

RISC Migration Guide

Siebel CRM

February 2012



Contents

1. Executive Summary.....	3
2. Introduction.....	4
2.1 Intended Audience.....	4
2.2 Oracle Siebel Customer Relationship Management Solution.....	4
2.3 Oracle Solaris Platform.....	4
2.4 Cisco Unified Computing System (UCS)	4
2.5 EMC VNX Storage Family.....	7
2.6 Why Migrate?.....	7
3. Planning for Migration.....	10
3.1 Migration Approach.....	10
3.2 Hardware Considerations	10
4. Migration Process.....	11
4.1 Source Setup Evaluation	11
4.2 Target(Cisco UCS) Enterprise Setup.....	12
4.3 Migration Activities	21
5. Performance Comparison	28
6. Conclusion	29
7. For More Information.....	30

1. Executive Summary

Migrating from proprietary technology to one based on industry standards will not only help reduce IT costs, but also allow the IT ecosystem to be more flexible.

Today, business demands are changing rapidly, requiring high IT responsiveness in order to keep pace with market dynamics. In this context, the Cisco Unified Computing System (UCS), the industry's first unified data center platform, can add great value to businesses looking to respond quickly.

Often, existing enterprise servers are not equipped to handle the ever-growing demands on IT. As a result, companies are unable to cope with market changes. In contrast, Cisco UCS delivers superior flexibility, compared to SPARC (which is built on RISC-based architecture). If a server fails, IT can provision any other available server in minutes by applying the Cisco UCS Manager service profile. This results in improved enterprise application availability and performance.

Moving Oracle's Siebel Customer Relationship Management (CRM) Applications (Release 8.1) to Cisco UCS will allow it to take full advantage of the stateless computing, unified fabric, and extended memory features of Cisco UCS. These features help optimize the application's potential to support business innovation.

The Cisco strategic migration procedure provides a road map to execute the migration from SPARC to the Cisco UCS platform securely and efficiently, with minimal disruption to business.

This document begins with a description of Cisco UCS and the drivers behind migration to this platform, such as architecture and performance. It then details the migration methodology, approach, hardware evaluation, and activities. The planning steps that should be taken when preparing for such a migration, as well as common implementation best practices, are included in this paper.

A thorough understanding of the migration environment is the critical first step to a successful migration. Likewise, understanding potential deployment scenarios will help to identify any roadblocks and help ensure a smooth transition.

There are references to related resources essential in understanding the existing configuration and application requirements to speed and simplify the migration.

2. Introduction

This guide shows how to migrate Siebel CRM from the Solaris platform, which is based on reduced instruction set computer (RISC) processor architecture, to Cisco UCS with x86-architecture hardware.

2.1 Intended Audience

This guide is for system administrators, database administrators, and data center architects who are tasked with migrating the Siebel Customer Relationship Management (CRM) application from the Solaris platform to Cisco UCS.

2.2 Oracle Siebel Customer Relationship Management Solution

Oracle Siebel CRM software was launched by Siebel Systems, Inc., and currently holds a significant market share in the CRM business.

Oracle Siebel CRM is a branch of Oracle's customer relations management program suite. It offers customer-centered applications for tracking, selling, billing, and analytics software for effectively tailoring use of the applications. Founded on a service-oriented architecture, the software allows businesses to build scalable, standards-based applications that can help an organization attract new business, increase customer loyalty, and improve profitability.

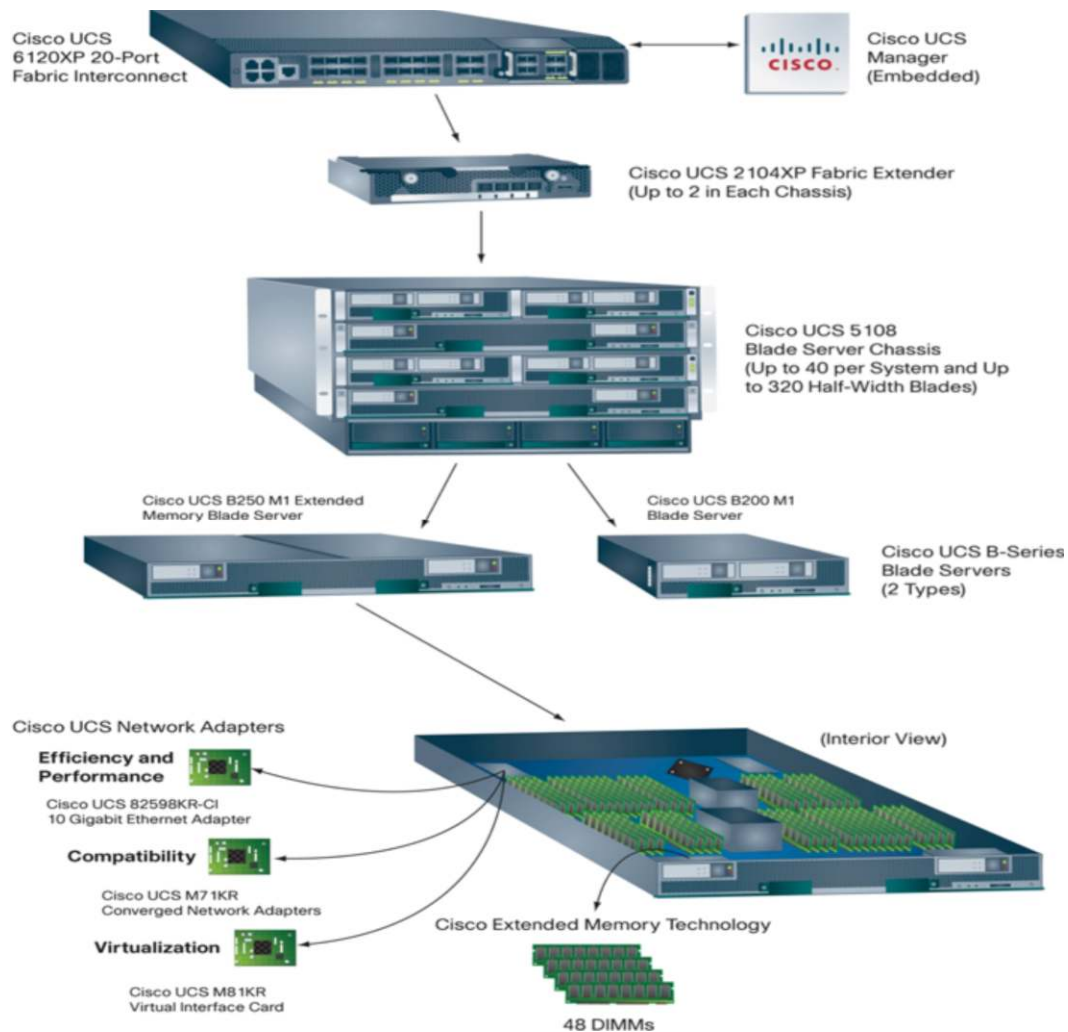
2.3 Oracle Solaris Platform

This Unix-based operating system environment supports multithreading, symmetric multiprocessing (SMP), and integrated TCP/IP networking, in addition to centralized network administration. In this migration exercise, Siebel Enterprise (version 8.1.1.4) running on Solaris 10 on Sun v240/v490/v880 servers, is considered.

2.4 Cisco Unified Computing System (UCS)

The Cisco Unified Computing System (UCS) is a next-generation data center platform that unites compute, network, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies, and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. It is an integrated, scalable, multichassis platform, in which all resources participate in a unified management domain.

Figure 1. The Cisco Unified Computing System



- **Cisco UCS 6100 Series Fabric Interconnects**

This family of line-rate, low-latency, lossless, 10-Gbps Ethernet interconnect switches consolidates I/O within the system. Both 20-port one-rack-unit (1RU) and 40-port 2RU versions accommodate expansion modules that provide Fiber Channel and 10 Gigabit Ethernet connectivity.

- **Cisco UCS 2100 Series Fabric Extenders**

These extenders bring unified fabric into the blade-server chassis, providing up to four 10-Gbps connections each between the blade server and the fabric interconnect, simplifying diagnostics, cabling, and management.

- **Cisco UCS Network Adapters**

These offer a range of options to meet application requirements, including adapters optimized for virtualization, converged network adapters (CNAs) for access to unified fabric and compatibility with existing driver stacks, Fiber Channel HBAs, and efficient, high-performance Ethernet adapters.

- **Cisco UCS Manager**

This embedded device-management software provides centralized management capabilities, creates a unified management domain, and serves as the central nervous system of Cisco UCS. It manages the entire system as a single logical entity through an intuitive GUI, command-line interface (CLI), or XML API.

Cisco UCS Manager implements role- and policy-based management, using service profiles and templates. This construct improves IT productivity and business agility. Now, infrastructure can be provisioned in minutes instead of days, shifting IT's focus from maintenance to strategic initiatives.

Cisco UCS resources are abstract, and use a just-in-time deployment model for programming:

- Identity and I/O configuration
- MAC addresses
- World Wide Names (WWNs), firmware versions
- BIOS boot order
- Network attributes (including quality-of-service [QoS] settings, access control lists [ACLs], pin groups, and threshold policies)

The manager stores this identity, connectivity, and configuration information in service profiles that reside on Cisco UCS 6100 Series Fabric Interconnects. A service profile can be applied to any resource, to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, providing flexibility in the use of system resources.

Cisco UCS Manager also offers role-based management that helps organizations make more efficient use of their limited administrator resources.

- **Cisco UCS 5100 Series Blade Server Chassis**

A logical part of the Cisco UCS fabric interconnects, this blade server chassis adds no management complexity to the system.

The Cisco UCS 5100 Series Blade Server Chassis fits on a standard rack and is six RUs in height. It physically houses blade servers and up to two Cisco UCS 2100 Series Fabric Extenders.

The chassis features eight cooling fans and four power supply units, which are hot swappable and redundant. It only requires two power supplies for normal operation; the additional power supplies are for redundancy. These highly efficient (in excess of 90 percent) power supplies, along with the simple chassis design (which incorporates front-to-back cooling), makes Cisco UCS very reliable and energy efficient.

- **Cisco UCS B200 M1 and M2 Blade Servers**

These blade servers balance simplicity, performance, and density for production-level virtualization and other mainstream data center workloads. The servers are half-width, two-socket blade servers.

The first-generation Cisco UCS B200 M1 uses the Intel Xeon 5500 Series processor, and the next-generation Cisco UCS B200 M2 uses the Intel Xeon 5600 processor.

Each Cisco UCS B200 server uses a CNA for consolidated access to the unified fabric. This design reduces the number of adapters, cables, and access-layer switches needed for LAN and SAN connectivity.

- **Cisco UCS B250 M1 and M2 Extended Memory Blade Servers**

These blade servers use patented Cisco Extended Memory Technology, which provides over twice as much industry-standard memory (384 GB) as traditional two-socket servers. This increases performance and capacity for demanding virtualization and large-data-set workloads. Alternatively, this technology offers a more cost-effective memory footprint for less-demanding workloads.

The first-generation Cisco UCS B250 M1 uses the Intel Xeon processor 5500 series, and the next-generation Cisco UCS B250 M2 uses the Intel Xeon 5600 processor series.

Each Cisco UCS B250 uses two CNAs for consolidated access to the unified fabric. This design reduces the number of adapters, cables, and access-layer switches needed for LAN and SAN connectivity.

2.5 EMC VNX Storage Family

These storage systems represent EMC's next generation of unified storage, optimized for virtual environments, while offering a cost-effective choice for deploying mission-critical enterprise applications such as Siebel. The massive virtualization and consolidation trends with servers demand a new storage technology that is dynamic and scalable. The EMC VNX series meets these requirements, and offers several software and hardware features for optimally deploying enterprise applications such as Siebel.

Figure 2. The EMC VNX Family of Unified Storage Platforms



A crucial 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 Fibre Channel (FC), iSCSI, Fibre Channel over Ethernet (FCoE), Network File System (NFS), and Common Internet File System (CIFS) network shared file access.

Furthermore, any 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. Additional information on Unisphere can be found on emc.com in the white paper titled: [Introducing EMC Unisphere: A Common Midrange Element Manager](#).

2.6 Why Migrate?

Siebel CRM can gain significant advantages by migrating from the current Solaris platform to UCS, mainly due to superior flexibility offered by UCS platforms compared to RISC-based platforms, as well as benefiting from reduced maintenance costs.

Cisco UCS represents a radical simplification of the traditional blade server deployment model, with simplified, stateless blades and a blade server chassis that is centrally provisioned, configured, and managed by Cisco UCS Manager. The result is a unified system that significantly reduces the number of components, while offering a just-in-time provisioning model that allows systems to be deployed or redeployed in minutes, rather than hours or days.

By uniting computing, networking, and storage access resources, combined with the single-view management of the complete solution, Cisco UCS is designed to deliver a cohesive, integrated system that is managed, serviced, and tested as a whole. Based on industry standards supported by a partner ecosystem of industry leaders, Cisco UCS also delivers:

- Reduced TCO at the platform, site, and organizational levels by simplifying data center resources
- Increased IT staff productivity and business agility, through just-in-time provisioning and mobility support
- Scalability through a design for up to 320 discrete servers and thousands of virtual machines
- The capability to scale I/O bandwidth to match demand
- Reduced number of devices requiring setup, management, power, cooling, and cabling

2.6.1 Cisco UCS Advantages

The benefits of innovative Cisco UCS technology are detailed below:

Exceptional Performance

- **Cisco Extended Memory Technology:** The Cisco UCS offers some of the industry's highest-density two-socket and four-socket rack and blade server platforms capable of delivering greater than 12 terabytes (TB) of memory in a single rack. This computing and memory density supports increased performance and capacity for demanding virtualization and large-data-set workloads. Alternatively, this technology can offer a more cost-effective memory footprint for less-demanding workloads, while still delivering outstanding performance.
- **Intel Xeon processor 5600 and 7500 series:** These processors offer advanced reliability, availability, and serviceability (RAS) features. When combined with the highly available Cisco UCS architecture with redundant, hot-swappable components, they improve data integrity and reduce downtime. They automatically and intelligently adjust server performance according to application needs, increasing performance as needed, and achieving substantial energy savings when performance requirements are low.

The Intel Xeon processor has become ubiquitous. It supports multiple operating systems, including Sun Solaris and all types of Linux, and Microsoft Windows systems, supporting a consistent deployment platform across the organization. With Intel's continued support of virtual server environments, the x86 architecture is the standard platform for virtualization, supporting efficient, cost-effective, and reliable computing platforms for the next-generation data center.

- **Broad enterprise application support:** Cisco UCS has been certified and has delivered industry-leading benchmarks on a wide variety of business applications, from vendors including Oracle, SAP, Microsoft, and VMware.

More Agility and Productivity

- **Stateless computing:** Cisco UCS provides built-in management that dynamically configures the entire hardware stack, ranging from network configurations to server firmware. Cisco service profiles apply state to every component in the stack, so that the system self-integrates and scales rapidly without adding complexity. Hardware failures no longer require IT to swap host bus adapters (HBAs) or network interface cards (NICs), rezone the SAN, alter network configurations, or reconfigure firmware. The system's single point of management increases IT staff productivity, improves compliance, and reduces the opportunity for errors that can cause downtime.

Less Cost and Complexity

- **Unified fabric:** The Cisco UCS server chassis revolutionizes the use and deployment of blade-based systems. By incorporating unified fabric and fabric extender technology, this chassis contains fewer physical components and requires no independent management, compared with traditional blade server chassis. This simplicity eliminates the need for dedicated chassis management and blade LAN and SAN switches, and reduces cabling and complexity.

Cisco integrates the system through a standards-based, high-bandwidth, low-latency, virtualization-aware unified fabric that supports the existing enterprise SAN, network-attached storage (NAS), and direct-attached storage (DAS) architectures.

The unified fabric is wired to support the desired bandwidth. It carries all network, storage, interprocess communication, and virtual machine traffic with security isolation, visibility, and control to individual virtual machines.

With 20 Gbps of I/O bandwidth per server slot and support for future 40 Gigabit Ethernet standards, the fabric meets the I/O demands of multicore processors. It eliminates costly redundancy, and increases workload agility, reliability, and performance. The fabric supports a wire-once deployment model, where changing I/O configurations no longer means installing adapters or recabling racks and switches, potentially saving thousands of dollars.

- **Cisco UCS Manager:** Embedded system management is uniquely integrated into all components of the system, so the entire solution can be managed as a single entity through Cisco UCS Manager. Whether managing one server or hundreds of servers, Cisco UCS Manager provides an intuitive GUI, a CLI, and a robust, open API for managing all system configuration and operations.

Cisco UCS Manager's flexible role-based management allows IT managers of storage, networking, and servers to communicate and collaborate easily on service profile definitions for applications. This can help increase personnel efficiency and reduce capital expenses (CapEx) and operating expenses (OpEx).

Virtualization—Aware Networking

- **Cisco VN-Link technology:** High-performance networking to virtual machines is achieved using Cisco VN-Link technology. It supports a consistent operational model, whether networks connect to physical servers or virtual machines. All links are centrally configured, secured, and managed without introducing additional switching layers into virtualized environments. I/O configurations and network policies move with virtual machines, increasing security and efficiency, while preserving a familiar management framework for network administrators.

Outstanding Energy Efficiency

- **Environmental footprint:** Cisco UCS incorporates unified fabric, fabric extender technology, and a single point of management. As a result, the chassis requires fewer physical components, consumes less power, and generates less heat than traditional blade server chassis, which use RISC- or x86-based architectures.

Applicability to Different Scenarios

Cisco UCS servers can easily fit the requirement of various configurations of application and database servers, to achieve scale-out and scale-up demands for Siebel Enterprise CRM deployment.

3. Planning for Migration

Appropriate due diligence needs to be carried out before beginning migration from Sun servers to Cisco UCS servers for the following considerations:

- **Siebel and operating system compatibility:** Cisco UCS servers support only Windows or Linux operating systems. Please refer to the “Siebel System Requirements and Supported Platforms” document, published at http://download.oracle.com/docs/cd/E11886_01/srsphomepage.html. Red Hat Linux (Version 5.4, 64 bit version) was used in this migration exercise.
- **Siebel upgrade path identification:** For example, Siebel Enterprise running on 7.x version on the Solaris operating system first needs to be upgraded to Siebel version 8.1.1.4 before migrating to the Linux operating system. For specific upgrade paths of Siebel versions, please refer to: http://download.oracle.com/docs/cd/E11886_01/V8/CORE/SRSP_81/SRSP_81_UpgradePaths2.html
- **Siebel customization:** Care should be taken while migrating custom components, as they may require additional efforts. Siebel-provided utilities, such as cfgmerge, do not support custom components. These components may also require dedicated development and testing efforts for operating system compatibility.

The migration of Siebel CRM from a RISC platform to Cisco UCS has been rigorously tested and validated in a lab environment. Customers can be assured of a smooth migration, based on the defined methodology.

3.1 Migration Approach

To minimize the downtime of the entire Siebel Enterprise, it is essential to follow a phased migration of components. For example, database servers and gateway servers should be migrated first, followed by application servers, web servers, etc. The migrated gateway server can register the existing RISC platform-based app/web servers, along with new app/web servers that are hosted on Cisco UCS servers until the migration is complete.

Also, it is recommended to follow the standard approach of migrating low-risk environments, such as development and QA, before production environment migration.

3.2 Hardware Considerations

The hardware capacity needs to be estimated before the Siebel migration, to make sure that the service-level agreements (SLAs) maintained in the current Siebel Enterprise are not impacted in the newer environment. The target platforms, such as Cisco UCS servers, are built to meet compute-intensive workloads. This is fully demonstrated by the world record benchmark results achieved by Cisco UCS (for more details, please refer to: http://www.cisco.com/en/US/prod/ps10265/industry_benchmarks.html).

4. Migration Process

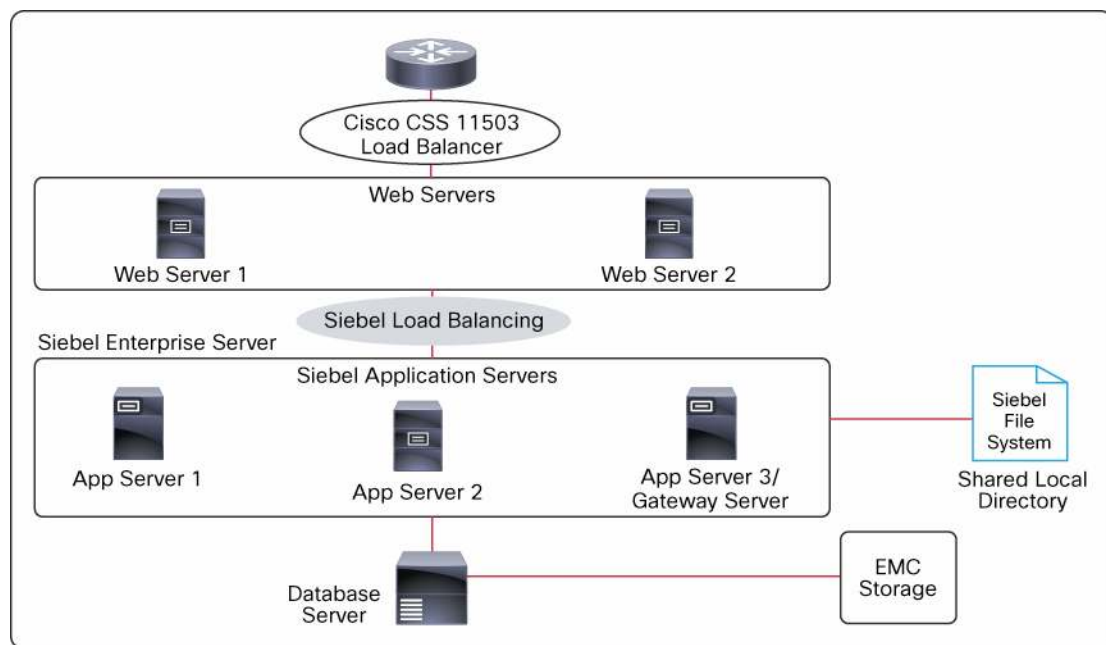
The source setup for the migration (Oracle Solaris environment) is understood as a first step in the migration process, in terms of hardware deployment, installed Siebel components in each, etc. The next step is the configuration of the Cisco UCS servers, along with the EMC storage connectivity and operating system installation, before the Siebel application is installed.

4.1 Source Setup Evaluation

The source setup for the migration (Oracle Solaris environment) is assessed in terms of hardware deployment, installed Siebel Components in each of them, etc. This helps ensure that the migrated environment matches it in terms of functionality and capacity.

4.1.1 Setup and Hardware Configuration

Figure 3. Initial Siebel Infrastructure on RISC Platforms



The hardware details and the components information are captured in the table below.

Table 1. RISC Siebel Infrastructure

Function	Qty	Server Model	CPU	Memory	Comments
Web Servers	2	Sun Fire V240	2x SPARC IIIi1.5GHz	32GB	These web servers are load-balanced by the Cisco CSS 11503 Content Services Switch.
Siebel Application Servers/ Gateway Server	3	Sun Fire V490	4x SPARC IV 1.5GHz	32GB	Three servers are dedicated to handle eCommerce/eSales/enterprise application integration (EAI) workflow client requests. They are load balanced using the Siebel load balancer. Gateway Server was installed on Application Server 3.
Siebel Database Server	1	Sun Fire V880	8x SPARC III 900GHz	32GB	Siebel database server.

Siebel File System	1	N/A	N/A	150GB	A RAID10 Logical Unit Number (LUN) carved from EMC storage array and mounted exported as a file system on to one of the application servers.
--------------------	---	-----	-----	-------	--

4.1.2 Workload Details and Running Information

This setup is modeled after a realistic workload of 500 peak concurrent users (of 800 registered users), expected to generate 10000 new orders in an 8-hour working day. The running status is captured here:

Figure 4. Current Setup: Application Server Status

State	Siebel Server	Server Group	Name	State	PID	Host Name	Start Time	End Time
Running	sunsiebas2			Running	7104	sunsiebas2	8/5/2011 03:53:00 I	
Running	sunsiebas1			Running	8701	sunsiebas1	8/5/2011 08:59:25 I	
Running	sunsiebas3			Running	3214	sunsiebas3	8/5/2011 03:52:54 I	

Figure 5. Current Setup: Siebel Server Management

State	Name	Running Normal	Max Tasks	Running MTS	Max MTS	Start Time	End Time	State
Online	Batch Assignment	0	20			8/5/2011 03:53:15 PM		Online
Online	EAI Object Manager (EAI)	0	120	4	4	8/5/2011 03:53:15 PM		Online
Online	Enterprise Integration Mgr	0	5			8/5/2011 03:53:15 PM		Online
Online	File System Manager	0	20	1	1	8/5/2011 03:53:05 PM		Online
Online	Siebel Product Configuration Object Manager (EAI)	0	20	1	1	8/5/2011 03:53:15 PM		Online
Online	Workflow Action Agent	0	5			8/5/2011 03:53:15 PM		Online
Online	Workflow Monitor Agent	0	5			8/5/2011 03:53:15 PM		Online
Online	Workflow Process Batch Manager	0	20	1	1	8/5/2011 03:53:15 PM		Online
Online	Workflow Process Manager	0	20	1	1	8/5/2011 03:53:15 PM		Online
Online	Workflow Recovery Manager	0	20	1	1	8/5/2011 03:53:15 PM		Online
Online	eSales Object Manager (EAI)	0	120	4	4	8/5/2011 03:53:15 PM		Online


4.2 Target(Cisco UCS) Enterprise Setup

Based on the planning as discussed in Section 3, the Cisco UCS hardware is chosen to migrate from the existing Siebel environment. In this exercise, Cisco UCS B200 M2(B-series) server is chosen for Web server, Gateway server and Application server (one per component) and Cisco UCS B230 blade server-hosted Oracle database.

4.2.1 Detailed Topology

The Cisco UCS blades are connected to fabric interconnects and Cisco Nexus 5548 switches before being connected to the EMC VNX5500 storage array, as shown below.

Figure 6. Detailed Topology



4.2.2 Storage Configuration

4.2.2.1 EMC VNX Storage Platforms

The new EMC VNX family of unified storage platforms continues the EMC tradition of providing among the highest rates of data reliability and availability in the industry. In addition, they include a boost in performance and bandwidth, to address sustained data access bandwidth rates. The new system design has also placed heavy emphasis on storage efficiencies and density, as well as crucial green storage factors, such as a smaller data center footprint, lower power consumption, and improvements in power reporting. The VNX5500 model was used in this RISC migration scenario with Siebel.

All models in EMC's new VNX storage family now support the 2.5" Serial Attached SCSI (SAS) drives in a 2U disk array enclosure (DAE) that can hold up to 25 drives, one of the densest offerings in the industry. For example, the older-generation technology required storing 15 x 600 GB worth of data using the 3.5" FC drives in a 3U DAE. However, the new DAE uses 25 x 600 GB drives in a 2U footprint, 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 this 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>.

Important efficiency features available with the VNX series include FAST Cache and FAST VP.

4.2.2.2 FAST Cache Technology

In traditional storage arrays, DRAM caches are too small to maintain hot data for long periods of time. Few storage arrays give an option to non-disruptively 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, subsequent access to that chunk of data is served at Flash latencies. Eventually, when the data temperature cools down, the data chunks are 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.

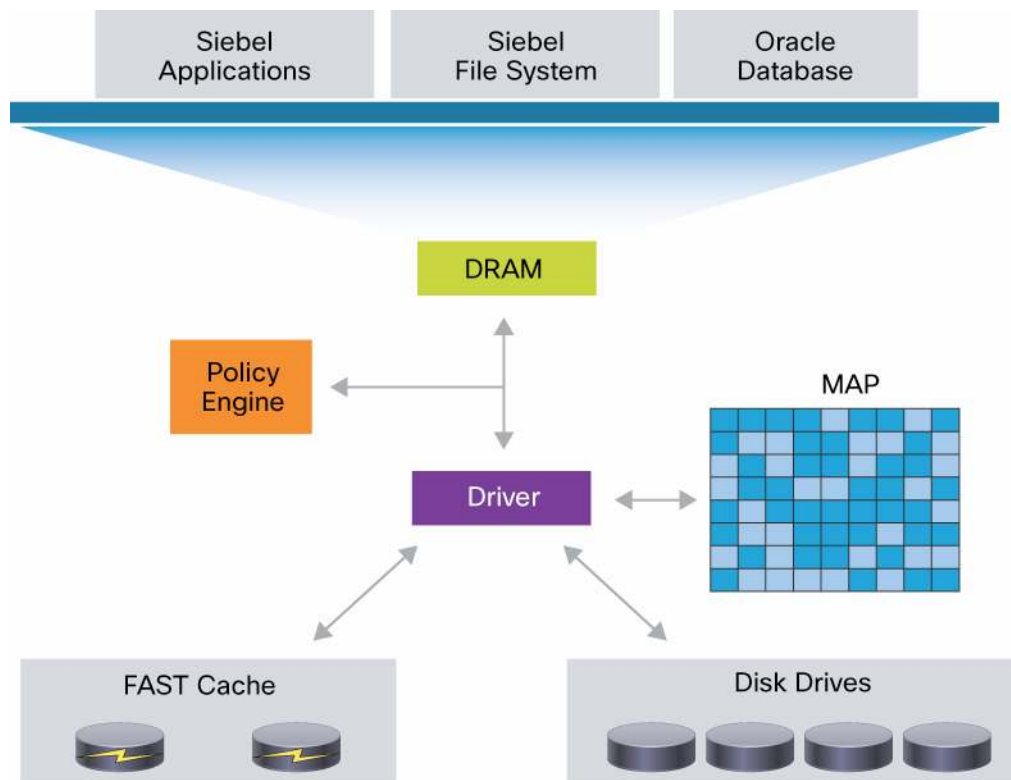
FAST Cache is built on the premise that the overall applications' latencies can improve when most frequently-accessed data is maintained on a relatively smaller-sized, but faster, storage medium, such as Flash drives. FAST Cache identifies the most frequently-accessed data that is temporal in nature, and copies it to Flash drives automatically and non-disruptively. The data movement is completely transparent to applications, thereby 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. This allows the data to be maintained on a faster medium for longer periods, irrespective of application read-to-write mix and data re-write 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 scratch. This can be performed 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 PCI-E slots inside the storage processor.

Figure 7. EMC FAST Cache



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-clarion-celerra-unified-fast-cache-wp.pdf>

4.2.2.3 FAST VP

VNX FAST VP is a policy-based auto-tiering solution for enterprise applications. FAST VP 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. This is accomplished by tiering colder slices of data to high-capacity drives, such as NL-SAS. FAST VP 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, which 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 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>.

FAST Cache and FAST VP are offered in a FAST Suite Package, as part of the VNX Total Efficiency Pack. This pack includes the FAST Suite which automatically optimizes for the highest system performance and lowest storage cost simultaneously. In addition, this pack includes the Security and Compliance Suite, which keeps data safe from changes, deletions, and malicious activity. For additional information on this Total Efficiency Pack as well as other offerings such as the Total Protection Pack, reference: <http://www.emc.com/collateral/software/data-sheet/h8509-vnx-software-suites-ds.pdf>.

The diagram below depicts the disk layout carved on the EMC VNX5500 storage, which is connected with the migration target: the Cisco UCS system.

Figure 8. Disk Layout: EMC VNX5500



4.2.3 Cisco UCS Configuration

This section details the Cisco UCS configuration that was performed as part of the infrastructure build out for deployment of Siebel platform. For the racking, power, and installation of the chassis described in the installation guide, please refer to:

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

One of the important aspects of configuring a physical blade in the Cisco UCS 5108 chassis is to develop a service profile through Cisco UCS Manager. A service profile is an extension of the virtual machine abstraction applied to physical servers. The definition has been expanded to include elements of the environment that span the entire data center. This encapsulates the server identity (LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and quality-of-service [QoS] policies) in logical service profiles. These can be dynamically created and associated with any physical blade in the system within minutes, faster than the conventional approach.

The association of service profiles with physical servers is performed as a simple, single operation. It supports migration of identities between servers in the environment, without requiring any physical configuration changes, and facilitates rapid bare metal provisioning of replacements for failed servers.

For more information on creating the service profile, please refer to the documentation “Create Service Profile for Cisco UCS Blade”

(http://www.cisco.com/en/US/products/ps10281/products_configuration_example09186a0080af7515.shtml).

As referenced in the Section 4.2.1 Detailed Topology, service profiles are created for each of the blades (Web Server/Gateway Server/Application Server and Database Server) and associated with the blades are shown below:

These profiles are created with 4 vHBA connections (2 vHBAs for OS traffic and other two for application traffic) with SAN LUNs in the boot policy.

Figure 9. UCS Manager: Service Profile Setup



4.2.4 BOOT from SAN Setup and OS Installation

4.2.4.1 Boot from SAN

This critical feature helps in moving towards stateless computing in which there is no static binding between a physical server and the operating system and applications it is supposed to run. The OS is installed on a SAN LUN, and boot from SAN policy is applied to the service profile template or the service profile. If the service profile was to be moved to another server, the port world wide name (pwwn) of the HBAs and the server policy also

moves along with it. The new server now takes the same exact view of the old server, the true stateless nature of the blade server.

The main benefits of booting from the SAN:

- Reduce server footprints: Boot from SAN removes the necessity for each server to have its own direct-attached disk, eliminating internal disks as a potential point of failure. Thin diskless servers also take up less facility space, require less power, and are generally less expensive because they have fewer hardware components.
- Disaster and server failure recovery: All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. If a disaster destroys functionality of the servers at the primary site, the remote site can take over with minimal downtime.
- Recovery from server failures is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of its image. As a result, boot from SAN can greatly reduce the time required for server recovery.
- High availability: A typical data center is highly redundant in nature, with redundant paths, redundant disks, and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for mechanical failure of a local disk.
- Rapid redeployment: Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may only need to be in production for hours or days, and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage very cost-effective.
- Centralized image management: When operating system images are stored on networked disks, all upgrades and fixes can be managed at a centralized location. Changes made to disks in a storage array are readily accessible by each server.

With boot from SAN, the image resides on the SAN, and the server communicates with the SAN through a host bus adapter (HBA). The HBAs BIOS contain the instructions that allow the server to find the boot disk. After power on self test (POST), the server hardware component fetches the boot device that is designated as the boot device in the hardware BIOS settings. Once the hardware detects the boot device, it follows the regular boot process.

In order to enable SAN booting, the Cisco Nexus 5548 Switch needs to be configured and zoned. The fabric interconnects are connected to the Cisco Nexus 5548 switches (refer to Figure 6. Detailed Topology), which is also connected to the EMC VNX storage.

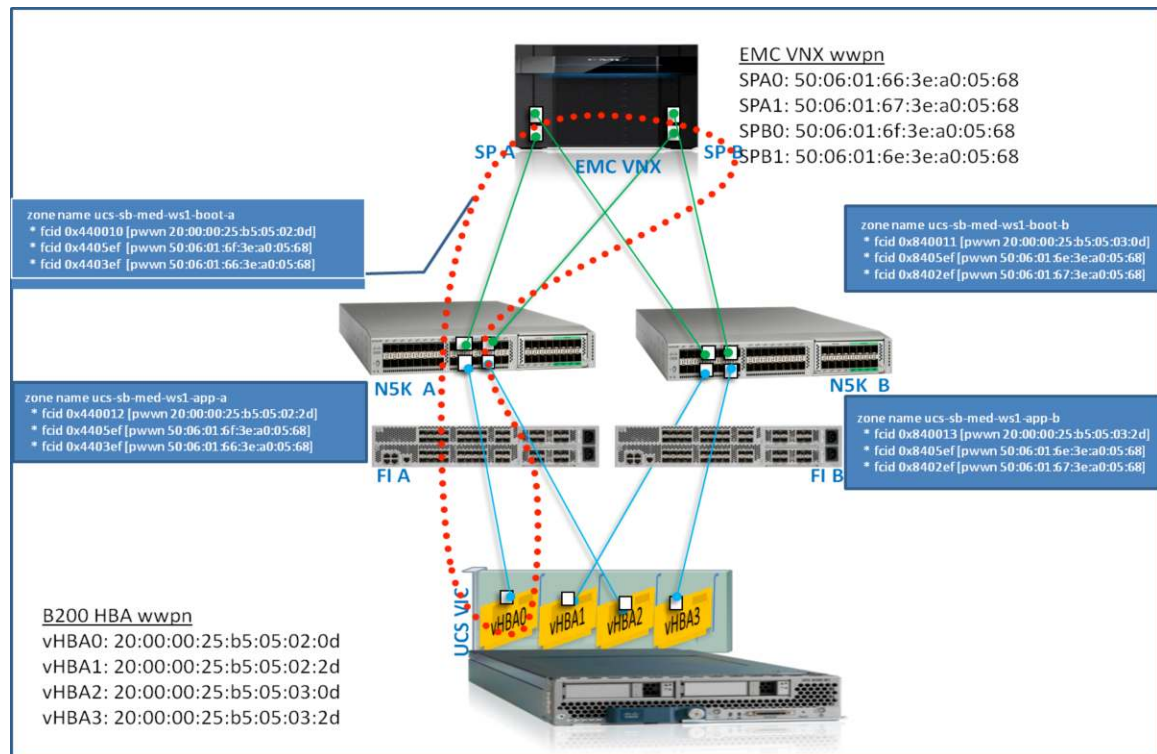
To configure the Cisco Nexus 5548 Switch, follow these steps

1. Make sure that the following configuration details are implemented:
 - The NPIV feature must be enabled on the Cisco Nexus 5548 Multilayer Fabric Switch
 - The 4-GB SPF+ modules must be connected to the Cisco UCS 6120 Series Fabric Interconnect with the port mode and speed set to AUTO
 - If you have created different VSANs, be sure to associate each FC uplink with the correct VSAN
2. Refer to established SAN and zoning best practices for your setup.

3. Complete the zoning.

Please refer to the SAN zoning configuration below (Figure 10). It captures the zones and their associated members that are used for one of the servers in the setup and zones for the Gateway, Application, and Database servers followed in the similar way.

Figure 10. SAN Zoning Configuration



4.2.4.2 Configuring Storage

Follow these steps to configure storage for Cisco UCS data center solution:

1. Make sure of host connectivity.
If each host has the EMC Unisphere Host Agent package installed, the agent should automatically registers the HBA initiators.
2. If the Unisphere Host Agent is not installed, make sure that all initiators are registered properly to complete the host registration.
3. Create the RAID groups and LUNs

It is extremely important that an appropriate storage processor is chosen as the default owner, so that the I/Os on the service processors are evenly balanced. The table below summarizes the LUN ownership and distribution for the setup, followed by additional recommendations.

LUN Configuration Data

RAID Group and Type	LUN	Size	Purpose	Owner (Storage Processor)
RAID Group 1 (RAID-1)	LUN 1	60GB	Boot LUNs for Application Server	SP-A

RAID Group 2 (RAID-1)	LUN2, 3	60GB	Boot LUNs for Database and Webserver	SP-B
RAID Group 3 (RAID-1/0)	LUN4, 5, 6, 7, 8, 14	80 GB	Application binary LUNs for Application Server, Gateway Server, and Database	SP-A
RAID Group 4 (RAID-1/0)	LUN 9	200 GB	Oracle DataLUN	SP-B
RAID Group 5 (RAID-1/0)	LUN11	200 GB	Redo log LUN	SP-A
RAID Group 6 (RAID-1/0)	LUN 11	150 GB	Siebel FileSystem	SP-B

Note that an important value proposition of using a VNX unified storage system in this deployment is highlighted in the table above. There, the Boot LUNs and Oracle data and log use the block capability of the VNX, while the Siebel File System is able to take advantage of the file capability of that same storage system.

Once the zoning is completed on the Cisco Nexus 5548 switch and the SAN connectivity is completed, the boot LUN would be visible on the host.

4.2.4.3 RedHat Linux OS Installation

After making sure that Boot LUN has been made visible to the host, installation of RHEL OS can be done by following the steps given in the document:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/b/os/linux/install/BSERIES-LINUX.pdf

4.2.4.4 EMC PowerPath

PowerPath is host-based software that provides automated data path management and load-balancing capabilities for heterogeneous server, network, and storage deployed in physical and virtual environments. A critical IT challenge is being able to provide predictable, consistent application availability and performance across a diverse collection of platforms. PowerPath is designed to address those challenges, helping IT meet service-level agreements and scale-out mission-critical applications.

This software supports up to 32 paths from multiple HBAs (iSCSI TCI/IP Offload Engines [TOEs] or FCoE CNAs) to multiple storage ports when the multipathing license is applied. Without the multipathing license, PowerPath will use only a single port of one adapter (PowerPath SE). In this mode, the single active port can be zoned to a maximum of two storage ports. This configuration provides storage port failover only, not host-based load balancing or host-based failover. It is supported, but not recommended, if the customer wants true I/O load balancing at the host and also HBA failover.

PowerPath balances the I/O load on a host-by-host basis. It maintains statistics on all I/O for all paths, then intelligently chooses the most underutilized path available. It makes this determination based on statistics and heuristics and the load-balancing and failover policy in effect.

In addition to the load balancing capability, PowerPath also automates path failover and recovery for high availability. If a path fails, I/O is redirected to another viable path within the set. This is transparent to the application, which is not aware of the error on the initial path, avoiding sending I/O errors to the application. Important features of PowerPath include standardized path management, optimized load balancing, and automated I/O path failover and recovery.

For more information reference: <http://www.emc.com/collateral/software/data-sheet/1751-powerpath-ve-multipathing-ds.pdf>

This environment tested with Siebel was performed using PowerPath Version 5.5 (Build 275).

Figure 11. PowerPath Output

```
Pseudo name=emcpowerd
CLARiiON ID=FNM00103100070 [small-sb-ws]
Logical device ID=60060160329029002E153BCA5CD2E011 [ucssmws-boot lun]
state=alive; policy=CLAROpt; priority=0; queued-IOs=0;
Owner: default=SP A, current=SP A   Array failover mode: 1
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode  State  Q-IOs  Errors
=====
  2 fnic              sda      SPA A6   active alive  0    0
  2 fnic              sdb      SP B7   active alive  0    0
  1 fnic              sdc      SPA A7   active alive  0    0
  1 fnic              sdd      SP B6   active alive  0    0
  0 fnic              sdg      SPA A6   active alive  0    0
  0 fnic              sdh      SP B7   active alive  0    0
  3 fnic              sdi      SPA A7   active alive  0    0
  3 fnic              sdj      SP B6   active alive  0    0
```

4.2.5 Siebel Installation

Siebel installation (Version 8.1.1.4) is performed as mentioned in the Siebel Bookshelf documentation:

http://docs.oracle.com/cd/E14004_01/books/SieblnstUNIX/SieblnstUNIXTOC.html.

For this installation, the order below is followed on the identified servers (refer to Figure 6. Detailed Topology):

- Oracle database installation
- Oracle client's installation on the gateway/application servers and database connectivity
- Oracle HTTP webserver installation
- Gateway Server or Siebel Application server installation and configuration
- Database configuration
- Installation and configuration of Siebel Web Server Extension (SWSE) on webserver

4.3 Migration Activities

As mentioned in Section 3.1: Migration Approach, component-by-component migration is followed in this exercise, beginning with database migration. The sections below describe migration steps in detail.

4.3.1 Database Migration

Overview

The existing data from Solaris database is migrated to the target database server, using Oracle's cross-platform transportable tablespace (XTTS) option. As a first step, the Endian format of the Oracle database is queried in the source (Solaris) database as well as on the target (Cisco UCS) database. As mentioned below, the Endian formats are different between the two.

Source Database

Operating System	Host Name	Endian Format	Storage
Solaris 10	Sunsiebdb	Big	SAN

Target Database

Operating System	Host Name	Endian Format	Storage
RedHat Linux 5.4	ucssmdb	Little	SAN

The data migration steps are followed as per Oracle's Maximum Availability Architecture document:

<http://www.oracle.com/technetwork/database/features/availability/maa-wp-10gr2-platformmigration-129296.pdf>

After installation of Oracle database software on the Target:

- Gathered information from the Source database required for migration
- Created the Target database using Database Configuration Assistant (DBCA)
- Created database link and directory for data pump use at Target database
- Created directory for data pump use at Source database
- Ran data pump import on the Target for metadata, using the database link (for the transportable import)
- Exported Source database metadata for importing into the Target database (to address the metadata which was not transported)

As a next phase, disconnected users and restricted access to Source database:

- Exported the user tablespace metadata from the source database
- Original data files were moved to target server by SCP commands (as Target server is more resourceful than the Source server)
- Converted the files using RMAN convert commands to address the Big Endian to Little Endian format difference
- Copied the dump files from the transportable import and from the metadata export of source database to target
- Imported tablespaces into target database as well as other metadata
- The system privileges and sequence values are fixed and the invalid objects are recompiled
- Verified the target database and backed up

4.3.2 Application Migration

Migrating Siebel application server involves moving the Repository data, Siebel files, and Enterprise configuration data to the target application server as detailed below.

Administrative Data (Database Objects)

Administrative data entities are created and modified by administrators, using administrative screens in the Siebel Web Client. Examples for administrative data are List of Values (LOVs), responsibilities and views, as well as positions and organizations.

Repository Migration

Typical repository migration would involve migrating the Siebel repository tables using the repimexp utility.

- Export the complete repository from the Source database to a flat file
- Import the data from the flat file to the Target database as the new Siebel Repository

- Rename the old Siebel Repository in the Target database
- Synchronize the physical schema of the Target database with the object definitions in the data layer of the new repository
- Update the schema version information in the Target database

Note: Since the database was migrated to the Target environment, the repository migration is not warranted.

Files on Siebel Servers and Siebel Web Server Extensions (SWSE)

- Siebel Repository Files (.srf)
- Siebel Web Templates (.swt)
- BI Publisher Report Templates and related files (.rtf and others)
- All file types residing in the public directory of the Siebel Web Server Extension (SWSE), including cascading stylesheets (.css), graphic files, and browser scripts

Example: Perform a file copy by staging all the important folders from the source env and unstage the same in the target env.

```
tar -cvf /siebel/siebsrvr/migration_file_Copy.tar objects/enu/siebel_sia.srf admin/*.ifb webmaster/images
sqltempl/*.sql webmaster/files/enu/*.css webmaster/siebel_build/scripts/*.jswebtempl/*.swt bin/enu/scomm.cfg
bin/enu/eai.cfg
```

Deploying Enterprise Configuration Data using the cfgmerge Utility

Siebel application configuration of enterprise parameters, servers, and components changes can be migrated through the server management screens and views in the Siebel Web Client. It can also be performed through the Siebel Server Manager (srvmgr) command line utility or cfgmerge command line utility.

The siebns.dat file (managed by the Siebel Gateway Name Server) contains some of the parameters refer to enterprise-specific paths or host names, such as the file system path or the chart server host name. As a result, the siebns.dat file cannot be simply copied from one enterprise to the other. Instead, cfgmerge command line utility can be used to achieve a merge of configuration stores.

cfgmerge Utility Features and Limitations:

The cfgmerge utility can produce input files on the enterprise level and the individual server level. It can only compare parameters, component definitions, and server components that are already present in both siebns.dat files and have the same names in both files.

If the cfgmerge utility is used to deploy new component definitions from the Source system to the Target system, they have to be created manually in the target system. This is done by simply copying and renaming an existing component definition in the GUI, or by using a script and then running the utility:

- Log in to the Siebel Server Manager (srvmgr) command line against the Source enterprise.
- Use the backup namesrvr command to create a backup of the current siebns.dat source file and rename it to siebns_source.dat.
- Log in to the Siebel Server Manager (srvmgr) command line against the Target enterprise.
- Use the backup namesrvr command to create a backup of the current siebns.dat target file and rename it to siebns_target.dat.

- Copy both files created in the previous steps to a temporary directory on a machine where a Siebel Gateway Name Server is installed.
- On the Siebel Gateway machine, from the command shell, navigate to the bin directory of the Siebel Gateway Name Server installation folder.
- Execute a command similar to the following to create the input file at the enterprise level:

```
Ex: cfgmerge -l ENU -i /cfgmerge/siebns_source.dat, /cfgmerge/ siebns_target.dat -e
SBA_81,SBA_81 -o /cfgmerge/SBA_81_ ENT.txt
```

Note: The above command invokes the `cfgmerge` utility. The `-l` parameter takes a three-letter language code. The `-i` switch is followed by a comma-separated list of the full paths to the file, representing the source `siebns.dat` and the file representing the target `siebns.dat`. The `-e` switch is followed by a comma-separated list of the names of the source enterprise and the target enterprise. The `-o` parameter value specifies the path where the output file should be written to.

- Execute a command similar to the following to create the input file at the server level and repeat for any additional Siebel servers in the Source or Target enterprise. It is not repeated in this current migration context, since there is only one Siebel server in the Target system.

```
Ex: cfgmerge -l ENU -i /cfgmerge/siebns_source.dat, /cfgmerge/ siebns_target.dat -e
SBA_81,SBA_81 -s siebserver1, ucssmas -o /cfgmerge/SBA_81_ SRVR.txt
```

- Review and modify the output files, if necessary. It is mandatory to thoroughly review the output files of the `cfgmerge` utility, in order to avoid unwanted changes being applied to the target enterprise configuration.
- Log in to the Siebel Server Manager (`srvmgr`) against the Target enterprise. Then, execute a command similar to the following to apply the changes (the `read` command opens the specified file and executes all commands in that file):

```
read /cfgmerge/ SBA_81_ ENT.txt
```

```
read /cfgmerge/ SBA_81_ SRVR.txt
```

The process is completed by shutting down and restarting all services of the Target enterprise. If end users are affected by the shutdown, they must be notified of the downtime well in advance.

4.3.3 File System Migration

Siebel typically stores reports, files generated during user navigations, and more, in the file system. As a result, migration of these files on the target servers is critical for business continuity. Below are the steps to be followed for file system migration:

- File system related batch jobs are terminated before migration
- Siebel Enterprise is shut down and File System is unmounted at the source
- Once the target file system is mounted and accessible, SAN copy options were used to copy the files
- File system is mounted at the target, edit, and other `fstab` entries and enterprise is restarted

4.3.4 Post-Migration Validation

- Run EVT (Enterprise Verification Tool Utility) from Oracle to check on the Target environment and resolve any differences found.
- Perform connectivity check from application server to database server, through sqlplus and odbcsql.
- Cross-verify that the file copy is successful.
- Start the enterprise and check for any errors:
 - a. Examine enterprise and components logs
 - b. Scan core files for any crashes
 - c. Perform UI navigation for any errors

As a part of validation, Siebel customizations were repeated on target, as reflected in Figure 12 below:

Figure 12. Post-Migration Validation

UI Changes: Opportunity List Applet Configured with New Field– Channel (RISC Platform)

Siebel Communications - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Home Search Favorites

Address http://10.104.111.14:61000/communications_enu/start.swe?SWECmd=Login&SWEPL=1&SWETS=1313667084066

File Edit View Navigate Query Tools Help

Opportunity:

Home Accounts Orders Contacts Opportunities Service Request

Opportunity Explorer | Manager's Explorer | Opportunities Home | Opportunities List | Opportunity Charts | Lead Sources Administration | Opportunities Administration | Opportun...

My Opportunities Menu New Delete Query

Opportunity Name Account Channel Sales Cycle Sales Method Sales Team Revenue Committed

Ugandaa001

133ACK1

Channel

Sales Cycle

Sales Method

Sales Team

Revenue

Committed

Ugandaa001

Noranda Inc.

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa002

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa003

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa004

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa005

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa006

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa007

Walk-In

ACAPS Application

SADMIN

\$0.00

Ugandaa008

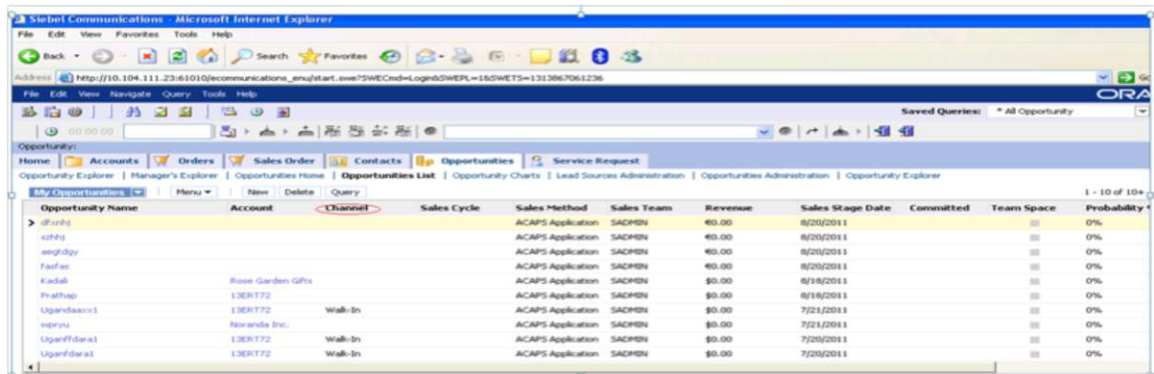
Walk-In

ACAPS Application

SADMIN

\$0.00

UI Changes: Opportunity List Applet Configured with New Field– Channel (Cisco UCS Platform)



Opportunity Name	Account	Channel	Sales Cycle	Sales Method	Sales Team	Revenue	Sales Stage Date	Committed	Team Space	Probability
Ugandaa001	133ACK1	Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa002	Noranda Inc.	Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa003		Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa004		Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa005		Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa006		Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa007		Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%
Ugandaa008		Walk-In	ACAPS Application	SADMIN	\$0.00	8/20/2011				0%

Custom Workflows (RISC Platform)

Siebel Communications - Microsoft Internet Explorer

Address: http://10.104.111.14:61000/ecomunications_enu/start.swe?SWECmd=Login&SWEPL=1&SWETS=1313867084066

Process:

Home Accounts Orders Contacts Opportunities Service Request Administration - Application Administration - Business Process

Workflow Policies | Workflow Policy Actions | Workflow Policy Explorer | Workflow Policy Groups | Workflow Policy Log | Workflow Deployment | Workflow Processes | Workflow Instance Admin

Repository Workflow Processes | Menu | Query | Activate | Query Results

Name	Business Object	Status	Group	Version	Mode
Cisco Opportunity Process	Opportunity	Completed		1	Service Flow
Cisco SR Process	Service Request	Completed		0	Service Flow

Custom Workflows (UCS Platform)

Siebel Communications - Microsoft Internet Explorer

Address: http://10.104.111.23:61010/ecomunications_enu/start.swe?SWECmd=Login&SWEPL=1&SWETS=1313867061236

Process:

Home Accounts Orders Sales Order Contacts Opportunities Service Request Administration - Application Administration - Business Process

Workflow Policies | Workflow Policy Actions | Workflow Policy Explorer | Workflow Policy Groups | Workflow Policy Log | Workflow Deployment | Workflow Processes | Workflow Instance Admin | Workflow Instance Admin

Repository Workflow Processes | Menu | Query | Activate | Query Results

Name	Business Object	Status	Group	Version	Mode
Cisco Opportunity Process	Opportunity	Completed		1	Service Flow
Cisco SR Process	Service Request	Completed		0	Service Flow

EAI Subsystem Changes (RISC Platform)

Enterprise Servers | Menu | Backup Enterprise

Enterprise Server Description

SBA_B1 Siebel Enterprise Server

Component Groups | Component Definitions | System Alerts | **Profile Configuration** | Parameters | Synchronize

Menu | New | Query | Query Results

Profile	Alias	Subsystem Type	Description
CISCOEAI	CISCOEAI	EAITransportDataHandlingSubsys	
CISCO EAI Order	CISCO EAI Order	EAITransportDataHandlingSubsys	CISCO EAI Order

EAI Subsystem Changes (UCS Platform)

Enterprise Servers | Menu | Backup Enterprise

Enterprise Server Description

SBA_B1 Siebel Enterprise Server

Component Groups | Component Definitions | System Alerts | **Profile Configuration** | Parameters | Synchronize


Menu | New | Query | Query Results

Profile	Alias	Subsystem Type	Description
CISCOEAI	CISCOEAI	EAITransportDataHandlingSubsys	To Handle Orders from Stub
CISCO EAI Order	CISCOEAIOrder	EAITransportDataHandlingSubsys	Update the Orders through Stub

5. Performance Comparison

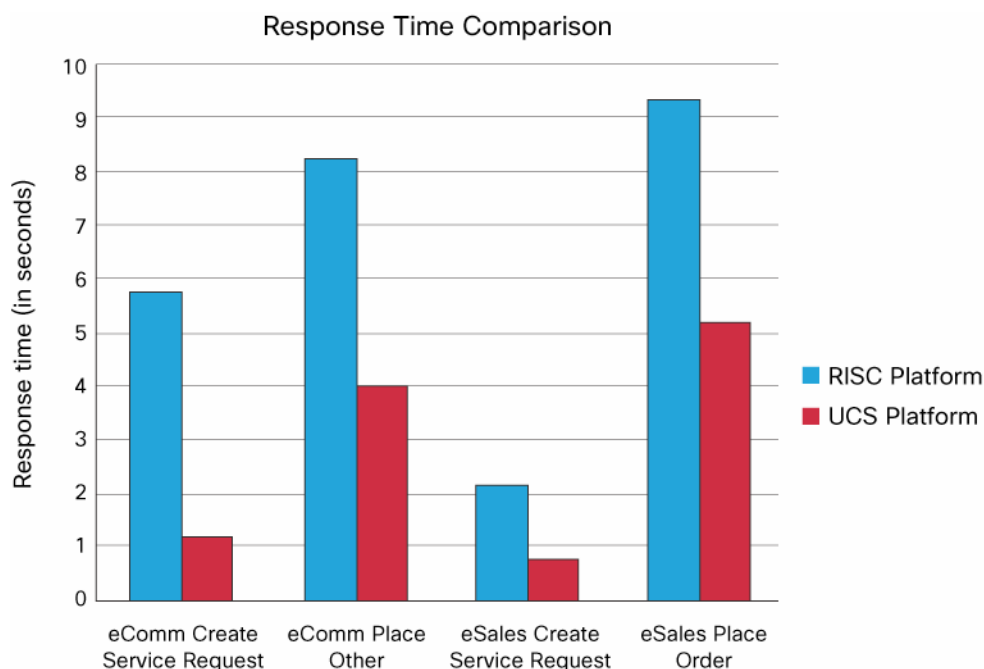
The performance of a Siebel application implemented on the RISC platform is compared with a migrated Siebel application on the Cisco UCS platform by executing performance tests using the LoadRunner tool. Siebel Enterprise on UCS servers has shown faster response times with reduced CPU usage, compared to Siebel Enterprise running on RISC-based servers, as captured below.

Figure 13. Comparison of CPU Usage: Sun with Solaris Compared to Cisco UCS with Linux



Response times for various Siebel services are compared in Figure 14 below.

Figure 14. Comparison of Response Times: Sun with Solaris Compared to Cisco UCS with Linux



6. Conclusion

Oracle's Siebel CRM applications, running on Cisco UCS, can reduce the TCO at the platform, site, and organizational levels. They can also increase IT staff productivity and business agility, through just-in-time provisioning and mobility support for both virtualized and non-virtualized environments.

Migrating Siebel CRM from Solaris to Cisco UCS is straightforward, and can be done with minimum downtime; however, it does require planning and coordination. The procedure has been tested and optimized by Cisco in the lab environment, to avoid surprises and inefficiencies in a real-life scenario. As a standard practice, it is recommended that the development and QA environments, which are low-risk, are migrated before moving the production environment.

This document provides guidance on the best way to plan and execute Cisco UCS migration, taking into account the best sequence of steps to be followed and preconditions necessary to help ensure a transparent transition.

The TCO, resource utilization, space, and power consumption levels in the existing environment need to be measured, so that the new environment can be compared on these parameters.

7. For More Information

Migrate from RISC Servers to Cisco UCS

http://www.cisco.com/en/US/partner/prod/ps10265/uc_risc.html

Cisco UCS Services: Accelerate Your Transition to a Unified Computing Architecture

http://www.cisco.com/en/US/services/ps2961/ps10312/Unified_Computing_Services_Overview.pdf.

Cisco UCS Delivers World-Record Application Server Performance

http://www.cisco.com/en/US/prod/collateral/ps10265/LE-212506_PB_jAppServer_B230.pdf.

x86 Blades: Shrinking the Branch Office

http://www.cisco.com/en/US/partner/prod/collateral/modules/ps10598/white_paper_c11-613183_ps11273_Products_White_Paper.html.

EMC VNX Series:

<http://www.emc.com/storage/vnx/vnx-series.htm>



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)

Copyright © 2012 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided "as is." EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

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

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