused on service profiles and templates. These mechanisms fully provision one or many servers and their network connectivity in minutes, rather than hours or days.

2.6 Cisco UCS Manager

Cisco UCS Manager is an embedded, unified manager that provides a single point of management for Cisco UCS. Cisco UCS Manager can be accessed through an intuitive GUI, a command-line interface (CLI), or the comprehensive open XML API. It manages the physical assets of the server and storage and LAN connectivity, and it is designed to simplify the management of virtual network connections through integration with several major hypervisor vendors. It provides IT departments with the flexibility to allow people to manage the system as a whole, or to assign specific management functions to individuals based on their roles as managers of server, storage, or network hardware assets. It simplifies operations by automatically discovering all the components available on the system and enabling a stateless model for resource use.

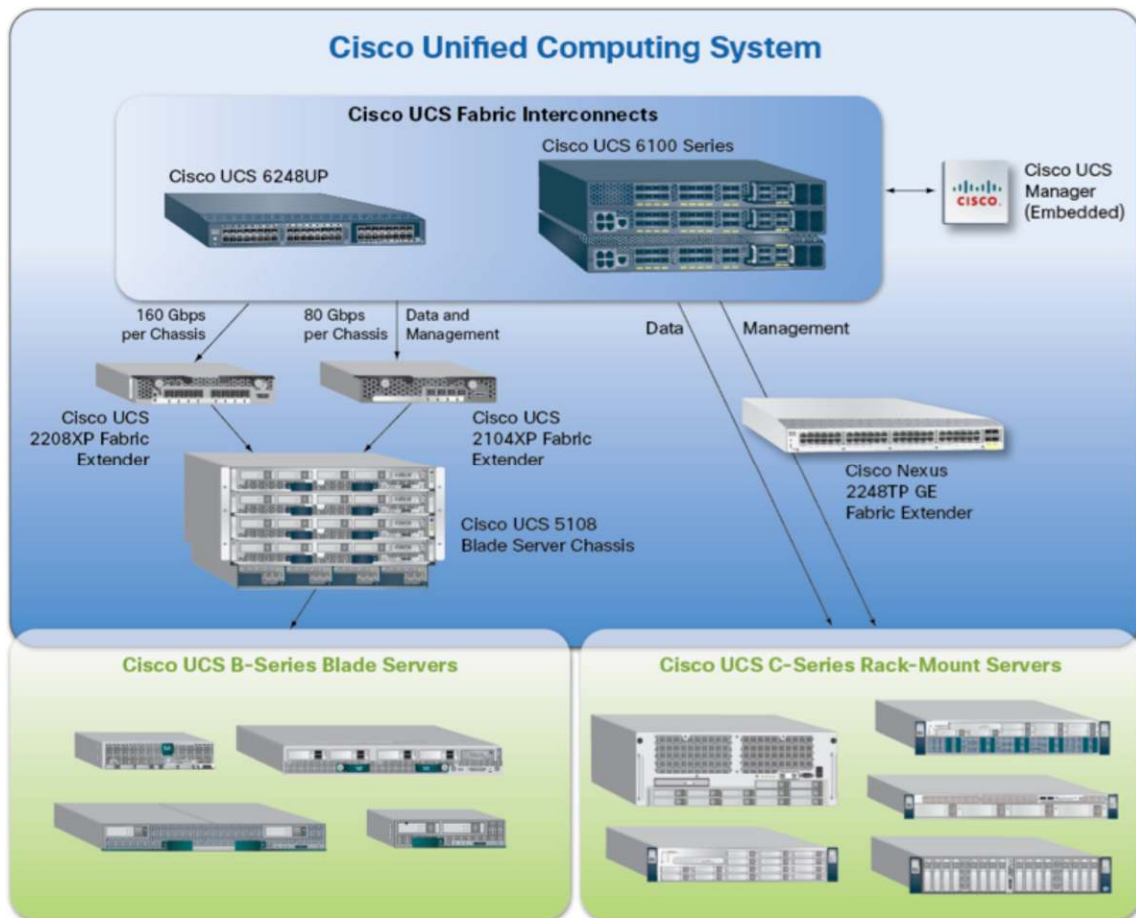
The elements managed by Cisco UCS Manager include:

- Cisco Integrated Management Controller (IMC) firmware
- RAID controller firmware and settings
- BIOS firmware and settings, including server universal user ID (UUID) and boot order
- Converged network adapter (can) firmware and settings, including MAC and worldwide name (WWN) addresses and SAN boot settings
- Virtual port groups used by virtual machines, using Cisco Data Center VM-FEX technology
- Interconnect configuration, including uplink and downlink definitions, MAC and WWN address pinning, VLANs, VSANs, quality of service (QoS), bandwidth allocations, Cisco Data Center VM-FEX settings, and EtherChannels to upstream LAN switches

Cisco Unified Computing System Components

Figure 1 shows the Cisco UCS components.

Figure 1. Cisco UCS Components



Cisco UCS is designed from the start to be programmable and self-integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards (VICs), even the number and type of I/O interfaces is programmed dynamically, making every server ready to power any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration and associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation speeds provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender Technology (FEX Technology) reduces the number of system components that need to be purchased, configured, managed, and maintained by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and virtual machines. Virtual networks are now managed exactly as physical networks are, but with

massive scalability. This approach represents a radical simplification compared to traditional systems, reducing capital and operating costs while increasing business agility, simplifying and accelerating deployment, and improving performance.

2.7 Cisco UCS Fabric Interconnects

Cisco UCS fabric interconnects create a unified network fabric throughout Cisco UCS. They provide uniform access to both networks and storage, eliminating the barriers to deployment of a fully virtualized environment based on a flexible, programmable pool of resources. Cisco fabric interconnects comprise a family of line-rate, low-latency, lossless 10 Gigabit Ethernet, IEEE Data Center Bridging (DCB), and FCoE interconnect switches. Based on the same switching technology as the Cisco Nexus® 5000 Series Switches, Cisco UCS 6100 Series Fabric Interconnects provide additional features and management capabilities that make them the central nervous system of Cisco UCS. The Cisco UCS Manager software runs inside the Cisco UCS fabric interconnects. The Cisco UCS 6100 Series Fabric Interconnects expand the Cisco UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS blade server chassis. All chassis and all blades that are attached to interconnects are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6100 Series provides the flexibility to support LAN and SAN connectivity for all blades within its domain at configuration time. Typically deployed in redundant pairs, Cisco UCS fabric interconnects provide uniform access to both networks and storage, facilitating a fully virtualized environment.

The Cisco UCS fabric interconnect portfolio currently consists of the Cisco 6100 and 6200 Series Fabric Interconnects.

Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a one-rack-unit (1RU), 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect providing more than 1 terabit per second (Tbps) throughput with low latency. It has 32 fixed ports of Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE Enhanced Small Form-Factor Pluggable (SFP+) ports.

One expansion module slot can provide up to 16 additional Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS U6120XP 20-Port Fabric Interconnect

The Cisco UCS U6120XP 20-Port Fabric Interconnect is a 1RU, 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect providing more than 500 Gbps throughput with very low latency. It has 20 fixed 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

One expansion module slot can be configured to support up to six additional 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS U6140XP 40-Port Fabric Interconnect

The Cisco UCS U6140XP 40-Port Fabric Interconnect is a 2RU, 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect built to provide 1.04 Tbps throughput with very low latency. It has 40 fixed 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Two expansion module slots can be configured to support up to 12 additional 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

2.8 Cisco UCS 2100 and 2200 Series Fabric Extenders

The Cisco UCS 2100 and 2200 Series Fabric Extenders multiplex and forward all traffic from blade servers in a chassis to a parent Cisco UCS fabric interconnect over 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis or virtual machines on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the fabric interconnect. At the core of the Cisco UCS fabric extender are application-specific integrated circuit (ASIC) processors developed by Cisco that multiplex all traffic.

Up to two fabric extenders can be placed in a blade chassis.

- The Cisco UCS 2104XP Fabric Extender has eight 10GBASE-KR connections to the blade chassis midplane, with one connection per fabric extender for each of the chassis' eight half slots. This configuration gives each half-slot blade server access to each of two 10-Gbps unified fabric-based networks through SFP+ sockets for both throughput and redundancy. It has four ports connecting the fabric interconnect.
- The Cisco UCS 2208XP is first product in the Cisco UCS 2200 Series. It has eight 10 Gigabit Ethernet, FCoE-capable, and Enhances Small Form-Factor Pluggable (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis.

2.9 Cisco UCS M81KR Virtual Interface Card

The Cisco UCS M81KR VIC is unique to the Cisco UCS blade system. This mezzanine adapter is designed around a custom ASIC that is specifically intended for VMware-based virtualized systems. It uses custom drivers for the virtualized host bus adapter (HBA) and the 10 Gigabit Ethernet NIC. As is the case with the other Cisco CNAs, the Cisco UCS M81KR VIC encapsulates Fibre Channel traffic within the 10 Gigabit Ethernet packets for delivery to the fabric extender and the fabric interconnect.

The Cisco UCS VIC is also unique in its ability to present up to 128 virtual PCI devices to the operating system on a given blade. Eight of those devices are used for management, leaving 120 virtual devices available for either storage or network use. The configurations can be changed as needed using Cisco UCS Manager. To the guest operating system, each virtualized device appears to be (from the viewpoint of the operating software that is running within VMware or other virtualized environments) a directly attached device. The adapter supports Cisco Data Center VM-FEX, which allows visibility all the way through to the virtual machine. This adapter is exclusive to Cisco and is not offered outside the Cisco UCS B-Series Blade Server product line.

2.10 Cisco UCS 5100 Series Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a 6RU blade chassis that accepts up to eight half-width Cisco UCS B-Series Blade Servers or up to four full-width Cisco UCS B-Series Blade Servers, or a combination of the two. The Cisco UCS 5108 Blade Server Chassis can accept four redundant power supplies with automatic load sharing and failover and two Cisco UCS 2100 or 2200 Series Fabric Extenders. The chassis is managed by Cisco UCS chassis management controllers, which are mounted in the Cisco UCS fabric extenders and work in conjunction with Cisco UCS Manager to control the chassis and its components.

A single Cisco UCS managed domain can theoretically scale to up to 40 individual chassis and 320 blade servers. At this time, Cisco UCS supports up to 20 individual chassis and 160 blade servers.

Basing the I/O infrastructure on a 10-Gbps unified network fabric allows Cisco UCS to have a streamlined chassis with a simple yet comprehensive set of I/O options. The result is a chassis that has only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two fabric extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel

2.11 Cisco UCS B200 M2 Blade Servers

The Cisco UCS B200 M2 Blade Server is a half-slot, 2-socket blade server. The system uses two Intel Xeon p5600 series processors, up to 192 GB of double-data-rate-3 (DDR3) memory, two optional Small Form Factor (SFF) SAS/SSD disk drives, and a single CNA mezzanine slot for up to 20 Gbps of I/O throughput. The Cisco UCS B200 M2 blade server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

2.12 Cisco UCS B250 M2 Extended Memory Blade Servers

The Cisco UCS B250 M2 Extended-Memory Blade Server is a full-slot, 2-socket blade server using Cisco Extended Memory Technology. The system supports two Intel Xeon 5600 series processors, up to 384 GB of DDR3 memory, two optional SFF SAS/SSD disk drives, and two CNA mezzanine slots for up to 40 Gbps of I/O throughput. The Cisco UCS B250 M2 blade server provides increased performance and capacity for demanding virtualization and large-data-set workloads, with greater memory capacity and throughput.

2.13 Cisco UCS B230 M2 Blade Servers

The Cisco UCS B230 M2 Blade Server is a full-slot, 2-socket blade server offering the performance and reliability of the Intel Xeon processor E7-2800 product family and up to 32 DIMM slots, which support up to 512 GB of memory. The Cisco UCS B230 M2 supports two SSD drives and one CNA mezzanine slot for up to 20 Gbps of I/O throughput. The Cisco UCS B230 M2 Blade Server platform delivers outstanding performance, memory, and I/O capacity to meet the diverse needs of virtualized environments with advanced reliability and exceptional scalability for the most demanding applications.

2.14 Cisco UCS B440 M2 High-Performance Blade Servers

The Cisco UCS B440 M2 High-Performance Blade Server is a full-slot, 2-socket blade server offering the performance and reliability of the Intel Xeon processor E7-4800 product family and up to 512 GB of memory. The Cisco UCS B440 M2 supports four SFF SAS/SSD drives and two CNA mezzanine slots for up to 40 Gbps of I/O throughput. The Cisco UCS B440 M2 blade server extends Cisco UCS by offering increased levels of performance, scalability, and reliability for mission-critical workloads.

3. EMC VNX Storage

3.1 EMC VNX Storage Platforms

The EMC VNX family of storage systems represents EMC's next generation of unified storage optimized for virtualized environments. The massive virtualization and consolidation trend of servers demands a new storage technology that is dynamic and scalable. The EMC VNX series offers several software and hardware features for optimally deploying mission-critical enterprise applications.

An important distinction of this new generation of platforms is support for both block- and file-based external storage access over a variety of access protocols, including FC, iSCSI, FCoE, Network File System (NFS), and Common Internet File System (CIFS) network shared file access. Furthermore, data stored in one of these systems, whether accessed as block- or file-based storage objects, is managed uniformly using Unisphere, a web-based interface window.

EMC's new VNX storage family now supports the 2.5" SAS drives in a 2 U disk array enclosure (DAE) that can hold up to 25 drives, one of the densest offerings in the industry. For example, older-generation technology calls for storing 15 x 600 GB worth of data using the 3.5" FC drives in a 3U DAE. However, the new DAE using 25 x 600 GB drives in a 2 U footprint results in an increase of 2.5 times.

The power efficiency of the new DAEs also makes it more cost-effective to store the increased data in a more compact footprint without the need to increase power consumption and cooling. Additional information on the VNX Series is available at: <http://www.emc.com/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf>

The data points discussed in this paper were generated on a VNX 5500 model.

3.2 EMC FAST Cache Technology

In traditional storage arrays, the DRAM caches are too small to maintain the hot data for long periods of time. Very few storage arrays give an option to nondisruptively expand DRAM cache, even if they support DRAM cache expansion. FAST Cache extends the cache available to customers by up to 2 TB using Flash drives.

FAST Cache tracks the data activity temperature at a 64 KB chunk size and copies the chunks to the Flash drives when its temperature reaches a certain threshold. After a data chunk gets copied to FAST Cache, the subsequent accesses to that chunk of data will be served at Flash latencies. Eventually, when the data temperature cools down, the data chunks get evicted from FAST Cache and will be replaced by newer hot data. FAST Cache uses a simple least-recently used (LRU) mechanism to evict the data chunks.

FAST Cache is built on the premise that the overall applications' latencies can improve when most frequently accessed data is maintained on a relatively smaller-sized, but faster, storage medium, like Flash drives. FAST Cache identifies the most frequently accessed data that is temporal in nature, and copies it to Flash drives automatically and non-disruptively.

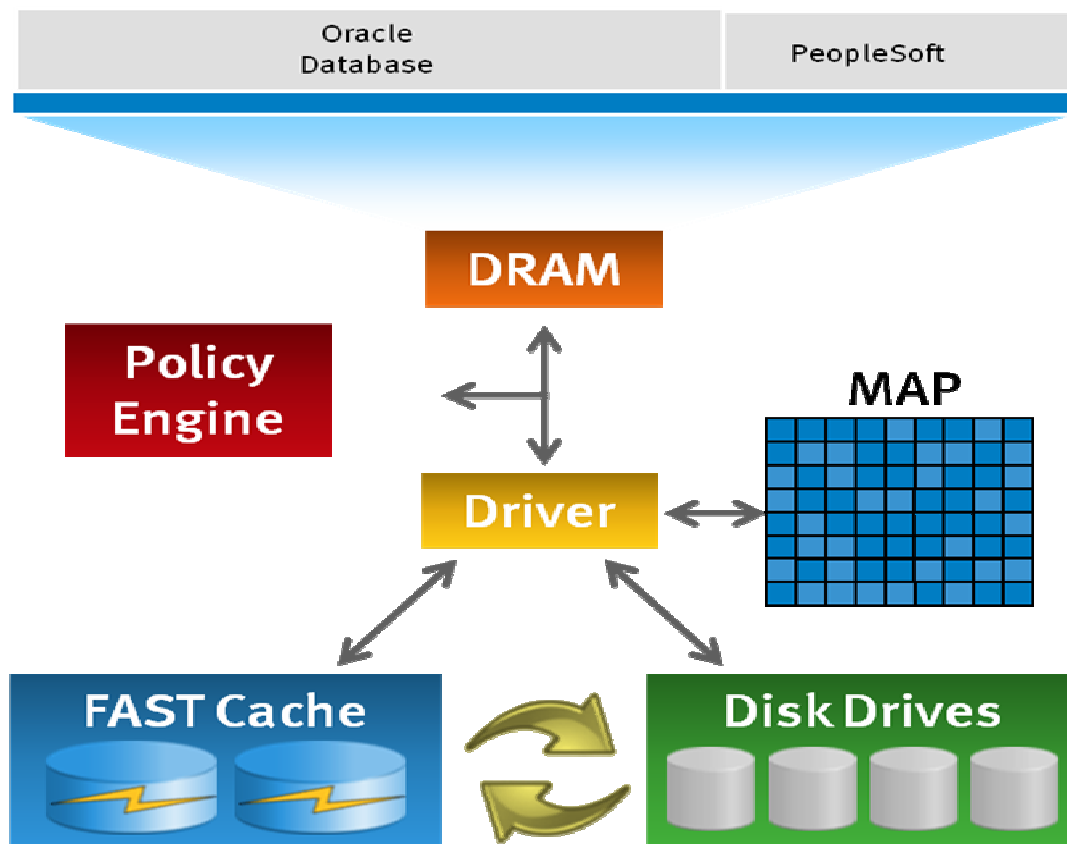
The data movement is completely transparent to applications, making this technology application-agnostic and management-free. For example, FAST Cache can be enabled or disabled on any storage pool simply by selecting or clearing the FAST Cache storage pool property in advanced settings.

FAST Cache can be selectively enabled on a few or all storage pools within a storage array, depending on application performance requirements and SLAs.

There are several distinctions to EMC FAST Cache:

- It can be configured in read/write mode, which allows the data to be maintained on a faster medium for longer periods, irrespective of application read-to-write mix and data rewrite rate.
- FAST Cache is created on a persistent medium like Flash drives, which can be accessed by both storage processors. In the event of a storage processor failure, the surviving storage processor can simply reload the cache, rather than repopulating it from the beginning, by observing the data access patterns again, which is a differentiator.
- Enabling FAST Cache is completely non-disruptive. It is as simple as selecting the Flash drives that are part of FAST Cache, and does not require any array disruption or downtime.
- Since FAST Cache is created on external Flash drives, adding FAST Cache will not consume any extra Peripheral Component Interconnect Express (PCI-E) slots inside the storage processor.

Figure 2.



Additional information on EMC Fast Cache is documented in the white paper titled “EMC FAST Cache—A Detailed Review,” which is available at: <http://www.emc.com/collateral/software/white-papers/h8046-clariion-celerra-unified-fast-cache-wp.pdf>

3.3 EMC FAST Virtual Pools

VNX FAST VP is a policy-based autotiering solution for enterprise applications. It operates at a granularity of 1 GB, referred to as a slice. The goal of FAST VP is to efficiently use storage tiers to reduce customer TCO by

tiering colder slices of data to high-capacity drives, such as Near-Line (NL)-SAS. It also increases performance by keeping hotter slices of data on performance drives, such as Flash drives. This occurs automatically and transparently to the host environment.

High locality of data is important to realize the benefits of FAST VP. When FAST VP relocates data, it will move the entire slice to the new storage tier. To successfully identify and move the correct slices, FAST VP automatically collects and analyzes statistics prior to relocating data. Customers can initiate the relocation of slices manually or automatically, by using a configurable, automated scheduler that can be accessed from the Unisphere management tool.

The multi-tiered storage pool allows FAST VP to fully use all three storage tiers: Flash, SAS, and NL-SAS. The creation of a storage pool allows for the aggregation of multiple RAID groups, using different storage tiers, into one object. The logical unit numbers (LUNs) created out of the storage pool can be either thickly or thinly provisioned. These pool LUNs are no longer bound to a single storage tier. Instead, they can be spread across different storage tiers within the same storage pool.

If you create a storage pool with one tier (Flash, SAS, or NL-SAS) then FAST VP has no impact on the performance of the system. To operate FAST VP, you need at least two tiers.

Additional information on EMC FAST VP for Unified Storage is documented in the white paper titled “EMC FAST VP for Unified Storage System—A Detailed Review,” which is available at:

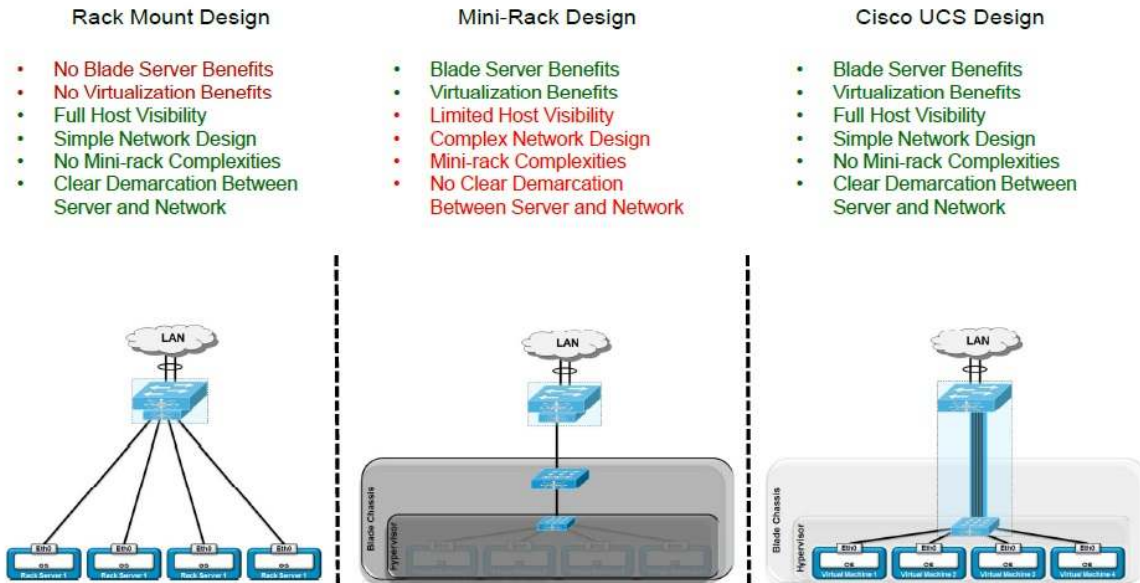
<http://www.emc.com/collateral/software/white-papers/h8058-fast-vp-unified-storage-wp.pdf>

4. Background—Why Migrate?

A massive shift is underway in the underlying computing architecture and platforms used to run enterprise applications. Traditional RISC/UNIX server platforms are not keeping pace with current demands for faster application deployments, flexible and simpler provisioning, and cost-effective licensing, support, and management.

Figure 3.

UCS Design Provides Best of Both Worlds



The racking, power, and installation of the chassis are described in the install guide (refer to http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html) and it is beyond the scope of this document. More details on each step can be found in the following documents:

- Cisco Unified Computing System CLI Configuration Guide
http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/1.4/b_UCSM_CLI_Configuration_Guide_1_4.html
- Cisco UCS Manager GUI Configuration guide
http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.4/b_UCSM_GUI_Configuration_Guide_1_4.html

The industry has moved on, as the diminished value of RISC/UNIX systems has been widely acknowledged. Cisco UCS provides innovative architectural advantages that simplify and accelerate deployment of enterprise-class applications running in bare metal, virtualized, and cloud computing environments. It offers alternative server architecture to RISC/UNIX, based on the lower-cost, high-performance x86 processor.

4.1 Hot Spare Servers: Cisco UCS sparing is another differentiated capability to improve resource use. In a typical data center, customers are required to keep one hot spare per blade. With Cisco UCS, customers can swap a failed blade with a cold spare in as little as few minutes. This eliminates the cost of having multiple spares, as well as the related software licenses required for hot spares.

4.2 Elimination of Active/Passive: With Oracle RAC for PeopleSoft, all nodes are Active/Active. Therefore when a node fails in an Oracle RAC cluster, the other nodes continue processing and Oracle clients can be configured to transparently failover to surviving nodes. Since Oracle RAC is scalable horizontally, additional nodes can be provisioned with Cisco UCS sparing.

4.3 Reduced downtime: Despite service level agreements (SLAs), downtime actually catches most organization by surprise when they measure the actual loss of unforeseen interruptions and hidden costs that cannot be measured. With the Active/Active setup, combined with UCS cold spare capability across the tiers in a Cisco UCS environment, better uptime for PeopleSoft applications can be guaranteed. In addition, a much more predictable environment can be provided for system administrators.

4.4 Horizontal scalability: Every tier in a Cisco UCS environment is horizontally scalable. Customers buy only what they need today, and add more nodes if they need faster processing in the future.

4.5 Reduced licensing cost: Organizations are increasingly looking at not only their server sprawl, but also their software licensing sprawl. Reducing the cost of licensing across tiers from the OS to the database has become a major goal, made possible by Cisco UCS. Higher memory density in x86 systems, such as Cisco UCS, allows cheaper memory modules to fulfill application memory requirements for greater memory per core. This supports larger memory-resident workloads that translate to lower licensing costs.

4.6 Maturity of Linux: For the past 5 years, with Linux being embraced by large IT organizations and the support that it has received from dedicated independent software vendors (ISVs), it is no longer unknown, and is a very stable and matured operating system. This has increased the penetration of Linux as an alternative to traditional proprietary operating systems.

4.7 Optimized footprint (power-cooling-rack space): Every data center customer has challenges in power, cooling, and space, and improvements to these features can greatly boost efficiency. The unique, end-to-end Cisco FCoE solution, extending from the server to the storage array and from access to the core, eliminates the need for a duplicate storage-networking layer. It reduces cabling, the number of ports and switches, power consumption, cooling, space requirements, and NICs/HBAs.

Another important feature Cisco offers is Unified Port capability. It allows a single port to act as an Ethernet, FC, or FCoE port. Customers can take advantage of a single switch with this capability, instead of separate FC and Ethernet switches.

5. The Migration: Solaris to Linux


Oracle provides many options to migrate their PeopleSoft environment from one operating system to another. Two very successful methods that the Oracle PeopleSoft production shops have adopted are:

- **Import-and-export:** Data was imported into the appropriate tier according to performance characteristics and the significance of the data. After the data was imported, a team validated the data.
- **Oracle Transportable Tablespaces (TTS):** This feature allows users to move a non-system tablespace across Oracle databases. It provides an efficient and much faster way to move bulk data between databases than an export-and-import. Transporting a tablespace requires only the copying of data files from the source to the destination, then integrating the tablespace structural information called the metadata.

The diagram below shows a PeopleSoft setup on Sun Servers running Solaris.

Figure 4.

PeopleSoft SUN Solaris Deployment



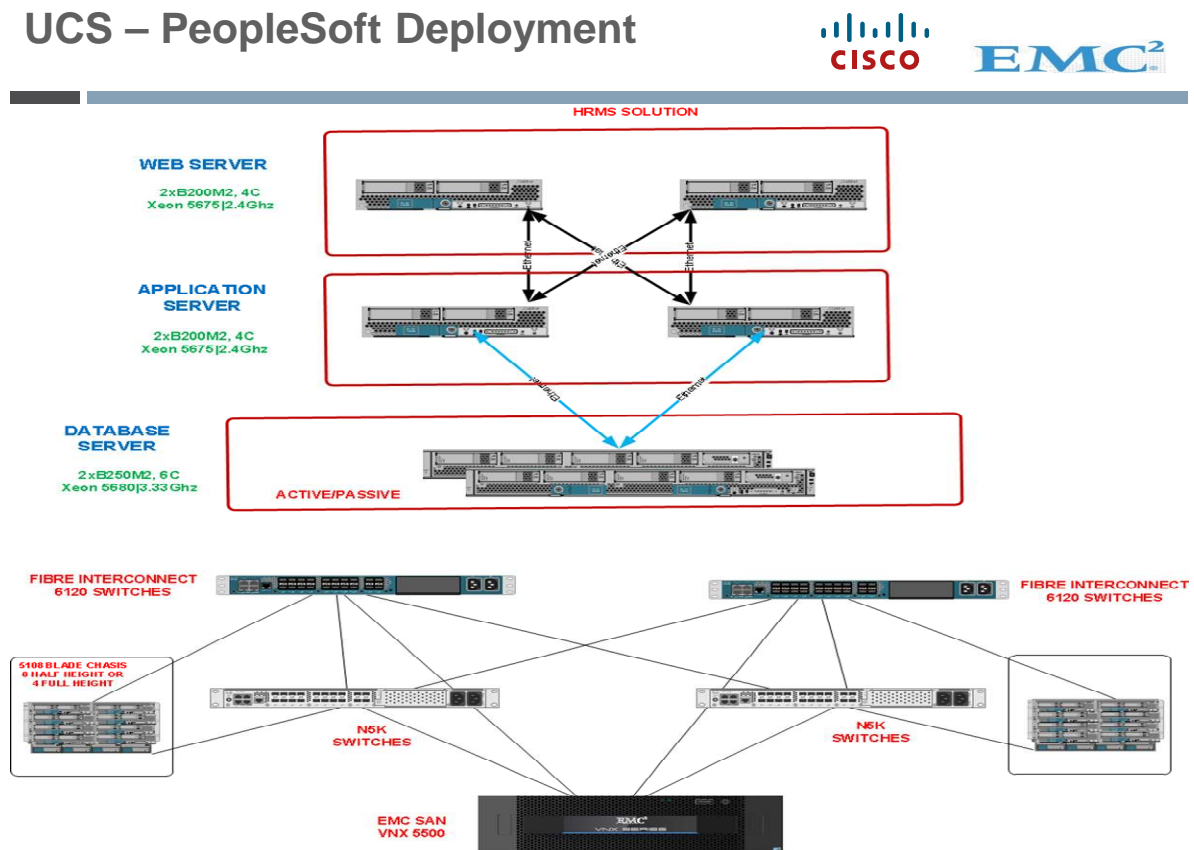
The following points show the actual steps involved in migrating a PeopleSoft Applications Database on SUN Solaris to UCS RHEL 5.6 utilizing the Oracle TTS option:

- Determine if Source (Solaris) and Target platforms (Red Hat Linux) are supported
- Determine the Endian format of Source
- Determine the support for the Target platform
- Install the Oracle Database 11g Release 2 (11.2) Software
- Purge recycle bin
- Verify objects in the SYSTEM or SYSAUX tablespaces
- Create a directory for Data Pump use
- Perform self-containment check and resolve violations
- Create database shell on Target system
- Verify database options and components used in the Source database are installed on the Target database
- Create Target database from the structure of the source database
- Create metadata required for TTS
- Drop user tablespaces

- Export Source database metadata
- Ready the Source database for transport
- Export tablespaces from Source database
- Convert and make Source data files available to Target database
- Copy Data Pump dump files to Target system
- Import tablespaces into Target database
- Make user tablespaces read/write on target database
- Import Source database metadata into Target database
- Fix sequence values
- Compile invalid objects

The diagram below shows a PeopleSoft setup on Cisco UCS Server running RHEL 5.6.

Figure 5.



In the following section, the above points are elaborated on at a much deeper technical level to help PeopleSoft database administrators. The actual flows involved in migrating a PeopleSoft Applications Database on SUN Solaris to UCS RHEL 5.6 using the Oracle TTS option are demonstrated.

5.1 Prerequisites

The following prerequisites were verified and completed to perform the cross-platform Tablespace Transport operation:

Determine if Source (Solaris) and Target Platforms (RHEL) are Supported

We determined if XTTS (Cross-platform transportable tablespace) is supported for both the Source and Target platforms, and determined the Endian format (Little or Big) of each platform.

We determined the Endian format of the Source.

```
SQL> select d.platform_name, endian_format from v$transportable_platform tp,
       v$database d where tp.platform_name = d.platform_name;
```

PLATFORM_NAME	ENDIAN_FORMAT
Solaris[tm] OE (64-bit)	Big

We determined support for the Target platform:

```
SQL> select platform_name, endian_format from v$transportable_platform;
```

PLATFORM_NAME	ENDIAN_FORMAT
Solaris[tm] OE (32-bit)	Big
Solaris[tm] OE (64-bit)	Big
Microsoft Windows IA (32-bit)	Little
Linux IA (32-bit)	Little
AIX-Based Systems (64-bit)	Big
HP-UX (64-bit)	Big
HP Tru64 UNIX	Little
HP-UX IA (64-bit)	Big
Linux IA (64-bit)	Little
HP Open VMS	Little
Microsoft Windows IA (64-bit)	Little
IBM zSeries Based Linux	Big
Linux x86 64-bit	Little
Apple Mac OS	Big
Microsoft Windows x86 64-bit	Little
Solaris Operating System (x86)	Little
IBM Power-Based Linux	Big
HP IA Open VMS	Little
Solaris Operating System (x86-64)	Little
Apple Mac OS (x86-64)	Little

Note: Target Endian was made Big, similar to Source, using the datafile conversion described later.

Install the Oracle Database 11g Release 2 (11.2) Software

Oracle 11gR2 software was installed on Target, the same as the Source system. Some of the parameters shown below are kernel-level settings for both Solaris and UCS RHEL server setup.

5.2 Target

```
Release                11.2.0.1.0
Physical Memory: Allocated 396193228 kB,
Swap Space:            31031288 kB
Disk:                  1.5 GB - 3.5 GB of disk space for Oracle software
Kernel:                2.6.18-238.el5
SELinux:               was disabled
Oracle base            "/u01/app/oracle"
Oracle home            "/u01/app/oracle/product/11.2.0/psftsmdb"
```

5.3 Solaris Kernel Setting:

```
# cat /etc/project
system:0::::
user.root:1::::
noproject:2::::
default:3::::
group.staff:10::::
user.oraclep:101::oraclep::process.max-file-
descriptor=(priv,65536,deny);process.max-sem-nsems=(priv,1024,deny);project.max-
sem-ids=(priv,1024,deny);project.max-shm-ids=(priv,1024,deny);project.max-shm-
memory=(priv,9663676416,deny)
```

5.4 Linux Kernel Setting:

```
# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 4294967296
fs.aio-max-nr = 1048576
fs.file-max = 6815744
kernel.shmmax = 64424509440
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
```

6. Preparing Source System

Purge Recycle Bin

The recycle bin was purged before export to improve the export/import performance, and to reduce required storage.

```
SQL> purge dba_recyclebin;
```

Verify Objects in the SYSTEM or SYSAUX Tablespaces

a) SYSTEM-owned objects residing in the SYSTEM or SYSAUX tablespaces

We verified that no application specific objects are in the tablespaces owned by SYSTEM.

b) User-owned tables residing in the SYSTEM or SYSAUX tablespaces

We ran the script below to verify if any user objects exist in SYSTEM or SYSAUX so that they could be moved separately.

```
SQL> @tts_system_user_obj.sql
```

We confirmed no user specific objects are found in SYSTEM or SYSAUX.

c) Gather information from the Source

The following information was gathered from the source database and used throughout this process. Following scripts are provided in the Appendix:

```
SQL> connect system/<password>
```

To drop tablespaces in the target database prior to the transport process:

```
SQL> @cr_tts_drop_ts.sql
```

To set all tablespaces to be transported to READ ONLY mode:

```
SQL> @cr_tts_tsro.sql
```

To set all tablespaces to READ WRITE mode after the transport process:

```
SQL> @cr_tts_tsrw.sql
```

To create GRANT commands to be run on the target database to give privileges that are not handled by Data Pump:

```
SQL> @cr_tts_sys_privs.sql
```

To reset the proper starting value for sequences on the target database:

```
SQL> @cr_tts_create_seq.sql
```

To Create Data Pump parameters files for:

- XTTS export (dp_ttsexp.par)
- XTTS import (dp_ttsimp.par)
- Test tablespace metadata-only export (dp_tsmeta_exp_TESTONLY.par)

```
SQL> @cr_tts_parfiles.sql
```

Create a Directory for Data Pump Use

```
SQL> connect system/<password>
```

```
SQL> create directory PUMP_DIR as '/solcrm/dump';
```

```
SQL> !mkdir /solcrm/dump
```

Perform Self-Containment Check and Resolve Violations

Make sure that all object references from the transportable set are contained in the transportable set. For example, the base table of an index must be in the transportable set, index-organized tables and their overflow

tables must both be in the transportable set, and a scoped table and its base table must be together in the transportable set.

```
SQL> @tts_check.sql
```

NOTE: After performing this step, no Data Definition Language (DDL) changes are to be made to the source database. DDL changes made to the database after the source database metadata export will not be reflected in the target database, unless handled manually.

7. Creating Target System

Created Database Shell on Target System

The target database shell was created using Database Configuration Assistant (DBCA).

When creating the Target database, the following points were covered:

- We created user tablespaces that were the same as Source, with smaller sizes as placeholders to be dropped in the process later

We edited the CreateDBFiles.sql script created during template creation, and changed the data file size for all permanent tablespaces to 1 M. For example:

```
CREATE SMALLFILE TABLESPACE "USERS" LOGGING DATAFILE SIZE 250M ...
```

Changed to:

```
CREATE SMALLFILE TABLESPACE "USERS" LOGGING DATAFILE SIZE 1M ...
```

- The sizes of the SYSTEM, SYSAUX, UNDO, and temporary tablespaces were made the same as the Source database
- The sizes of log files and number of members per log file group in the new target database were made same as the Source database
- We verified that the Source and Target database have the same character set and national character set

```
SQL> select * from database_properties  
where property_name like '%CHARACTERSET';
```

Source :

```
NLS_CHARACTERSET : AL32UTF8
```

```
NLS_NCHAR_CHARACTERSET : AL16UTF16
```

Target :

```
NLS_CHARACTERSET: AL32UTF8
```

```
NLS_NCHAR_CHARACTERSET: AL16UTF16
```

We verified the database options and components used in the source database were installed on the target database

- Query V\$OPTION to get currently installed database options.
- Query DBA_REGISTRY to get currently installed database components.

Created the Target Database from the Structure of the Source Database

We launched DBCA and clicked Next to continue to the Operations window. On the Operations window, we selected Manage Templates and clicked Next to continue to the Template Management window. We selected from an existing database (structure only), and followed the remaining windows to create a template of the existing source database.

Created Database Link and Directory for Data Pump

On the Target database, we created a database link from the Target system to the Source system, plus a directory for Data Pump use.

```
SQL> connect system/<password>
SQL> create database link ttslink using 'mig';
SQL> create directory PUMP_DIR as '/solcrm/dump';
SQL> !mkdir /solcrm/dump
```

Created Metadata Required for XTTS

We ran Data Pump on the Target system to import the database metadata necessary for the transportable import.

```
$ impdp system/password DIRECTORY= PUMP_DIR LOGFILE=dp_userimp.log
NETWORK_LINK=mig FULL=y INCLUDE=USER,ROLE,ROLE_GRANT,PROFILE
```

Drop User Tablespaces

We dropped the placeholder tablespaces in the Target database that were created when the Target database was initially created by DBCA. Tablespace USERS was the default permanent tablespace, so we changed the database default permanent tablespace.

```
SQL> select property_value from database_properties where
property_name='DEFAULT_PERMANENT_TABLESPACE';
PROPERTY_VALUE
-----
USERS

SQL> alter database default tablespace SYSTEM;
Database altered.
Dropped all user tablespaces, running the tts_drop_ts.sql
SQL> @tts_drop_ts.sql
```

8. Export Source Database Metadata

We exported all metadata from the source database, and made sure no DDL was performed after this step.

```
$ expdp system/password DIRECTORY= PUMP_DIR LOGFILE=dp_fullexp_meta.log
DUMPFILE=dp_full.dmp FULL=y CONTENT=METADATA_ONLY
EXCLUDE=USER,ROLE,ROLE_GRANT,PROFILE
```

9. Perform the Transport

Make Sure that the Source Database is Ready for Transport

Disconnect the Users and Restrict Access to Source Database

```
SQL> alter system enable restricted session;
SQL> alter system disconnect session '<SID>,<SERIAL#>';
```

Make All User Tablespaces READ ONLY

```
SQL> @tts_tsro.sql
```

Gather Sequence Information

Proper sequence starting values need to be captured from the source database, which is used to recreate sequences in the target database with the correct starting values.

```
SQL> @cr_tts_create_seq.sql
```

Transport the User Tablespaces

We followed the steps below to perform the tablespace transport:

Export Tablespaces from Source Database

We exported the user tablespace metadata from the Source database.

```
$ expdp system/password PARFILE=dp_ttsexp.par
```

Convert and Make Source Datafiles Available to Target Database

Once the Source tablespaces were placed in READ ONLY mode, the datafiles were made available to the Target database.

Note: As Endian was different on the Source and the Target, the data file conversion was made on the Target. Also the data files were converted into Automatic Storage Management (ASM) format on the Target.

1. We transferred the original data files to a staging area on the target system
Location : solcrm/dump (NFS mount point shared between both systems.)
2. We ran RMAN CONVERT DATAFILE on the target system to convert the datafiles to the new Endian format, and place the converted copy in the final destination on the Target system. The data files of all tablespaces being transported were specified.

```
RMAN> CONVERT DATAFILE
'/solcrm/dump/SBL_DATA_16.dbf',
'/solcrm/dump/sbl_data_81_3.dbf',
'/solcrm/dump/sbl_data_81_4.dbf',
'/solcrm/dump/sbl_data_81_1.dbf',
'/solcrm/dump/sbl_data_81.dbf',
'/solcrm/dump/sbl_data_81_2.dbf',
'/solcrm/dump/sbl_indx_81_3.dbf',
'/solcrm/dump/sbl_indx_81_1',
'/solcrm/dump/sbl_indx_81.dbf',
'/solcrm/dump/sbl_indx_81_2.dbf',
'/solcrm/dump/sbl_indx_81_4.dbf',
'/solcrm/dump/SBL_IND_16.dbf',
'/solcrm/dump/users01.dbf'
FROM PLATFORM 'Solaris[tm] OE (64-bit)'
PARALLELISM 4
DB_FILE_NAME_CONVERT '/solcrm/dump/', '+DATA/psftsmdb/' ;
```

Run script cr_rman_df_convert.sql

Moving the Datafiles

NFS mount is already shared between Source and Target. Converted ASM files are within the NFS mountpoint only.

Copy Data Pump Dump Files to Target System

We copied the dump to the Target system in the shared folder solcrm/dump, which is also used by the Target.

10. Import Tablespaces into Target Database

We imported the user tablespaces into the Target database.

Note: The dp_ttsimp.par file contains a list of datafiles to be transported into the Target database. The contents of the file were generated from the Source database, including datafile names. The datafile paths specified in the file must be changed to reflect the location where the datafiles exist on the Target database.

```
$ impdp system/password PARFILE=dp_ttsimp.par
Parfile look like, modify the ASM path of datafiles
directory=PUMP_DIR
dumpfile=dp_tts.dmp
logfile=dp_ttsimp.log
transport_datafiles= '+DATA/psftsmdb/aaapp.dbf',
'+DATA/psftsmdb/aalarge.dbf',
'+DATA/psftsmdb/adapp.dbf',
'+DATA/psftsmdb/amapp.dbf',
'+DATA/psftsmdb/avapp.dbf',
'+DATA/psftsmdb/bdapp.dbf',
'+DATA/psftsmdb/bnapp.dbf',
'+DATA/psftsmdb/bnlarge.dbf',
'+DATA/psftsmdb/ccapp.dbf',
'+DATA/psftsmdb/coapp.dbf',
'+DATA/psftsmdb/cuaudit.dbf',
'+DATA/psftsmdb/cularg1.dbf',
'+DATA/psftsmdb/cularg2.dbf',
'+DATA/psftsmdb/cularg3.dbf',
'+DATA/psftsmdb/cularge.dbf',
'+DATA/psftsmdb/diapp.dbf',
'+DATA/psftsmdb/dtapp.dbf',
'+DATA/psftsmdb/eoapp.dbf',
'+DATA/psftsmdb/eobfapp.dbf',
'+DATA/psftsmdb/eocfapp.dbf',
'+DATA/psftsmdb/eocmapp.dbf',
'+DATA/psftsmdb/eocmlrg.dbf',
'+DATA/psftsmdb/eocmwrk.dbf',
'+DATA/psftsmdb/eocuapp.dbf',
'+DATA/psftsmdb/eoculrg.dbf',
'+DATA/psftsmdb/eodsapp.dbf',
'+DATA/psftsmdb/eodslrg.dbf',
'+DATA/psftsmdb/eoecapp.dbf',
```

'+DATA/psftsmdb/eoeclrg.dbf',
'+DATA/psftsmdb/eocwrk.dbf',
'+DATA/psftsmdb/eoeiapp.dbf',
'+DATA/psftsmdb/eoeilrg.dbf',
'+DATA/psftsmdb/eoewapp.dbf',
'+DATA/psftsmdb/eoewlrg.dbf',
'+DATA/psftsmdb/eoewwrk.dbf',
'+DATA/psftsmdb/eoiuapp.dbf',
'+DATA/psftsmdb/eoiulrg.dbf',
'+DATA/psftsmdb/eoiuwrk.dbf',
'+DATA/psftsmdb/eolarge.dbf',
'+DATA/psftsmdb/eoltapp.dbf',
'+DATA/psftsmdb/eoppapp.dbf',
'+DATA/psftsmdb/eopplrg.dbf',
'+DATA/psftsmdb/eotpapp.dbf',
'+DATA/psftsmdb/eotplrg.dbf',
'+DATA/psftsmdb/epapp.dbf',
'+DATA/psftsmdb/eplarge.dbf',
'+DATA/psftsmdb/erapp.dbf',
'+DATA/psftsmdb/erlarge.dbf',
'+DATA/psftsmdb/erwork.dbf',
'+DATA/psftsmdb/faapp.dbf',
'+DATA/psftsmdb/falarge.dbf',
'+DATA/psftsmdb/fgapp.dbf',
'+DATA/psftsmdb/fglarge.dbf',
'+DATA/psftsmdb/fsapp.dbf',
'+DATA/psftsmdb/giapp.dbf',
'+DATA/psftsmdb/gpapp.dbf',
'+DATA/psftsmdb/gpdeapp.dbf',
'+DATA/psftsmdb/hpapp.dbf',
'+DATA/psftsmdb/hrapp.dbf',
'+DATA/psftsmdb/hrapp1.dbf',
'+DATA/psftsmdb/hrapp2.dbf',
'+DATA/psftsmdb/hrapp3.dbf',
'+DATA/psftsmdb/hrapp4.dbf',
'+DATA/psftsmdb/hrapp5.dbf',
'+DATA/psftsmdb/hrapp6.dbf',
'+DATA/psftsmdb/hrapp7.dbf',
'+DATA/psftsmdb/hrimage.dbf',
'+DATA/psftsmdb/hrlarg1.dbf',
'+DATA/psftsmdb/hrlarge.dbf',
'+DATA/psftsmdb/hrsapp.dbf',
'+DATA/psftsmdb/hrsarch.dbf',
'+DATA/psftsmdb/hrslarge.dbf',
'+DATA/psftsmdb/hrswork.dbf',
'+DATA/psftsmdb/hrwork.dbf',
'+DATA/psftsmdb/htapp.dbf',

```
'+DATA/psftsmdb/inapp.dbf',
'+DATA/psftsmdb/paapp.dbf',
'+DATA/psftsmdb/palarge.dbf',
'+DATA/psftsmdb/pcapp.dbf',
'+DATA/psftsmdb/pclarge.dbf',
'+DATA/psftsmdb/piapp.dbf',
'+DATA/psftsmdb/pilarge.dbf',
'+DATA/psftsmdb/piwork.dbf',
'+DATA/psftsmdb/poapp.dbf',
'+DATA/psftsmdb/psdefault.dbf',
'+DATA/psftsmdb/psimage.dbf',
'+DATA/psftsmdb/psimgr.dbf',
'+DATA/psftsmdb/psindex.dbf',
'+DATA/psftsmdb/pswork.dbf',
'+DATA/psftsmdb/ptamsg.dbf',
'+DATA/psftsmdb/ptapp.dbf',
'+DATA/psftsmdb/ptappe.dbf',
'+DATA/psftsmdb/ptaudit.dbf',
'+DATA/psftsmdb/ptcmstar.dbf',
'+DATA/psftsmdb/ptlock.dbf',
'+DATA/psftsmdb/ptprc.dbf',
'+DATA/psftsmdb/ptprjwk.dbf',
'+DATA/psftsmdb/ptrpts.dbf',
'+DATA/psftsmdb/pttbl.dbf',
'+DATA/psftsmdb/pttlrg.dbf',
'+DATA/psftsmdb/pttree.dbf',
'+DATA/psftsmdb/ptwork.dbf',
'+DATA/psftsmdb/pvapp.dbf',
'+DATA/psftsmdb/py0lrg.dbf',
'+DATA/psftsmdb/pyapp.dbf',
'+DATA/psftsmdb/pylarge.dbf',
'+DATA/psftsmdb/pywork.dbf',
'+DATA/psftsmdb/saapp.dbf',
'+DATA/psftsmdb/sacapp.dbf',
'+DATA/psftsmdb/salarge.dbf',
'+DATA/psftsmdb/srapp.dbf',
'+DATA/psftsmdb/stapp.dbf',
'+DATA/psftsmdb/stlarge.dbf',
'+DATA/psftsmdb/stwork.dbf',
'+DATA/psftsmdb/users01.dbf',
'+DATA/psftsmdb/tlapp.dbf',
'+DATA/psftsmdb/tllarge.dbf',
'+DATA/psftsmdb/tlwork.dbf',
'+DATA/psftsmdb/waapp.dbf;
```

Perform Post-Transport Actions on Target Database

Make User Tablespaces READ WRITE on Target Database

```
SQL> @tts_tsrw.sql
```

11. Import Source Database Metadata into Target Database

After the tablespaces were imported into the Target database, the remaining database metadata from the Source database was imported.

```
$ impdp system/password DIRECTORY= PUMP_DIR LOGFILE=dp_fullimp.log
DUMPFILE=dp_full.dmp FULL=y
```

We reviewed the tts_dpnet_fullimp.log file for errors. No errors were found.

Create System Privileges in Target Database

```
SQL> @tts_sys_privs.sql
```

Fix Sequence Values

Sequences may have values in the Target database that do not match the Source database, because the sequences were referenced after the dictionary export was created. The supported method of resetting a sequence to a different starting value is to drop and recreate the sequence. The script tts_create_seq.sql, created in an earlier step in Phase 3, IS be used to drop and recreate sequences based on the values in the Source database.

```
SQL> @tts_create_seq.sql
```

```
SQL> @tts_create_seq.sql
```

Compile Invalid Objects

```
SQL> @?/rdbms/admin/utlrp.sql
```

```
Compiled all invalid objects
```

Once the transport process is finished, we verified that the Target database was complete and functional. The Target database is now open and available.

12. Migrating the Web and Application Layer

Few tools are available to migrate application code from Solaris to Red Hat Enterprise Linux. All test runs had been performed and a listing of all patches applied to the PeopleSoft application on the Solaris environment. As a result, it was decided to do a fresh install and reapply all the patches, bringing the application to the current required patch level.

For deploying PeopleSoft on Cisco UCS Server and RHEL, please refer to the “PeopleSoft Deployment Guide on UCS.”

13. Post-Migration Activities

Below are the post migration activities that need to be performed after the database migration from Solaris to UCS RHEL servers:

1. Update the DB name from PSHRSOL to PSFTSMDB. The result should be as below.

```
SQL> select * from ps.psdowner;
DBNAME          OWNERID
-----
PSFTSMDB        SYSADM
```

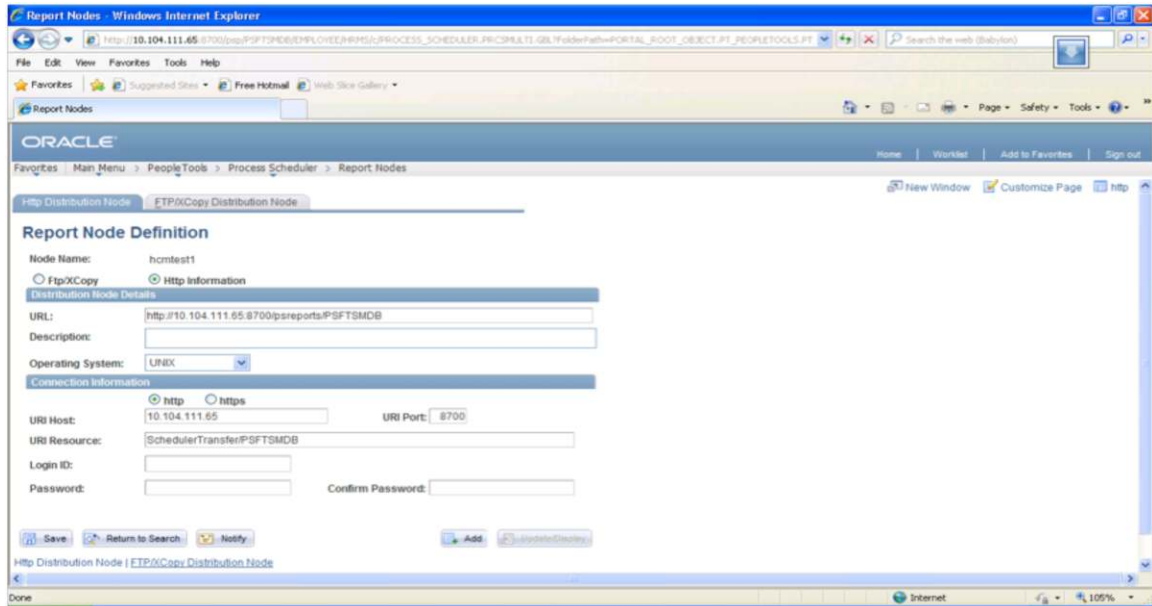
2. Change the report nodes configuration as shown below:

Login to the application using PeopleSoft user credentials.

Go to the navigation:

PeopleTools->process scheduler->report nodes

Change the settings as shown in the screen shot below:



14. Validation of the Migrated Environment

A production shop environment was simulated by writing Top-25 business (critical) scripts and checking to see if the response time matches or exceeds the old setup. The same scripts were used to check the data integrity.

Additionally we ran the queries below in both Solaris and Cisco UCS environments, and the output should be same from both environments.

```
Select count(*) from dba_objects where OWNER='SYSADM';
SELECT COUNT(*) FROM DBA_USERS;
SELECT COUNT(*) FROM DBA_TABLESPACES;
SELECT COUNT(*) FROM PS_JOB;
SELECT COUNT(*) FROM PS_PERSONAL_DATA;
etc
;;
;;;
;;;;
;;;;;
```

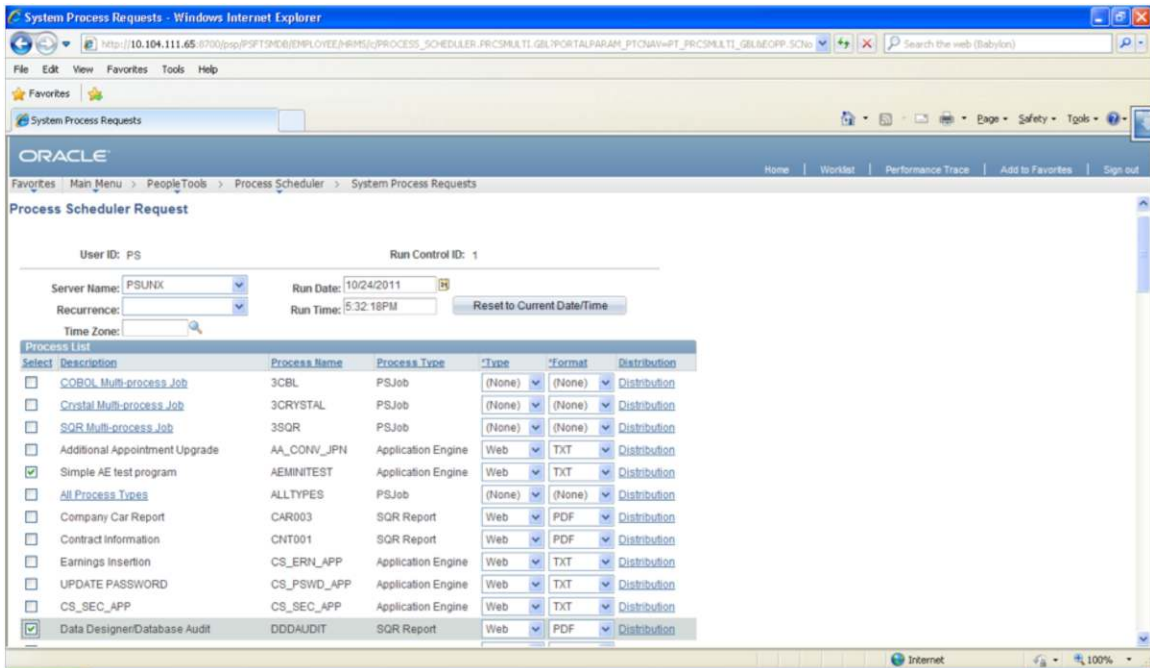
Login to the migrated environment. Run some sample reports, such as dddaudit.sqr, sysaudit.sqr, etc.

Go to the navigation:

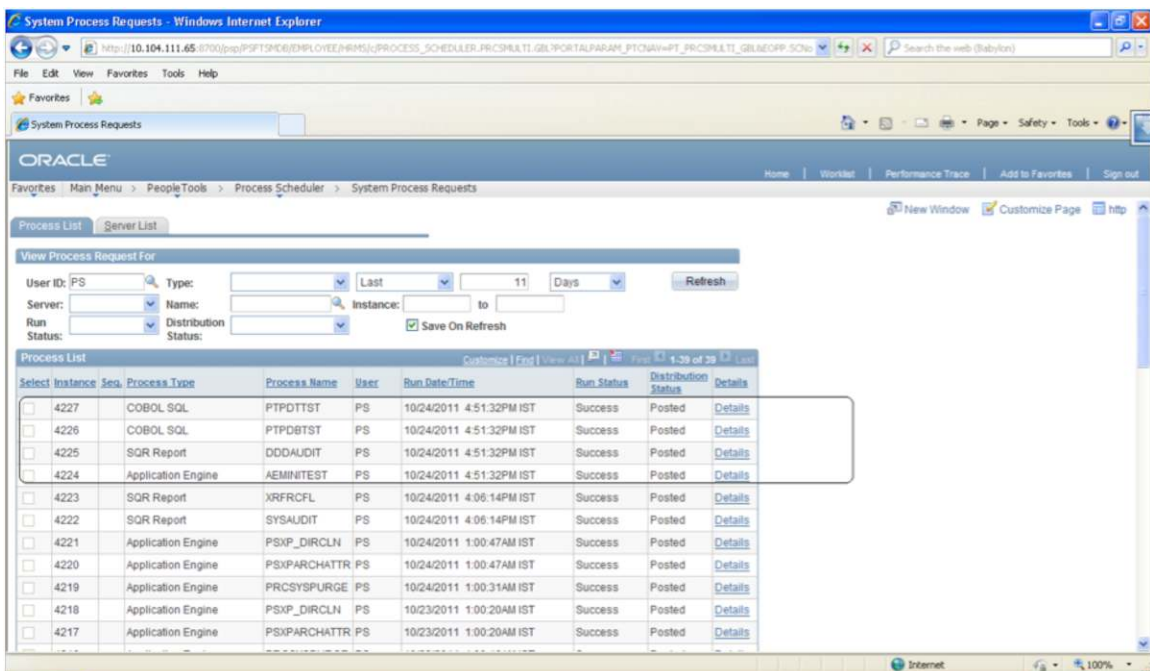
PeopleTools-> process scheduler-> system process requests->

select any run control id-> run-> now select the reports->

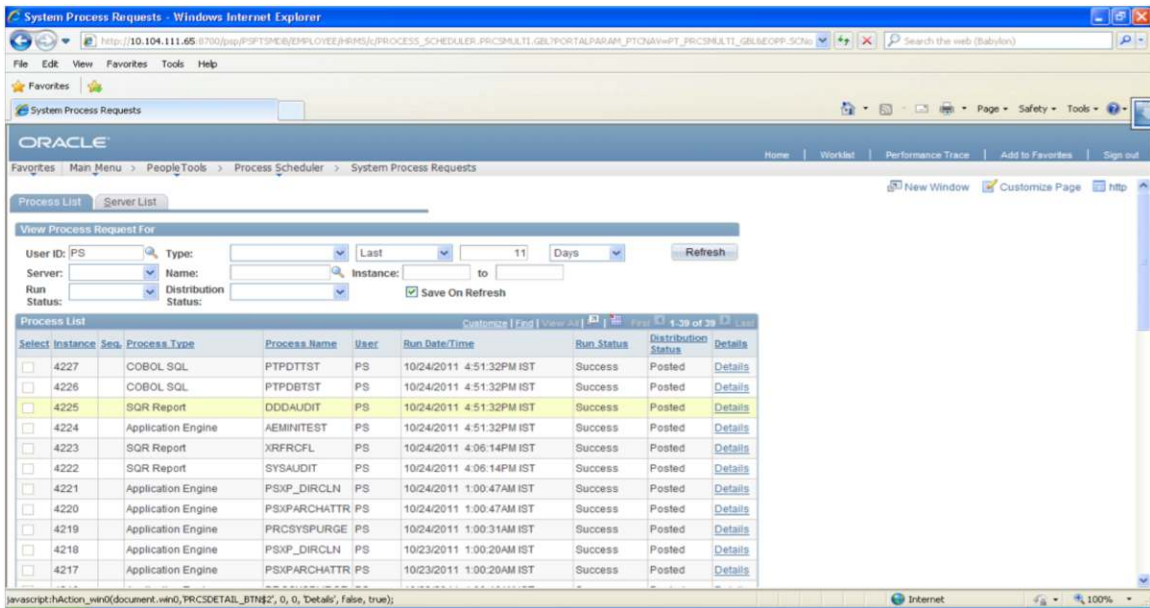
simple ae test program, data designer/database audit , simple COBOL test program, PTPDDTTST



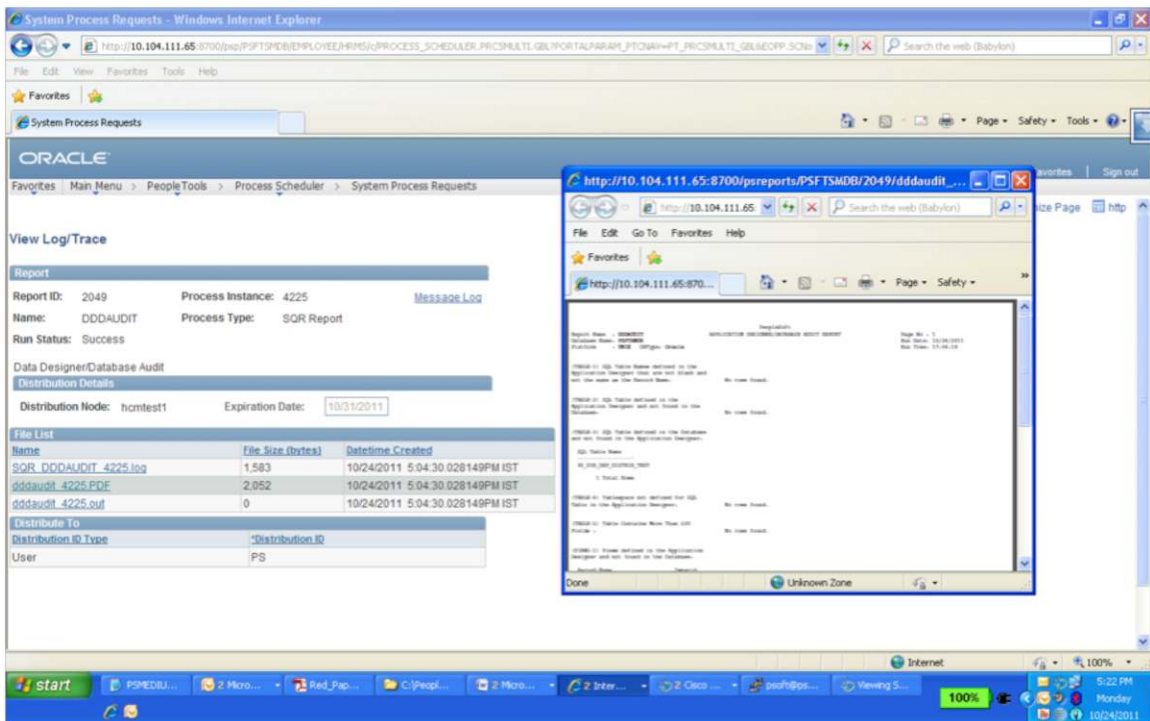
Next, go to the process monitor page and check the reports status. It should show Success and Posted.



Open a report by clicking on the details of the corresponding report. It should open without asking any userID, password, or other issue.



The report should be opened as follows:



Appendix A

Hardware and Software

For PeopleSoft Server Environment on Solaris Operating System:

Server Function	IP Address/ Hostname	Quantity	Server Model	CPU	Memory	Function/ Components	Comments
PeopleSoft Database Server	10.104.111.19 sunsiebdb	1	Sun Fire V880	8XuSPARC III+ 900 MHz	32 GB	DB Server	PeopleSoft Database server
PeopleSoft Web Server	10.104.111.99 ssr-savbu	2	Sun Fire V480R	4XuSPARC III+ 1.2 GHz	16 GB	PeopleSoft Web Server	PeopleSoft Web Server
PeopleSoft Application Server	10.104.111.16 sunsiebas1	2	Sun Fire V490	4XuSPARC IV 1.35 GHz	32 GB	PeopleSoft Application Server	PeopleSoft Application Server

For PeopleSoft Server Environment on UCS RHEL Operating System:

Server Function	Qty	Server Model	CPU	Memory	Function/ Components	Comments
PeopleSoft Database Server	1	B250 M2	2x Intel Xeon 6C X5675	96 GB	DB Server	PeopleSoft Database Server
PeopleSoft Web Server	2	B200 M2	2x Intel Xeon 4C 8T E5620	12 GB	PeopleSoft Web Server	PeopleSoft Web Server
PeopleSoft Application Server	3	B200 M2	2x Intel Xeon 4C 8T E5620	24 GB	PeopleSoft Application Server	PeopleSoft Application Server

Supported Operating System:

Operating System and Product	Minimum Patch Level	Current Environment
Solaris 10	Oracle Solaris 10 9/10 s10s_u9wos_14a SPARC	Solaris 10
Red Hat Enterprise Linux 5	Linux X86-64 > 5.6	RHEL 5.6 64-Bit

Supported Oracle People Tools and Application Releases

PeopleTools and Application Product	Operating System and Minimum Patch Level	Current Environment
PeopleTools: 8.51 (Minimum PeopleTools patch version supported is 8.51.02)	Oracle Solaris 10 9/10 s10s_u9wos_14a SPARC	PeopleTools: 8.51 with Patch 8.51.11
Application: HRMS 9.1 feature pack December 2010		Application: HRMS 9.1 feature pack December 2010
Microfocus Server Express: Microfocus Server Express 5.1 Wrap Pack4		Microfocus Server Express: Microfocus Server Express 5.1 Wrap Pack4
PeopleTools 8.51	RHEL 5.6	PeopleTools 8.51.11
HRMS Application 9.1	RHEL 5.6	HRMS 9.1 Dec 2010
Microfocus Cobol Server Express 5.1	RHEL 5.6	Micro Focus Server Express 5.1 64-bit Wrap Pack 4

Supported Web Servers

Web Server and Product	Operating System and Minimum Patch Level	Current Environment
Oracle Weblogic 10.3.4.0.0 (JDK or Jrockit needs to be installed first) Sun Java 6 Update 17 or higher; 64 bit JDK for Solaris SPARC Java Version 1.6.0_20	Oracle Solaris 10 9/10 s10s_u9wos_14a SPARC	Oracle WebLogic 10.3.4.0.0 (JDK or Jrockit needs to be installed first) Sun Java 6 Update 17 or higher; 64 bit JDK for Solaris SPARC Installed the jrockit version jrockit-jdk1.6.0_26-R28.1.4-4.0.1

Web Server and Product	Operating System and Minimum Patch Level	Current Environment
Weblogic 10g WebSphere JRE	RHEL 5.6	Oracle Weblogic 10.3.4.0.0 Installed JROCKIT jrockit28.1.4 (p12706519_2814_Linux-x86-64.zip) Java Version 1.6.0_20 Java(TM) SE Runtime Environment (build 1.6.0_20-b02) Java HotSpot(TM) 64-Bit Server VM (build 16.3-b01, mixed mode)

Supported Application Server (Tuxedo)

Web Server and Product	Operating System and Minimum Patch Level	Current Environment
Oracle Tuxedo 10gR3 minimum patch level RP031 64-bit (the patch RP065 has been installed)	Oracle Solaris 10 9/10 s10s_u9wos_14a SPARC	Oracle Tuxedo 10.3.0.0 (64 Bit)
Tuxedo 10gR3 RP031 64-bit	RHEL 5.6	Oracle Tuxedo 10.3.0.0.0

Supported Database Server

Database Server/Client and Product	Operating System and Minimum Patch Level	Current Environment
Vendor and Product: Oracle 11g Enterprise Server Version: 11.2.0.1.0 Connectivity Software: Oracle 11g Client Version: 11.1.0.6 or above	Oracle Solaris 10 9/10 s10s_u9wos_14a SPARC	Oracle 11g 11.2.0.1.0 (64 -Bit)
Oracle 11g Enterprise Server Version: 11.2.0.1.0 Connectivity Software: Oracle 11g Client	RHEL 5.6	Oracle 11.2.0.2.0

Appendix B

Installation of Oracle PeopleSoft on Sun Solaris

The sequence below shows the installation of PeopleSoft Enterprise Server:

Installations in Web Server

- Installation of JDK or jrockit jrockit-jdk1.6.0_26-R28.1.4-4.0.1 in Web Server box before doing the WebLogic server installation
- Installation of Oracle WebLogic 10.3.4.0 (64-bit mode) in Web Server box
- Installation of PeopleTools 8.51 in Web Server box
- Installation of PeopleTools patch 8.51.11 in Web Server box

Installations in Application Server

- Installation of Oracle Client 11.2.0.1.0 in Application Server box
- Installation of Oracle Tuxedo 10.3.0.0 in Application Server box
- Installation of Oracle Tuxedo Patch RP061 in Application Server box
- Installation of PeopleTools 8.51 in Application Server box
- Installation of PeopleTools Patch 8.51.11 in Application Server box

- Installation of HRMS9.1 feature pack December 2010 in Application Server box

Installations in Database Server

- Installation of Oracle Server 11.2.0.1.0 binaries in Database Server box
- Creation of Oracle/PeopleSoft database in Database Server box.
- Running PeopleSoft-delivered scripts on Oracle Database in Database Server box
- Running Datamover Setup to load PeopleSoft-delivered data in Database Server box

Configuration Sequence

- Configure Application Server in Application Server box
- Configure Web Server in Web Server box
- Configure Process Scheduler Server in Application Server box
- Configure Report Nodes in Peripheral Interface Adapter (PIA)

Appendix C

Some of the scripts that were called and used in this migration activity are listed below:

<p>TTS_SYSTEM_USER_OBJ.sql</p> <pre>select owner, segment_name, segment_type from dba_segments where tablespace_name in ('SYSTEM', 'SYSAUX') and owner not in ('SYS', 'SYSTEM', 'DBSNMP', 'SYSMAN', 'OUTLN', 'MDSYS', 'ORDSYS', 'EXFSYS', 'DMSYS', 'WMSYS', 'WKSYS', 'CTXSYS', 'ANONYMOUS', 'XDB', 'WKPROXY', 'ORDPLUGINS', 'DIP', 'SI_INFORMTN_SCHEMA', 'OLAPSYS', 'MDDATA', 'WK_TEST', 'MGMT_VIEW', 'TSMSYS');</pre>	<p>CR_TTS_DROP_TS.sql</p> <pre>set heading off feedback off trimspool on linesize 500 spool tts_drop_ts.sql prompt /* ===== */ prompt /* Drop user tablespaces */ prompt /* ===== */ select 'DROP TABLESPACE ' tablespace_name ' INCLUDING CONTENTS AND DATAFILES;' from dba_tablespaces where tablespace_name not in ('SYSTEM','SYSAUX') and contents = 'PERMANENT'; spool off</pre>
<p>CR_TTS_TSRO.sql</p> <pre>set heading off feedback off trimspool on linesize 500 spool tts_tsro.sql prompt /* ===== */ prompt /* Make all user tablespaces READ ONLY */ prompt /* ===== */ select 'ALTER TABLESPACE ' tablespace_name ' READ ONLY;' from dba_tablespaces where tablespace_name not in ('SYSTEM','SYSAUX') and contents = 'PERMANENT'; spool off</pre>	<p>CT_TTS_TSRW.sql</p> <pre>set heading off feedback off trimspool on linesize 500 spool tts_tsrw.sql prompt /* ===== */ prompt /* Make all user tablespaces READ WRITE */ prompt /* ===== */ select 'ALTER TABLESPACE ' tablespace_name ' READ WRITE;' from dba_tablespaces where tablespace_name not in ('SYSTEM','SYSAUX') and contents = 'PERMANENT'; spool off</pre>
<p>CT_TTS_SYS_PRIVS.sql</p> <pre>set heading off feedback off trimspool on escape off set long 1000 linesize 1000 col USERDDL format A150 spool tts_sys_privs.sql prompt /* ===== */ prompt /* Grant privs */ prompt /* ===== */ select 'grant ' privilege ' on "' owner '."' table_name '" to "' grantee '" </pre>	<p>CR_TTS_CREATE_SEQS.sql</p> <pre>set heading off feedback off trimspool on escape off set long 1000 linesize 1000 pagesize 0 col SEQDDL format A300 spool tts_create_seq.sql prompt /* ===== */ prompt /* Drop and create sequences */ prompt /* ===== */ select regexp_replace(dbms_metadata.get_ddl('SEQUENCE',sequence_name,sequence_own</pre>

<pre> decode(grantable,'YES',' with grant option ') decode(hierarchy,'YES',' with hierarchy option ') '; from dba_tab_privs where owner in ('SYS', 'SYSTEM', 'DBSNMP', 'SYSMAN', 'OUTLN', 'MDSYS', 'ORDSYS', 'EXFSYS', 'DMSYS', 'WMSYS', 'KSYS', 'CTXSYS', 'ANONYMOUS', 'XDB', 'WKPROXY', 'ORDPLUGINS', 'DIP', 'SI_INFORMTN_SCHEMA', 'OLAPSYS', 'MDDATA', 'WK_TEST', 'MGMT_VIEW', 'TSMSYS') and grantee in (select username from dba_users where username not in ('SYS', 'SYSTEM', 'DBSNMP', 'SYSMAN', 'OUTLN', 'MDSYS', 'ORDSYS', 'EXFSYS', 'DMSYS', 'WMSYS', 'KSYS', 'CTXSYS', 'ANONYMOUS', 'XDB', 'WKPROXY', 'ORDPLUGINS', 'DIP', 'SI_INFORMTN_SCHEMA', 'OLAPSYS', 'MDDATA', 'WK_TEST', 'MGMT_VIEW', 'TSMSYS')); spool off </pre>	<pre> er), '^.*(CREATE SEQUENCE.*CYCLE).*\$', 'DROP SEQUENCE "' sequence_owner '."' sequence_name '";' chr(10) \1;') SEQDDL from dba_sequences where sequence_owner not in ('SYS', 'SYSTEM', 'DBSNMP', 'SYSMAN', 'OUTLN', 'MDSYS', 'ORDSYS', 'EXFSYS', 'DMSYS', 'WMSYS', 'KSYS', 'CTXSYS', 'ANONYMOUS', 'XDB', 'WKPROXY', 'ORDPLUGINS', 'DIP', 'SI_INFORMTN_SCHEMA', 'OLAPSYS', 'MDDATA', 'WK_TEST', 'MGMT_VIEW', 'TSMSYS'); spool off </pre>
<pre> CR_TTS_PAR_FILES.sql REM REM Create TTS Data Pump export and import PAR files REM set feedback off trimspool on set serveroutput on size 1000000 REM REM Data Pump parameter file for TTS export REM spool dp_ttsexp.par declare tsname varchar(30); i number := 0; begin dbms_output.put_line('directory=PUMP_DIR'); dbms_output.put_line('dumpfile=dp_tts.dmp'); dbms_output.put_line('logfile=dp_ttsexp.log'); dbms_output.put_line('transport_full_check=no'); dbms_output.put_line('transport_tablespaces='); for ts in (select tablespace_name from dba_tablespaces where tablespace_name not in ('SYSTEM','SYSAUX') and contents = 'PERMANENT' order by tablespace_name) loop if (i!=0) then dbms_output.put_line(tsname ','); end if; i := 1; tsname := ts.tablespace_name; end loop; dbms_output.put_line(tsname); dbms_output.put_line(''); end; / spool off REM REM Data Pump parameter file for TTS import REM spool dp_ttsimp.par </pre>	<pre> TTS_CHECK.sql declare checklist varchar2(4000); i number := 0; begin for ts in (select tablespace_name from dba_tablespaces where tablespace_name not in ('SYSTEM','SYSAUX') and contents = 'PERMANENT') loop if (i=0) then checklist := ts.tablespace_name; else checklist := checklist ',' ts.tablespace_name; end if; i := 1; end loop; dbms_tts.transport_set_check(checklist,TRUE,TRUE); end; / select * from transport_set_violations; </pre>

```

declare
fname varchar(513);
i number := 0;
begin
dbms_output.put_line('directory=PUMP_DIR');
dbms_output.put_line('dumpfile=dp_tts.dmp');
dbms_output.put_line('logfile=dp_ttsimp.log');
dbms_output.put('transport_datafiles=+DATA/ucssmdb/');
for df in
(select file_name from dba_tablespaces a, dba_data_files b
where a.tablespace_name = b.tablespace_name
and a.tablespace_name not in ('SYSTEM','SYSAUX')
and contents = 'PERMANENT'
order by a.tablespace_name)
loop
if (i!=0) then
dbms_output.put_line('"'||'+DATA/ucssmdb'||fname||'",');
end if;
i := 1;
fname := df.file_name;
end loop;
dbms_output.put_line('"'||'+DATA/ucssmdb'||fname||'"');
dbms_output.put_line("");
end;
/
spool off
REM
REM Data Pump parameter file for tablespace metadata export
REM Only use this to estimate the TTS export time
REM
spool dp_tsmeta_exp_TESTONLY.par
declare
tsname varchar(30);
i number := 0;
begin
dbms_output.put_line('directory=PUMP_DIR');
dbms_output.put_line('dumpfile=dp_tsmeta_TESTONLY.dmp');
dbms_output.put_line('logfile=dp_tsmeta_exp_TESTONLY.log');
dbms_output.put_line('content=metadata_only');
dbms_output.put('tablespaces=');
for ts in
(select tablespace_name from dba_tablespaces
where tablespace_name not in ('SYSTEM','SYSAUX')
and contents = 'PERMANENT'
order by tablespace_name)
loop
if (i!=0) then
dbms_output.put_line(tsname||',');
end if;
i := 1;
tsname := ts.tablespace_name;
end loop;
dbms_output.put_line(tsname);
dbms_output.put_line("");
end;
/
spool off

```

CR_TTS_CR_SEQS.sql

```
set heading off feedback off trimspool on escape off
set long 1000 linesize 1000 pagesize 0
col SEQDDL format A300
spool tts_create_seq.sql
prompt /* ===== */
prompt /* Drop and create sequences */
prompt /* ===== */
select regexp_replace(
dbms_metadata.get_ddl('SEQUENCE',sequence_name,sequence_owner),
'^.*(CREATE SEQUENCE.*CYCLE).*$',
'DROP SEQUENCE "'||sequence_owner||"."||sequence_name
||";'||chr(10)||chr(1);') SEQDDL from dba_sequences
where sequence_owner not in ('SYS', 'SYSTEM', 'DBSNMP', 'SYSMAN',
'OUTLN', 'MDSYS',
'ORDSYS', 'EXFSYS', 'DMSYS', 'WMSYS', 'KSYS', 'CTXSYS',
'ANONYMOUS', 'XDB', 'WKPROXY', 'ORDPLUGINS', 'DIP',
'SI_INFORMTN_SCHEMA', 'OLAPSYS', 'MDDATA', 'WK_TEST',
'MGMT_VIEW', 'TSMSYS');
spool off
```

CR_RMAN_DF_CONVERT.sql

```
REM
REM Create RMAN CONVERT DATAFILE script for cross platform TTS
REM Use for target system conversion only
REM
set feedback off trimspool on
set serveroutput on size 1000000
spool df_convert.rman
declare
fname varchar(513);
i number := 0;
begin
dbms_output.put_line('# Sample RMAN script to perform file conversion
on all user datafiles');
dbms_output.put_line('# Datafile names taken from
DBA_DATA_FILES');
dbms_output.put_line('# Please review and edit before using');
dbms_output.put_line('CONVERT DATAFILE ');
for df in
(select substr(file_name,instr(file_name,/, -1)+1) file_name
from dba_tablespaces a, dba_data_files b
where a.tablespace_name = b.tablespace_name
and a.tablespace_name not in ('SYSTEM','SYSAUX')
and contents = 'PERMANENT'
order by a.tablespace_name)
loop
if (i!=0) then
dbms_output.put_line('"/solcrm/dump/"||fname||"');
end if;
i := 1;
fname := df.file_name;
end loop;
dbms_output.put_line('"/solcrm/dump/"||fname||"');
dbms_output.put_line('FROM PLATFORM "<Enter source platform
here>");
dbms_output.put_line('PARALLELISM 4');
dbms_output.put_line('DB_FILE_NAME_CONVERT
"/solcrm/dump/" ,">+DATA/ucssmdb/");
dbms_output.put_line(';');
end;
/
spool off
```

Note: Please check the “Cisco UCS Deployment Guide for PeopleSoft” for more information on how to setup and install PeopleSoft on Cisco UCS servers.

Appendix D

Installation of PeopleSoft on Cisco UCS

The PeopleSoft Enterprise Server requires that you create a standard UNIX system user account; for example: psoft. This account must be available on each PeopleSoft server in your enterprise under which PeopleSoft processes and components operate. Use the following guidelines to create the PeopleSoft service owner account:

- The PeopleSoft service owner account must be defined or available on each applicable server machine: on each Application Server, Web Server, and Process Scheduler

-
- The PeopleSoft owner account password must not require a change on next logon and must be set not to expire
 - The PeopleSoft owner account name or password cannot contain any spaces

Prerequisites

- A root user and a non-root user (psft) to perform any installation tasks on the Linux boxes
- Root user credentials to create the required PeopleSoft or Oracle users
- Required file systems configured to install Tuxedo, WebLogic, PeopleTools, and Oracle Database
- Required software and versions from following the PeopleSoft Certification Matrix
- All required software downloaded
- Network connectivity established between all the machines involved
- JRE version installed 1.6.0_20 on each server where PeopleSoft components will be installed
- Oracle database software installed on the machine that will host the PeopleSoft database
- Oracle client installed on all machines that will host PeopleSoft servers
- Database connectivity between machines which host PeopleSoft Servers and the Oracle server
- Umask set to 027 on the installation directory before installing

Creating PeopleSoft Service Account

The PeopleSoft Enterprise Server requires that you create a standard UNIX system user account; for example: psft. This account must be available on each PeopleSoft server in your enterprise under which PeopleSoft processes and components operate. Use the following guidelines to create the PeopleSoft service owner account:

- The PeopleSoft service owner account must be defined or available on each applicable server machine: Application Server, Web Server, and Process Scheduler
- The PeopleSoft owner account password must not require a change on next logon and must be set so that it does not expire
- The PeopleSoft owner account name or password cannot contain any spaces

General Install Requirements

Below are the general requirements to follow before installing the PeopleSoft Enterprise Server.

- Determine your load balancing strategy
- Make sure there is sufficient disk space for the installation
- Database Server software must be installed in PeopleSoft Database Server
- Database Client software must be installed in other PeopleSoft Servers, such as Application Servers and Web Servers
- Database Server and Client must be installed prior to the installation of PeopleSoft Server components, as required
- Install all third-party software required for the PeopleSoft, such as JRE, JDK, etc.
- Create Directories for the PeopleSoft Software and PeopleSoft File System in the Linux Machine
- Make sure enough temporary disk space is available for the installers and wizards

- If you are installing the PeopleSoft Products in GUI mode, you must set the DISPLAY variable to display the Java Installer User Interface on your machine
- If you are installing in console mode, the mode=-console parameter need to be specified during the installation procedure

Unix- and Linux-Specific Install Requirements

- Installation can be performed either as a root or as a non-root user. In most cases, it is recommended that installation be performed by a non-root user for simpler administration and maintenance.
- If the PeopleSoft Application Server is installed by root, then only root can stop and start the server. To avoid this requirement, you can use an account other than root that has the correct authorizations to install. All future patch releases must be installed as the same user who installed the base installation being patched.
- If web server is installed by root user, then only root can stop and start the server. To avoid this requirement, you can use an account other than root that has the correct authorizations to install. All future patch releases must be installed as the same user who installed the base installation being patched.
- Use VNC Viewer, Xterm, or Xmanager, which is a third-party software or remote access to the Linux machine for installing PeopleSoft products in GUI mode
- Installation can be done in console mode by specifying the parameter, mode=-console during the installation
- Set the user .profiles and environment variables for Oracle and PeopleSoft

Oracle PeopleSoft Web Server Installation

- Installation of JDK or jrockit jrockit-jdk1.6.0_26-R28.1.4-4.0.1 in Web Server box before doing the WebLogic Server installation
- Installation of Oracle WebLogic 10.3.4.0 (64-bit mode) in Web Server box
- Installation of PeopleTools 8.51 in Web Server box
- Installation of PeopleTools patch 8.51.11 in Web Server box

Note: Refer to Appendix C for a complete PeopleSoft Web Server installation listing in the “Cisco UCS Deployment Guide for Oracle PeopleSoft.”

Oracle PeopleSoft Application Server Installation

- Installation of Oracle Client 11.2.0.10 in Application Server box
- Installation of Oracle Tuxedo 10.3.0.0 in Application Server box
- Installation of Oracle Tuxedo Patch RP065 in Application Server box
- Installation of PeopleTools 8.51 in Application Server box
- Installation of PeopleTools Patch 8.51.11 in Application Server box
- Installation of HRMS9.1 feature pack December 2010 in Application Server box.

Note: Refer to Appendix D for a complete PeopleSoft Application server installation listing in the “Cisco UCS Deployment Guide for Oracle PeopleSoft.”

Oracle PeopleSoft Database Server installation

- Installation of Oracle Server 11.2.0.1.0 binaries in Database Server box
- Installation of Oracle Client 11.2.0.10 in Database Server box
- Installation of Oracle Tuxedo 10.3.0.0 in Database Server box
- Installation of Oracle Tuxedo Patch RP065 in Database Server box
- Installation of PeopleTools 8.51 in Database Server box
- Installation of PeopleTools Patch 8.51.11 in Database Server box
- Installation of HRMS9.1 feature pack December 2010 in Database Server box
- Installation of Microfocus Server Express Cobol compiler 5.1wp4

Note: Refer to Appendix E for a complete PeopleSoft Database server installation listing in the “Cisco UCS Deployment Guide for Oracle PeopleSoft.”

Oracle PeopleSoft Install Port Numbers

We have used the port numbers below for WSL, JSL, HTTP ports, and more:

WSL	:	7700
JSL	:	9700
HTTP	:	8700
HTTPS	:	443
JRAD	:	9100
PSDBGSRV	:	9500
SMTP	:	25

Appendix E

Storage Configuration

Oracle PeopleSoft on Cisco UCS server storage requirements were carved out of the EMC VNX 5500 .

[illegible]

14.1 PowerPath/VE:

Increasingly, deployments are using virtualization for consolidation and scale-out of mission-critical applications. PowerPath/VE manages the complexity of large virtual environments, which could contain hundreds or thousands of independent virtual machines running intensive I/O applications.

PowerPath/VE manages this complexity, adjusting the I/O path usage to changes in IO loads coming from the VMs. Simply assign all devices to all paths and let PowerPath/VE do the work, optimizing the overall I/O performance for the virtual environment.

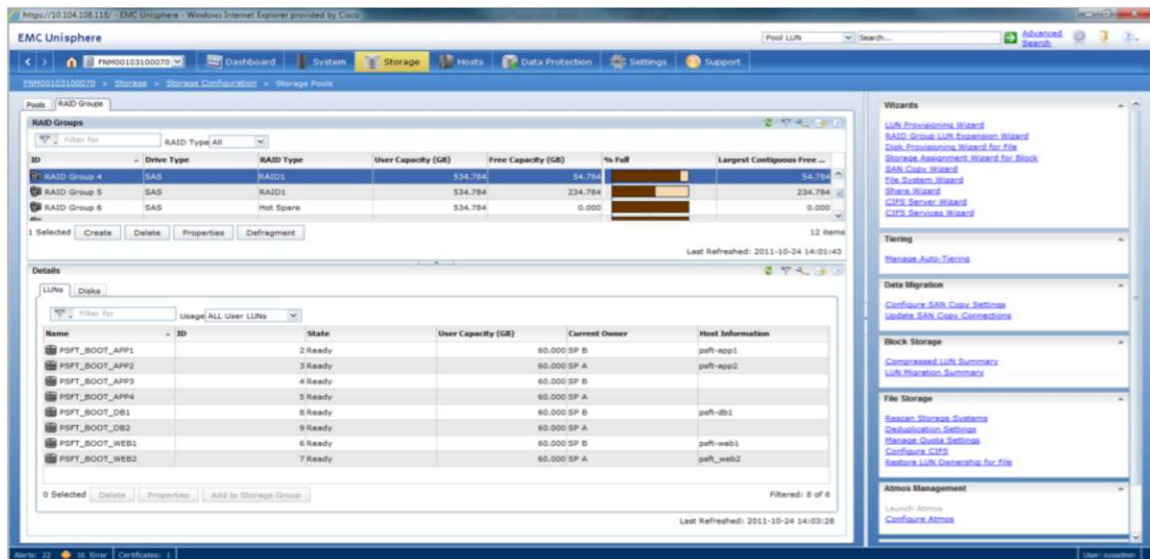
© 2012 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public.

- The ability to manage large environments
- Support for increased performance through optimum use of resources
- Providing for HA and automating I/O path failover
- Support for recovery in the event of a path failure

Additional information on EMC PowerPath/VE is available at: <http://www.emc.com/collateral/software/data-sheet/751-powerpath-ve-multipathing-ds.pdf>

14.2 Configuring LUNs with Unisphere

Create Raid Group and Boot LUNS (RAID 1)



Create Raid Groups for Oracle Application Binaries

The screenshot shows the EMC Unisphere interface for Storage Configuration. The 'RAID Groups' section displays a table with the following data:

ID	Drive Type	RAID Type	User Capacity (GB)	Free Capacity (GB)	% Full	Largest Contiguous Free...
RAID Group 9	SAS	Hot Spare	534.784	0.000		0.000
RAID Group 10	NL SAS	Hot Spare	1831.851	0.000		0.000
RAID Group 11	SAS	RAID5	2241.336	41.336		41.336

The 'Details' section shows a table of LUNs:

Name	ID	State	User Capacity (GB)	Current Owner	Host Information
PSFT-App1	41	Ready	300.000 SP B	psft-app1	
PSFT-App2	55	Ready	300.000 SP A	psft-app2	
PSFT-App3	57	Ready	300.000 SP A		
PSFT-App4	58	Ready	300.000 SP B		
PSFT-Ors	42	Ready	300.000 SP B	psft-db1	
PSFT-Web1	40	Ready	300.000 SP B	psft-web1	
PSFT-Web2	56	Ready	300.000 SP B	psft_web2	

Create a Storage Pool for the PSFT Database (Raid 5)

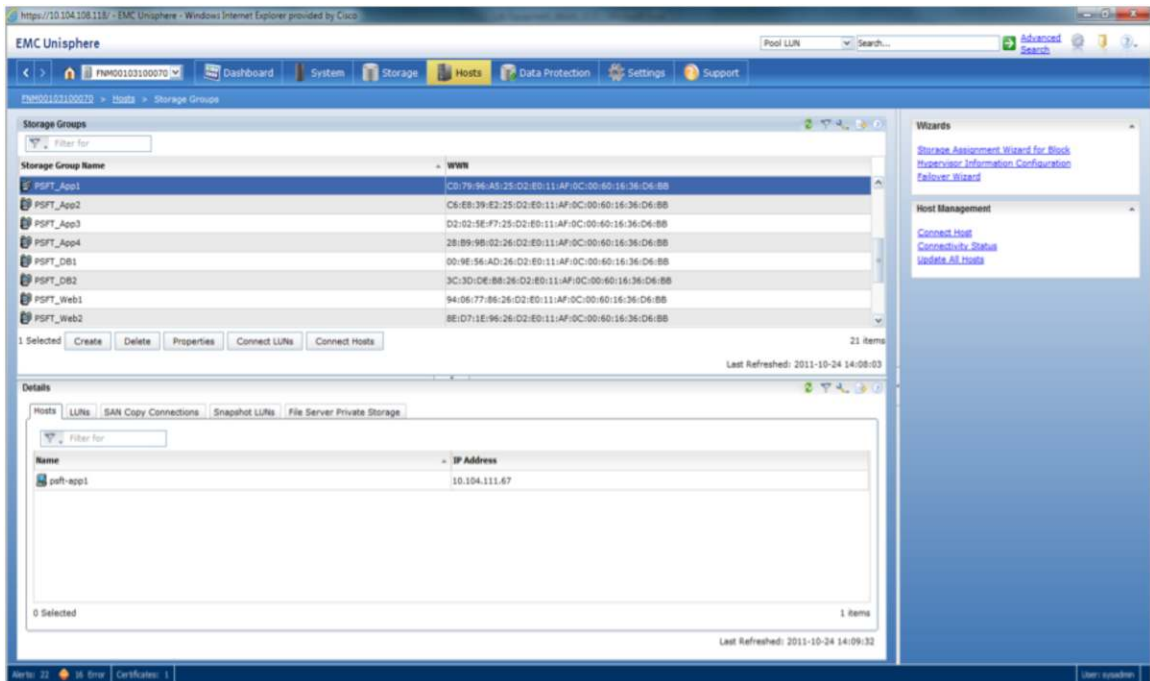
The screenshot shows the EMC Unisphere interface for Storage Configuration. The 'Storage Pools' section displays a table with the following data:

Name	State	RAID Type	Drive Type	User Capacity...	Free Capacity...	Allocated (GB)	%Consumed	Subscribed (GB)	%Subscribed
Pool 0 - Oracle Data+Application LUN	Ready	RAID5	Mixed	7159.957	4437.650	2722.307		2.693.25	37.611
Pool 1 - Siebel Filesystem	Ready	RAID1/0	SAS	1066.078	909.773	136.305		155.303	14.568

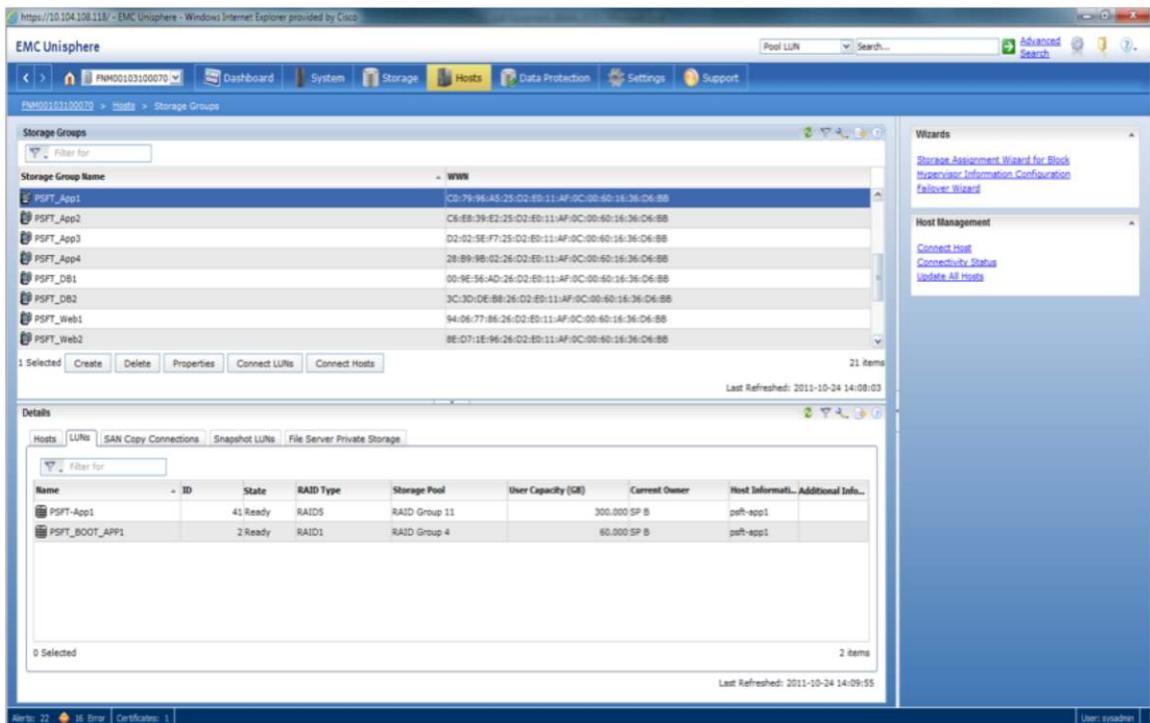
The 'Details' section shows a table of LUNs:

Name	ID	State	User Capacity (GB)	Current Owner	Host Information
Grid_1	45	Ready	3.000 SP A	ucamedb2.localdomain	
Grid_2	46	Ready	3.000 SP B	ucamedb2.localdomain	
Grid_3	47	Ready	3.000 SP A	ucamedb2.localdomain	
Grid_4	48	Ready	3.000 SP B	ucamedb2.localdomain	
Grid_5	49	Ready	3.000 SP A	ucamedb2.localdomain	
PSFT-Database	39	Ready	800.000 SP A	psft-db1	
UoMedA1-app lun	31	Ready	80.000 SP A	ucamedas1 - /siebel	
UoMedA1-ora client lun	32	Ready	20.000 SP A	ucamedas1 - /u01c	
UoMedA2-app lun	33	Ready	80.000 SP A	ucamedas2 - /siebel	

Create Storage Group and Attach HOST



Create Storage Group and Attach LUNs



Appendix F

Configuring Cisco Nexus Switches

These switches are connected to the blade chassis, with the I/O modules (2x4 cables) brought out and connected to each switch.

From the fabric interconnect switch, the FCoE cables are drawn out and patched to the Cisco Nexus 5000 CNA switch, which is also connected to the VNx. The uplink to the external world is from Cisco Nexus 5000.

Creating Zone and Zoneset

Logon to the Cisco Nexus 5000 Switch - IP : 10.1xx.1xx.2xx and 10.1xx.1xx.2xx

Log on to the Cisco UCS Manager (Capitola).

List the WWPN Listed in the Cisco Nexus 5000 Database

VSAN 4 is the zone under which the PeopleSoft Servers have been accommodated.

rk4-n5k8-a# sh flogi database

```
-----
INTERFACE      VSAN  FCID      PORT NAME      NODE NAME
-----
fc1/21          4    0x4203ef  50:06:01:6d:3e:a0:05:68 50:06:01:60:be:a0:05:6 - VNx to N5K
fc1/22          4    0x4202ef  50:06:01:65:3e:a0:05:68 50:06:01:60:be:a0:05:6 - VNx to N5K
fc1/23          4    0x42000e  20:43:00:05:73:a2:c2:80 20:04:00:05:73:a2:c2:8 - N5K to FI
fc1/23          4    0x42000f  20:00:00:25:b5:00:00:0c 20:00:00:25:b5:00:00:0 - N5K to FI
fc1/23          4    0x420011  20:01:00:25:b5:00:00:06 20:00:00:25:b5:00:00:1 - N5K to FI
fc1/23          4    0x420012  20:00:00:25:b5:00:00:0e 20:00:00:25:b5:00:00:0 - N5K to FI
fc1/24          4    0x42000d  20:42:00:05:73:a2:c2:80 20:04:00:05:73:a2:c2:8 - N5K to FI
fc1/24          4    0x420010  20:00:00:25:b5:00:00:17 20:00:00:25:b5:00:00:0 - N5K to FI
fc1/24          4    0x420013  20:00:00:25:b5:00:00:09 20:00:00:25:b5:00:00:0 - N5K to FI
fc1/24          4    0x420014  20:01:00:25:b5:00:00:07 20:00:00:25:b5:00:00:1 - N5K to FI
fc1/27          2    0x4405ef  50:06:01:6f:3e:a0:05:68 50:06:01:60:be:a0:05:6
fc1/28          2    0x4403ef  50:06:01:66:3e:a0:05:68 50:06:01:60:be:a0:05:6
```

Create a New Volume to Install Oracle Software:

rk4-n5k8-a# conf t

Enter configuration commands, one per line. End with CNTL/Z.

rk4-n5k8-a(config)# zone name psft_app1_data_vhba0 vsan 4

rk4-n5k8-a(config-zone)# member pwwn 20:00:00:25:b5:00:00:09

rk4-n5k8-a(config-zone)# member pwwn 50:06:01:6d:3e:a0:05:68

rk4-n5k8-a(config-zone)# member pwwn 50:06:01:65:3e:a0:05:68

rk4-n5k8-a(config-zone)# exit

Show the Zone Names of ZoneSet

```
rk4-n5k8-a(config)# sh zone

zone          zone-attribute-group  zoneset
rk4-n5k8-a(config)# sh zoneset active vsan 4

zoneset name psft_zoneset vsan 4

zone name psft_web1_vhba0 vsan 4
* fcid 0x4203ef [pwwn 50:06:01:6d:3e:a0:05:68]
* fcid 0x4202ef [pwwn 50:06:01:65:3e:a0:05:68]
* fcid 0x42000f [pwwn 20:00:00:25:b5:00:00:0c]

zone name psft_app1_vhba0 vsan 4
* fcid 0x4203ef [pwwn 50:06:01:6d:3e:a0:05:68]
* fcid 0x4202ef [pwwn 50:06:01:65:3e:a0:05:68]
* fcid 0x420010 [pwwn 20:00:00:25:b5:00:00:17]

zone name psft_db1_vhba0 vsan 4
* fcid 0x4203ef [pwwn 50:06:01:6d:3e:a0:05:68]
* fcid 0x4202ef [pwwn 50:06:01:65:3e:a0:05:68]
* fcid 0x420011 [pwwn 20:01:00:25:b5:00:00:06]
```

Add New Zone Name to ZoneSet

```
rk4-n5k8-a(config)# zoneset name psft_zoneset vsan 4
rk4-n5k8-a(config-zoneset)# member psft_app1_data_vhba0
rk4-n5k8-a(config-zoneset)# zoneset activate name psft_zoneset vsan 4
Zoneset activation initiated. check zone status
```

Appendix G

Discover LUNs and Install RHEL

Note: Refer to “Appendix G” for a RHEL install on Cisco UCS Server in “Cisco CVD for Oracle PeopleSoft.”

Discovering All Attached Disks on the Server

```
login as: root
```

root@10.104.111.71's password:

Last login: Mon Sep 5 14:10:51 2011 from dhcp-72-163-185-123.cisco.com

```
[root@psft-db1 ~]# echo "1" > /sys/class/fc_host/host0/issue_lip
```

```
[root@psft-db1 ~]# echo "1" > /sys/class/fc_host/host1/issue_lip
```

```
[root@psft-db1 ~]# echo "1" > /sys/class/fc_host/host2/issue_lip
```

```
[root@psft-db1 ~]# echo "1" > /sys/class/fc_host/host3/issue_lip
```

Creating a File System Partition on the PeopleSoft Servers

List out the connected hard disks and their partitions:

```
[root@psft-db1 ~]# fdisk -l
```

Partition the disk with fdisk:

fdisk /dev/sdk and input the following in the order shown:

(n->p->1->default value ->default value -> p -> w)

```
[root@psft-db1 ~]# fdisk /dev/sdk
```

Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel.

Build a new DOS disklabel. Changes will remain in memory only, until you decide to write them. After that, of course, the previous content will not be recoverable.

The number of cylinders for this disk is set to 39162.

Since this number is larger than 1024, it could, in certain setups; cause problems with:

- 1) Software that runs at boot time (for example, old versions of LILO).
- 2) Booting and partitioning software from other OSs
(for example, DOS FDISK, OS/2 FDISK).

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

Command (m for help): n

Command action

e extended

p primary partition (1-4)

p

Partition number (1-4): 1

First cylinder (1-39162, default 1):

Using default value 1

Last cylinder or +size or +sizeM or +sizeK (1-39162, default 39162):

Using default value 39162

Command (m for help): w

The partition table has been altered!

Calling ioctl() to re-read partition table.

Syncing disks.

Create a File System

After creating the partition, we can overlay a file system (format) using mkfs command as shown below.

```
[root@psft-db1 ~]# mkfs -t ext3 /dev/sdk
```

Add the file system created to FSTAB

```
[root@psft-db1 ~]# vi /etc/fstab
```

```
Add "/dev/sdk    /db2-data ext3 defaults    0 0"
```

LABEL=/	/	ext3	defaults	1 1
LABEL=/boot	/boot	ext3	defaults	1 2
devpts	/dev/pts	devpts	gid=5,mode=620	0 0
/dev/sdk	/db2-data	ext3	defaults	0 0
tmpfs	/dev/shm	tmpfs	defaults	0 0

proc	/proc	proc	defaults	0 0
sysfs	/sys	sysfs	defaults	0 0

Create a Dir

```
[root@psft-db1 ~]# mkdir /db2-data
```

Mount the Dir

```
[root@psft-db1 ~]# mount /db2-data
```

```
[root@psft-web1 host4]# mkfs -t ext3 /dev/sdk
```

```
mke2fs 1.39 (29-May-2006)
```

```
Filesystem label=
```

```
OS type: Linux
```

```
Block size=4096 (log=2)
```

```
Fragment size=4096 (log=2)
```

```
39321600 inodes, 78642183 blocks
```

```
3932109 blocks (5.00%) reserved for the super user
```

```
First data block=0
```

```
Maximum filesystem blocks=4294967296
```

```
2400 block groups
```

```
32768 blocks per group, 32768 fragments per group
```

```
16384 inodes per group
```

```
Superblock backups stored on blocks:
```

```
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
```

```
4096000, 7962624, 11239424, 20480000, 23887872, 71663616
```

```
Writing inode tables: done
```

```
Creating journal (32768 blocks): done
```

```
Writing superblocks and filesystem accounting information: done
```

This file system will be automatically checked every 27 months or 180 days, whichever comes first. Use `tune2fs -c` or `-i` to override.

Remember:

Now -> `vi /etc/fstab`

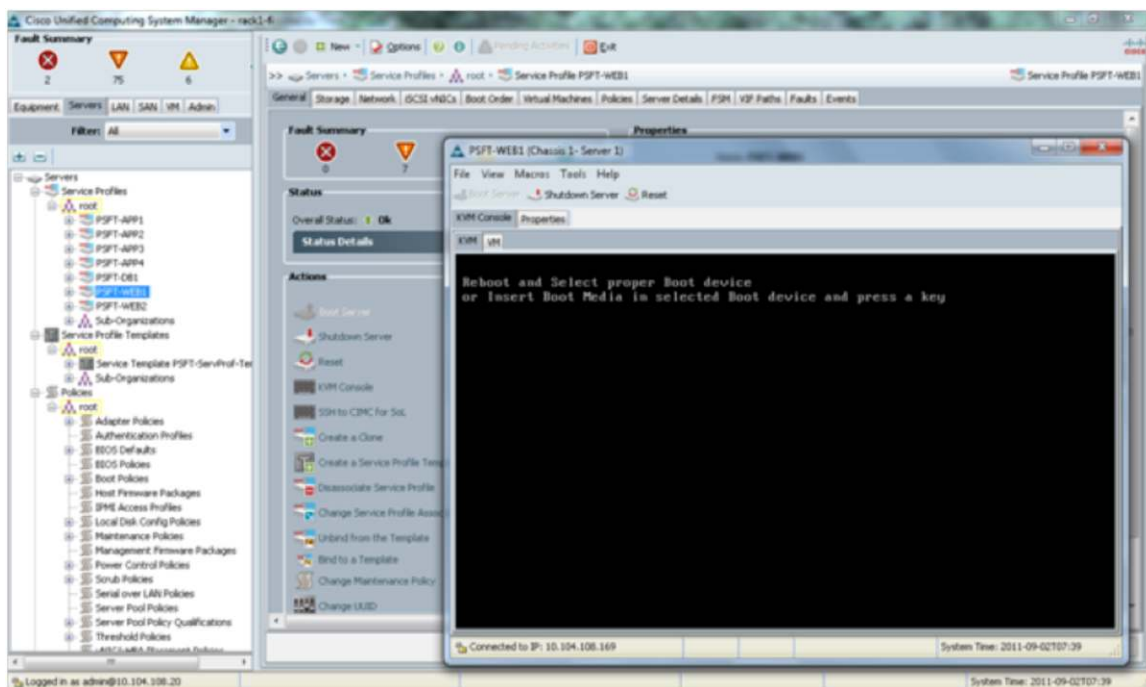
Add `"/dev/sdc1 /db2-data ext3 defaults 0 0"` in `fstab`


Then -> `mkdir /db2-data`

then -> `mount /db2-data`


Logon to Cisco UCS Manager (Capitola).

Remember to click on the VM instead of KVM in earlier versions of UCS Manager to get to load the OS boot image.

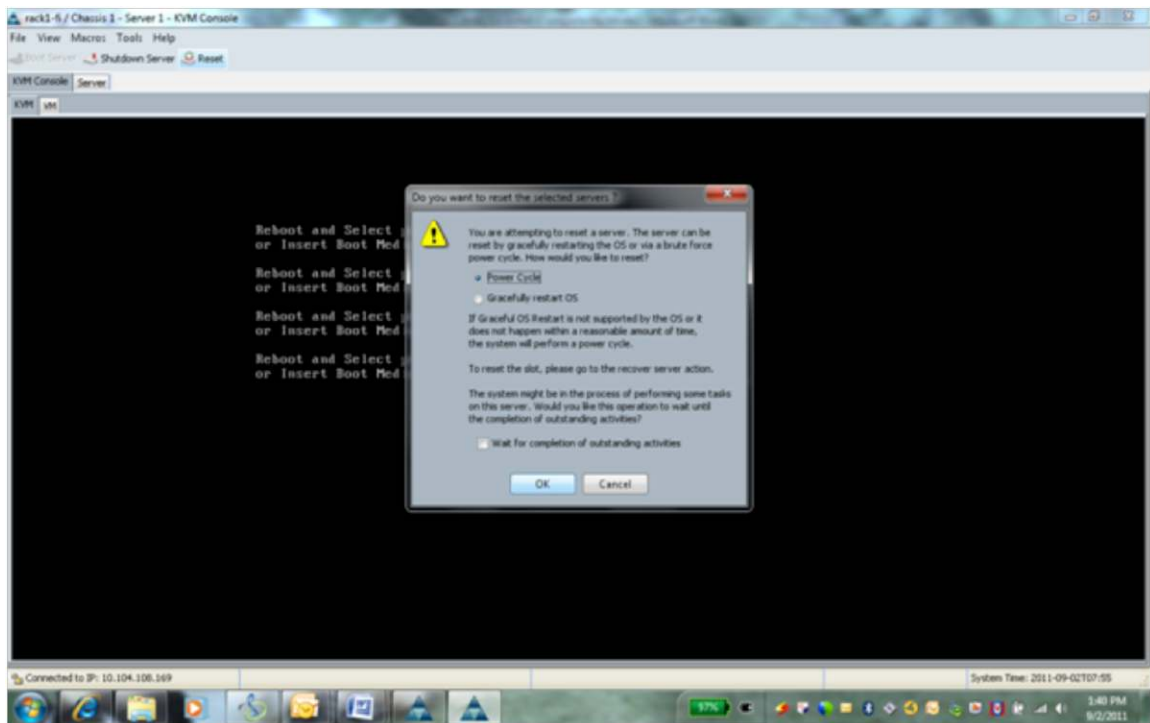





Add the image and map it.

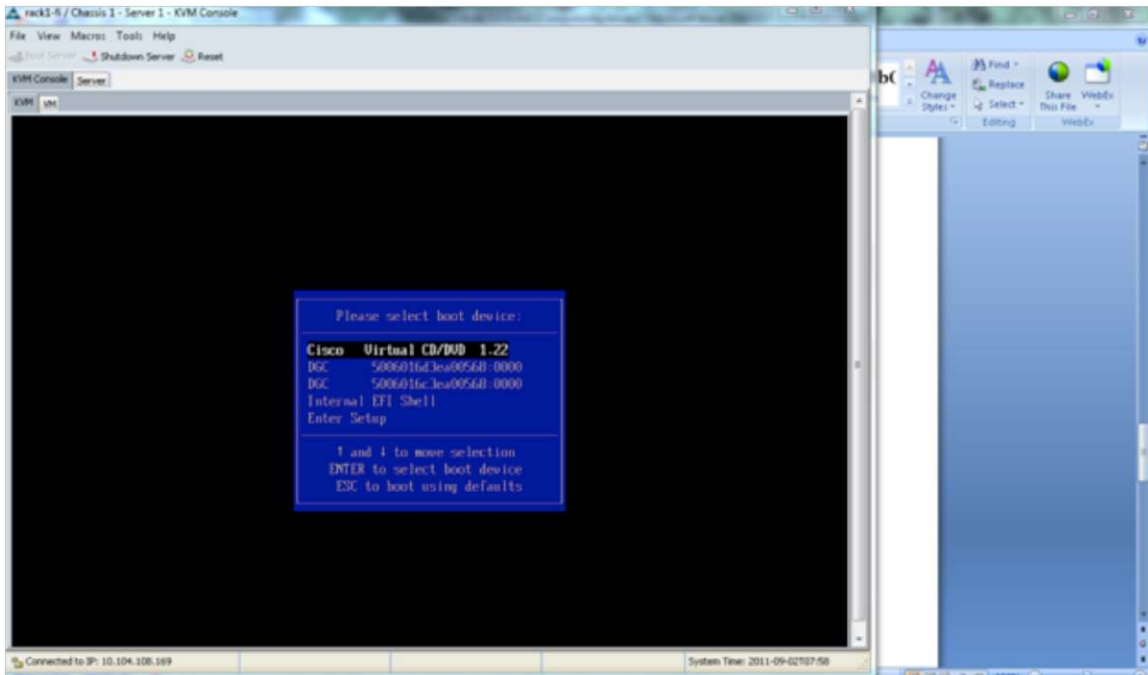


Click Reset to restart the server.

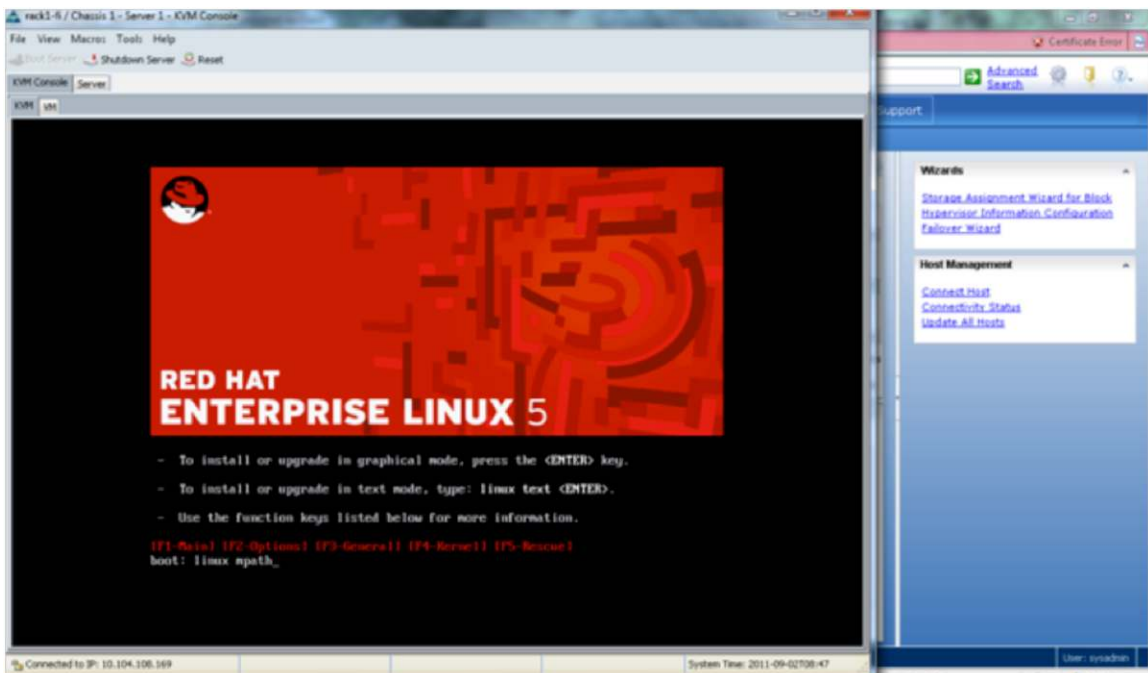


Press F6 to enter the boot media.





Remember to type **linux mpath**.



References Links

The racking, power, and installation of the chassis are described in the Install Guide:

http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html

Cisco Unified Computing System CLI Configuration Guide:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/1.4/b_UCSM_CLI_Configuration_Guide_1_4.html

Cisco UCS Manager GUI Configuration Guide:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/1.4/b_UCSM_GUI_Configuration_Guide_1_4.html

"VP for Unified Storage System – A Detailed Review," is available at:

<http://www.emc.com/collateral/software/white-papers/h8058-fast-vp-unified-storage-wp.pdf>

"EMC FAST Cache – A Detailed Review," is available at:

<http://www.emc.com/collateral/software/white-papers/h8046-clariion-celerra-unified-fast-cache-wp.pdf>

"EMC FAST VP for Unified Storage System – A Detailed Review," is available at:

<http://www.emc.com/collateral/software/white-papers/h8058-fast-vp-unified-storage-wp.pdf>

Additional information on EMC PowerPath/VE is available at:

<http://www.emc.com/collateral/software/data-sheet/l751-powerpath-ve-multipathing-ds.pdf>

Additional information on the VNX Series is available at:

<http://www.emc.com/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf>

References Documents

- Oracle PeopleSoft PeopleTools 8.51 Certification Matrix on Linux x86-64 Red Hat Enterprise Linux 5
<https://support.oracle.com/CSP/ui/flash.html>
- Cisco Hardware and Software Interoperability Matrix Release 1.4.3
http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html
- Oracle PeopleSoft PeopleTools 8.51 Release Notes
[https://support.oracle.com/CSP/ui/flash.html - tab=KBHome\(page=KBHome&id=\(\)\),\(page=KBNavigator&id=\(viewingMode=1143&bmDocTitle=PeopleTools 8.51 Release Notes&bmDocDsrc=KB&bmDocType=REFERENCE&bmDocID=1203023.1&from=BOOKMARK\)\)](https://support.oracle.com/CSP/ui/flash.html-tab=KBHome(page=KBHome&id=()),(page=KBNavigator&id=(viewingMode=1143&bmDocTitle=PeopleTools%208.51%20Release%20Notes&bmDocDsrc=KB&bmDocType=REFERENCE&bmDocID=1203023.1&from=BOOKMARK)))
- Hardware and Software Guide for Oracle PeopleTools 8.51
http://docs.oracle.com/cd/E18373_01/psft/acrobat/PeopleTools_8.51_HardwareSoftwareGuide.pdf
- 733205.1: Migration of Oracle Database Instances Across OS Platforms
[https://support.oracle.com/CSP/ui/flash.html - tab=KBHome\(page=KBHome&id=\(\)\),\(page=KBNavigator&id=\(bmDocDsrc=KB&bmDocID=733205.1&from=BOOKMARK&viewingMode=1143&bmDocTitle=Migration of an Oracle Database Across OS Platforms&bmDocType=HOWTO\)\)](https://support.oracle.com/CSP/ui/flash.html-tab=KBHome(page=KBHome&id=()),(page=KBNavigator&id=(bmDocDsrc=KB&bmDocID=733205.1&from=BOOKMARK&viewingMode=1143&bmDocTitle=Migration%20of%20an%20Oracle%20Database%20Across%20OS%20Platforms&bmDocType=HOWTO)))
- 243304.1-10g: Transportable Tablespaces Across Different Platforms
[https://support.oracle.com/CSP/ui/flash.html - tab=KBHome\(page=KBHome&id=\(\)\),\(page=KBNavigator&id=\(bmDocDsrc=KB&bmDocID=243304.1&from=BOOKMARK&viewingMode=1143&bmDocTitle=10g : Transportable Tablespaces Across Different Platforms&bmDocType=BUht](https://support.oracle.com/CSP/ui/flash.html-tab=KBHome(page=KBHome&id=()),(page=KBNavigator&id=(bmDocDsrc=KB&bmDocID=243304.1&from=BOOKMARK&viewingMode=1143&bmDocTitle=10g%20:%20Transportable%20Tablespaces%20Across%20Different%20Platforms&bmDocType=BUht)
- 371556.1: How to Move Tablespaces Across Platforms Using Transportable Tablespaces with RMAN
<https://support.oracle.com/CSP/ui/flash.html>

[tab=KBHome\(page=KBHome&id=\(\)\),\(page=KBNavigator&id=\(bmDocDsrc=KB&bmDocID=371556.1&from=BOOKMARK&viewingMode=1143&bmDocTitle=How to Move Tablespaces Across Platforms Using Transportable Tablespacht](#)

Disclaimer

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2010 Cisco Systems, Inc. All rights reserved.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)