

# Oracle E-Business Suite R12 on Cisco UCS with EMC VNX 5500

## Contents

<b>Introduction.....</b>	<b>3</b>
<b>Leadership from Cisco.....</b>	<b>3</b>
<b>Introducing Cisco UCS.....</b>	<b>3</b>
<b>Overview of Cisco UCS .....</b>	<b>4</b>
<b>Cisco Nexus 5548UP Switch.....</b>	<b>10</b>
<b>EMC VNX Unified Storage System .....</b>	<b>11</b>
<b>Why Unified Infrastructure?.....</b>	<b>13</b>
<b>Overview of the Solution.....</b>	<b>14</b>
<b>Oracle E-Business Suite R12 on Cisco UCS 2.1 and EMC VNX 5500 .....</b>	<b>15</b>
<b>Configuring Cisco Unified Computing System for Oracle E-Business Suite.....</b>	<b>15</b>
Configuring Fabric Interconnects.....	15
Configuring Server Ports.....	16
Configuring SAN and LAN on UCS Manager.....	16
Configure Ethernet Port Channels.....	19
Preparatory Steps Before Creating Service Templates .....	21
Create Service Profile Template: Database.....	24
Create Service Profile Template: Apps.....	30
Create Service Profiles from Service Profile Templates .....	33
<b>Setting Up EMC VNX Storage .....</b>	<b>35</b>
Hardware Storage Processors Configuration .....	36
<b>Configure SAN Zoning on Nexus 5548UP Switches.....</b>	<b>37</b>
Fibre Channel Zoning .....	37
Setting Up vPC on the Nexus 5548UP Switches .....	39
<b>Configure the VNX 5500 for NFS .....</b>	<b>41</b>
Create a Storage Pool for NFS (RAID 0/1) .....	44
Assign Host IDs to the LUNs .....	47
Create Apps File Systems .....	49
Create NFS Exports.....	51
<b>Install Operating System, Additional RPMs, and Prepare the System for the Oracle Database and Application Tier Servers.....</b>	<b>52</b>
Prepare to Install the Boot LUNs .....	52
Configure Storage Boot LUNs .....	52
Install Oracle Linux 5.8 from Image .....	55
Configure PowerPath.....	60
Configuring the Boot LUN.....	61
Reconfigure Zoning and Boot Policies.....	62
Miscellaneous Post-Installation Steps .....	63
Configure Oracle ASM on Database Server .....	65

Configure ASM LUNs and Create Disks .....	66
Installing the Operating System on Oracle Application-Tier Servers .....	67
<b>Rapid Install of Oracle Apps R12 .....</b>	<b>67</b>
Invoking Rapid Install.....	69
Moving the Database Files to ASM.....	76
Sanity Tests on Oracle Applications .....	77
Adding Another Web/Apps Node to the Infrastructure.....	79
<b>Configure a Load Balancer .....</b>	<b>81</b>
Configure a Hardware Load Balancer.....	81
Run Autoconfig After Configuring the Load Balancer .....	81
<b>Install EM Grid Control 12c.....</b>	<b>82</b>
<b>Using Cisco UCS Service Profiles for Failover to a Spare Blade .....</b>	<b>82</b>
<b>Performance and Destructive Tests.....</b>	<b>85</b>
Performance Tests .....	85
Details Collected from Tools .....	85
Destructive Tests .....	87
<b>Appendix .....</b>	<b>87</b>
Appendix A: Cisco UCS Service Profiles .....	87
Appendix B: Nexus 5548UP Zoning Details .....	89
Appendix C: Cisco ACE Configuration.....	90
Appendix D: ASM parameters .....	92
Appendix E: Database Instance Parameters.....	93
Appendix F: Oracle Apps Context File Parameters .....	93

---

## Introduction

This document describes how the Cisco Unified Computing System™ (Cisco UCS®) can be used in conjunction with EMC VNX unified storage systems to implement an Oracle E-Business Suite solution. Cisco UCS provides the compute, network, and storage access components of the cluster, deployed as a single, cohesive system. The document also showcases the boot-over-SAN capabilities of Cisco UCS that reduce the mean time to recover from hardware failures with little or no interruption to the Oracle E-Business system, depending on the topology. The result is an implementation that addresses many of the challenges that mission-critical environments in data centers face today, including the need for a simplified deployment and operational model, high performance with flexibility, and lower total cost of ownership (TCO).

## Leadership from Cisco

Cisco is the undisputed leader in providing network connectivity in enterprise data centers. With the introduction of the Cisco Unified Computing System, Cisco is now equipped to provide the entire clustered infrastructure for Oracle Database, E-business Suite, fusion middleware, and several other Oracle software products. Cisco UCS provides compute, network, virtualization, and storage access resources that are centrally controlled and managed as a single, cohesive system. With the capability to centrally manage both blade and rack-mount servers, Cisco UCS provides an ideal foundation for Oracle deployments.

### Target Audience

This document is intended to assist solution architects, system and storage administrators, database administrators, sales engineers, field engineers, and consultants in the planning, design, deployment, and migration of Oracle E-Business systems to Cisco UCS servers. It assumes that the reader has an architectural understanding of Cisco UCS servers, Oracle E-Business Suite, and storage and networking concepts.

### Purpose

The purpose of this document is to demonstrate best practices in setting up Oracle E-Business Suite R12 with multiple application and web servers using shared APPL\_TOP concepts and database software. While the paper focuses mainly on multiple web servers and SAN boot capabilities to achieve stateless computing, it can be further extended as needed by deploying Oracle RAC and Ebusiness Suite R12 configurations to further reduce any single point of failure in the system.

### Configuration

All components in an Oracle E-Business Suite implementation must work together flawlessly, and Cisco has worked closely with EMC and Oracle to create, test, and certify a configuration of Oracle E-Business Suite on the Cisco UCS. This paper provides an implementation of Oracle Database consistent with industry best practices. For back-end SAN storage, the environment included an EMC VNX storage system. Also, VNX capabilities were harnessed to use NFS in a multitier Oracle applications install.

## Introducing Cisco UCS

Cisco UCS addresses many of the challenges faced by database administrators and their IT departments, making it an ideal platform for Oracle Real Application Clusters (RAC) implementations.

### Comprehensive Management

The system uses an embedded, end-to-end management system that uses a high-availability active-standby configuration. Cisco UCS Manager uses role and policy-based management that allows IT departments to continue to use subject-matter experts to define server, network, and storage access policy. After a server and its

---

identity, firmware, configuration, and connectivity are defined, the server, or a number of servers like it, can be deployed in minutes, rather than the hours or days that it typically takes to move a server from the loading dock to production use. This capability relieves database administrators from tedious, manual assembly of individual components and makes scaling an Oracle RAC configuration a straightforward process.

#### Radical Simplification

Cisco UCS represents a radical simplification compared to the way servers and networks are deployed today. It reduces network access-layer fragmentation by eliminating switching inside the blade server chassis. It integrates compute resources on a unified I/O fabric that supports standard IP protocols as well as Fibre Channel through Fibre Channel over Ethernet (FCoE) encapsulation. The system eliminates the limitations of fixed I/O configurations with an I/O architecture that can be changed through software on a per-server basis to provide needed connectivity using a just-in-time deployment model. The result of this radical simplification is fewer switches, cables, adapters, and management points, helping reduce cost, complexity, power needs, and cooling overhead.

#### High Performance

The system's blade servers are based on the Intel® Xeon® 5670 and 7500 series processors. These processors adapt performance to application demands, increasing the clock rate on specific processor cores as workload and thermal conditions permit. The system is integrated within a 10 Gigabit Ethernet-based unified fabric that delivers the throughput and low-latency characteristics needed to support the demands of the cluster's public network, storage traffic, and high-volume cluster messaging traffic.

### Overview of Cisco UCS

Cisco UCS unites computing, networking, storage access, and virtualization resources into a single, cohesive system. When used as the foundation for Oracle Database and E-Business Suite software, the system brings lower total cost of ownership (TCO), greater performance, improved scalability, increased business agility, and Cisco's hallmark investment protection.

The system represents a major evolutionary step away from the current traditional platforms in which individual components must be configured, provisioned, and assembled to form a solution. Instead, the system is designed to be stateless. It is installed and wired once, with its entire configuration—from RAID controller settings and firmware revisions to network configurations—determined in software using integrated, embedded management.

The system brings together server resources powered by Intel Xeon processors on a 10-Gbps unified fabric that carries all IP networking and storage traffic, eliminating the need to configure multiple parallel IP and storage networks at the rack level. It uses dramatically fewer components compared to other implementations, reducing TCO, simplifying and accelerating deployment, and reducing the complexity that can be a source of errors and downtime.

Cisco UCS is designed to be form-factor neutral. The core of the system is a pair of fabric interconnects that link all the computing resources together and integrate all system components into a single point of management. Today, blade server chassis are integrated into the system through fabric extenders that bring the system's 10-Gbps unified fabric to each chassis.

---

The FCoE protocol collapses Ethernet-based networks and storage networks into a single common network infrastructure, thus reducing capital expenditures (CapEx) by eliminating redundant switches, cables, networking cards, and adapters, and reducing operating expenses (OpEx) by simplifying administration of these networks (Figure 1). Other benefits include:

- I/O and server virtualization
- Transparent scaling of all types of content, either block or file based
- Simpler and more homogeneous infrastructure to manage, enabling data center consolidation

#### Fabric Interconnects

The Cisco® fabric interconnect is a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. It offers line-rate, low-latency, lossless 10 Gigabit Ethernet, FCoE, and Fibre Channel functions.

The fabric interconnect provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the fabric interconnect become part of a single, highly available management domain. In addition, by supporting unified fabric, the fabric interconnect supports both LAN and SAN connectivity for all blades within their domain. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from a blade server through an interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

The Cisco UCS 6120XP 20-Port Fabric Interconnect provides low-latency, lossless, 10-Gbps unified fabric connectivity for the cluster. The interconnect provides connectivity to blade server chassis and the enterprise IP network.

#### Fabric Extenders

The Cisco fabric extenders multiplex and forward all traffic from blade servers in a chassis to a parent Cisco UCS fabric interconnect from 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis, 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.

The Cisco UCS 2104XP Fabric Extender brings the unified fabric into each blade server chassis. The fabric extender is configured and managed by the fabric interconnects, eliminating the complexity of blade-server-resident switches. Two fabric extenders are configured in each of the cluster's two blade server chassis.

Each fabric extender on either side of the chassis is connected through 10 Gigabit Ethernet links to the fabric interconnects and offers:

- Connection of the Cisco UCS blade chassis to the fabric interconnect
- Four 10 Gigabit Ethernet, FCoE-capable SFP+ ports
- Built-in chassis management function to manage the chassis environment (the power supply and fans as well as the blades) along with the fabric interconnect, eliminating the need for separate chassis management modules
- Full management by Cisco UCS Manager through the fabric interconnect

- Support for up to two fabric extenders, enabling increased capacity as well as redundancy
- Up to 160 Gbps of bandwidth per chassis

#### Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of Cisco UCS, delivering a scalable and flexible blade server chassis.

#### Cisco UCS Manager

Cisco UCS Manager provides unified, embedded management of all software and hardware components of Cisco UCS across multiple chassis, rack-mount servers, and thousands of virtual machines. It manages Cisco UCS as a single entity through an intuitive GUI, a command-line interface (CLI), or an XML API for comprehensive access to all Cisco UCS Manager functions.

#### Cisco UCS Virtual Interface Card 1280

Cisco UCS Virtual Interface Card (VIC) 1280 is the second generation of mezzanine adapters from Cisco. The VIC 1280 supports up to 256 PCIe devices and up to 80 Gbps of throughput. Compared with its earlier generation of Palo adapters, it has double the capacity in throughput and PCIe devices and is compatible with many OS and storage vendors. Cisco VIC 1280 card was used in the database server B440-M2.

#### Cisco UCS Virtual Interface Card 1240

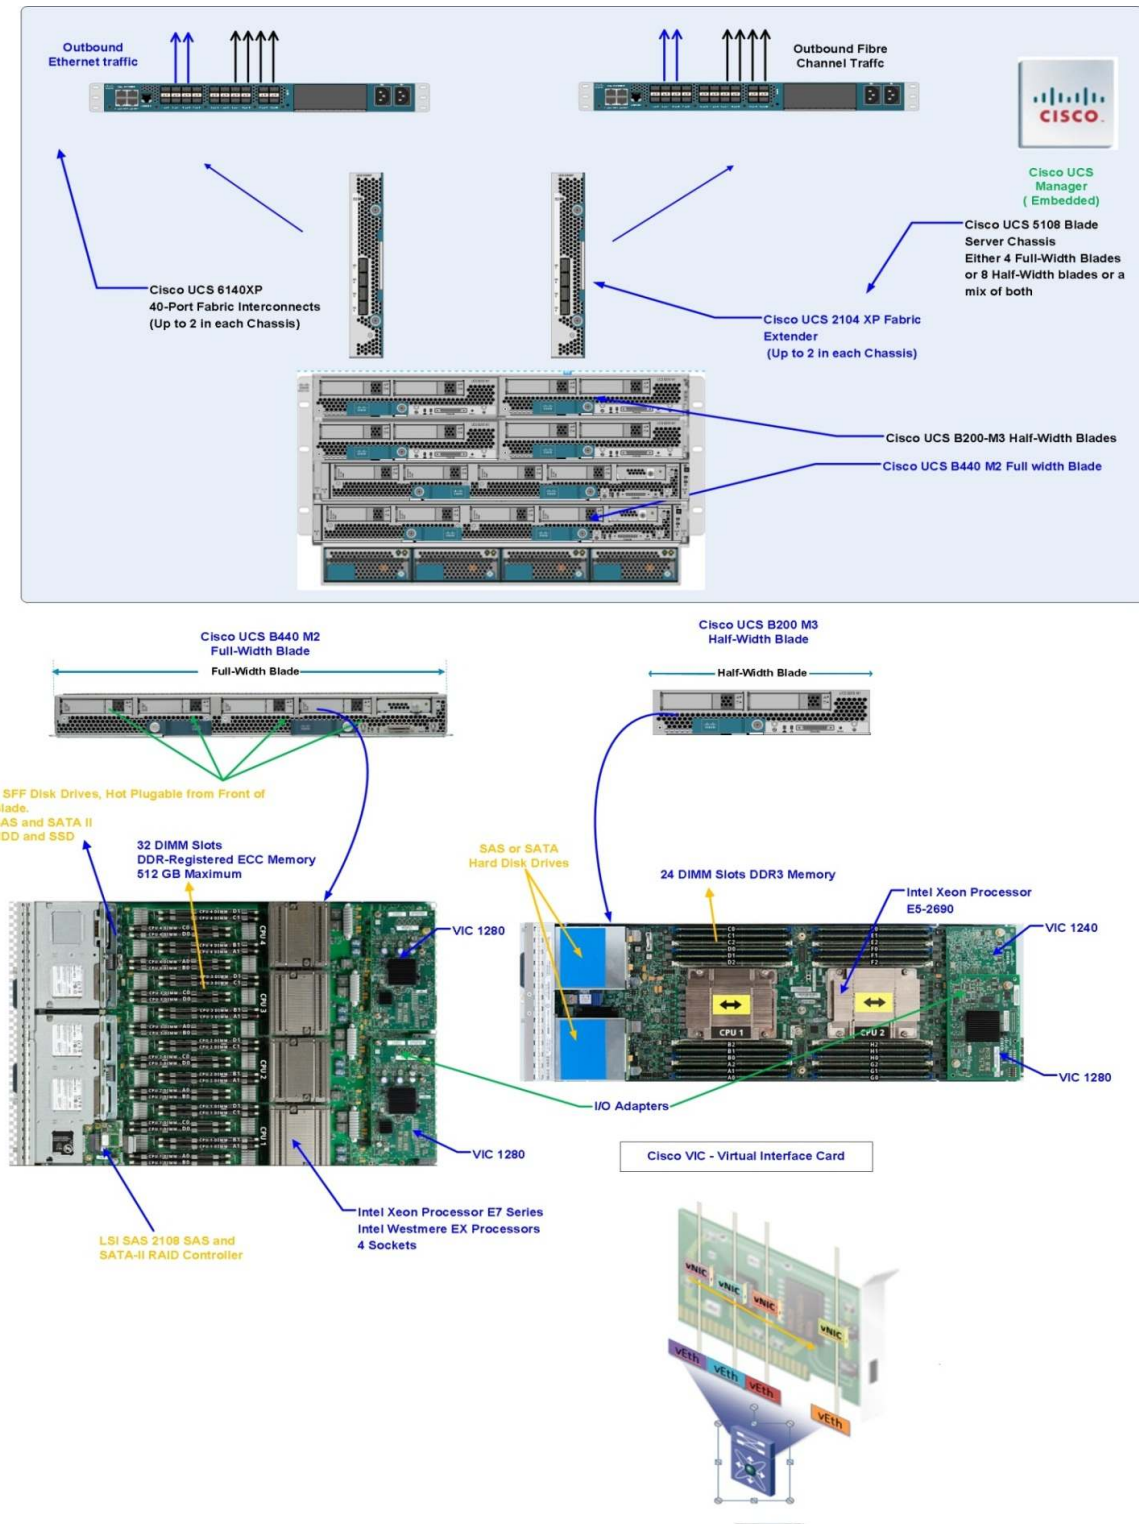
Based on second-generation Cisco VIC technology, the VIC 1240 is a modular LAN on motherboard (LOM) that is designed specifically for the M3 generation of Cisco UCS B-Series Servers. The Cisco UCS VIC 1240 offers industry leading performance, flexibility, and manageability. The VIC 1240 is capable of aggregate 8 x 10 Gbps speed to the half-width blade slot when used with the Port Expander Card. Without the Port Expander Card, the VIC 1240 enables four ports of 10 Gbps network I/O to each half-width blade server. Cisco VIC 1240 card was used in the Oracle Ebusiness Suite Server B200-M3.

#### Cisco UCS B440 M2 High-Performance Blade Servers

The Cisco UCS B440 M2 High-Performance Blade Servers are full-slot, 4-socket, high-performance blade servers 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 supports four Small Form Factor (SFF) SAS and SSD drives and two converged network adapter (CNA) mezzanine slots up to 80 Gbps of I/O throughput. The Cisco UCS B440 blade server extends Cisco UCS by offering increased levels of performance, scalability, and reliability for mission-critical workloads.

The Cisco UCS components are shown in [Figure 1](#).

**Figure 1.** Cisco UCS Components





## Service Profiles: Cisco Unified Computing System Foundation Technology

Cisco UCS resources are abstract in the sense that their identity, I/O configuration, MAC addresses and worldwide names (WWNs), firmware versions, BIOS boot order, and network attributes (including quality-of-service [QoS] settings, pin groups, and threshold policies) are all programmable using a just-in-time deployment model. The manager stores this identity, connectivity, and configuration information in service profiles that reside on the Cisco UCS 6100 or 6200 Series Fabric Interconnects. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. Service profiles allow server and network definitions to move within the management domain, enabling flexibility in the use of system resources. Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools.

### Service Profile Description, Overview, and Elements

#### Service Profile Description

Conceptually, 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, encapsulating 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 that can be dynamically created and associated with any physical server in the system within minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation. It enables 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. Service profiles also include operational policy information, such as information about firmware versions.

This highly dynamic environment can be adapted to meet rapidly changing needs in today's data centers with just-in-time deployment of new computing resources and reliable movement of traditional and virtual workloads. Data center administrators can now focus on addressing business policies and data access on the basis of application and service requirements, rather than physical server connectivity and configurations. In addition, using service profiles, Cisco UCS Manager provides logical grouping capabilities for both physical servers and service profiles and their associated templates. This pooling or grouping, combined with fine-grained role-based access, allows businesses to treat a farm of compute blades as a flexible resource pool that can be reallocated in real time to meet their changing needs, while maintaining any organizational overlay on the environment that they want.

#### Service Profile Overview

A service profile typically includes four types of information:

- **Server definition:** Defines the resources (for example, a specific server or a blade inserted into a specific chassis) that are required to apply to the profile.
- **Identity information:** Includes the universally unique identifier (UUID), MAC address for each virtual NIC (vNIC), and WWN specifications for each HBA.
- **Firmware revision specifications:** Used when a certain tested firmware revision is required to be installed or if a specific firmware is used for some other reason.
- **Connectivity definition:** Used to configure network adapters, fabric extenders, and parent interconnects; however, this information is abstract, as it does not include the details of how each network component is configured.



---

A service profile is created by the UCS server administrator. This service profile uses configuration policies that were created by the server, network, and storage administrators. Server administrators can also create a service profile template, which can later be used to create service profiles in an easier way. A service template can be derived from a service profile, with server and I/O interface identity information abstracted. Instead of specifying exact UUID, MAC address, and WWN values, a service template specifies where to get these values. For example, a service profile template might specify the standard network connectivity for a web server and the pool from which its interface's MAC addresses can be obtained. Service profile templates can be used to provision many servers with the same simplicity as creating a single one.

### **Service Profile Elements**

In summary, service profiles represent all the attributes of a logical server in the Cisco UCS data model. These attributes have been abstracted from the underlying attributes of the physical hardware and physical connectivity. Using logical servers that are disassociated from the physical hardware removes many limiting constraints involving how servers are provisioned. Using logical servers also makes it easy to repurpose physical servers for different applications and services.

### **Understanding Service Profile Template**

A lot of time can be lost between the point when a physical server is in place and when that server begins hosting applications and meeting business needs. Much of this lost time is due to delays in cabling, connecting, configuring, and preparing the data center infrastructure for a new physical server. In addition, provisioning a physical server requires a large amount of manual work that must be performed individually on each server. In contrast, the Cisco UCS Manager uses service profile templates to significantly simplify logical (virtual) server provisioning and activation. The templates also allow standard configurations to be applied to multiple logical servers automatically, which reduces provisioning time to just a few minutes.

Logical server profiles can be created individually or as a template. Creating a service profile template allows rapid server instantiation and provisioning of multiple servers. The Cisco UCS data model (with pools, policies, and isolation security methods) also creates higher-level abstractions such as vNICs and virtual HBAs (vHBAs). Ultimately, these service profiles are independent of the underlying physical hardware. One important aspect of the Cisco UCS data model is that it is highly referential. This means you can easily reuse and refer to previously defined objects and elements in a profile without having to repeatedly redefine their common attributes and properties.

The Cisco system used in the setup is based on Cisco UCS B-Series blade servers; however, the breadth of Cisco's server and network product line suggests that similar product combinations will meet the same requirements. It was built from the following hierarchy of components.

- The Cisco UCS 6120XP Fabric Interconnect provides low-latency, lossless, 10-Gbps unified fabric connectivity for the cluster. The fabric interconnect provides connectivity to blade server chassis and the enterprise IP network. Two fabric interconnects are configured in the cluster and providing the capability to securely takeover the other in the event of a failure.
- The Cisco UCS 2104XP Fabric Extender brings the unified fabric into each blade server chassis. The fabric extender is configured and managed by the fabric interconnects, eliminating the complexity of blade-server-resident switches. Two fabric extenders are configured in each of the cluster's two blade server chassis.

- The Cisco UCS 5108 Blade Server Chassis houses the fabric extenders, up to four power supplies, and up to four full-width blade servers. As part of the system's radical simplification, the blade server chassis is also managed by the fabric interconnects, eliminating another point of management.
- The blade server form factor supports a range of mezzanine-format Cisco UCS network adapters, including a 40 Gigabit MLOM adapter designed for efficiency and performance, the Cisco UCS VIC 1240, designed to deliver outstanding performance and full compatibility with existing Ethernet and Fibre Channel environments. These adapters present both an Ethernet NIC and a Fibre Channel HBA to the host operating system. They make the existence of the unified fabric transparent to the operating system, passing traffic from both the NIC and the HBA on to the unified fabric. The database server B440-M2 had two Cisco UCS VIC 1280s per blade that provided 80 Gbps of performance per blade server.

## Cisco Nexus 5548UP Switch

**Figure 2.** Cisco Nexus 5548UP Switch



The Cisco Nexus<sup>®</sup> 5548UP Switch (Figure 2) delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages, including:

### Architectural Flexibility

- Includes unified ports that support traditional Ethernet, Fibre Channel, and FCoE
- Synchronizes system clocks with accuracy of less than one microsecond, based on IEEE 1588
- Supports secure encryption and authentication between two network devices, based on Cisco TrustSec<sup>®</sup> IEEE 802.1AE
- Offers converged fabric extensibility, based on emerging standard IEEE 802.1BR, with the fabric extender (FEX) technology portfolio, including:
  - Cisco Nexus 2000 Series FEX
  - Cisco Adapter FEX
  - Cisco VM-FEX

### Infrastructure Simplicity

- Provides a common high-density, high-performance, data center-class, fixed-form-factor platform
- Consolidates LAN and storage
- Supports any transport over an Ethernet-based fabric, including Layer 2 and Layer 3 traffic
- Supports storage traffic, including iSCSI, NAS, Fibre Channel, etc.
- Reduces management points with FEX technology

#### Business Agility

- Enables diverse data center deployments on one platform
- Provides rapid migration and transition for traditional and evolving technologies
- Offers performance and scalability to meet growing business needs

#### Specifications at a Glance

- A 1RU, 1/10 Gigabit Ethernet switch
- 32 fixed unified ports on base chassis and one expansion slot, totaling 48 ports
- The slot can support any of the three modules: unified ports, 1/2/4/8 native Fibre Channel, and Ethernet or FCoE
- Throughput of up to 960 Gbps

### EMC VNX Unified Storage System

The EMC VNX Series Unified Storage Systems (Figure 3) deliver uncompromising scalability and flexibility for the midtier while providing market-leading simplicity and efficiency to reduce TCO.

Based on the powerful family of Intel Xeon 5600 processors, the EMC VNX Series implements a modular architecture that integrates hardware components for block, file, and object with concurrent support for native NAS, iSCSi, Fibre Channel, and FCoE protocols. The unified configuration includes the following rack-mounted enclosures:

- Disk processor enclosure (holds disk drives) or storage processor enclosure (requires disk drive tray) plus standby power system to deliver block protocols.
- One or more data mover enclosures to deliver file protocols (required for file and unified configurations)
- Control station (required for file and unified configurations)

A robust platform designed to deliver five-9s availability, the EMC VNX Series enables organizations to dynamically grow, share, and cost-effectively manage multiprotocol file systems and multiprotocol block storage access. The EMC VNX Series has been expressly designed to take advantage of the latest innovation in flash drive technology, increasing the storage system's performance and efficiency while reducing the cost per GB.

Finally, Cisco and EMC are collaborating on solutions and services to help build, deploy, and manage IT infrastructures that adapt to changing needs. Industry-leading EMC information infrastructure and intelligent Cisco networking products, including Cisco UCS, will reduce the complexity of data centers.

Together, EMC and Cisco provide comprehensive solutions that can benefit customers now and in the future, including:

- High-performance storage and SANs that reduce TCO
- Disaster recovery to protect data and improve compliance
- Combined computing, storage, networking, and virtualization technologies

Using EMC software creates additional benefits, which can be derived when using products such as:

- **FAST Cache:** Dynamically absorbs unpredicted spikes in system workloads.
- **Fully Automated Storage Tiering for Virtual Pools (FAST VP):** Tiers data from high-performance to high-capacity drives in 1-GB increments, resulting in overall lower costs, regardless of application type or data age.
- **FAST Suite:** Automatically optimizes for the highest system performance and the lowest storage cost simultaneously (includes FAST VP and FAST Cache). For additional information, see [www.emc.com/collateral/hardware/white-papers/h8242-deploying-oracle-vnx-wp.pdf](http://www.emc.com/collateral/hardware/white-papers/h8242-deploying-oracle-vnx-wp.pdf).
- **EMC PowerPath:** Provides automated data path management and load-balancing capabilities for heterogeneous server, network, and storage deployed in physical and virtual environments. For additional information, see [www.emc.com/collateral/software/data-sheet/l751-powerpath-ve-multipathing-ds.pdf](http://www.emc.com/collateral/software/data-sheet/l751-powerpath-ve-multipathing-ds.pdf).
- **EMC Unisphere:** Delivers simplified management via a single management framework for all NAS, SAN, and replication needs. For additional information, see [www.emc.com/collateral/software/data-sheet/h7303-unisphere-ds.pdf](http://www.emc.com/collateral/software/data-sheet/h7303-unisphere-ds.pdf).

For additional information on the EMC VNX Series, visit [www.emc.com/storage/vnx/vnx-series.htm](http://www.emc.com/storage/vnx/vnx-series.htm).

For additional detail regarding the EMC VNX Series Software Suites and the resulting value in performance, protection, and TCO, see [www.emc.com/collateral/software/data-sheet/h8509-vnx-software-suites-ds.pdf](http://www.emc.com/collateral/software/data-sheet/h8509-vnx-software-suites-ds.pdf).

**Figure 3.** EMC VNX Storage Systems




To learn more about the features available in the VNX product line that enable value in your Oracle deployment environment, see [www.emc.com/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf](http://www.emc.com/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf).

## Why Unified Infrastructure?

Today, most organizations have to run parallel network infrastructures for their LANs and SANs, with separate switches and separate HBAs. For Fibre Channel SANs, one or more HBAs must be purchased for each server, which adds considerably to equipment costs. For mission-critical applications (and often others), most organizations provide redundant connectivity, increasing costs even more. With Fibre Channel SANs, separate networks must be operated for the LAN and SAN environments, as shown in Figure 4.

**Figure 4.** Separate LAN and SAN Networks, Requiring a Dedicated Fibre Channel Infrastructure




These separate networks involve added expense due to requirements such as increased number of network interfaces, additional cabling and switch ports, and more complex support needs. Another factor in the increased number of network adapters is server virtualization. Server virtualization, such as that provided by VMware, requires multiple adapters to carry traffic for LAN, SAN, hypervisor management, and virtual infrastructure services.

As the environment grows over time, those expenses become even greater. Currently, the server environment and the server access layer of the network are particular areas of focus. Because of the scale of the server environment, with hundreds or even thousands of servers, small changes can have significant effects. Cisco Nexus 5000 Series Switches are the devices in the Cisco data center switch portfolio that deliver on the promise of a unified fabric for the data center. This unified fabric will have the operational characteristics to concurrently handle LAN, SAN, and server clustering traffic. In addition, the Cisco Nexus 5000 Series has a cut-through design architecture that can deliver a consistent port-to-port latency of 3.2 microseconds, independent of the packet size. The Cisco Nexus 5000 Series' high-performance, line-rate 10 Gigabit Ethernet throughput, combined with its low latency, enable an outstanding efficiency and performance for storage networks. Low latency and low jitter are essential requirements for the high-performance computing applications that the Cisco Nexus 5000 Series consolidates over the same unified fabric.

## Overview of the Solution

At a high level, the solution consists of Cisco UCS B440 M2 blade for the database, while a multinode setup was used for the Web Apps and Concurrent Manager servers using shared APPL\_TOP. Cisco Nexus 5000 Series Switches carried the Ethernet and Fibre Channel traffic to the load balancer and the EMC VNX storage (Figure 5).

**Figure 5.** Solution Overview



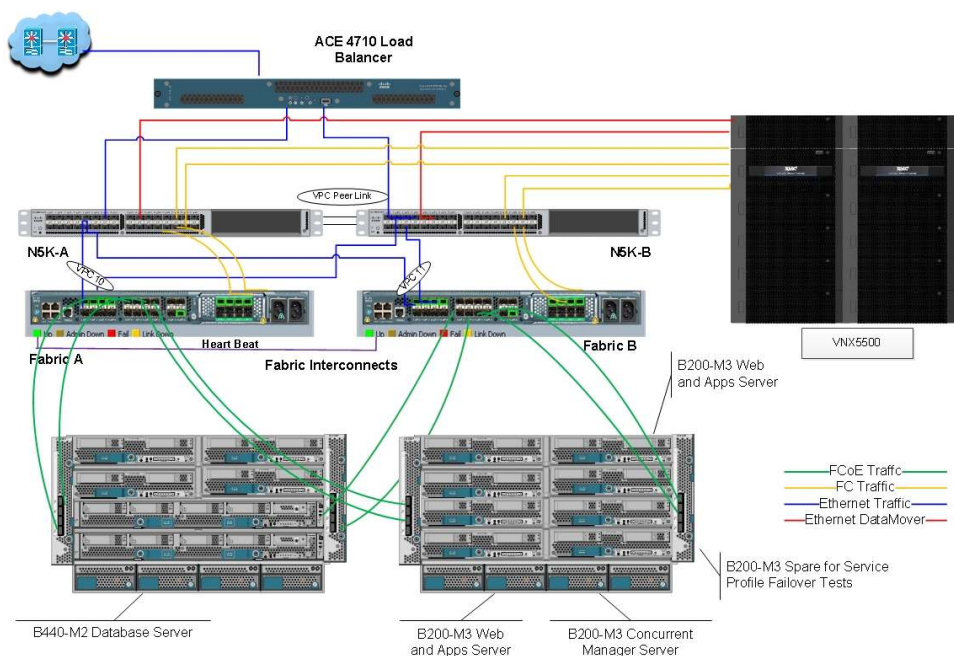
Oracle E-Business Suite is installed as a multi-node environment. This consists of three application-tier nodes and a node for the database tier. A spare node was kept ready to simulate failovers and using service profiles. The application tier is installed using Oracle shared APPL\_TOP. The shared APPL\_TOP file systems are created on NFS shared folders on the VNX unified storage array, and shared folders are accessed over IP. The load-balancing feature is provided by Cisco Application Control Engine (ACE) load balancer. The database is Oracle 11gR2 implemented on Fibre Channel SAN to Automatic Storage Management (ASM) disk groups created on the storage. All files, such as data files, control files, and online redo log files, are deployed on the ASM disk groups. This is a typical configuration that can be deployed in a customer's environment, and the use cases, best practices, and setup recommendations are described in subsequent sections of this document. In addition to the above, Enterprise Manager (EM) Grid Control 12c was installed on one of the blades to monitor the database and application tier nodes. R12 Apps Plugin for EM 12c was installed on this blade, and the application and database nodes were discovered in EM Grid Control.



## Oracle E-Business Suite R12 on Cisco UCS 2.1 and EMC VNX 5500

Figure 6 depicts the deployment architecture.

**Figure 6.** Oracle E-Business Suite Deployment Architecture



## Configuring Cisco Unified Computing System for Oracle E-Business Suite

Detailed information about configuring the Cisco UCS system is available at [www.cisco.com/en/US/products/ps10281/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html).

It is beyond the scope of this document to cover all of these. However, we have included as much information as possible.

### Configuring Fabric Interconnects

Cisco UCS 6120XP 20-Port Fabric Interconnects are configured for redundancy. This provides resiliency in case of failures.

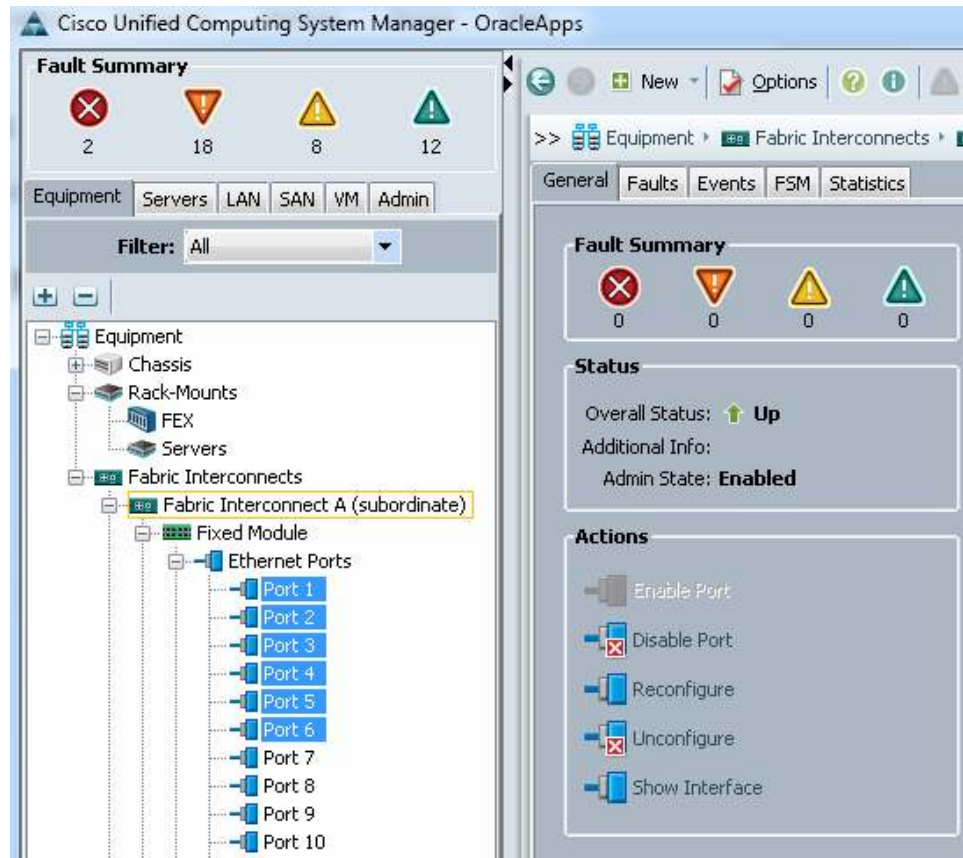
The first step is to establish connectivity between the blades and fabric interconnects. As shown in Figure 6 four 8-Gb links were used from each chassis. Each of the I/O modules is connected to either of the fabric interconnects, providing failover capabilities. These take care of both IOM and fabric failures. Configurations may vary, depending on the distribution of the database and middle-tier servers across the chassis. For simplicity, the database server B440 M2 and EM servers were hosted on one chassis, while the Web Apps and Concurrent Manager servers were on another.



## Configuring Server Ports

Figure 7 shows a screen shot with the configuration of the server ports.

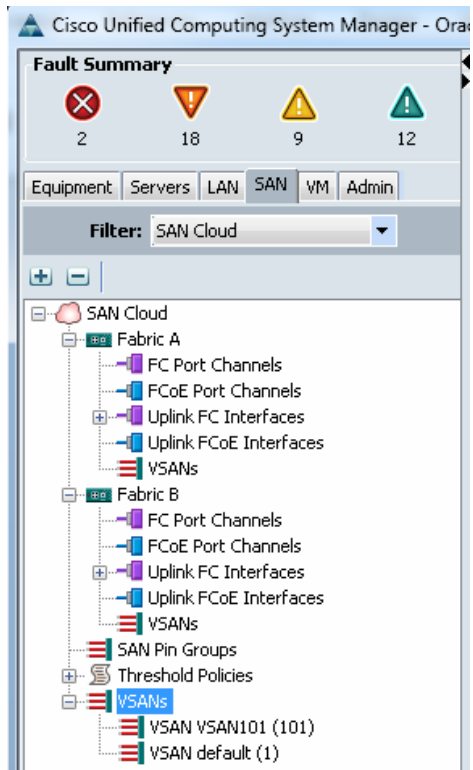
**Figure 7.** How the Server Ports Are Configured



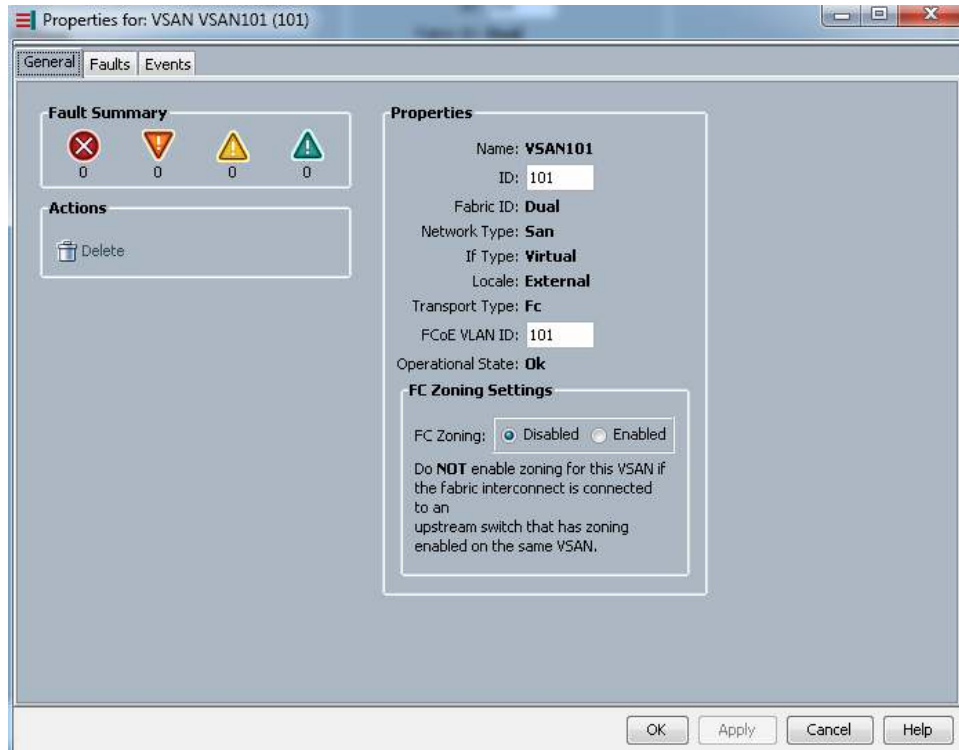
## Configuring SAN and LAN on UCS Manager

On the SAN tab, create and configure the VSANs to be used for database as shown in Figures 8, 9, and 10. In our setup, we used VSAN 101.


**Figure 8.** The SAN Tab



**Figure 9.** Properties for VSAN 101



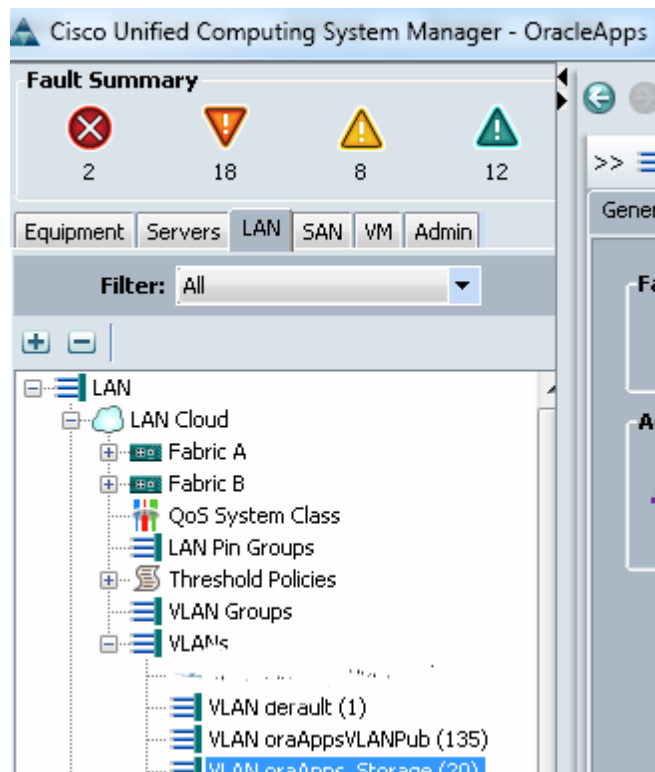
**Figure 10.** Displaying the VSANs



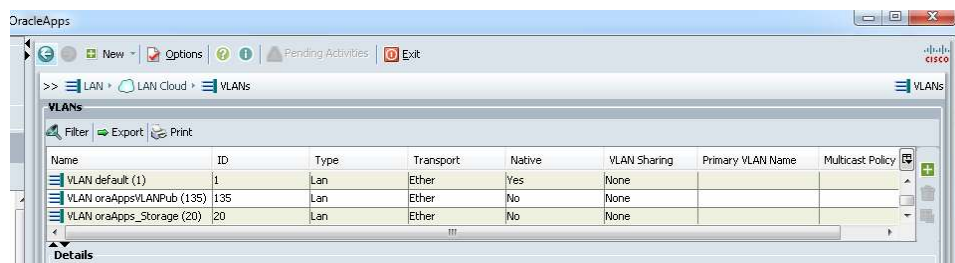
### Configure the LAN

On the LAN tab, create VLANs that will be used later for virtual NICs, for public network communication and also with storage (Figures 11 and 12). You can also set up MAC address pools for assignment to vNICs. For this setup, we used VLAN 135 for public interfaces and VLAN 20 for E-Business Suite storage traffic. It is also very important that you create both VLANs as global across both fabric interconnects. That way, VLAN identity is maintained across the fabric interconnects in case of failover.

**Figure 11.** The LAN Tab



**Figure 12.** Displaying the VLANs

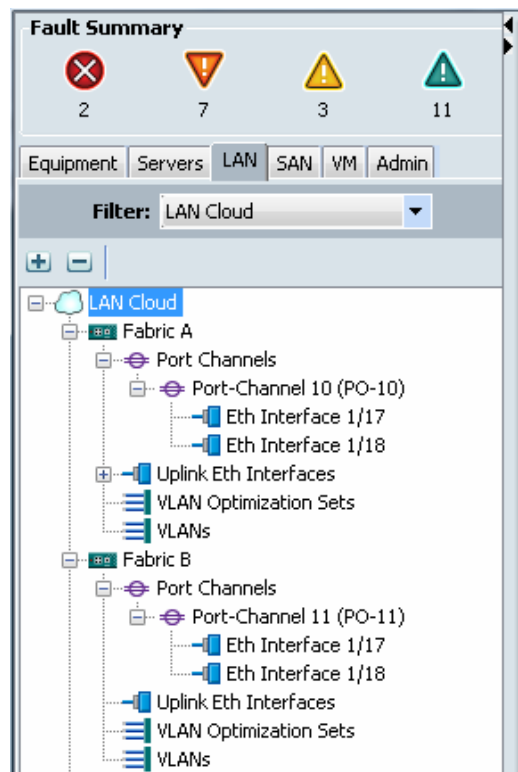


## Configure Ethernet Port Channels

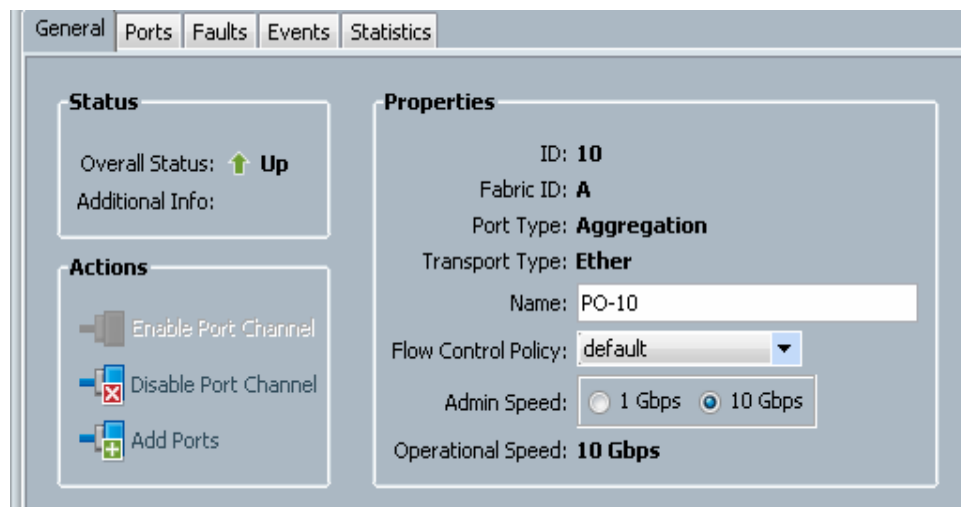
To configure the port channels, log in to Cisco UCS Manager, display the LAN tab, and filter on LAN Cloud. Select Fabric A, right-click Port Channels, and create a port channel. In the current setup, ports 17 and 18 on Fabric A were selected to be configured as port channel 10.

Similarly, ports 17 and 18 on Fabric B with port channel 11 (Figures 13 through 16).

**Figure 13.** Configuring the Port Channel for Fabric A



**Figure 14.** Port Channel 10 Details



**Figure 15.** Port Channels on Fabric A

Port Channels				
Name	Fabric ID	If Type	If Role	Transport
Port-Channel 10 (PO-10)	A	Aggregation	Network	Ether
Eth Interface 1/17	A	Physical	Network	Ether
Eth Interface 1/18	A	Physical	Network	Ether

**Figure 16.** Port Channels on Fabric B

Port Channels				
Name	Fabric ID	If Type	If Role	Transport
Port-Channel 11 (PO-11)	B	Aggregation	Network	Ether
Eth Interface 1/17	B	Physical	Network	Ether
Eth Interface 1/18	B	Physical	Network	Ether

The next step is to set up a virtual port channel (vPC) on the Cisco Nexus 5000 Series. This is covered in a later section.


### Preparatory Steps Before Creating Service Templates

First create the UUID, IP, MAC, worldwide node name (WWNN) and worldwide port name (WWPN) pools and keep them handy in case they are not pre-created. If pre-created, make sure you have enough of them free and unallocated.

#### UUID Pool

Click the Servers tab, and filter on Pools. Expand the UUID suffix pools and create a new pool (Figure 17).


**Figure 17.** Creating a New UUID Suffix Pool



#### IP and MAC Pools

Click the LAN tab, filter on Pools, and create IP and MAC pools (Figure 18).

**Figure 18.** Creating a MAC Pool

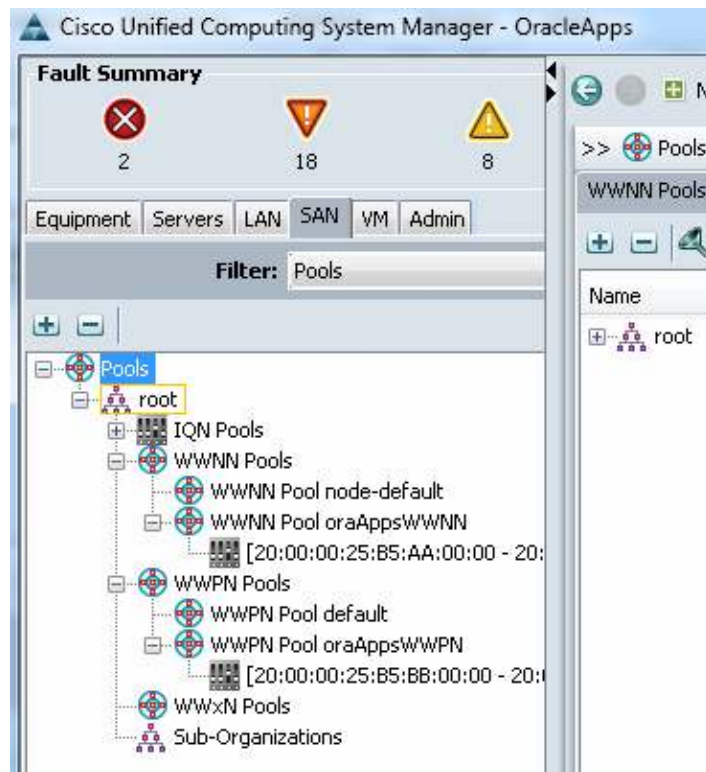


The IP pools will be used for console management, while the MAC addresses are for the vNICs being carved out later.

#### WWNN and WWPN pools

Click the SAN tab, filter on Pools, and create the pools as shown in Figure 19.

**Figure 19.** WWNN and WWPN Pools

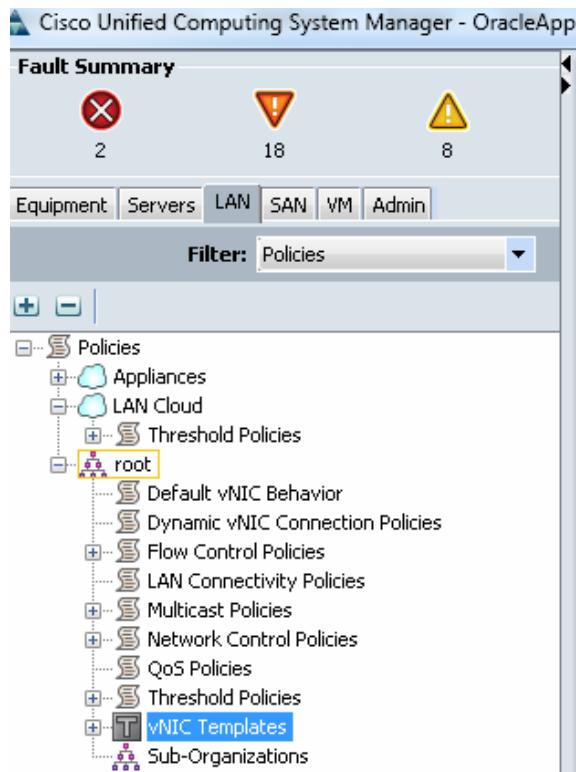


#### Configure vNIC Templates

Click the LAN tab, filter on Policies, and select vNIC Templates (Figure 20). Two templates are created, one for the public network and one for the storage network.



**Figure 20.** Configuring the vNIC Templates

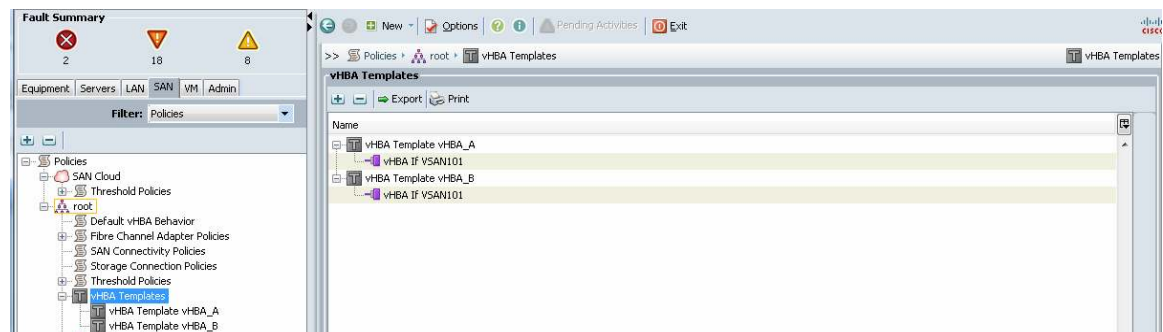


The vNIC template for storage was configured with 9000 MTU.

### Create HBA Templates

Click the SAN tab, filter on Policies, right-click vHBA Templates, and create a template (Figures 21 through 23). The HBA templates will be used only for the Oracle Database server, in our case Cisco UCS B440 M2.

**Figure 21.** Creating the HBA Templates



**Figure 22.** Properties for the vHBA\_A Template

The screenshot shows the Cisco UCS Manager web interface. The breadcrumb navigation is: Policies > root > vHBA Templates > vHBA Template vHBA\_A. The 'vHBA Template vHBA\_A' tab is selected. The 'General' tab is active, showing the 'Properties' section. On the left, there is an 'Actions' panel with a 'Delete' button. The 'Properties' section contains the following fields:

- Name: vHBA\_A
- Description: (empty text box)
- Fabric ID: ☒ A ☐ B
- VSAN: VSAN101 (dropdown menu)
- Target: Adapter
- Template Type: ☒ Initial Template ☐ Updating Template
- Max Data Field Size: 2048
- Policies section:
  - WWPN Pool: oraAppsWWPN (dropdown menu)
  - QoS Policy: <not set> (dropdown menu)
  - Pin Group: <not set> (dropdown menu)
  - Stats Threshold Policy: default (dropdown menu)

**Figure 23.** Properties for the vHBA\_B Template

The screenshot shows the Cisco UCS Manager web interface. The breadcrumb navigation is: Policies > root > vHBA Templates > vHBA Template vHBA\_B. The 'vHBA Template vHBA\_B' tab is selected. The 'General' tab is active, showing the 'Properties' section. On the left, there is an 'Actions' panel with a 'Delete' button. The 'Properties' section contains the following fields:

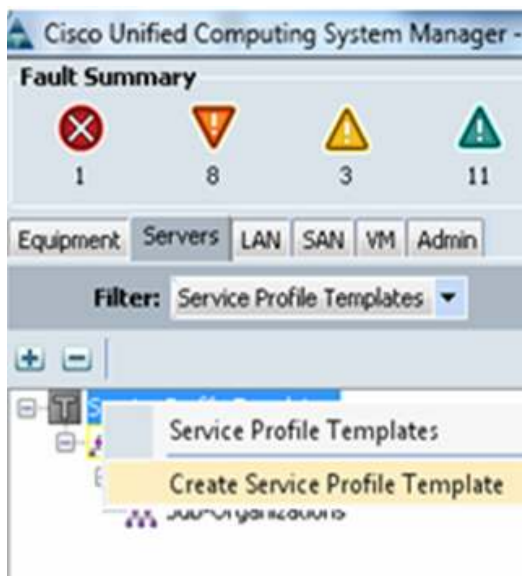
- Name: vHBA\_B
- Description: (empty text box)
- Fabric ID: ☐ A ☒ B
- VSAN: VSAN101 (dropdown menu)
- Target: Adapter
- Template Type: ☒ Initial Template ☐ Updating Template
- Max Data Field Size: 2048
- Policies section:
  - WWPN Pool: (empty dropdown menu)
  - QoS Policy: (empty dropdown menu)
  - Pin Group: <not set> (dropdown menu)
  - Stats Threshold Policy: (empty dropdown menu)

Once the above preparatory steps are complete, you can create a service template from which the service profiles can be easily created.

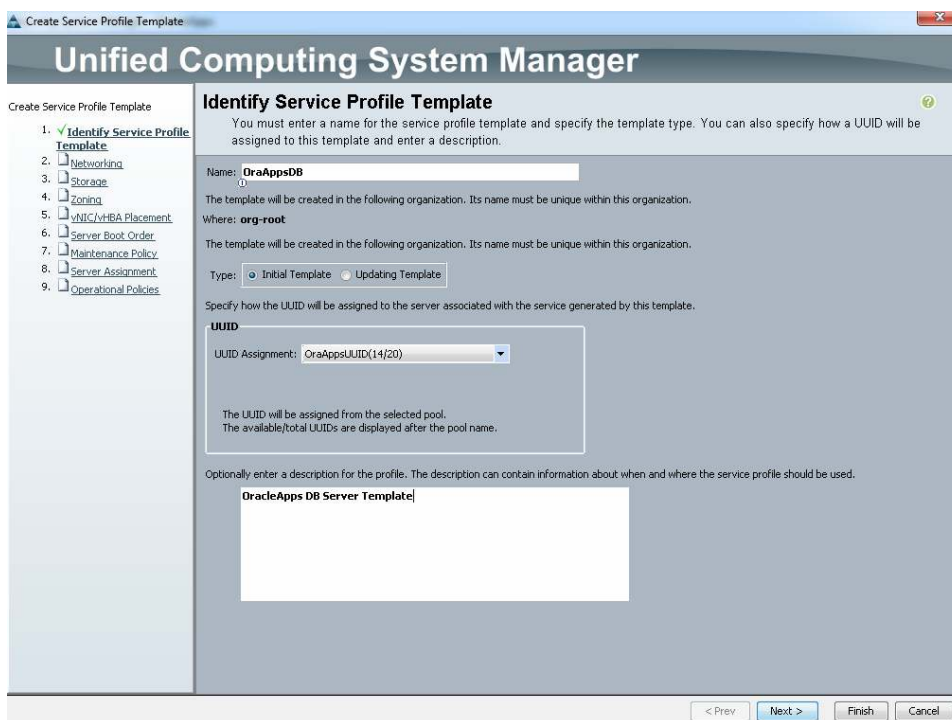
### Create Service Profile Template: Database

Create a service profile template before forking service profiles to be allocated to the servers later. Click the Servers tab in Cisco UCS Manager, filter on Service Profile Templates, and select Create Service Profile Template (Figure 24).

**Figure 24.** Creating a Service Profile Template



**Figure 25.** Identifying the Database Service Profile Template



Enter the name, pick up the default UUID created earlier, and move on to the next screen (Figure 25).

On the Networking page, create vNICs for the public network and associate them with the VLAN policies created earlier.

Select Expert mode, and click Add in the section that specifies one or more vNICs that the server should use to connect to the LAN (Figures 26 through 28).

**Figure 26.** Adding a vNIC

The screenshot shows the 'Unified Computing System Manager' window with the 'Networking' tab selected. On the left, a sidebar lists steps for 'Create Service Profile Template', with 'Networking' highlighted. The main area is titled 'Networking' and includes the instruction 'Optionally specify LAN configuration information.' Below this, there's a section for 'Dynamic vNIC Connection Policy' with a dropdown menu set to 'Select a Policy to use (no Dynamic vNIC Policy by default)' and a '+ Create Dynamic vNIC Connection Policy' button. A radio button group asks 'How would you like to configure LAN connectivity?' with options: Simple, Expert (selected), No vNICs, and Use Connectivity Policy. Below the radio buttons, a text prompt says 'Click Add to specify one or more vNICs that the server should use to connect to the LAN.' This is followed by a table with columns: Name, MAC Address, Fabric ID, and Native VLAN. Below the table are buttons for Delete, Add, and Modify. Another text prompt says 'Click Add to specify one or more iSCSI vNICs that the server should use.' followed by a table with columns: Name, Overlay vNIC Name, iSCSI Adapter Policy, and MAC Address. Below this table are buttons for Add, Delete, and Modify. At the bottom right are navigation buttons: < Prev, Next >, Finish, and Cancel.

**Figure 27.** The Create vNIC Page

The screenshot shows the 'Create vNIC' window. It has a title bar 'Create vNIC' and a question mark icon. The main area is titled 'Create vNIC'. It contains a 'Name' field with 'eth0' entered. Below it is a 'Use vNIC Template' checkbox which is checked, with a '+ Create vNIC Template' button. A 'vNIC Template' dropdown menu is set to 'oraApps\_A'. Below this is an 'Adapter Performance Profile' section with an 'Adapter Policy' dropdown set to 'Linux' and a '+ Create Ethernet Adapter Policy' button. At the bottom right are navigation buttons: < Prev, Next >, Finish, and Cancel.

**Figure 28.** After Adding the vNIC

Create Service Profile Template

## Unified Computing System Manager

Create Service Profile Template

1. ☒ Identify Service Profile Template
2. ☒ **Networking**
3. ☐ Storage
4. ☐ Zoning
5. ☐ vNIC/vHBA Placement
6. ☐ Server Boot Order
7. ☐ Maintenance Policy
8. ☐ Server Assignment
9. ☐ Operational Policies

### Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

How would you like to configure LAN connectivity? ☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
vNIC eth0	Derived	derived	

Click **Add** to specify one or more iSCSI vNICs that the server should use.

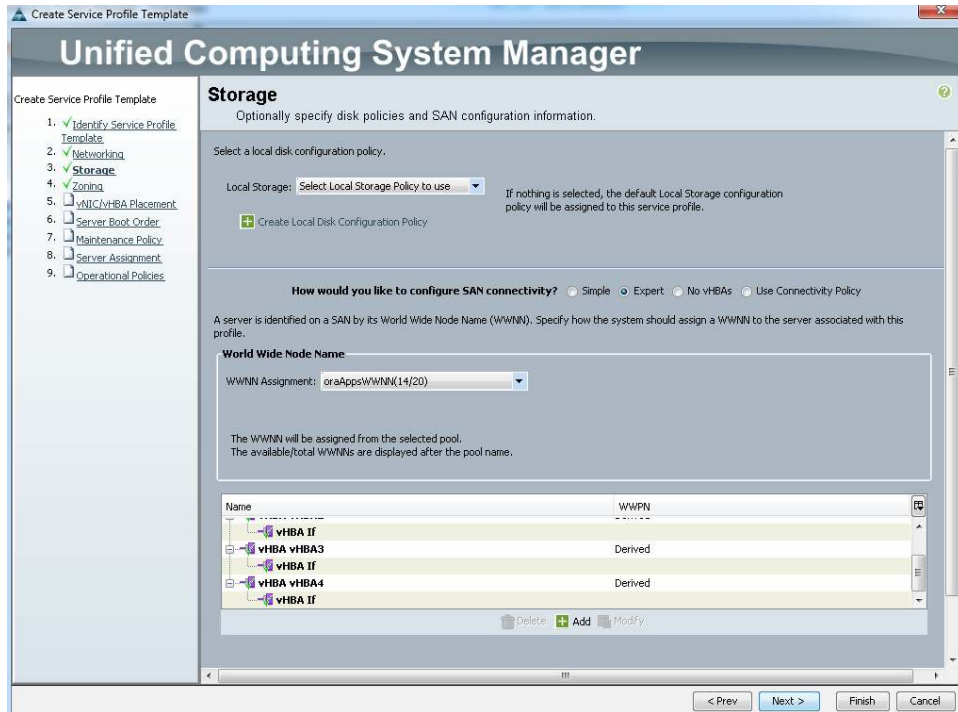
Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
------	-------------------	----------------------	-------------

< Prev Next > Finish Cancel


On the Storage page, as you did for the vNICs, select Expert mode in the adapter, choose the WWNN pool created earlier, and click the Add button to create vHBAs (Figures 29 and 30). We selected the following four vHBAs:

- Create vHBA1 using template vHBA\_A.
- Create vHBA2 using template vHBA\_B.
- Create vHBA3 using template vHBA\_A.
- Create vHBA4 using template vHBA\_B.

**Figure 29.** Adding a vHBA



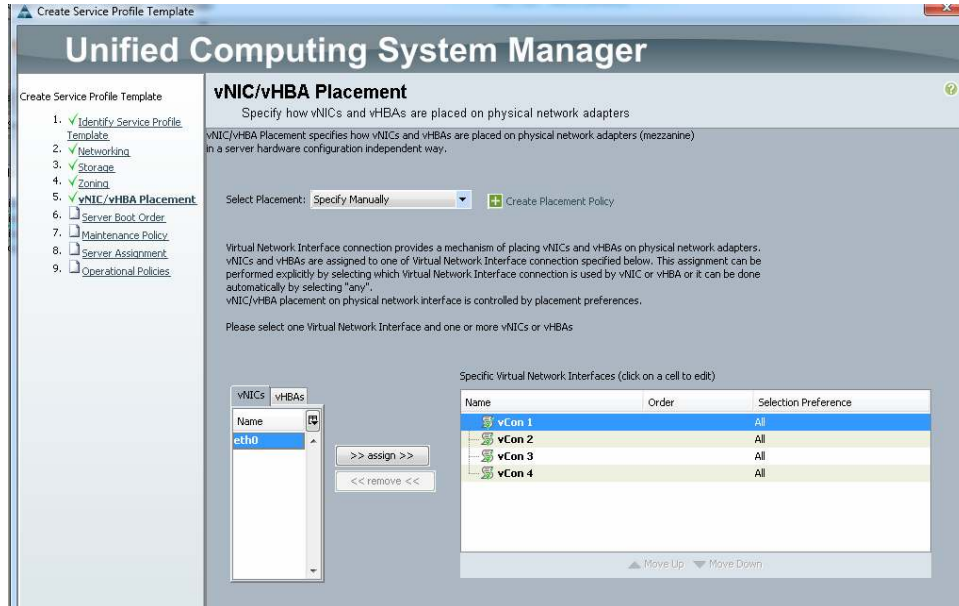
**Figure 30.** The Create vHBA Page



Skip the Zoning section and go to vNIC/vHBA Placement.

While you can leave this to the system defaults, you can also specify the vNIC and vHBA placement manually, as shown in Figure 31.

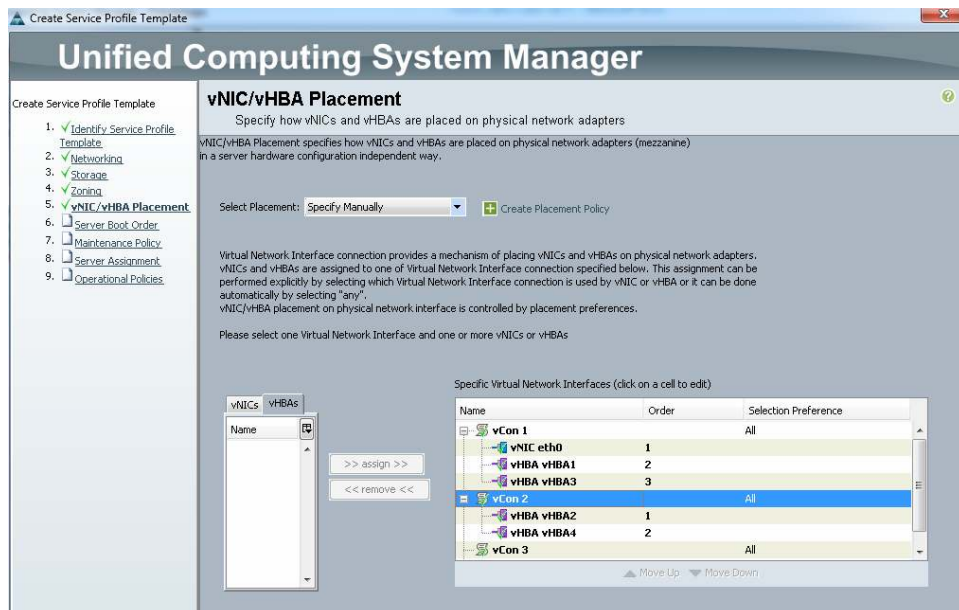
**Figure 31.** Specifying vNIC and vHBA Placement



Highlight eth0 and click assign to vCon1.

Highlight the vHBAs one after another and assign them respectively to vCons, as shown in Figure 32.

**Figure 32.** Assigning the vHBAs to vCons





Here we allocated vNIC1, vHBA1, and vHBA3 to the first Cisco UCS VIC 1280, and allocated vNIC2, vHBA2, and vHBA4 to the second VIC 1280 of the database server Cisco UCS B440 M2.

Server Boot Policy:

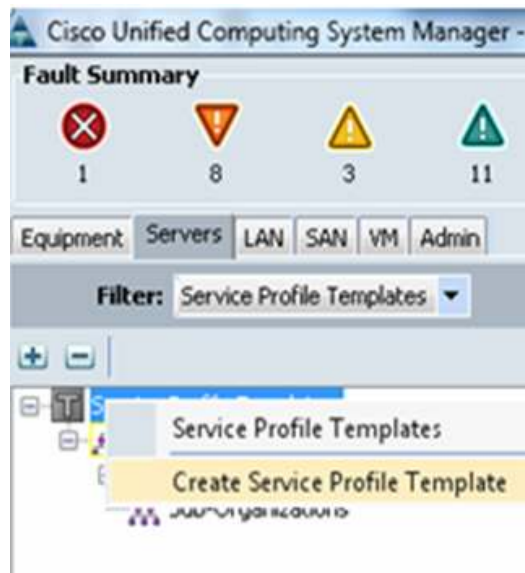
Leave this at the default, as the initiators may vary from one server to the other.

We left the rest of the maintenance and assignment policies at the default settings in the test bed. But they may be selected and may vary from site to site, depending on your workloads, best practices, and policies.

### Create Service Profile Template: Apps

Similar to the procedure in the previous section, create a service profile template for the Oracle application tier (Figures 33 through 38). The service profiles for the application tiers can be forked from this template.

**Figure 33.** Creating a Service Profile Template



**Figure 34.** Identifying the Apps Service Profile Template

The screenshot shows the 'Create Service Profile Template' window in the Unified Computing System Manager. The left sidebar lists steps: 1. Identify Service Profile Template (selected), 2. Networking, 3. Storage, 4. Zoning, 5. vNIC/vHBA Placement, 6. Server Boot Order, 7. Maintenance Policy, 8. Server Assignment, and 9. Operational Policies. The main area is titled 'Identify Service Profile Template' and contains the following fields and instructions:

- Name:** oraAppsMT
- Where:** org-root
- Type:** Initial Template (selected), Updating Template
- UUID Assignment:** OraAppsUUID(14/20)
- Description:** Template for Oracle Apps Server

Buttons at the bottom include '< Prev', 'Next >', 'Finish', and 'Cancel'.

**Figure 35.** Creating a vNIC in the Networking Section

The screenshot shows the 'Create vNIC' window in the Unified Computing System Manager. The main area contains the following fields and instructions:

- Name:** eth0
- Use vNIC Template:** [checked]
- vNIC Template:** oraApps\_A
- Adapter Performance Profile:** Adapter Policy: Linux

Buttons at the bottom include '+ Create vNIC Template' and '+ Create Ethernet Adapter Policy'.

**Figure 36.** After Adding the vNIC

The screenshot shows the 'Networking' configuration page in the Unified Computing System Manager. The left sidebar lists the steps of the 'Create Service Profile Template' wizard, with 'Networking' selected. The main area is titled 'Networking' and includes the instruction 'Optionally specify LAN configuration information.' Below this, there is a dropdown for 'Dynamic vNIC Connection Policy' and a '+ Create Dynamic vNIC Connection Policy' button. A section titled 'How would you like to configure LAN connectivity?' has four radio buttons: 'Simple' (selected), 'Expert', 'No vNICs', and 'Use Connectivity Policy'. Below this, a table lists vNICs. The first row is 'vNIC eth0' with 'Derived' for both MAC Address and Fabric ID. Below the table are 'Delete', 'Add', and 'Modify' buttons. Another section titled 'Click Add to specify one or more iSCSI vNICs that the server should use.' contains an empty table with columns 'Name', 'Overlay vNIC Name', 'iSCSI Adapter Policy', and 'MAC Address', and 'Add', 'Delete', and 'Modify' buttons. At the bottom are '< Prev', 'Next >', 'Finish', and 'Cancel' buttons.

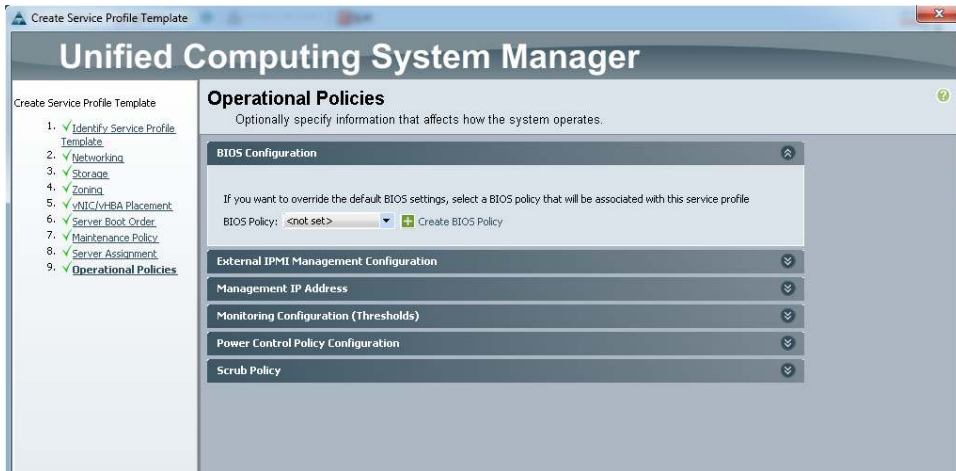
Name	MAC Address	Fabric ID	Native VLAN
vNIC eth0	Derived	derived	

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
------	-------------------	----------------------	-------------

**Figure 37.** Specifying No vHBAs

The screenshot shows the 'Storage' configuration page in the Unified Computing System Manager. The left sidebar lists the steps of the 'Create Service Profile Template' wizard, with 'Storage' selected. The main area is titled 'Storage' and includes the instruction 'Optionally specify disk policies and SAN configuration information.' Below this, there is a dropdown for 'Local Storage' and a '+ Create Local Disk Configuration Policy' button. A section titled 'How would you like to configure SAN connectivity?' has four radio buttons: 'Simple', 'Expert', 'No vHBAs' (selected), and 'Use Connectivity Policy'. Below this, a message states: 'This server associated with this service profile will not be connected to a storage area network.' At the bottom are '< Prev', 'Next >', 'Finish', and 'Cancel' buttons.

**Figure 38.** Viewing the Operational Policies



The rest of the entries can be left at the default settings for this template.

### Create Service Profiles from Service Profile Templates

Click the Servers tab, right-click on the root, and select Create Service Profile from Template (Figure 39).

**Figure 39.** Creating a Service Profile from a Template



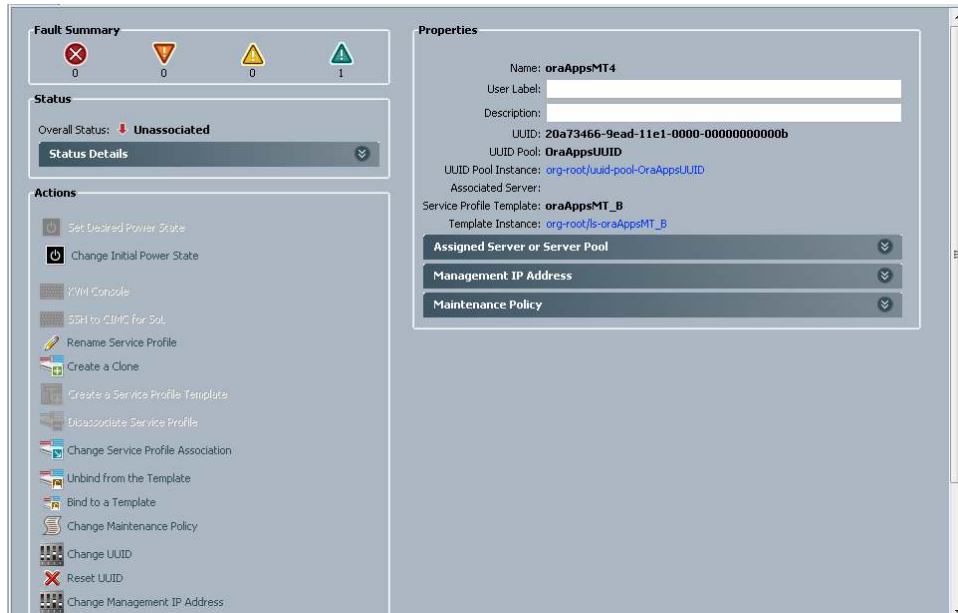
We created three templates in the system, two for the Apps/Web server and one for Concurrent Management server.

### Associating Service Profile with the Servers

To associate a service profile with a server, perform the following steps.

On the Servers tab, select the desired service profile and select Change Service Profile Association (Figure 40).

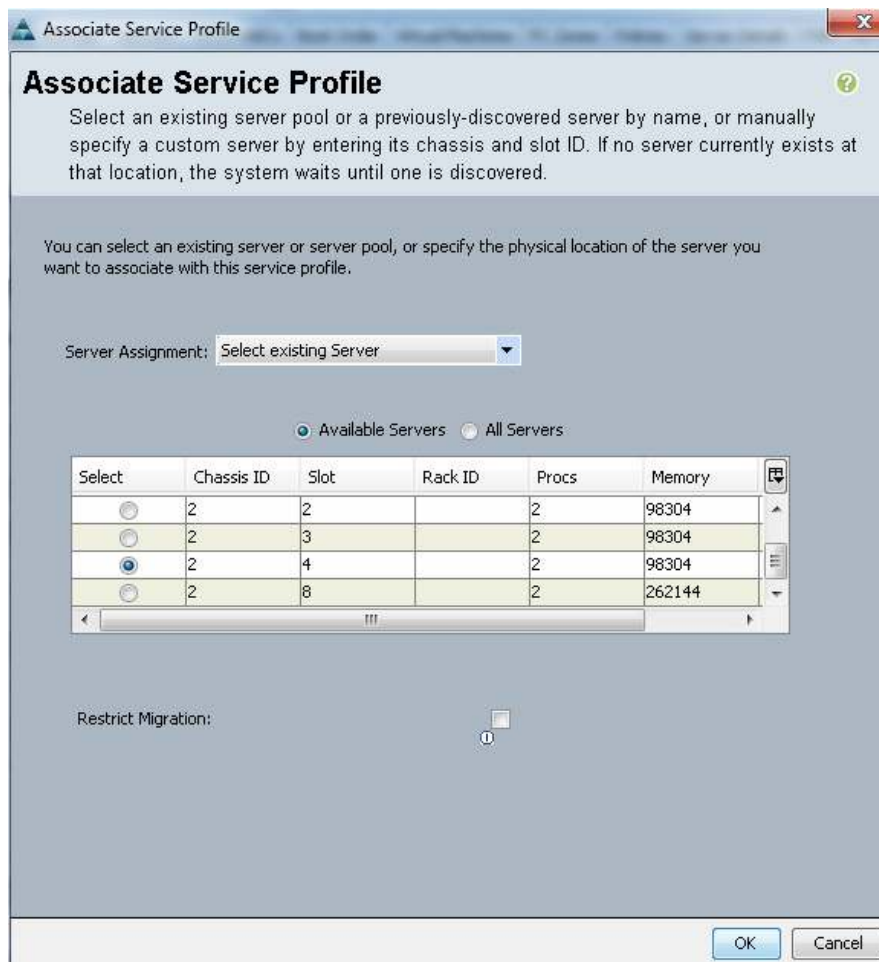
**Figure 40.** A Service Profile



The service profile is unassociated as of now and can be assigned to a server in the pool.

Click Change Service Profile Association, and from the Server Assignment drop-down, select the existing server that you would like to assign. Click OK (Figure 41).

**Figure 41.** Associating the Service Profile with a Server



## Setting Up EMC VNX Storage

This section provides a general overview of the storage configuration for the database and Apps layout. However, it is beyond the scope of this document to provide full details about host connectivity and logical unit numbers (LUNs) in RAID configuration and Data Mover connectivity. For more information about Oracle Database best practices for deployments with EMC VNX storage, refer to [www.emc.com/oracle](http://www.emc.com/oracle).

The following are some generic recommendations for EMC VNX storage configuration. Turn off the read and write caches for flash drive-based LUNs. In most situations, it is better to turn off both the read and write caches on all the LUNs that reside on flash drives, for the following reasons:

The flash drives are extremely fast. When the read cache is enabled for the LUNs residing on them, the read cache lookup for each read request adds more overhead, compared to SAS drives. This scenario occurs in an application profile that is not expected to get many read cache hits at any rate. It is generally much faster to directly read the block from the flash drives.

Typically, the storage array is also shared by several other applications, along with the database. In some situations, the write cache may become fully saturated, placing the flash drives in a force-flush situation. This adds unnecessary latency. This typically occurs when storage deploys mixed drives and consists of slower near line SAS drives. Therefore, it is better in these situations to write the block directly to the flash drives than to the write cache of the storage system.

Tables 1 and 2 illustrate the distribution of LUNs carved out from a VNX 5500 for the setup.

**Table 1.** LUNs in the Apps Database

Purpose	Apps Database Data and Temp Files	Redo Log Files for Apps Database
Disk type	SAS	SAS
RAID type	RAID 5 RAID groups	RAID 1/0 RAID groups
SAS disks	50	16
Flash disks	0	0
Total LUNs	10	8
LUN size	262 GB	20 GB

**Table 2.** Boot LUNs

Purpose	Boot LUNs
Disk type	SAS
RAID type	RAID 5
SAS disks	5
Flash disks	0
Total LUNs	4 boot LUNs
LUN size	Boot LUNs: 100 GB

## Hardware Storage Processors Configuration

A total of four ports were used from storage processors and were equally distributed between the storage processors A and B and were connected to the respective Cisco Nexus 5000s (Table 3).

**Table 3.** Distribution of Ports Between Storage Processors

Processor	Slot/Port	WWPN
SP A	A2P0	50:06:01:60:47:20:2c:af
	A2P2	50:06:01:62:47:20:2c:af
	A3P0	50:06:01:64:47:20:2c:af
	A3P2	50:06:01:66:47:20:2c:af
SP B	B2P0	50:06:01:68:47:20:2c:af
	B2P2	50:06:01:6a:47:20:2c:af
	B3P0	50:06:01:6c:47:20:2c:af
	B3P2	50:06:01:6e:47:20:2c:af



## Configure SAN Zoning on Nexus 5548UP Switches

Two Cisco Nexus 5548UP Switches were configured.

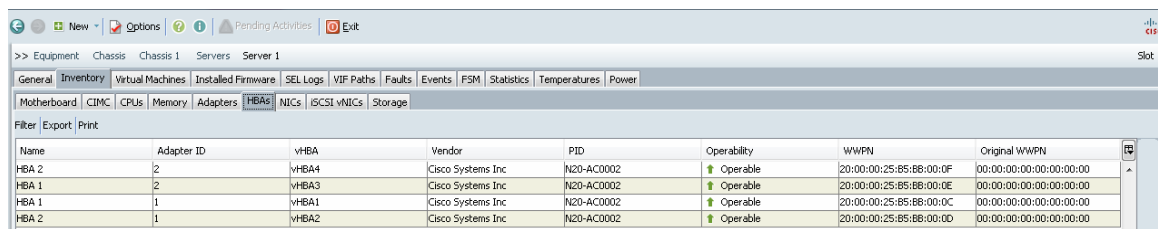
### Fibre Channel Zoning

Before going into the zoning details, decide how many paths are needed for each LUN and extract the WWPN numbers for each of the HBAs.

To see the WWPNs for each of the HBAs, log in to Cisco UCS Manager.

Click Equipment, Chassis, Servers, and the desired server. Click the Inventory tab and the HBAs subtab, as shown in Figure 42.

**Figure 42.** Displaying the WWPNs for the HBAs



Name	Adapter ID	vHBA	Vendor	PID	Operability	WWPN	Original WWPN
HBA 2	2	vHBA4	Cisco Systems Inc	N20-AC0002	Operable	20:00:00:25:b5:bb:00:0f	00:00:00:00:00:00:00:00
HBA 1	2	vHBA3	Cisco Systems Inc	N20-AC0002	Operable	20:00:00:25:b5:bb:00:0e	00:00:00:00:00:00:00:00
HBA 1	1	vHBA1	Cisco Systems Inc	N20-AC0002	Operable	20:00:00:25:b5:bb:00:0c	00:00:00:00:00:00:00:00
HBA 2	1	vHBA2	Cisco Systems Inc	N20-AC0002	Operable	20:00:00:25:b5:bb:00:0d	00:00:00:00:00:00:00:00

Figure 42 shows the WWPNs for all four HBAs for server 1. In the current setup, it was decided to have a total of two paths from each fabric and Nexus 5548UP to the storage.

Therefore, the zoning for Server1, HBA1 can be set up as follows:

```
* fcid 0x380007 [pwwn 20:00:00:25:b5:bb:00:0c] [OraAppsDB_vHBA1]
* fcid 0x3806ef [pwwn 50:06:01:60:3d:e0:21:f6] [OraAppsSPA2]
* fcid 0x3807ef [pwwn 50:06:01:68:3d:e0:21:f6] [OraAppsSPB2]
* fcid 0x380007 [pwwn 20:00:00:25:b5:bb:00:0d] [OraAppsDB_vHBA2]
* fcid 0x3804ef [pwwn 50:06:01:61:3d:e0:21:f6] [OraAppsSPA3]
* fcid 0x3805ef [pwwn 50:06:01:69:3d:e0:21:f6] [OraAppsSPB3]
```

Effectively, the HBAs are distributed to both Nexus 5548UP switches.

The WWPNs from storage are distributed between both storage processors, providing distribution and redundancy in case of failures.

Table 4 shows an example for the Database server.

**Table 4.** Example of WWPNs for Database Server

Switch A		
	zone OraAppsDB_1	[pwwn 20:00:00:25:b5:bb:00:0c]
		[pwwn 50:06:01:60:3d:e0:21:f6]
		[pwwn 50:06:01:68:3d:e0:21:f6]
	zone OraAppsDB_3	[pwwn 20:00:00:25:b5:bb:00:0e]
		[pwwn 50:06:01:60:3d:e0:21:f6]

<b>Switch B</b>		[pwwn 50:06:01:68:3d:e0:21:f6]
	zone OraAppsDB_2	[pwwn 20:00:00:25:b5:bb:00:0d]
		[pwwn 50:06:01:61:3d:e0:21:f6]
		[pwwn 50:06:01:69:3d:e0:21:f6]
	zone OraAppsDB_4	[pwwn 20:00:00:25:b5:00:00:2f]
		[pwwn 50:06:01:61:3d:e0:21:f6]
		[pwwn 50:06:01:69:3d:e0:21:f6]

Log in through SSH and issue the following commands.

Here is an example for one zone on one Nexus 5548UP:

```
conf term
zoneset name OraAppsZoneset vsan 101
zone name OraAppsDB_1
member pwwn 50:06:01:60:3d:e0:21:f6
member pwwn 50:06:01:68:3d:e0:21:f6
member pwwn 20:00:00:25:b5:bb:00:0c
exit
```

Add other zones:

```
exit
zoneset activate name OraAppsZoneset vsan 101
copy running-config startup-config
```

Repeat the above for all the HBAs. A detailed list of zones added during setup is provided in Appendix B.

Set Up VLAN and VSANs on Both Nexus 5548UP Switches

```
conf term
vlan 135
name Oracle_Traffic
exit
vlan 20
name Oracle_NFS_Traffic
exit
vsan database
vsan 101
exit
```

## Setting Up vPC on the Nexus 5548UP Switches

**Figure 43.** vPC Setup on the Nexus 5548UP Switches

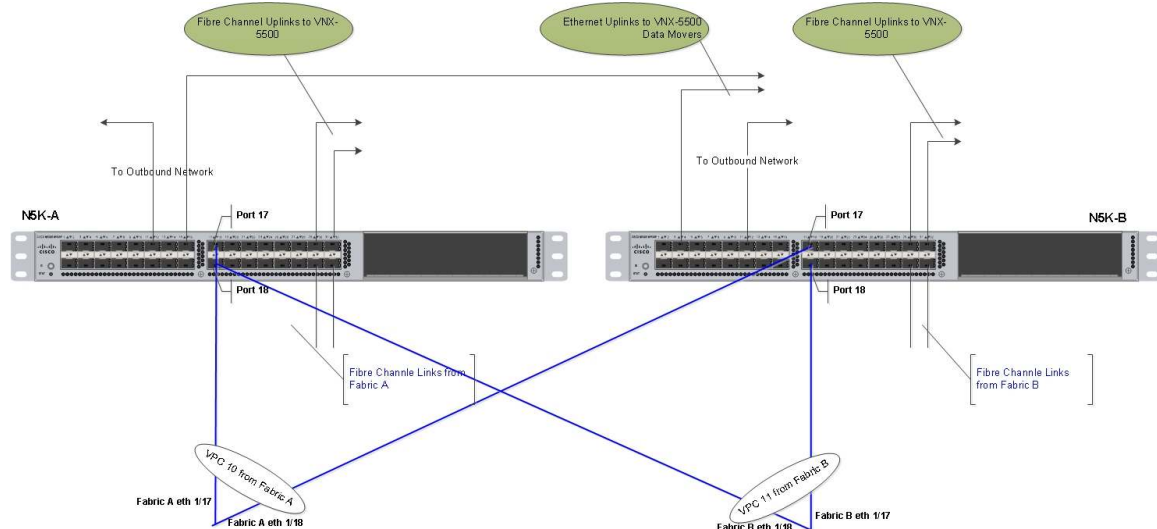


Figure 43 diagrammatically represents how the Nexus 5548UP switches are connected to the northbound switches and storage while connected to the underlying Cisco UCS fabrics. The Nexus 5548UP switches form a core group in controlling SAN zoning.

In the figure, port 17 on both Nexus 5548UP switches receives traffic from UCS Fabric A, which has port-channel 10 defined. Similarly port 18 on both switches receives traffic from UCS Fabric B, which has port-channel 11 configured.

```
Log in to switch A as admin.
conf term
feature vpc

vpc domain 1
peer-keepalive destination <IP Address of peer-N5K>
exit

interface port-channel 1
switchport mode trunk
vpc peer-link
switchport trunk allowed vlan 1,20,135
spanning-tree port type network
exit
```

```

interface port-channel 10
  switchport mode trunk
  vpc 10
  switchport trunk allowed vlan 1,20,135
  spanning-tree port type edge trunk
exit

interface port-channel 11
  switchport mode trunk
  vpc 11
  switchport trunk allowed vlan 1,20,135
  spanning-tree port type edge trunk
exit

interface eth 1/17
  switchport mode trunk
  switchport trunk allowed vlan 1,20,135
  channel-group 10 mode active
  no shut

interface eth 1/18
  switchport mode trunk
  switchport trunk allowed vlan 1,20,135
  channel-group 11 mode active
  no shut

copy running-config startup-config

```

Repeat the above on both Nexus 5548UP switches.

The vPC status should show the following for successful configuration.

```

vPC Peer-link status
-----
id Port   Status  Active vlans
--  ---
1   Po1    up      1,20,135

vPC status
-----
Id Port   Status  Consistency  Reason  Active vlans
--  ---
10 Po10    up      success    success  1,20,135
11 Po11    up      success    success  1,20,135
show interface port-channel 10-11 brief
-----
Port-channel  VLAN Type  Mode  Status Reason  Speed Protocol

```

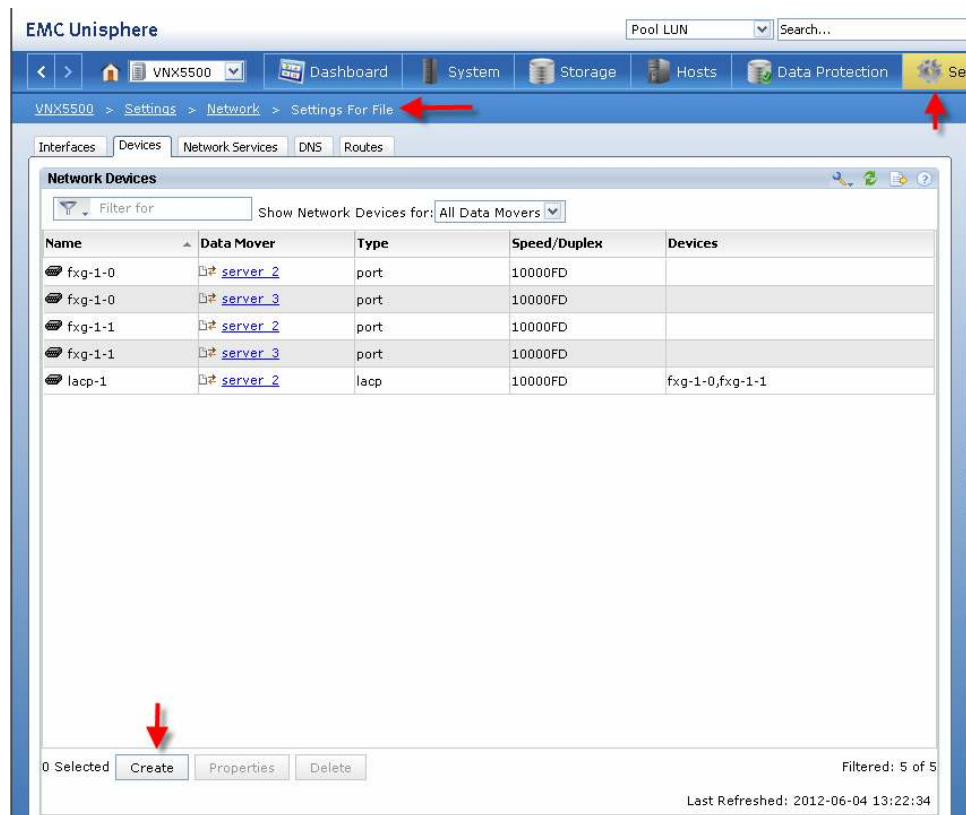
## Interface

Po10	1	eth	trunk	up	none	a-10G(D)	lacp
Po11	1	eth	trunk	up	none	a-10G(D)	lacp

## Configure the VNX 5500 for NFS

Open the Unisphere Console and navigate to Settings>Network>Settings for File>Devices tab (Figure 44).

**Figure 44.** The Devices Tab



Click Create.

**Figure 45.** Creating a Device

**Data Mover:** server\_2

**Type:**

- ☐ Ethernet Channel
- ☒ Link Aggregation
- ☐ Fail Safe Network

**Device Name:** larp-1

**10/100 ports:** None available

**Gigabit ports:** None available

**10/100/1000 ports:** None available

**10 Gigabit ports:** ☒ fxg-1-0 ☒ fxg-1-1

**Speed/Duplex:**

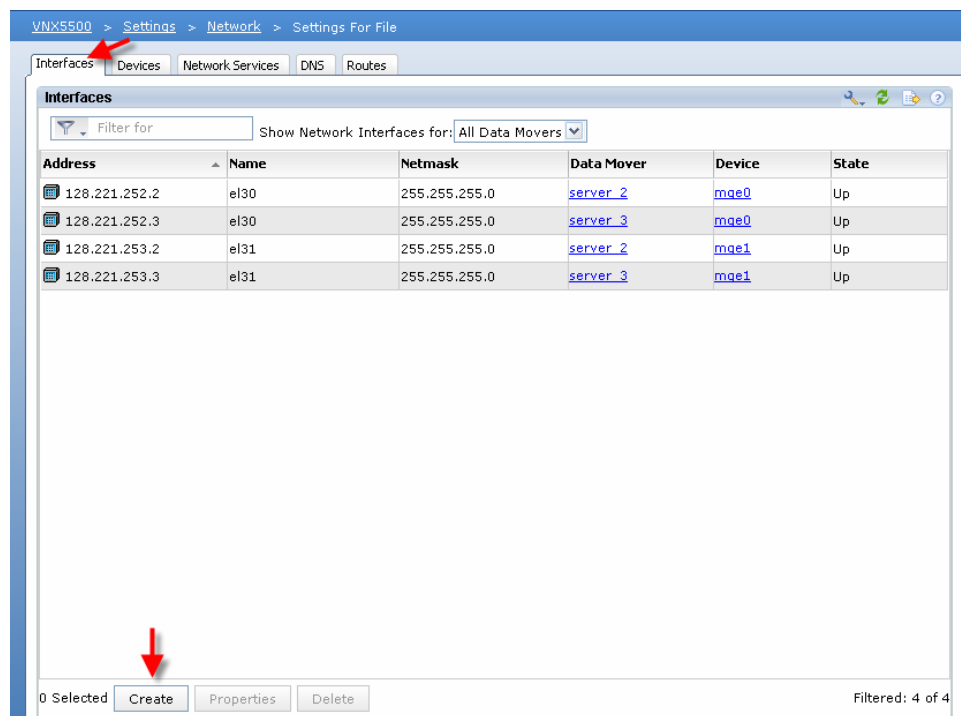
**Buttons:** OK Apply Cancel Help

**Status Bar:** Internet | Protected Mode: Off 125%

Choose the Data Mover, select Link Aggregation as the Type, provide a unique Device Name, select both 10 GE ports, then click OK (Figure 45).

Click the Interfaces tab, then click Create (Figure 46).

**Figure 46.** The Interfaces Tab



Specify a Data Mover, Device Name, IP Address, Name (it should NOT be the same as the device name), Netmask, MTU, and VLAN\* (Figure 47).

**Figure 47.** Creating an Interface

Data Mover:

Device Name:

Address:

Name:

Netmask:

Broadcast Address:

MTU:

VLAN ID:

[OK](#) [Apply](#) [Cancel](#) [Help](#)

\*Populate VLAN only if the ports that the fxx ports/lacp are connected to are trunked. If they are not trunked, leave the VLAN field blank.

Confirm that the Link Aggregation Control Protocol (LACP) configuration is working by running the following command at the VNX command line:

```
server_sysconfig server_2 -v -i lacp1
```

**Figure 48.** Confirming the LACP Configuration

```
[nasadmin@VNX5500 bin]$ server_sysconfig server_2 -v -i lacp-1
server_2 :
*** Trunk lacp-1: Link is Up ***
*** Trunk lacp-1: Timeout is Short ***
*** Trunk lacp-1: Statistical Load Balancing is IP ***
Device      Local Grp  Remote Grp  Link  LACP  Duplex Speed
-----
fxg-1-0     10003      2944       Up   Up    Full   10000 Mbs
fxg-1-1     10003      2944       Up   Up    Full   10000 Mbs
[nasadmin@VNX5500 bin]$
```

Link and LACP should be up on both 10 GE interfaces (Figure 48).

If there are problems with LACP configuration or communications, use the following commands to troubleshoot:

```
server_ifconfig server_2 -a
server_sysconfig server_2 -pci
server_sysconfig server_2 -v -i lacp-1
watch server_sysconfig server_2 -virtual -info lacp1
server_ping server_2 -interface lacp-1-a 10.29.166.62
```

### Create a Storage Pool for NFS (RAID 0/1)

Navigate to Storage>Storage Configuration>Storage Pools, and click Create (Figure 49).



**Figure 49. The Pools Tab**

The screenshot displays the EMC Unisphere web interface for storage management. The top navigation bar includes tabs for Dashboard, System, Storage, Hosts, and Data Protection. The 'Storage' tab is active, showing the 'Storage Configuration' > 'Storage Pools' path. The 'Pools' section contains a table of storage pools and a 'Create' button highlighted by a red arrow.

Name	State	RAID Type	Drive Type	User Capa...	Free Capa...	Allocated (...)	%Consum...	Subscribed...	%Subscri
Pool 2	Ready	RAID1/0	SAS	13319.965	24.047	13295.918		13,245.82	99.4
Pool 1	Ready	RAID6	SAS	3205.248	2151.193	1054.055		1,052.051	32.4
Pool 0	Ready	RAID5	SAS	4286.355	535.043	3751.312		3,661.137	85.4

Below the table, there are buttons for 'Create', 'Delete', 'Properties', and 'Expand'. The 'Create' button is highlighted with a red arrow. The 'Details' section below shows 'Pool LUNs' and 'Disks' tabs, with a 'Usage' dropdown set to 'ALL User LUNs'. The 'Pool LUNs' table is currently empty.

Select the RAID Type (10 for NFS) (Figure 50).

**Figure 50.** Creating a Storage Pool

**VNX5500 - Create Storage Pool**

General Advanced

**Storage Pool Parameters**

Storage Pool Type: ☒ Pool ☐ RAID Group

☒ Scheduled Auto-Tiering

Storage Pool ID: 3

Storage Pool Name: Pool 3

RAID Type: RAID1/0

Number of Disks:

**Extreme Performance**

SSD Disks: 0

**Performance**

SAS Disks: 10

**Distribution**

Performance : 5368.076 GB (100.00%)

**Disks**

☒ Automatic ☐ Use Power Saving Eligible Disks

☐ Manual  Total Raw Capacity: 5368.0...

Disk	Capacity	Drive Type	Model	State
Bus 0 Enclosure 1 Disk 8	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 3 Disk 5	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 3 Disk 4	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 3 Disk 3	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 3 Disk 2	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 3 Disk 1	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 3 Disk 0	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 0 Disk 14	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 0 Disk 13	536.808...	SAS	STE600...	Unbo...
Bus 0 Enclosure 0 Disk 12	536.808...	SAS	STE600...	Unbo...

Apply Cancel Help

Create LUNs for Use with NFS

Right-click the new pool that was just created, and click Create LUNs (Figure 51).

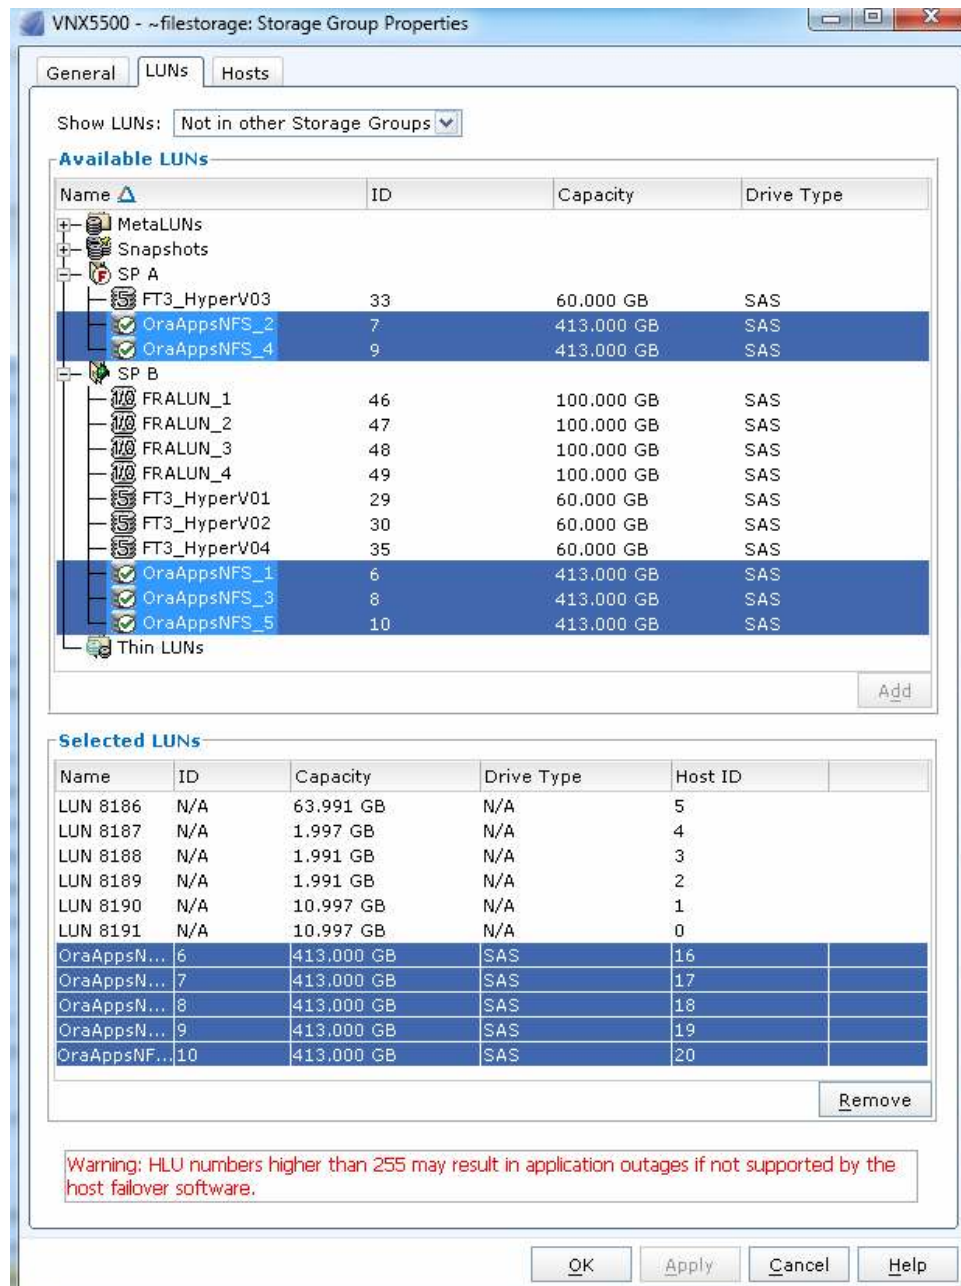
**Figure 51.** Creating a LUN

Set the User Capacity to MAX and the Number of LUNS equal to the number of disks used in the pool, then click Apply.

#### Assign Host IDs to the LUNs

Navigate to Hosts>Storage. Right-click ~filestorage, and select the LUNs that you just created (Figure 52).

**Figure 52.** Assigning Host IDs



Create the LUNs, assign host IDs, and distribute them to both of the storage processors.

Then, from the VNX command line, verify that the volumes were created:

```
nas_disk -list (shows dvols)
nas_diskmark -mark -all
nas_pool -list (you should see the pool created earlier)
```

## Create Apps File Systems

Navigate to Storage>Storage Configuration>File Systems. Click Create.

Configure the Name, Capacity, and Data Mover, then click OK (Figure 53).

**Figure 53.** Configuring the File System

Create from: ☒ Storage Pool ☐ Meta Volume

File System Name:

Storage Pool:

Storage Capacity:

Auto Extend Enabled: ☐

Thin Enabled: ☐

Slice Volumes: ☒

Deduplication Enabled: ☐

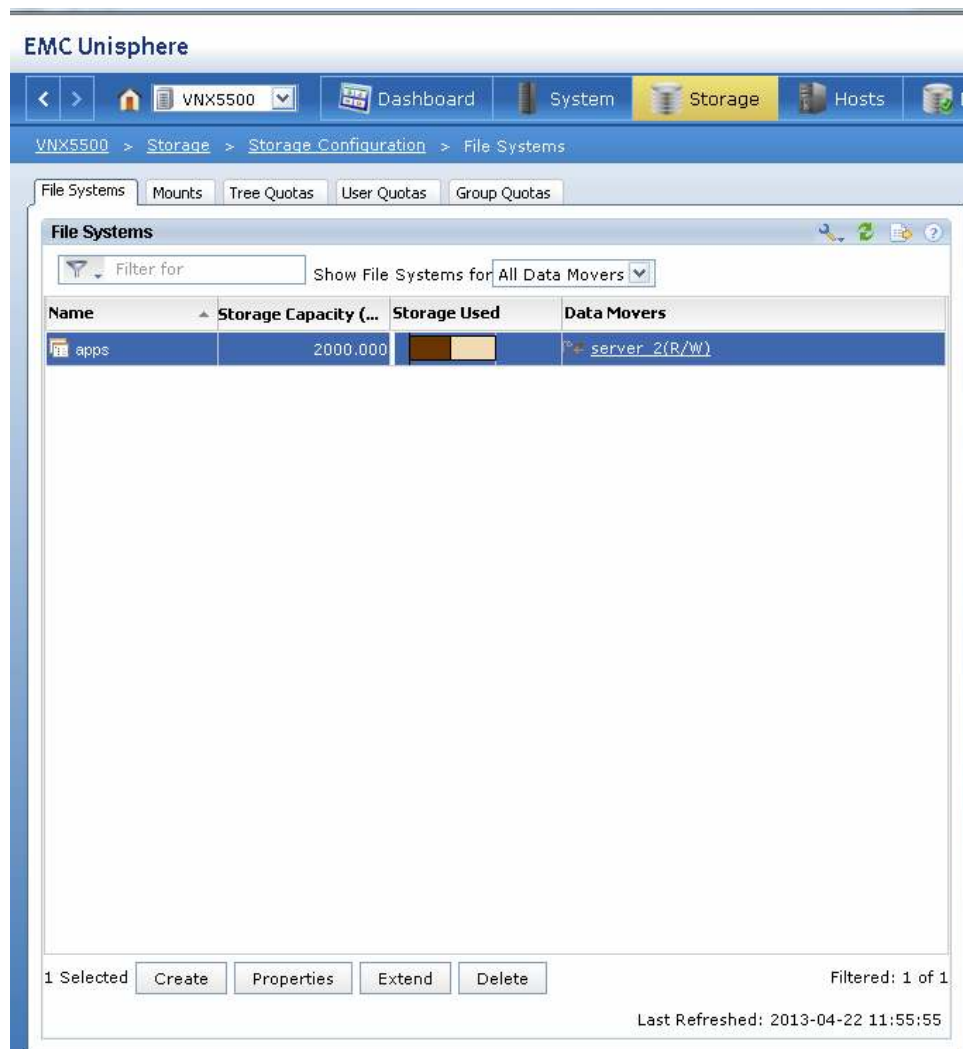
Data Mover (R/W):

Mount Point: ☒ Default ☐ Custom

OK Apply Cancel Help

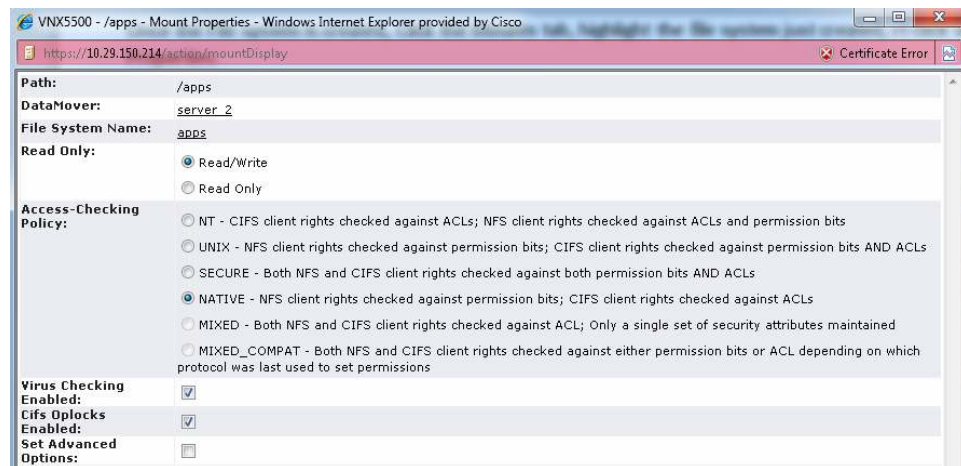
When you click the Properties tab, should appear as shown in Figure 54.

**Figure 54.** Displaying the File System Properties



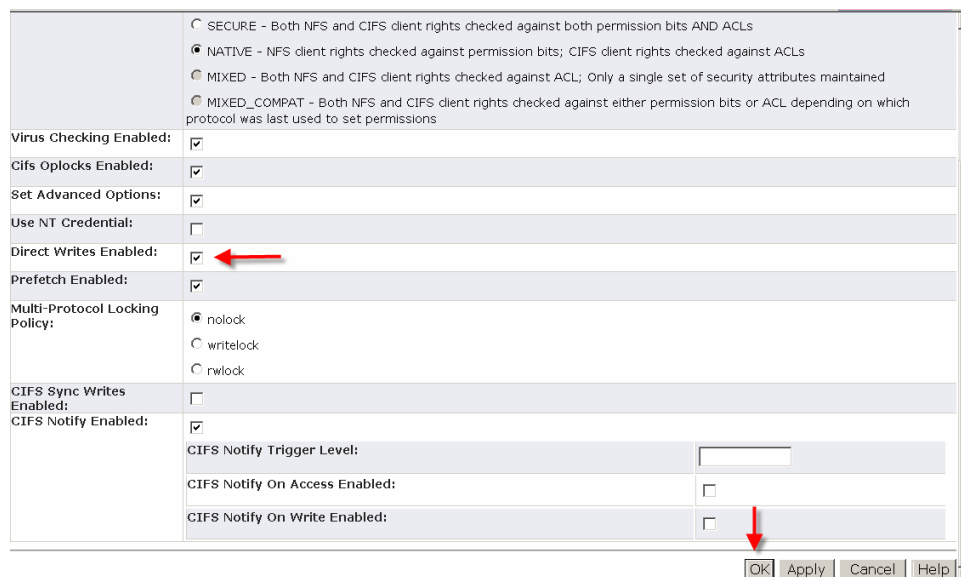
Once the file system is created, click the Mounts tab, highlight the file system just created, right-click, and choose Properties (Figure 55).

**Figure 55.** Choosing File System Properties



Click Set Advanced Options (Figure 56).

**Figure 56.** Enabling Direct Writes



Check Direct Writes Enabled, then click OK.

## Create NFS Exports

Map NFS Export and assign to it to a storage group.

Navigate to Storage>Shared Folders>NFS. Click Create (Figure 57).

**Figure 57.** Creating NFS Exports

SECURE - Both NFS and CIFS client rights checked against both permission bits AND ACLs  
NATIVE - NFS client rights checked against permission bits; CIFS client rights checked against ACLs  
MIXED - Both NFS and CIFS client rights checked against ACL; Only a single set of security attributes maintained  
MIXED\_COMPAT - Both NFS and CIFS client rights checked against either permission bits or ACL depending on which protocol was last used to set permissions

Virus Checking Enabled: ☒

Cifs Oplocks Enabled: ☒

Set Advanced Options: ☒

Use NT Credential: ☐

Direct Writes Enabled: ☒

Prefetch Enabled: ☒

Multi-Protocol Locking Policy:  
nolock  
writelock  
rlock

CIFS Sync Writes Enabled: ☐

CIFS Notify Enabled: ☒

CIFS Notify Trigger Level:

CIFS Notify On Access Enabled: ☐

CIFS Notify On Write Enabled: ☐

OK Apply Cancel Help

The IP addresses here are the host's NFS vNIC IP addresses, as listed below. These are covered as part of the installation of the hosts, discussed .

```
eth1 Link encap:Ethernet HWaddr 00:25:B5:AA:00:0D
      inet addr:192.168.20.21 Bcast:192.168.20.255 Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST MTU:9000 Metric:1
```

This should complete the EMC VNX setup, and the mounts are ready to be exported. For the sake of simplicity, only one mount point is carved out of the system. However, depending on policies and standards, you can create multiple mount points, instead of /apps alone, such as one for software, one for logs, etc.

## Install Operating System, Additional RPMs, and Prepare the System for the Oracle Database and Application Tier Servers

Oracle Linux 5.8 was used; however, it was the Red Hat Compatible kernel that was used in the setup.

### Prepare to Install the Boot LUNs

You may have to make a few changes to the storage and the Cisco Nexus 5548UP Switches before installing the kernel with boot LUNs, configured with EMC PowerPath. More detailed steps are provided in the EMC PowerPath for Linux version 5.7 Installation and Administration guide.

Cisco UCS Manager allows you to define boot policies for each server that can be configured to present the boot LUN.

### Configure Storage Boot LUNs

- Make sure that the boot LUN for the server is presented to the host first from the storage side. Four storage groups were defined: one for the database server, and one each for the middle-tier servers. Also make a note of the host ID (preferably 0, as this is the first LUN presented to the host) before going further.



## SAN Zoning Changes on the Nexus Switches for boot

Change the zoning policy on the Nexus 5548UP switches so that only one path is available during the boot time. Disable the zones on, say, switch B and enable them only on switch A. Also make sure that only one path is available before installation. Once the installation is complete and PowerPath is completely set up, this can be reverted back to its full paths. As an example, for server 1 (Database server), only one zone is made available before installation, as follows:

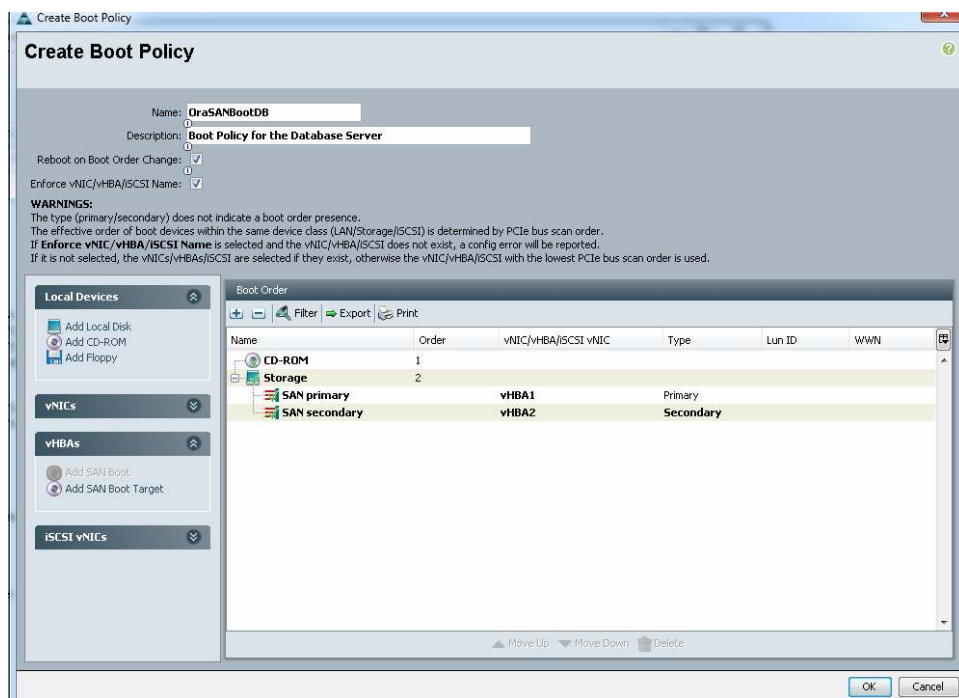
```
* fcid 0x380007 [pwwn 20:00:00:25:b5:bb:00:0c] [OraAppsDB_vHBA1]
* fcid 0x3806ef [pwwn 50:06:01:60:3d:e0:21:f6] [OraAppsSPA2]
```

## Configure Boot Policies on UCS Servers

- Define a boot policy for each server.

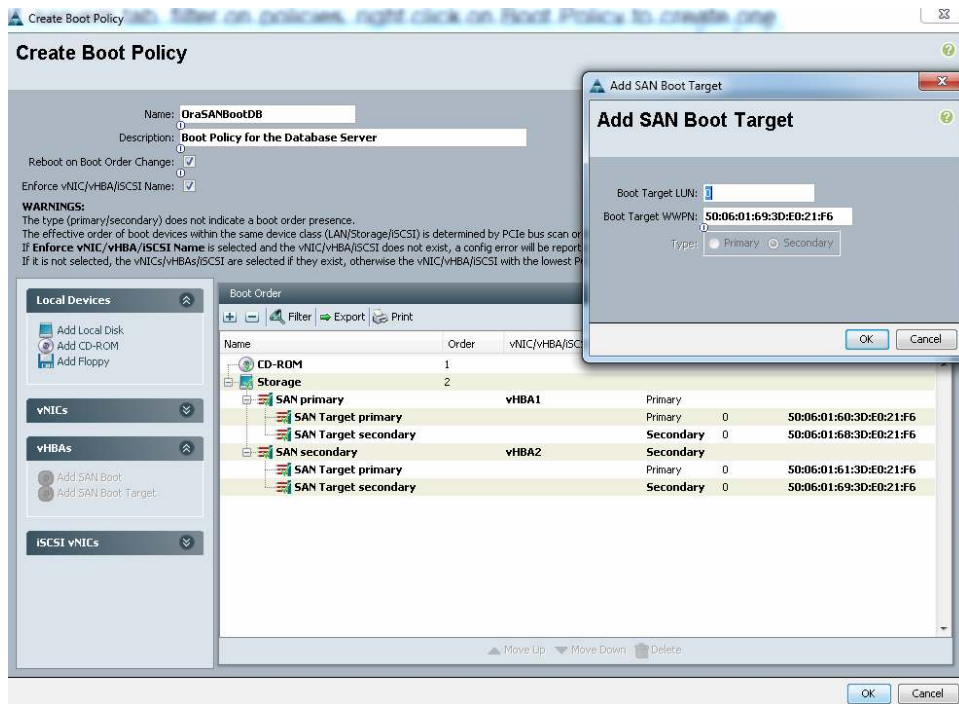
Log in to UCS Manager, display the Servers tab, filter on Policies, and right-click Boot Policy to create one (Figure 58).

**Figure 58.** Creating a Boot Policy



For both SAN primary and SAN secondary, add the SAN boot targets as shown in Figure 59. The Boot target LUN ID should match the host ID from VNX, as mentioned earlier.

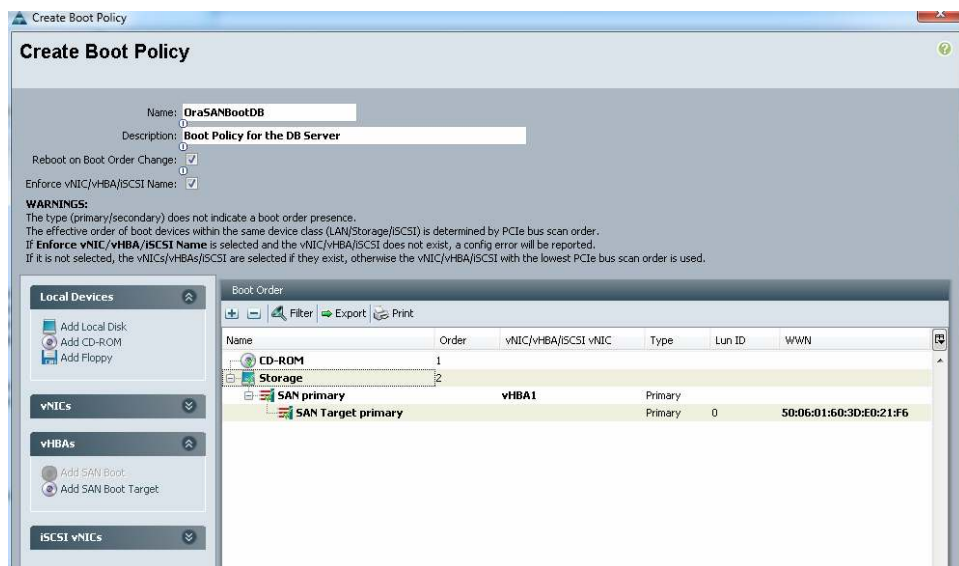
**Figure 59.** Adding SAN Boot Targets



Click OK to create the boot policy for the server. This has to be repeated for all the servers.

To be doubly sure that you do not have multiple paths during installation, temporarily disable all the paths and enable only one, as shown in Figure 60.

**Figure 60.** Disabling All But One Path



This is only for the installation. Once Linux is installed and PowerPath is fully configured, this must be reverted to the earlier settings, with both SAN primary and SAN secondary, each with a SAN boot target.

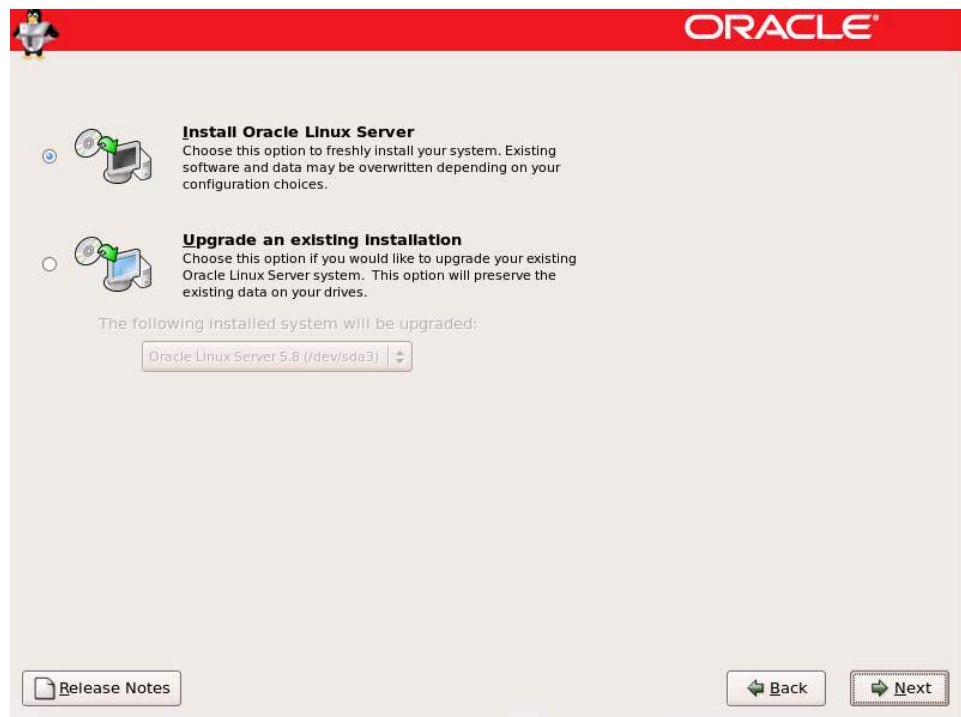
### Install Oracle Linux 5.8 from Image

Download the Oracle Linux 5.8 images from <https://edelivery.oracle.com/linux>, or as appropriate. Mount the image and launch the installer.

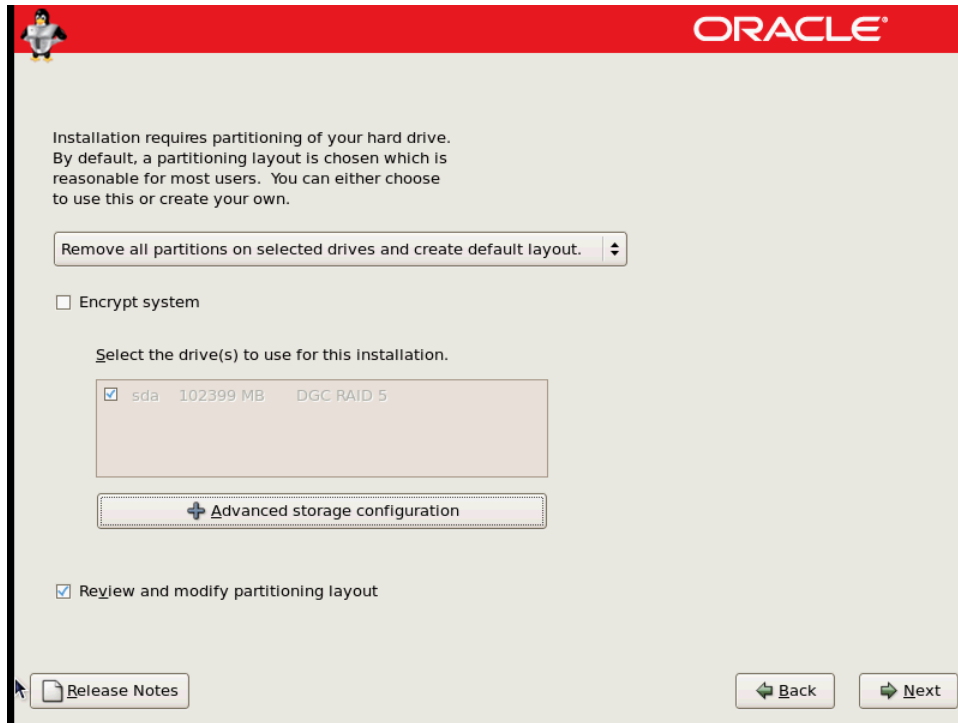
Launch the KVM console for the desired server, click on virtual media, add the image and reset the server. When the server comes up, it launches the Oracle Linux Installer (Figure 61).

Only a few of the screen shots for the installation are provided here.

**Figure 61.** The Oracle Linux Installer



**Figure 62.** Specifying a Partitioning Layout



The image shows a screenshot of the Oracle Linux installation partitioning screen. The window has a red header bar with the Oracle logo on the right and a small penguin icon on the left. The main content area is light gray. At the top, there is a text block explaining that installation requires partitioning and that a default layout is chosen. Below this is a dropdown menu with the text "Remove all partitions on selected drives and create default layout." and a small up/down arrow. Underneath is a checkbox labeled "Encrypt system" which is currently unchecked. Then, there is a section titled "Select the drive(s) to use for this installation." which contains a table with one row: a checked checkbox, the text "sda", "102399 MB", and "DGC RAID 5". Below the table is a button with a plus icon and the text "Advanced storage configuration". At the bottom of the main area is a checkbox labeled "Review and modify partitioning layout" which is checked. In the bottom left corner, there is a button with a document icon and the text "Release Notes". In the bottom right corner, there are two buttons: "Back" with a left arrow and "Next" with a right arrow.

Installation requires partitioning of your hard drive. By default, a partitioning layout is chosen which is reasonable for most users. You can either choose to use this or create your own.

Remove all partitions on selected drives and create default layout.

☐ Encrypt system

Select the drive(s) to use for this installation.

<input checked="" type="checkbox"/>	sda	102399 MB	DGC RAID 5
-------------------------------------	-----	-----------	------------

+ Advanced storage configuration

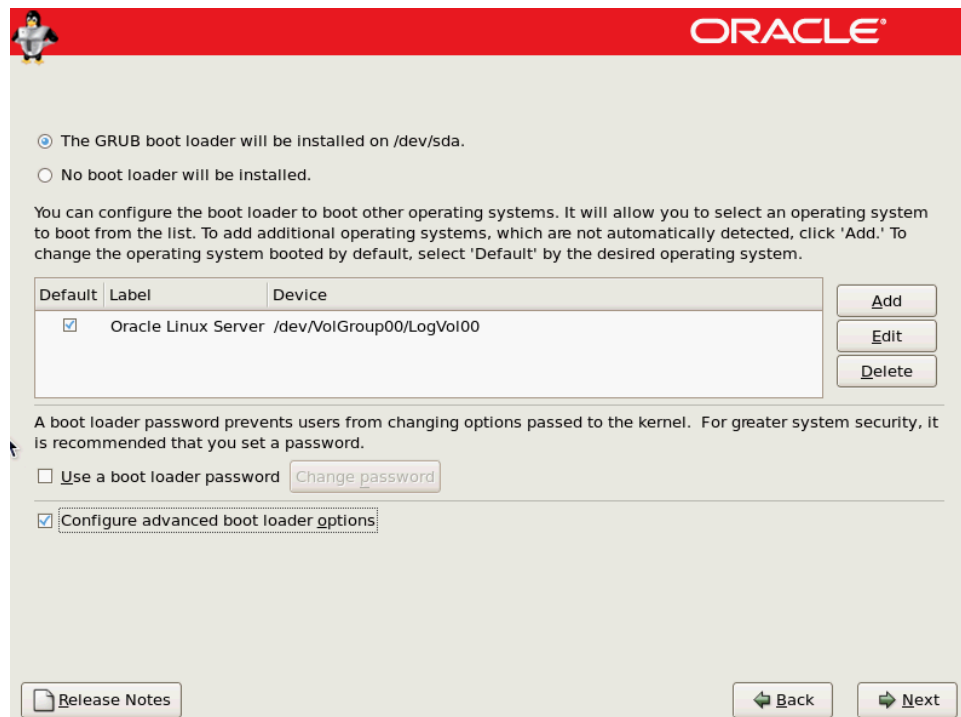
☒ Review and modify partitioning layout

Release Notes

Back Next

Make sure that you check the box for modifying the partitioning layout. Modify your layout as desired and click "Configure advanced boot loader options" (Figure 63).

**Figure 63.** Configuring the Boot Loader



The screenshot shows the 'Configuring the Boot Loader' screen in the Oracle Linux Server installer. The top header is red with the Oracle logo and a penguin icon. The main content area is light gray. It starts with two radio buttons: 'The GRUB boot loader will be installed on /dev/sda.' (selected) and 'No boot loader will be installed.' Below this is a paragraph explaining that the boot loader can be configured to boot other operating systems and that additional systems can be added via the 'Add' button. A table with three columns: 'Default', 'Label', and 'Device' is shown. The first row has a checked 'Default' box, the label 'Oracle Linux Server', and the device '/dev/VolGroup00/LogVol00'. To the right of the table are three buttons: 'Add', 'Edit', and 'Delete'. Below the table is a paragraph about boot loader passwords, followed by a checkbox 'Use a boot loader password' (unchecked) and a 'Change password' button. Below that is a checkbox 'Configure advanced boot loader options' (checked). At the bottom are a 'Release Notes' button and 'Back' and 'Next' navigation buttons.

**ORACLE**

☒ The GRUB boot loader will be installed on /dev/sda.  
☐ No boot loader will be installed.

You can configure the boot loader to boot other operating systems. It will allow you to select an operating system to boot from the list. To add additional operating systems, which are not automatically detected, click 'Add.' To change the operating system booted by default, select 'Default' by the desired operating system.

Default	Label	Device
<input checked="" type="checkbox"/>	Oracle Linux Server	/dev/VolGroup00/LogVol00

[Add](#)  
[Edit](#)  
[Delete](#)

A boot loader password prevents users from changing options passed to the kernel. For greater system security, it is recommended that you set a password.

☐ Use a boot loader password [Change password](#)

☒ [Configure advanced boot loader options](#)

[Release Notes](#) [Back](#) [Next](#)

Modify the Install Boot Loader record to the partition (Figure 64).

**Figure 64.** Specifying the Boot Loader Partition



The screenshot shows the 'Specifying the Boot Loader Partition' screen in the Oracle Linux Server installer. The top header is red with the Oracle logo and a penguin icon. The main content area is light gray. It starts with the text 'Install Boot Loader record on:' followed by two radio buttons: '/dev/sda Master Boot Record (MBR)' and '/dev/sda1 First sector of boot partition' (selected). Below the radio buttons is a 'Change Drive Order' button. Below that is a checkbox 'Force LBA32 (not normally required)' (unchecked). Below this is a paragraph explaining that default options can be added to the boot command by entering them into the 'General kernel parameters' field. Below that is a text input field for 'General kernel parameters'. At the bottom are a 'Release Notes' button and 'Back' and 'Next' navigation buttons.

**ORACLE**

Install Boot Loader record on:

☐ /dev/sda Master Boot Record (MBR)  
☒ /dev/sda1 First sector of boot partition

[Change Drive Order](#)

☐ Force LBA32 (not normally required)

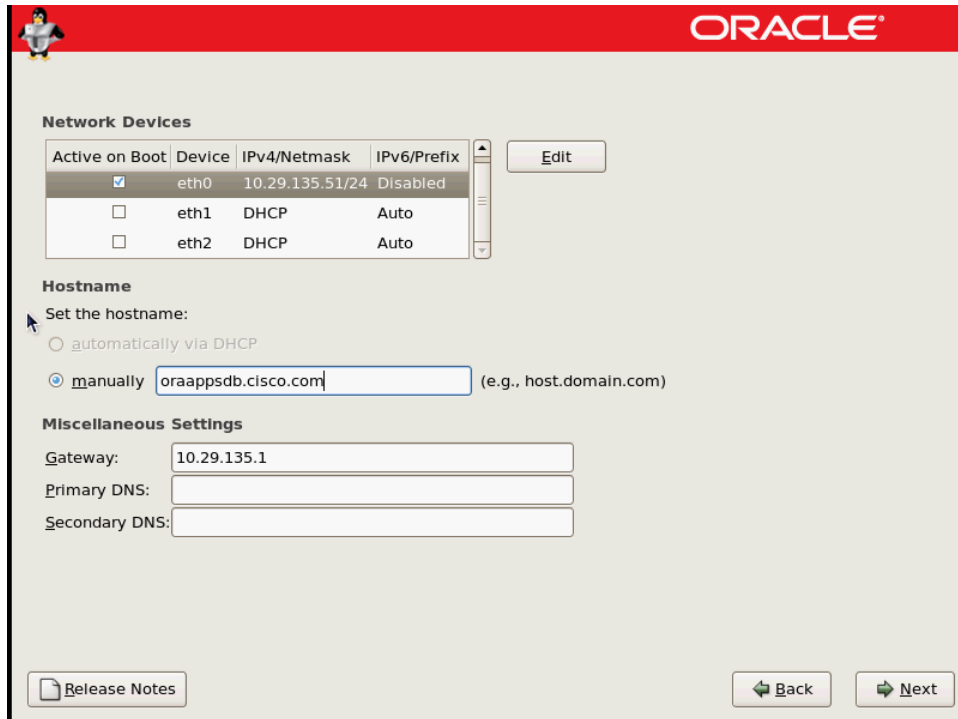
If you wish to add default options to the boot command, enter them into the 'General kernel parameters' field.

General kernel parameters

[Release Notes](#) [Back](#) [Next](#)

Configure the network as needed (Figures 65 and 66).

**Figure 65.** Configuring the Network



The image shows the Oracle Linux Network Configuration GUI. At the top is a red header with the Oracle logo and a Tux penguin icon. The main content area is titled "Network Devices" and contains a table with columns: "Active on Boot", "Device", "IPv4/Netmask", and "IPv6/Prefix". There is an "Edit" button to the right of the table. Below the table is the "Hostname" section, which includes a "Set the hostname:" label and two radio buttons: "automatically via DHCP" and "manually". The "manually" option is selected, and a text box contains "oraappsdb.cisco.com" with a hint "(e.g., host.domain.com)". Below this is the "Miscellaneous Settings" section, which includes fields for "Gateway:", "Primary DNS:", and "Secondary DNS:". At the bottom left is a "Release Notes" button, and at the bottom right are "Back" and "Next" buttons.

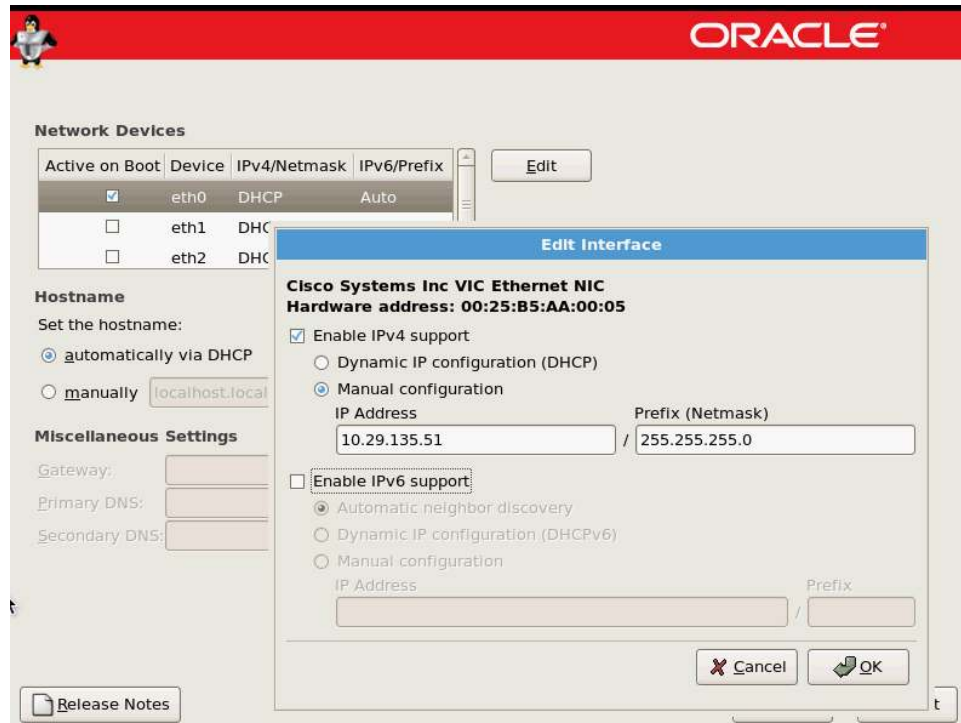
Active on Boot	Device	IPv4/Netmask	IPv6/Prefix
<input checked="" type="checkbox"/>	eth0	10.29.135.51/24	Disabled
<input type="checkbox"/>	eth1	DHCP	Auto
<input type="checkbox"/>	eth2	DHCP	Auto

**Hostname**  
Set the hostname:  
☐ automatically via DHCP  
☒ manually  (e.g., host.domain.com)

**Miscellaneous Settings**  
Gateway:   
Primary DNS:   
Secondary DNS:

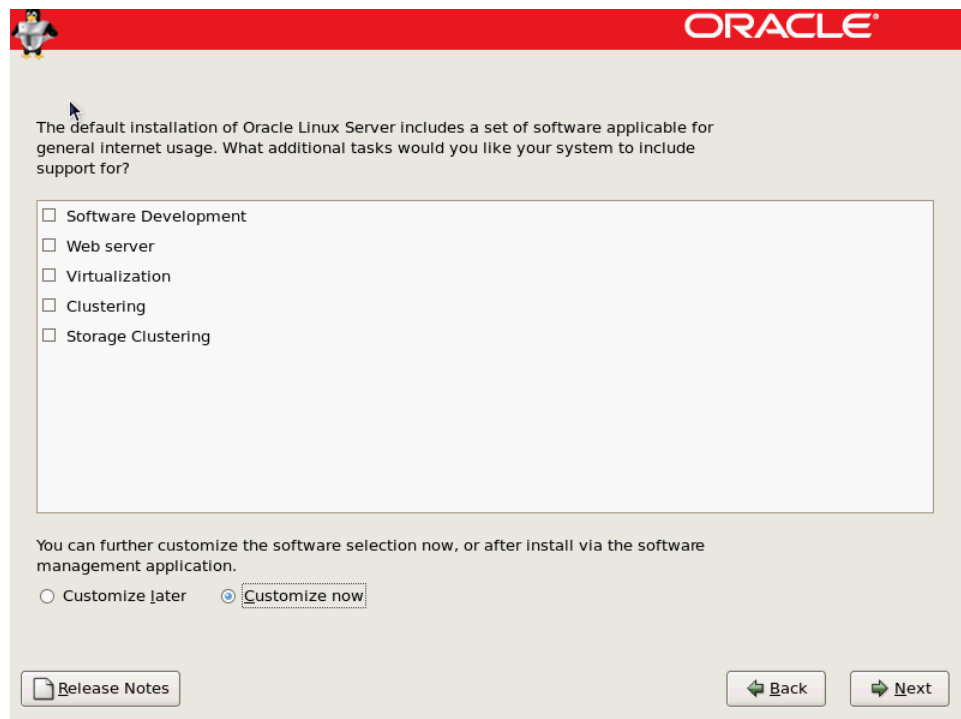
[Release Notes](#) [Back](#) [Next](#)

Figure 66. Editing the Network Interface



Select the desired time zone and move ahead.

Figure 67. Customizing the Software Selection



Click “Customize now” (Figure 67).

Under System Tools, ASM packages were selected, as the database was installed with Oracle ASM.

When prompted, reboot the server.

Accept the license information, register the system as needed, and synchronize the time with Network Time Protocol (NTP).

This completes the OS installation.

### Configure PowerPath

After reboot, configure PowerPath, as it has only a single path now. Contact EMC for the appropriate version of PowerPath for the operating system.

The version of PowerPath used in the setup was EMCpower.LINUX-5.7.1.00.00-029.el5 for Linux 5.8, 2.6.18-308.el5.

Make sure that multipath is not running, as follows:

```
# service -status-all | grep multipath
# multipath -ll
-bash: multipath: command not found

# rpm -ivh HostAgent-Linux-64-x86-en_US-1.0.0.1.0474-1.x86_64.rpm
# rpm -ivh EMCPower.LINUX-5.7.1.00.00-029.RHEL5.x86_64.rpm
If it complains about Oracle Linux, move /etc/oracle-release to /etc/oracle-
release.org

# service hostagent start
Starting Navisphere agent: [ OK ]

# service PowerPath start
Starting PowerPath: done

# powermt check_registration
There are no license keys now registered.

# emcpreg -add < power path key here >
1 key(s) successfully added.

# powermt set policy=co
# powermt config
# powermt save
#

[root@oraappsmt1 etc]# powermt display dev=all
Pseudo name=emcpowera
VNX ID=APM00121402725 [OraAppsMT1_Boot]
```



```

Logical device ID=600601606BB330009041DDD14DB5E111 [OraAppsMT1_Boot]
state=alive; policy=CLAROpt; queued-IOS=0
Owner: default=SP B, current=SP B Array failover mode: 4
=====
-----Host-----Stor---I/O Path-- --Stats--
### HW Path   I/O Paths   Interf. Mode   State   Q-IOS   Errors
=====
      1 fnic     sdc           SP B1      active  alive   0       0

```

Only one path is active right now.

## Configuring the Boot LUN

Follow the instructions in the EMC PowerPath Install and Administration guide. A few of the steps are mentioned below.

- Capture the partitions from /proc/partitions:

```

[root@oraappsmt1 etc]# cat /proc/partitions | grep emcpower
120 0 62914560 emcpowera
120 1 104391 emcpowera1 ← Boot Partition
120 2 62806117 emcpowera2

```

- Backup /etc/fstab:

```

[root@oraappsmt1 by-uuid]# ls -l /dev/disk/by-uuid/*
lrwxrwxrwx 1 root root 16 Feb 14 19:48 /dev/disk/by-uuid/bff7a2da-6583-48a5-8169-147643452cb8 -> ../../emcpowera1. This is the boot partition by uuid.
Update the boot LABEL in fstab file.

/dev/VolGroup00/LogVol00/ext3 defaults 1 1
#LABEL=/boot/boot ext3 defaults 1 2
/dev/disk/by-uuid/bff7a2da-6583-48a5-8169-147643452cb8 /boot ext3 defaults 1 0
/dev/VolGroup00/LogVol01 swap swap defaults 0 0

```

- Make logical volume manager (LVM) changes:

```

# filter = [ "a./." ]
filter = [ "a/sda[1-9]$/", "r/emcpowera2/", "r/sd.*/", "r/disk.*/", "a./." ]

```

Here emcpowera2 is for the root logical volume.

- Create a new ramdisk image file.

```

Take a backup of the existing image.
cp /boot/initrd-2.6.18-308.el5.img /boot/initrd-2.6.18-308.el5.img.org
mkinitrd /boot/initrd-2.6.18-308.el5.img 2.6.18-308.el5

```

- Modify the entry in grub.conf and reboot the host.

## Reconfigure Zoning and Boot Policies

Once PowerPath is installed, make necessary changes both in boot policies and zoning info, as mentioned earlier, to revert back to all the paths.

The zoning attribute for each HBA needs to be reverted back to what was planned earlier.

### Switch A

```
* fcid 0x380007 [pwwn 20:00:00:25:b5:bb:00:0c] [OraAppsDB_vHBA1]
* fcid 0x3806ef [pwwn 50:06:01:60:3d:e0:21:f6] [OraAppsSPA2]
* fcid 0x3807ef [pwwn 50:06:01:68:3d:e0:21:f6] [OraAppsSPB2]
```

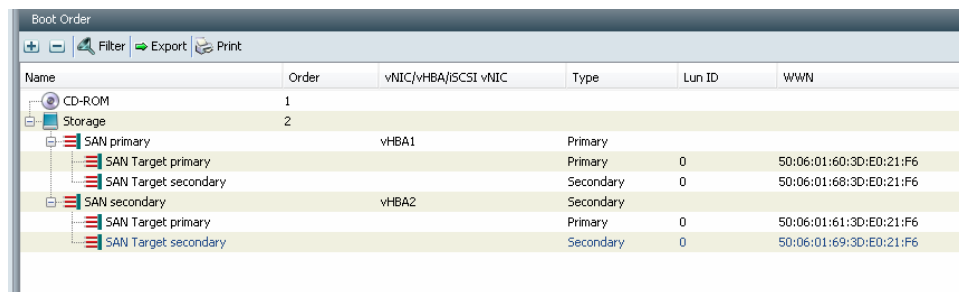
### Switch B

```
* fcid 0x380007 [pwwn 20:00:00:25:b5:bb:00:0d] [OraAppsDB_vHBA2]
* fcid 0x3804ef [pwwn 50:06:01:61:3d:e0:21:f6] [OraAppsSPA3]
* fcid 0x3805ef [pwwn 50:06:01:69:3d:e0:21:f6] [OraAppsSPB3]
```

Make the changes to all the HBAs for the SAN LUNs for the database server and to the boot LUNs for all the servers.

Modify the boot policies on the servers, too (Figure 68).

**Figure 68.** The Boot Order



Name	Order	vNIC/vHBA/SCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		vHBA1	Primary		
SAN Target primary			Primary	0	50:06:01:60:3D:E0:21:F6
SAN Target secondary			Secondary	0	50:06:01:68:3D:E0:21:F6
SAN secondary		vHBA2	Secondary		
SAN Target primary			Primary	0	50:06:01:61:3D:E0:21:F6
SAN Target secondary			Secondary	0	50:06:01:69:3D:E0:21:F6

After you reboot, all the paths should be active.

After activating, powermt should display something like the following.

This is an example of the boot LUN on one of the middle-tier servers:

```
[root@oraappsm1 etc]# powermt display dev=all
Pseudo name=emcpowera
VNX ID=APM00121402725 [OraAppsMT1_Boot]
Logical device ID=600601606BB330009041DDD14DB5E111 [OraAppsMT1_Boot]
state=alive; policy=CLAROpt; queued-IOS=0
Owner: default=SP B, current=SP B Array failover mode: 4
```

```

=====
-----Host-----Stor-----I/O Path-- --Stats--
### HW Path I/O Paths Interf. Mode State Q-I/Os Errors
=====
    0 fnic sda SP B0 active alive 0 0
    0 fnic sdb SP A0 active alive 0 0
    1 fnic sdc SP B1 active alive 0 0
    1 fnic sdd SP A1 active alive 0 0
=====

```

### Miscellaneous Post-Installation Steps

Note that not all of the following may have to be changed on your setup. Validate and change as needed. The following changes were made on the test bed on which the Oracle RAC install was done.

#### Disable selinux

```

It is recommended that you disable selinux.
Edit /etc/selinux/config and change to:
SELINUX=disabled
#SELINUXTYPE=targeted

```

#### Modify/Create the dba Group If Needed

```

groupmod -g 500 dba

```

#### Disable Firewalls

```

service iptables stop
service ip6tables stop
chkconfig iptables off
chkconfig ip6tables off
Make sure /etc/sysconfig/network has an entry for hostname. Preferably add
NETWORKING_IPV6=no

```

#### Set Up yum.repository

```

cd /etc/yum.repos.d
wget http://public-yum.oracle.com/public-yum-el5.repo
edit the downloaded file public-yum-ol6.repo and change status as enabled=1

```

Run yum update.

You may have to set up an http\_proxy environment variable in case the server accesses the Internet via a proxy.

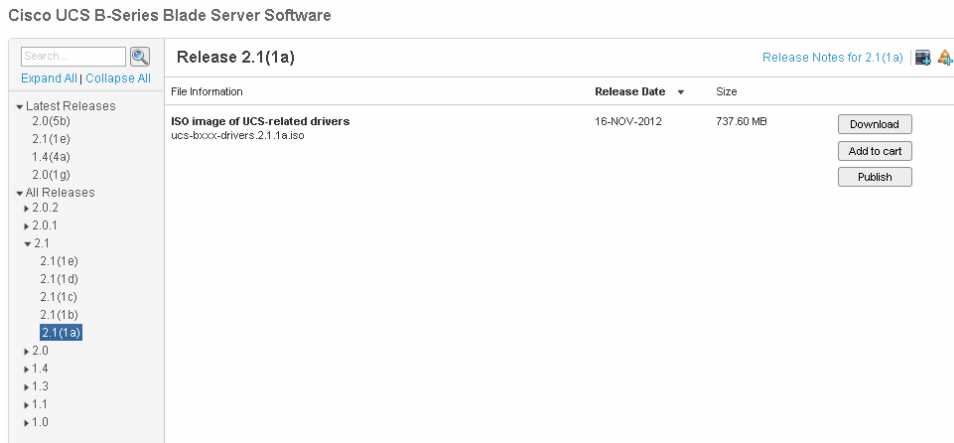
#### Install the Latest Linux Drivers for the Cisco VICs

Go to <http://software.cisco.com/download/navigator.html>

On the download page, select servers-Unified computing. On the right menu select your class of servers, for example, Cisco UCS B-Series Blade Server Software, and then select Unified Computing System (UCS) Drivers on the following page.

Select your firmware version under All Releases, such as 2.1, and download the ISO image of the UCS drivers for your matching firmware, for example, ucs-bxxx-drivers.2.1.1a.iso (Figure 69).

**Figure 69.** Downloading the UCS Drivers



Download and extract the ISO file for the drivers.

Extract the fnic/enic RPMs from the ISO.

Alternatively, you can also mount the ISO file. You can use a KVM console too and map the ISO.

After mapping the virtual media, log in to the host to copy the RPM.

For storage drivers, navigate to

```
/mnt/Linux/Storage/Cisco/MLOM/RHEL/RHEL5.8 ← This is for MLOM
/mnt/Linux/Storage/Cisco/1280/RHEL/RHEL5.8 ← This is for the VIC 1280
```

Extract the fnic drivers and install, following the instructions in the readme files.

For network drivers, navigate to

```
/mnt/Linux/Network/Cisco/MLOM/RHEL/RHEL5.8
/mnt/Linux/Network/Cisco/1280/RHEL/RHEL5.8
```

Depending on the VICs being used in the blades, you may have to navigate to the exact version and install the enic and fnic drivers.

Do a modinfo on the enic/fnic drivers to validate:

```
filename: /lib/modules/2.6.18-308.el5/extra/enic-rhel5u8/enic.ko
version: 2.1.1.41
license: GPL
author: Scott Feldman <scofeldm@cisco.com>
description: Cisco VIC Ethernet NIC Driver
srcversion: 051C239D11B943D7B04439E
alias: pci:v00001137d00000071sv*sd*bc*sc*i*
alias: pci:v00001137d00000044sv*sd*bc*sc*i*
```

```
alias: pci:v00001137d00000043sv*sd*bc*sc*i*
depends: 8021q
vermagic: 2.6.18-308.el5 SMP mod_unload gcc-4.1
```

#### Update grub.conf

By default, Oracle Linux installs with the uek kernel. Update grub.conf entries to Red Hat Compatible Kernel, 2.6.18-308.el5.

Reboot the host after making the changes, and verify.

#### Configure Oracle ASM on Database Server

Download Oracle ASMLib for RHEL5 from [www.oracle.com/technetwork/server-storage/linux/downloads/rhel5-084877.html](http://www.oracle.com/technetwork/server-storage/linux/downloads/rhel5-084877.html)

Check for the appropriate version of the kernel and download the driver.

At a minimum, you should have the following drivers:

```
oracleasm-support-2.1.8-1.el5
oracleasm-lib-2.0.4-1.el5
oracleasm-2.6.18-308.el5-2.0.5-1.el5
```

```
# /etc/init.d/oracleasm configure
```

#### Configuring the Oracle ASM library driver.

This will configure the on-boot properties of the Oracle ASM library driver. The following questions will determine whether the driver is loaded on boot and what permissions it will have. The current values are shown in brackets ([""]).

Pressing Enter without typing an answer will keep the current value. Ctrl-C will abort.

```
Default user to own the driver interface []: oracle
Default group to own the driver interface [dba]: dba
Start Oracle ASM library driver on boot (y/n) [y]: y
Scan for Oracle ASM disks on boot (y/n) [y]: y
Writing Oracle ASM library driver configuration: done
Initializing the Oracle ASMLib driver: [ OK ]
Scanning the system for Oracle ASMLib disks: [ OK ]
```

```
# cat /etc/sysconfig/oracleasm | grep -v '^#'
ORACLEASM_ENABLED=true
ORACLEASM_UID=oracle
ORACLEASM_GID=dba
ORACLEASM_SCANBOOT=true
ORACLEASM_SCANORDER="emcpower" ← Add this entry
ORACLEASM_SCANEXCLUDE="sd" ← Add this entry
```

This should create a mount point /dev/oracleasm/disks.

## Configure ASM LUNs and Create Disks

### Mask the LUNs and Create Partitions

#### Configure Storage LUNs

Add the necessary LUNS to the storage groups and provide connectivity to the hosts. Reboot the hosts so that SCSI is scanned and the LUNS are visible.

Is `/dev/emcpower*` or `powermt display dev=all` should reveal that all devices are now visible on the host.

#### Partition LUNs

Partition the LUNs with an offset of 1 MB. Although it is necessary to create partitions on disks for Oracle ASM (to prevent any accidental overwrite), it is equally important to create an aligned partition. Setting this offset aligns host I/O operations with the back-end storage I/O operations.

Use host utilities such as `fdisk` to create a partition on the disk.

Create an input file, `fdisk.input` as follows.

```
d
n
p
1

x
b
1
2048
p
w
```

Execute as `fdisk /dev/emcpower[name] < fdisk.input`. This makes a partition of 2048 cylinders. In fact, this can be scripted for all the LUNs.

All the pseudo-partitions should now be available in `/dev` as `emcpowera1`, `emcpowerb1`, `emcpowerab1`, etc.

#### Create ASM Disks

Once the partitions are created, create ASM disks with `oracleasm` APIs.

```
oracleasm createdisk -v DSS_DATA_1 /dev/emc[partition name ]
```

This will create a disk label of `DSS_DATA_1` on the partition. This label can be queried with Oracle-supplied `kfed/kfod` tools, covered later.

Repeat the process for all the partitions, and create ASM disks for all your database and RAC files.

---

Scan the disks with oracleasm; they should be visible under the /dev/oracleasm/disks mount point created earlier by oracleasm, as follows:

```
[root@rac1 disks]# oracleasm scandisks
Reloading disk partitions: done
Cleaning any stale ASM disks...
Scanning system for ASM disks...
```

The Database node is now ready.

### Installing the Operating System on Oracle Application-Tier Servers

To install the operating system on the Oracle Application servers, follow the same procedure as for the Database server: Create service profiles from the service templates, and assign the profiles to two blades, one for Web/Apps and the other for Concurrent Manager. To install Oracle Linux, move to Red Hat Compatible Kernel and set up the boot LUNs and zoning policies for the boot LUNs. These procedures are almost the same as those for the Database server.

As the system accesses the Oracle Apps middle-tier files over NFS, an entry in fstab should suffice, provided the exports from VNX are done as mentioned earlier.

#### Additional Steps on Application-Tier Servers Before Attempting a Rapid Install

- Make sure that the executables such as ar, gcc, g++, ld, ksh, make, and x-display are configured. Follow metalink note 761566.1 to check for any missing packages from the system.
- Install the following additional packages on the application-tier server boxes:
  - openmotif21-2.1.30-11.EL5.i386
  - xorg-x11-libs-compat-6.8.2-1.EL.33.0.1.i386

The above can be downloaded from [https://oss.oracle.com/projects/compat-oracle/files/Enterprise\\_Linux/](https://oss.oracle.com/projects/compat-oracle/files/Enterprise_Linux/)

- Create applmgr, oracle, and dba accounts as necessary.
- Download and apply patch 6078836. For more information on this patch, refer to metalink notes 1329085.1 and 1325822.1.
- Create a mount point /apps after making an entry in fstab as follows:

```
192.168.20.20:/apps /apps nfs
rw,hard,intr,bg,rsiz=32768,wsiz=32768,tcp,vers=3 0 0
```

Here the IP address is the storage NIC configured on the host and also exported from the VNX.

### Rapid Install of Oracle Apps R12

Only a brief overview of Rapid Install is being provided here. You may have to refer to metalink notes for detailed steps and configuration. The modules you may have to consider for install, licensing options, etc. are beyond the scope of this document.

### Download Oracle E-Business Suite

Download the Oracle E-Business Suite media packs from <https://edelivery.oracle.com>. For the product pack, select E-Business Suite, and for the platform, select x86\_64. In the current configuration, Oracle E-Business Suite Release 12.1.1 Media Pack for Linux x86-64 was selected.

Download the files to a staging location.

The files could be downloaded as follows:

```
[root@oraappsmt1 Stager12]# ls
oraAppDB oraApps oraAS oraDB startCD
```

### Create a Temporary Mount Point for Database Files

Rapid Install installs the database files onto a local file system. Hence, a default install is done to a temporary file system, followed by moving this to ASM later.

Create four LUNs of 128 GB each in VNX and mask them to the Database server:

```
cat /proc/partitions
120 16 134217728 emcpowerb
120 32 134217728 emcpowerc
120 48 134217728 emcpowerd
120 64 134217728 emcpowere
```

Volume created with the four LUNS as below.

```
[root@OraAppsDB ~]# pvcreate /dev/emcpowerb1
Physical volume "/dev/emcpowerb1" successfully created
[root@OraAppsDB ~]# pvcreate /dev/emcpowerc1
Physical volume "/dev/emcpowerc1" successfully created
[root@OraAppsDB ~]# pvcreate /dev/emcpowerd1
Physical volume "/dev/emcpowerd1" successfully created
[root@OraAppsDB ~]# pvcreate /dev/emcpowere1
Physical volume "/dev/emcpowere1" successfully created
[root@OraAppsDB ~]# vgcreate vgRIpool /dev/emcpowerb1 /dev/emcpowerc1
/dev/emcpowerd1 /dev/emcpowere1
```

```
Volume group "vgRIpool" successfully created
```

```
lvcreate -L 500G -n local_db vgRIpool
```

```
Logical volume "local_db" created
```

```
mkfs -t ext3 /dev/vgRIpool/local_db
OS type: Linux
Block size=4096 (log=2)
```

```
.....
Writing inode tables: done
```



Creating journal (32768 blocks): done

Writing superblocks and filesystem accounting information:

```
=====
Mount the LVM in /etc/fstab and reboot the host for confirmation before starting
the Rapid Install.
=====
```

Rapid Install will write the database files in this location, which later will be moved to ASM.

### Invoking Rapid Install

The NFS mount point for the staging was also mounted on the Database server, and Rapid Install was invoked from there.

```
cd /apps/StageR12/startCD/Disk1/rapidwiz
[root@OraAppsDB ]# ls
adautostg.pl ClientWiz.cmd etc images jre rapidwiz RapidWiz.ini
RapidWizVersion.cmd unzip
bin driver File jlib oui RapidWiz.cmd RapidWizVersion template
```

Set up your display, logging through VNC, Cygwin, etc., and launch the Rapidwiz wizard.

The figures that follow show a few of the screen shots of installing on the Database server, followed by both Apps nodes, as listed below.

```
[root@OraAppsDB ]# pwd
/apps/StageR12/startCD/Disk1/rapidwiz
[root@OraAppsDB ]# ./rapidwiz
```

**Figure 70.** Starting the Install

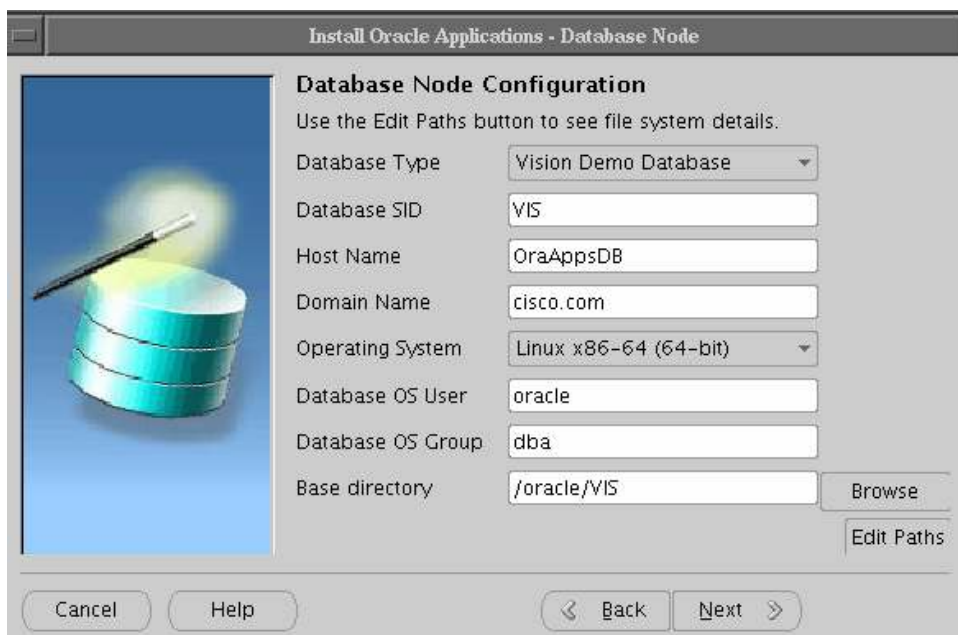


**Figure 71.** Creating a New Configuration

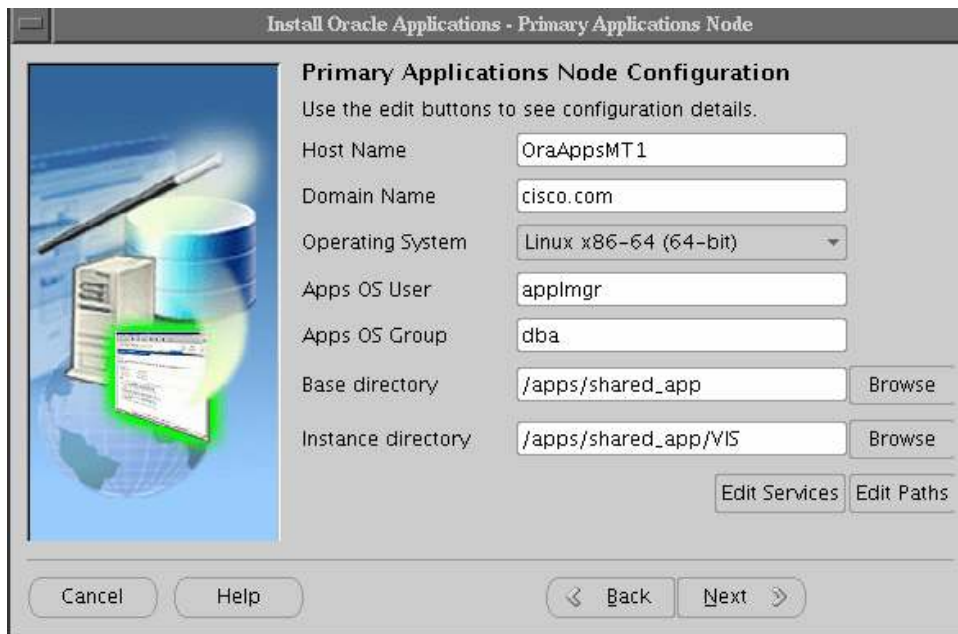


**Select default port pool 0 or edit the setting as required.**

**Figure 72.** Configuring the Database Node



**Figure 73.** Configuring the Primary Applications Node



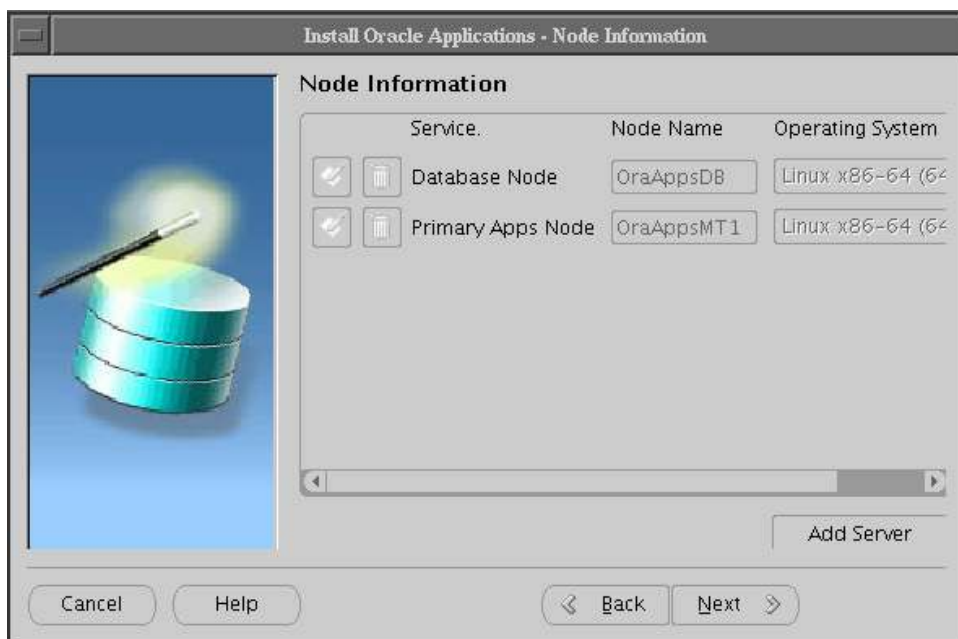
**Install Oracle Applications - Primary Applications Node**

**Primary Applications Node Configuration**  
Use the edit buttons to see configuration details.

Host Name: OraAppsMT1  
Domain Name: cisco.com  
Operating System: Linux x86-64 (64-bit)  
Apps OS User: applmgr  
Apps OS Group: dba  
Base directory: /apps/shared\_app [Browse](#)  
Instance directory: /apps/shared\_app/VIS [Browse](#)  
[Edit Services](#) [Edit Paths](#)

[Cancel](#) [Help](#) [Back](#) [Next](#)

**Figure 74.** Viewing Node Information



**Install Oracle Applications - Node Information**

**Node Information**

Service	Node Name	Operating System
<a href="#">Add Node</a> Database Node	OraAppsDB	Linux x86-64 (64-bit)
<a href="#">Add Node</a> Primary Apps Node	OraAppsMT1	Linux x86-64 (64-bit)

[Add Server](#)

[Cancel](#) [Help](#) [Back](#) [Next](#)

Click Add Server (Figure 74).

**Figure 75.** Specifying Additional Apps Node Information

The screenshot shows a window titled "Additional Applications Node Configuration". It contains the following fields and controls:

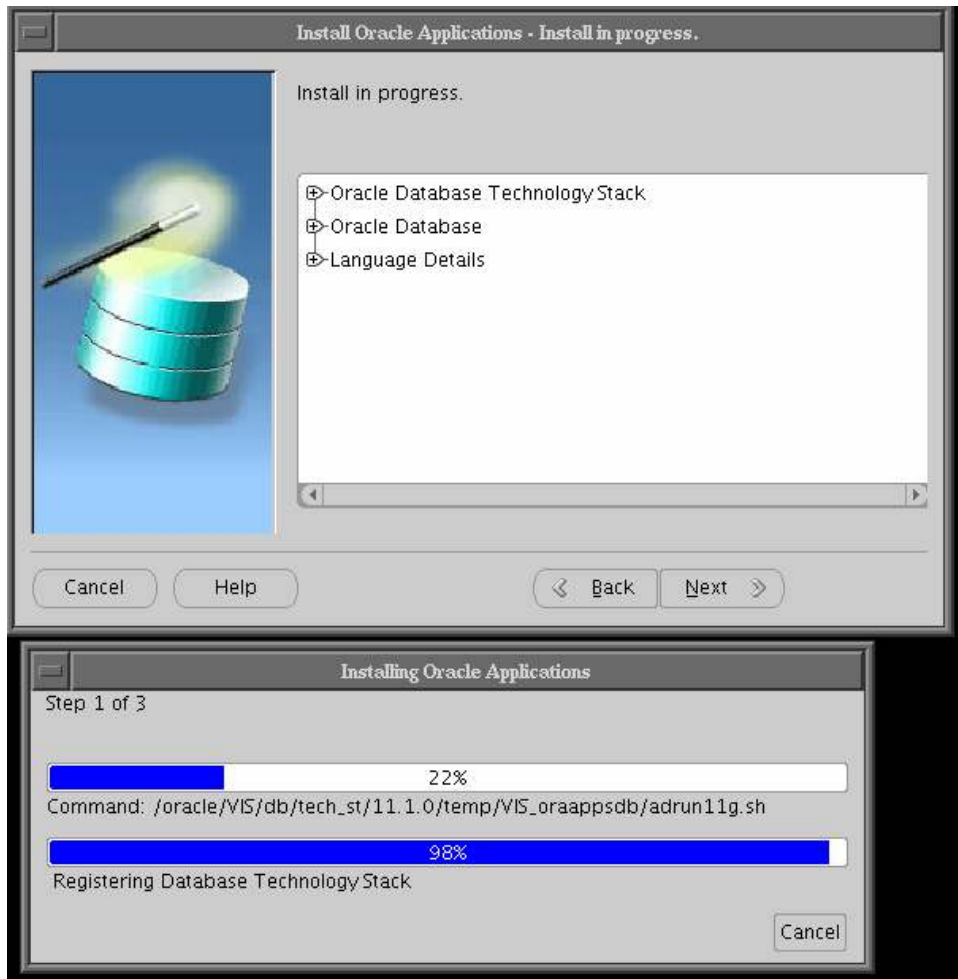
- Host Name: OraAppsMT2
- Domain Name: cisco.com
- Operating System: Linux x86-64 (64-bit) (dropdown menu)
- Apps OS User: applmgr
- Apps OS Group: dba
- Base directory: /apps/shared\_app (with a "Browse" button)
- Instance directory: /apps/shared\_app/VIS
- ☒ Shared File System: OraAppsMT1 (dropdown menu)
- Buttons: "Edit Paths", "Edit Services", "OK", and "Cancel".

Click the Shared File System check box, and select the server name from the drop-down menu (Figure 75).

You can optionally click the Edit Services button to enable or disable services here. You can edit services directly in this screen and/or edit the auto configuration files later to enable or disable the desired services on the node.

Start the installation (Figure 76).

**Figure 76.** Viewing the Installation Progress



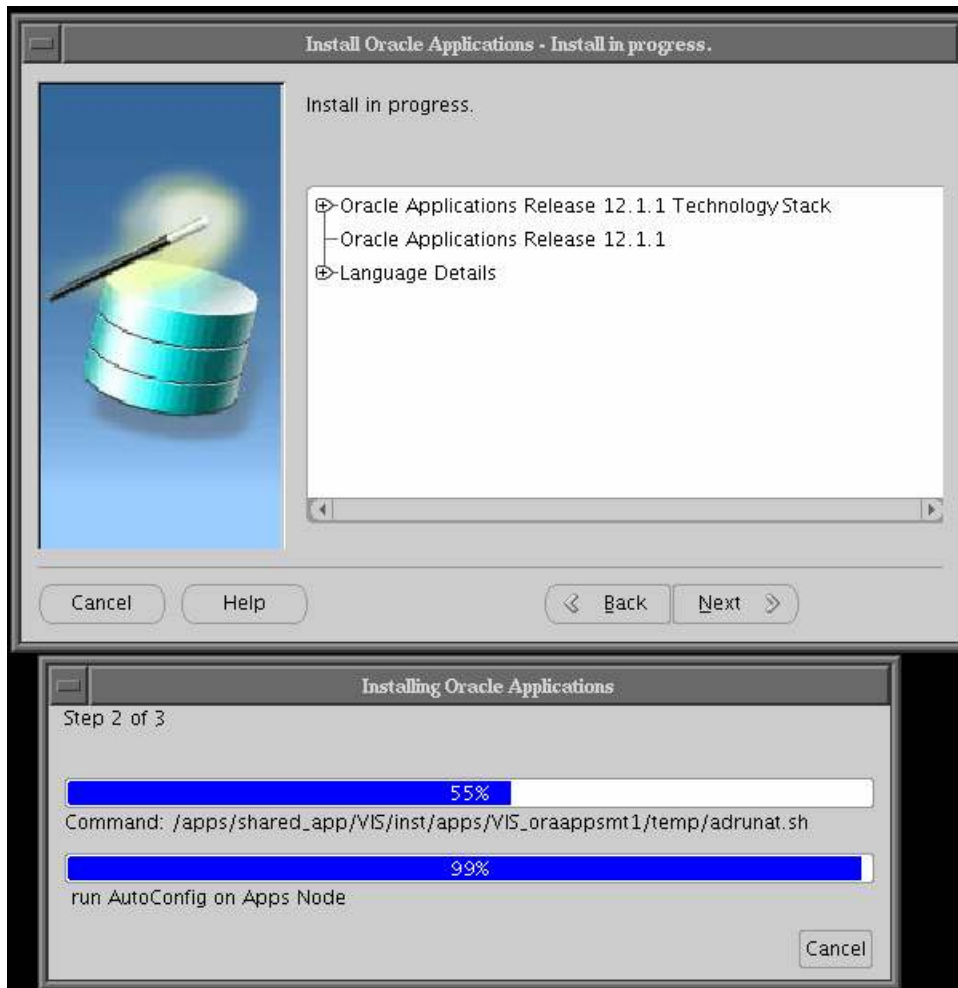
**Complete the installation and check the logfiles for any errors.**

Go to Middle Tier Box 1 and restart as ./rapidwiz, then restart as root user again.

Give the connectivity as oraappsdb.cisco.com:VIS:1521.

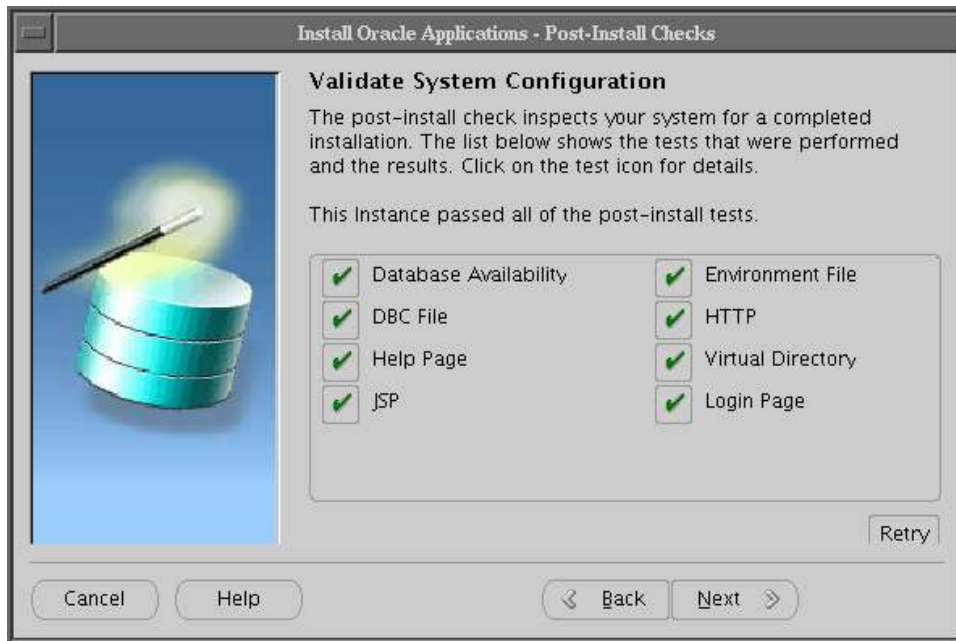
Replace the above with the thin Java Database Connectivity (JDBC) parameters that were used when running on the database node: <DBNODE>SID>:<Port#>

**Figure 77.** Installing on the First Middle Tier



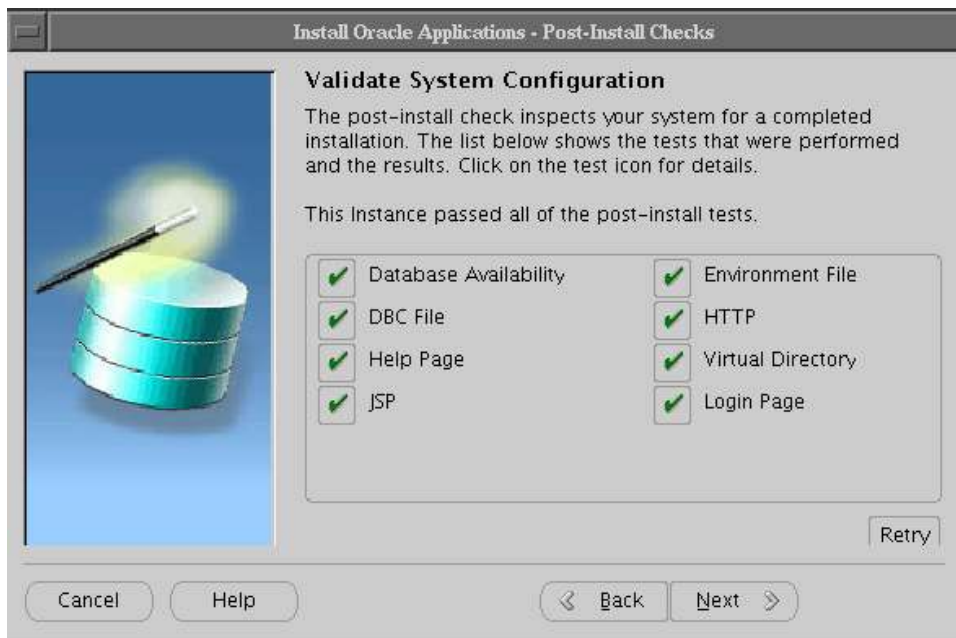
The installation of the Apps node completes (Figure 78).

**Figure 78.** Post-Installation Checks for the first Apps node.



Install the second Apps node (Figure 79).

**Figure 79.** Post-Installation Checks for the Second Apps Node



Once the installation is complete, log in to the database to validate that services are registered correctly.

```
SQL> select node_name, support_db as "DB", SUPPORT_CP as "Conc", SUPPORT_FORMS as
"Forms", SUPPORT_WEB as "Web",
2 SUPPORT_ADMIN as "Admin"
3 from apps.fnd_nodes
4 order by 1;
```

**Table 5.** Services Assignments

NODE_NAME	DB	Conc	Forms	Web	Admin
AUTHENTICATION	N	N	N	N	N
ORAAPPSDB	Y	N	N	N	N
ORAAPPSMT1	N	N	Y	Y	Y
ORAAPPSMT2	N	Y	N	N	Y

From this we can conclude that MT1, the first middle tier, is the web and forms server, while the second node is for the Concurrent Manager servers.

### Moving the Database Files to ASM

The next step is to install ASM on the database node and move the database files to ASM.

The detailed steps of installing ASM on a standalone node are not covered here. Whether you plan to use a single node or a RAC, install and configure the HAS or CRS and have the ASM processes up and running on the node(s) before attempting this step.

While there are many documented methods for migrating the database files from a file system to ASM, such as using RMAN, we used asmcp on the test bed. However, a few parameters from the ASM side had to be bumped up in order to use this method. The Cisco UCS B440 we used had ample memory and cores to support simultaneous copies from the file system. The bottleneck observed on the test bed was due to the read capacity of the local\_db file system, which was created with four LUNs, where Rapid Install created the database files.

- Create three ASM diskgroups, one for DATA, one for REDOCTL, and one for FRA (optional). Follow ASM best practices here. Copy all the data files to the DATA diskgroup, and the REDO and Control files to the REDOCTL diskgroup.
- Issue 'Alter database backup controlfile to trace' to capture all the files. Take a note of the database files, redo log files, and tempfiles.
- Shut down the database.

```
cd /oracle/VIS/db/tech_st/11.1.0/appsutil/scripts/VIS_oraappsdb/
./addbctl.sh stop VIS
```

- Set up your ASM environment.

```
[oracle@oraappsdb ~]$ . oraenv
ORACLE_SID = [VIS] ? +ASM
```

The Oracle base for ORACLE\_HOME=/oracle/ASM is /oracle/VIS/db/tech\_st/11.1.0



- Prepare a script either spooled from database earlier or from control file trace created.  
`asmcmd cp /local_db/VIS/data/<filename1.dbf> +DATA/VIS/<filename1.dbf>`
- Similarly copy the redo and control files to +REDOCTL diskgroup.

Once the script is prepared, run the script with around 50 copy commands in parallel. The parallelism depends on the server CPU, memory, and ASM processes parameters. Appendix D lists some of the ASM parameters used.

This will migrate all the database, temp, and redo log files into ASM.

Update the create control file script, and create the control file. A sample is given below.

```
STARTUP NOMOUNT
CREATE CONTROLFILE REUSE DATABASE "VIS" NORESETLOGS NOARCHIVELOG
  MAXLOGFILES 192
  MAXLOGMEMBERS 3
  MAXDATAFILES 1024
  MAXINSTANCES 32
  MAXLOGHISTORY 1460
LOGFILE
  GROUP 4 '+REDOCTL/vis/log1.dbf' SIZE 512M BLOCKSIZE 512,
  GROUP 5 '+REDOCTL/vis/log2.dbf' SIZE 512M BLOCKSIZE 512,
  GROUP 6 '+REDOCTL/vis/log3.dbf' SIZE 512M BLOCKSIZE 512
DATAFILE
  '+DATA/vis/sys1.dbf',
  '+DATA/vis/tx_data11.dbf',
  '+DATA/vis/tx_idx11.dbf',
  .....
```

Once the control file is created and the database is opened, do the sanity checks on the database.

### Sanity Tests on Oracle Applications

Login in to Oracle Applications and do the sanity tests.

<http://oraappsmt1.cisco.com:8000>

This should redirect you to the login page. Enter the default credentials sysadmin/sysadmin and select sysadmin responsibility.

If the login fails for any reasons, you can start from

[http://oraappsmt1.cisco.com:8000/OA\\_HTML/jsp/fnd/aoljtest.jsp](http://oraappsmt1.cisco.com:8000/OA_HTML/jsp/fnd/aoljtest.jsp). Enter apps username, apps password, and Database host, sid, and port data to debug further.

The page also should take you to the aolj\_setup\_test, which can shed light on the issues (Figure 80).

**Figure 80.** Testing Options

- Connection Test
  - [Locate DBC File](#)
  - [Verify DBC Settings](#)
  - [Verify Guest User](#)
  - [AOL/J Connection Test](#)
- [Virtual Directory settings](#)
- [APPS\\_WEB\\_AGENT](#)
  - [Virtual Directory Settings](#)
  - [FND\\_WEB.PING](#)
  - [Custom Authentication](#)
  - [GFM](#)
- [APPS\\_SERVLET\\_AGENT](#)
  - [Virtual Directory Settings](#)
  - [Servlet Ping](#)
  - [Jsp Ping](#)
- [APPS\\_FRAMEWORK\\_AGENT](#)
  - [Virtual Directory Settings](#)
  - [Servlet Ping](#)
  - [Jsp Ping](#)
  - [Cabo Setup Tests](#)
  - [X Server Accessibility](#)
  - [OA Framework System Info](#)
  - [Versions for Loaded Classes](#)
- [Online Help](#)
- Tool Launcher Profile Settings
  - [ICX\\_REPORT\\_LAUNCHER](#)
  - [ICX\\_DISCOVERER\\_LAUNCHER](#)
- Application Login
  - [Login Page](#)
  - [Show Responsibilities](#)(Must login first)
  - [Show Session Properties](#)(Must login first)
- End aoljtest session
  - [End current AOL/J Test session](#)

If successfully logged into the R12 sysadmin responsibility, you can do few sanity tests from Oracle Applications Manager, as well as on workflow, Concurrent Managers, etc. (Figure 81).

**Figure 81.** Performing Sanity Tests

Filter	Any	contains		Go
Select Host and ...	View Status	View Configuration	Online	Go
Select Details	Name	State	Status	Platform
<input checked="" type="radio"/> <a href="#">Show</a>	ORAAPPSDB	Online	✓	Linux x86-64 (64-bit)
<input type="radio"/> <a href="#">Show</a>	ORAAPPSMT1	Online	✓	Linux x86-64 (64-bit)
<input type="radio"/> <a href="#">Show</a>	ORAAPPSMT3	Online	✓	Linux x86-64 (64-bit)

## Adding Another Web/Apps Node to the Infrastructure

We can optionally add another apps node to the apps infrastructure. Follow metalink note 384248.1 for complete details.

- On the existing web node, run adpreclone.

```
cd $INST_TOP/admin/scripts
perl adpreclone.pl appsTier
```

- Mount the shared file system that's on the NFS mount point onto the new host. Create a new service profile, apply to a blade, install the OS with all the prerequisites as above, export the NFS file system from VNX, make changes to /etc/fstab, and mount the partition.
- As applmgr user, run adclonctx.pl to clone the context file.

```
cd /apps/shared_app/VIS/apps/apps_st/comm/clone/bin
[applmgr@oraappsmt3 bin]$ perl adclonctx.pl addnode
contextfile=/apps/shared_app/VIS/inst/apps/VIS_oraappsmt1/appl/admin/VIS_oraapp
smt1.xml
```

Running:

```
/apps/shared_app/VIS/apps/apps_st/comm/clone/bin/../../jre/bin/java -Xmx600M -
classpath
/apps/shared_app/VIS/apps/apps_st/comm/clone/bin/../../jlib/ojdbc14.jar:/apps/shar
ed_app/VIS/apps/apps_st/comm/clone/bin/../../jlib/xmlparserv2.jar:/apps/shared_app
/VIS/apps/apps_st/comm/clone/bin/../../jlib/java
oracle.apps.ad.context.CloneContext -e
/apps/shared_app/VIS/inst/apps/VIS_oraappsmt1/appl/admin/VIS_oraappsmt1.xml -
addnode
```

Enter the APPS password : apps

Log file located at

/apps/shared\_app/VIS/apps/apps\_st/comm/clone/bin/CloneContext\_0709141022.log

Provide the values required for creation of the new APPL\_TOP Context file.

Target System Hostname (virtual or normal) [oraappsmt3] :

It is recommended that your inputs are validated by the program.

However you might choose not to validate them under following circumstances:

-If cloning a context on source system for a remote system.

-If cloning a context on a machine where the ports are taken and you do not want to shut down the services at this point.

-If cloning a context but the database it needs to connect is not available.

Do you want the inputs to be validated (y/n) [n] ? : y

Target System Root Service [enabled] : yes

Target System Web Entry Point Services [enabled] : yes

Target System Web Application Services [enabled] : yes

Target System Batch Processing Services [enabled] : no

Target System Other Services [disabled] :

Do you want to preserve the Display [oraappsmt1:0.0] (y/n) ? : n

Target System Display [oraappsmt3:0.0] :

Database port is 1521

Creating the new APPL\_TOP Context file from :

/apps/shared\_app/VIS/apps/apps\_st/appl/ad/12.0.0/admin/template/adxmlctx.tmp

The new APPL\_TOP context file has been created :

/apps/shared\_app/VIS/inst/apps/VIS\_oraappsmt3/appl/admin/VIS\_oraappsmt3.xml

- Run adconfig.pl to run the autoconfig.

```
cd $AD_TOP/bin
```

```
./adconfig.pl
```

```
contextfile=/apps/shared_app/VIS/inst/apps/VIS_oraappsmt3/appl/admin/VIS_oraappsmt3.xml
```

Enter the APPS user password:

AutoConfig is configuring the Applications environment...

AutoConfig will consider the custom templates if present.

Using CONFIG\_HOME location : /apps/shared\_app/VIS/inst/apps/VIS\_oraappsmt3

Classpath :

/apps/shared\_app/VIS/apps/apps\_st/comn/java/lib/appsborg2.zip:/apps/shared\_app/VIS/apps/apps\_st/comn/java/classes

Using Context file :

/apps/shared\_app/VIS/inst/apps/VIS\_oraappsmt3/appl/admin/VIS\_oraappsmt3.xml

Context Value Management will now update the Context file

Updating Context file...COMPLETED

Attempting upload of Context file and templates to database...COMPLETED

Configuring templates from all of the product tops...

Configuring AD\_TOP.....

Configuring FND\_TOP.....COMPLETED

Configuring ICX\_TOP.....COMPLETED

Configuring MSC\_TOP.....COMPLETED

Configuring IEO\_TOP.....COMPLETED

.....

Autoconfig completed successfully.

```
SQL> select node_name, support_db as "DB", SUPPORT_CP as "Conc", SUPPORT_FORMS  
as "Forms", SUPPORT_WEB as "Web",
```

```
2 SUPPORT_ADMIN as "Admin" from apps.fnd_nodes order by 1;
```

NODE_NAME	DB	Conc	Forms	Web	Admin
AUTHENTICATION	N	N	N	N	N
ORAAPPSDB	Y	N	N	N	N
ORAAPPSMT1	N	N	Y	Y	Y
ORAAPPSMT2	N	Y	N	N	Y
ORAAPPSMT3	N	N	Y	Y	Y

```
./adconfig.pl <new context file >
```

The above will create a new instance top directory structure.

```
[applmgr@oraappsmt1 apps]$ ls
```

```
VIS_oraappsmt1 VIS_oraappsmt2 VIS_oraappsmt3
```

- Run Auto Config on all the nodes again as new host is added.
- `cd $INST_TOP/admin/scripts/; ./adautoCfg.sh`
- Select from apps.fnd\_nodes to make sure that the new host is seeded in the database.

Note that in case of multiple middle-tier servers in the system, the order in which autoconfig is run in the earlier step, the context variable `s_external_url` is updated. Hence, if oraappsmt3 is the last node where autoconfig was run, this is what is seeded in the system as the login URL.

<http://oraappsmt3.cisco.com:8000/>.

In order for both web nodes to load-balance the web transactions, a hardware load balancer needs to be configured and autoconfig has to be rerun after updating the context variables. A few of the steps in configuring the ACE load balancer are shown below. However, any of the load balancers that are certified with Oracle E-Business Suite can be used. It's beyond the scope of this document to cover all the details of setting up load balancers with Oracle E-Business Suite. Please refer to the metalink notes for full details.

## Configure a Load Balancer

To load-balance between two web nodes, you need to configure a load balancer. In the test bed, a Cisco Application Control Engine (ACE) load balancer was used. While a simple configuration was attempted as a test case here, for more details on how to configure a load balancer, refer to the metalink notes or get the details directly from the manufacturer. Metalink note 727171.1 provides details on Oracle certified load balancers for Oracle E-Business Suite. Metalink note 380489.1 provides details on using load balancers with R12 Oracle E-Business Suite.

### Configure a Hardware Load Balancer

The ACE load balancer comes with both a GUI and a CLI. The load balancer can be configured with either of these. Please refer to metalink note 603325.1 for information on configuring an ACE load balancer. We do not present detailed steps for doing so here. Appendix C lists some of the configuration options used. Session persistence was configured with active insert cookies in the setup. Optionally, you may set SSL termination and service policies as desired in your configuration.

### Run Autoconfig After Configuring the Load Balancer

After configuring the load balancer, it is necessary to rerun autoconfig across all the nodes.

- Run `exec fnd_conc_clone.setup_clean` as apps users; this will clean up the existing configurations.
- Run `autoconfig.sh` first on the database node, and then on all the middle tier nodes.  
  
In the setup, we had one database node, two web nodes, and one Concurrent Manager node. As the last run was done on the Concurrent Manager node, the `tnsnames.ora` and other files were populated correctly only on the last node. Hence you may have to rerun autoconfig again on the database and two apps nodes.
- Before running autoconfig, we had to update the context variables shown in Table 6 in the context file.

**Table 6.** Updated Context Variables

<b>s_webentryurlprotocol</b> □	http(s)
<b>s_webentryhost</b> □	Name of the host
<b>s_webentrydomain</b>	cisco.com
<b>s_active_webport</b> □	80 or 8000
<b>s_login_page</b>	<a href="http://ebizoncus.cisco.com:8000/OA_HTML/AppsLogin">http://ebizoncus.cisco.com:8000/OA_HTML/AppsLogin</a> □
<b>s_external_url</b>	<a href="http://ebizonucs.cisco.com:8000">http://ebizonucs.cisco.com:8000</a> □

Here it is assumed that the 8000 port is configured as the entry point on the load balancer.

## Install EM Grid Control 12c

This is an optional setup done on the test bed. Oracle Enterprise Manager Grid Control was installed, followed by Oracle Applications Management Suite for Oracle E-Business Suite (12.1.0.1.0). The agents were installed on all the hosts with the Apps Plugin patch and discovered the R12 infrastructure to monitor the system while doing some stress and performance tests. For more information, see *Getting Started with Oracle Application Management Pack for Oracle E-Business Suite, Release 12.1.0.1.0* (My Oracle Support Note 1434392.1).

Figure 82 provides a glimpse of a few targets in R12 E-Business Suite in EM Grid Control.

**Figure 82.** Targets in Grid Control

Name	Type	Status	Member Status Summary	Incidents	Members
EMSC_GCDomain/GCDomain(EMSC_OMS1)emsc	Application Deployment	OK	- - - - -	-	
EMSC_GCDomain/GCDomain(EMSC_OMS1)emsc	Application Deployment	OK	- - - - -	-	
EMSC_GCDomain/GCDomain(EMSC_OMS1)OC4Jspacore	Application Deployment	OK	- - - - -	-	
ebiz-cv-system	Generic System	OK	- 25 - - -	-	OC4J JRM(5), Host(3), OC4J(3), ... More
VIS	Database System	OK	- 3 - - -	-	Oracle Home(2), Database Instance(1), Listener(1), ... More
VIS-Core Managers for Concurrent Processing	Oracle Concurrent Manager	OK	- 1 - - -	-	Internal Concurrent Manager(1)
VIS-Forms Applications Infrastructure/VIS_orapoint1	Oracle Apps Forms System	OK	- 4 - - -	-	OC4J(1), OC4J JRM(1), Oracle HTTP Server(1), ... More
VIS_orapoint1.orapoint1.cisco.com_forms	OC4J	OK	- 1 - - -	-	OC4J JRM(1)
VIS-Infrastructure/VIS_orapoint1-Database Context	Oracle E-Business Suite Infrastructure	OK	- 1 - - -	-	Database Instance(1)
VIS-Infrastructure/VIS_orapoint1-APPL_TOP Context	Oracle E-Business Suite Infrastructure	OK	- 10 - - -	-	OC4J(3), OC4J JRM(3), Oracle HTTP Server(1), ... More
VIS_orapoint1.orapoint1.cisco.com	Oracle Application Server	OK	- 7 - - -	-	OC4J(3), OC4J JRM(3), Oracle HTTP Server(1)
VIS-Redundancy Group for Self Service Applications Infrastructure	Oracle Apps SSA Redundancy Group	OK	- 4 - 1 -	-	Oracle Apps SSA System(2), OC4J(1), OC4J JRM(1), ... More
VIS-Self Service Applications Infrastructure/VIS_orapoint1	Oracle Apps SSA System	OK	- 3 - - -	-	OC4J(1), OC4J JRM(1), Oracle HTTP Server(1)
VIS_orapoint1.orapoint1.cisco.com_pacore	OC4J	OK	- 1 - - -	-	OC4J JRM(1)
VIS-Self Service Applications Infrastructure/VIS_orapoint2	Oracle Apps SSA System	OK	- - - 1 -	-	OC4J(3), OC4J JRM(3), Oracle HTTP Server(1)
VIS_orapoint1.orapoint1.cisco.com	Oracle Application Server	OK	- 7 - - -	-	OC4J JRM(1)
VIS_orapoint1.orapoint1.cisco.com_forms	OC4J	OK	- 1 - - -	-	OC4J JRM(1)
VIS_orapoint1.orapoint1.cisco.com_pacore	OC4J	OK	- 1 - - -	-	OC4J JRM(1)
VIS_orapoint1.orapoint1.cisco.com_safm	OC4J	OK	- 1 - - -	-	OC4J JRM(1)

## Using Cisco UCS Service Profiles for Failover to a Spare Blade

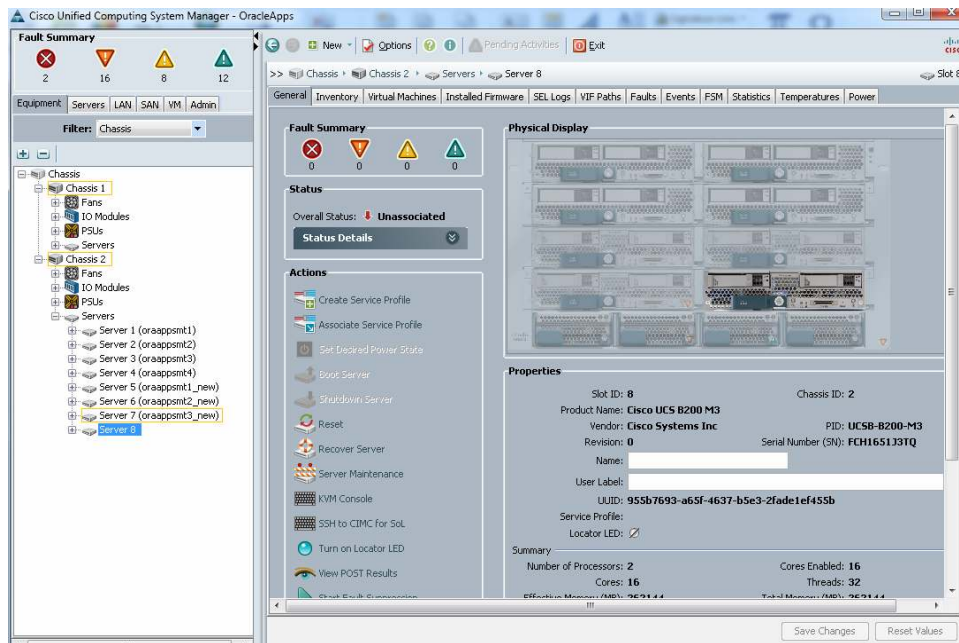
An attempt was made to test UCS blade failovers through UCS service profiles by associating a service profile with a spare blade in the UCS domain.

The service profile of the second web server was associated with a spare blade.

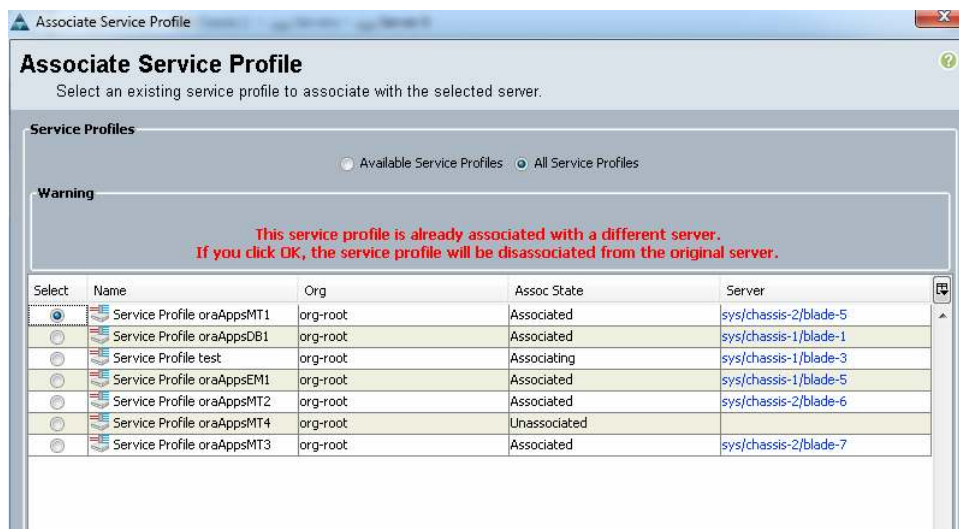
Here are the steps.

Click the Equipment tab, filter on Chassis, and select the empty blade with which you would like associate the profile (Figure 83).

**Figure 83.** Associating an Empty Blade with a Service Profile



**Figure 84.** Selecting a Service Profile



Click Associate Service Profile, then select All profiles and the service profile that you want to move (Figure 84). If it is already associated, it may display a warning. Click OK to associate the Service Profile.

This association was done with web transactions happening on the system.

A script was added under the /etc/rc2.d directory to start up the services as part of the server boot. In fact the script was to set up the apps environment as applmgr (application user), navigate to \$INST\_TOP/admin/scripts, and start adstrall.sh. No other changes were made to the system for this exercise.

Fault injection and end-to-end time taken for the business continuity were as shown in Table 7.

**Table 7.** Time Required to Restore Services

Activity	Time taken
Service profile association	1' 07"
Server uptime	6' 50"
Restoration of Oracle services	2'20"
Total downtime	10'17"

The total time taken was around 10 minutes. Because of the presence of the load balancer, the transactions shifted to OC4J running on the surviving blade during this short period.

The process was monitored in EM Grid Control. Figure 85 shows a snippet of the Oracle E-Business Suite services that were up and running after the association.

**Figure 85.** Some of the E-Business Suite Services Running After the Association

Name	Type	Member Status Summary	Incidents	Members
ebiz2	Application Deployment	25	0	OC4J JMXS, Host(3), OC4J(3), ... More
VTS	Generic System	3	0	Oracle Home(2), Database Instance(1), Listener(1), ... More
VTS-Core Manager for Concurrent Processing	Oracle Concurrent Manager	1	0	Internal Concurrent Manager(1)
VTS-Forms Applications Infrastructure VTS_orappem1	Oracle Apps Forms System	4	0	OC4J(1), OC4J JMX(1), Oracle HTTP Server(1), ... More
VTS-Infrastructure VTS_orappem1-cisco.com_forms	OC4J	1	0	OC4J JMX(1)
VTS-Infrastructure VTS_orappem1-Database Context	Oracle E-Business Suite Inf...	1	0	Database Instance(1)
VTS-Infrastructure VTS_orappem1-APPL_TOP Context	Oracle E-Business Suite Inf...	10	0	OC4J(3), OC4J JMX(3), Oracle HTTP Server(1), ... More
VTS_orappem1-cisco.com	Oracle Application Server	7	0	OC4J(3), OC4J JMX(3), Oracle HTTP Server(1)
VTS-Self Service Applications Infrastructure	Oracle Apps SSA Redundanc...	4	1	Oracle Apps SSA System(2), OC4J(1), OC4J JMX(1), ... More
VTS-Self Service Applications Infrastructure VTS_orappem1	Oracle Apps SSA System	3	0	OC4J(1), OC4J JMX(1), Oracle HTTP Server(1)
VTS-Self Service Applications Infrastructure VTS_orappem2	OC4J	1	0	OC4J JMX(1)
VTS_orappem1	Oracle Apps SSA System	1	0	OC4J JMX(1)
VTS_orappem1-cisco.com	Oracle Application Server	7	0	OC4J(3), OC4J JMX(3), Oracle HTTP Server(1)
VTS_orappem1-cisco.com_forms	OC4J	1	0	OC4J JMX(1)
VTS_orappem1-cisco.com_gaconn	OC4J	1	0	OC4J JMX(1)
VTS_orappem1-cisco.com_nafm	OC4J	1	0	OC4J JMX(1)

EM recorded around 11 minutes of downtime (Figure 86).

**Figure 86.** Downtime as Measured by Enterprise Manager




Figure 86 was extracted from EM Grid Control, which monitored the complete failover. It took few more seconds for the EM agent running on the system to upload the data to the central console; hence the total duration of 11 minutes.



While the presence of the load balancer makes it transparent to the end user, the load on the surviving Web and Apps nodes increases, resulting in performance degradation. The use of service profiles in a SAN boot environment reduces this window. The time it takes to fix the hardware failure or to configure another blade with similar hardware characteristics, such as MAC and IP addresses, etc., is much simplified by taking advantage of Cisco UCS service profiles.

## Performance and Destructive Tests

### Performance Tests

In order to do performance and destructive tests, we needed a tool that continuously keeps a load on the system at the web layer and also does batch processing through Concurrent Manager. At the time of testing, only the batch processing OATS (Oracle Application Testing Suite) kit was feasible. The web transaction toolkit still had issues and could not be used. Whenever and wherever required, a few synthetic transactions were posted through EM Grid Control, but that was not a true load-generation kit.

The load-generation toolkit was provided by Oracle. The kit comes with its own database and apptop. This was installed in parallel with the vision database and apptop, was converted to shared apptop, and was set up with the load balancer.

The purpose of using the toolkit was not to benchmark on the existing setup, but to create sufficient load on the system before running any destructive tests.

For details on benchmark data on Cisco UCS B200 M3 server, see [www.oracle.com/us/solutions/benchmark/apps-benchmark/results-166922.html](http://www.oracle.com/us/solutions/benchmark/apps-benchmark/results-166922.html).

### Details Collected from Tools


#### Automatic Workload Repository (AWR) Report Events

**Table 8.** Results of Performance Tests in AWR

Event	Waits	Time(s)	Average Wait (ms)	% DB Time	Wait Class
DB CPU		46,743		84.64	
Buffer busy waits	4,446,406	2,918	1	5.28	Concurrency
Enq: TX - index contention	1,004,789	2,648	3	4.79	Concurrency
DB file scattered read	136,822	939	7	1.7	User I/O
DB file sequential read	169,972	819	5	1.48	User I/O

## EM Graphs

**Figure 87.** Performance Test Results: Active Sessions

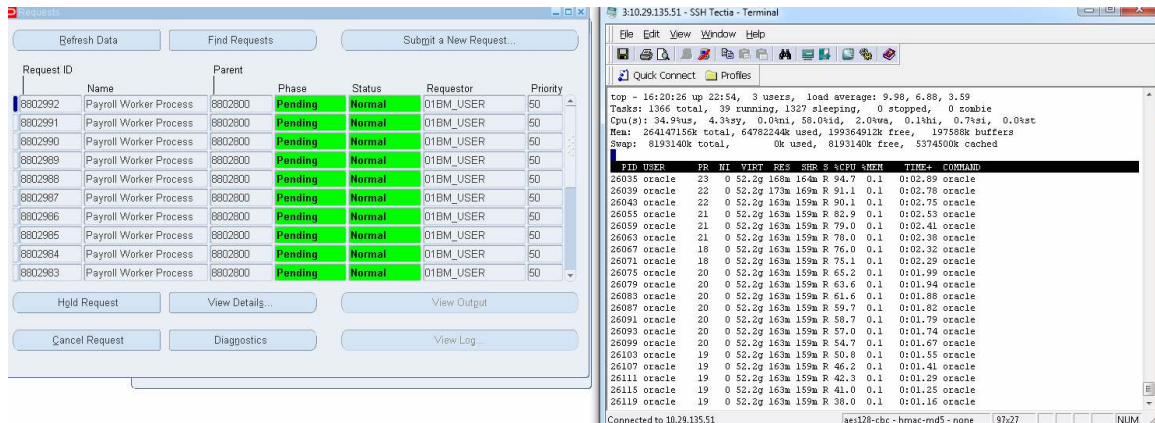


**Figure 88.** Performance Test Results: Runnable Processes and Active Sessions



## Payroll Run Data

**Figure 89.** Performance Test Results: Payroll and Order to Cash Tests



Both Payroll and Order to Cash tests were done on the system (Figure 89).

## Destructive Tests

The destructive tests were done on the Apps server. The database was on single instance. For details on RAC failures, refer to the 11gR2 white papers and/or CVDs at <http://www.cisco.com/go/oracle>. The RAC failures are not covered here.

The load was generated as mentioned in the Performance Tests section and the fault was induced (Table 9).

**Table 9.** Results of Destructive Tests on the Apps Server

Test Case	Fault Injection	Expected Result	Status
<b>Reboot one of the Web tier nodes</b>	The web node was rebooted	System should continue working.	The load balancer in front of the Apps nodes continued transactions. The OC4J processes on the surviving node picked up the load.
<b>Reboot Concurrent Manager node</b>	The Concurrent Manager node was rebooted	The web transactions should continue	The batch processing failed. The test bed did not include PCP setup. Otherwise, continuity could have been expected.
<b>Reboot one of the fabrics</b>	Rebooted one of the fabrics	No interruption	The fabric joined back in 10 to 15 minutes. The NICs failed over to the surviving fabric, seamlessly without any interruption.
<b>Associate service profile with a spare node</b>	One of the web nodes was disassociated and associated with a spare blade	System should continue working	The spare blade was back in the pool and shared the load within 10 minutes. No interruption.

## Appendix

### Appendix A: Cisco UCS Service Profiles

#### Fabric Interconnect

ID: A  
Product Name: Cisco UCS 6120XP  
HW Revision: 0  
Total Memory (MB): 3548  
OOB IP Addr: 10.29.135.4  
OOB Gateway: 10.29.135.1  
OOB Netmask: 255.255.255.0  
Operability: Operable  
Thermal Status: Ok

ID: B  
Product Name: Cisco UCS 6120XP  
HW Revision: 0  
Total Memory (MB): 3548  
OOB IP Addr: 10.29.135.6  
OOB Gateway: 10.29.135.1  
OOB Netmask: 255.255.255.0  
Operability: Operable  
Thermal Status: Ok

## Server Inventory

Server	Equipped	PID	Equipped VID	Equipped Serial (SN)	Ackd Slot Status	Ackd Memory (MB)	Ackd Cores
-----	-----	-----	-----	-----	-----	-----	-----
1/1	B440-BASE-M2	V01	FCH16177CZP	Equipped		262144	40
1/2			Equipped Not Pri				
1/5	B230-BASE-M2	V01	FCH16017E1E	Equipped		262144	20
1/8			Equipped Not Pri				
2/5	UCSB-B200-M3	V01	FCH164579EP	Equipped		262144	16
2/6	UCSB-B200-M3	V01	FCH1651J48Y	Equipped		262144	16
2/7	UCSB-B200-M3	V01	FCH1651J42C	Equipped		262144	16
2/8	UCSB-B200-M3	V01	FCH1651J3TQ	Equipped		262144	16

## Show interface brief

Ethernet Interface	VLAN	Type	Mode	Status	Reason	Speed Ch #	Port
-----	-----	-----	-----	-----	-----	-----	-----
Eth1/1	1	eth	fabric	up	none	10G(D)	-
Eth1/2	1	eth	fabric	up	none	10G(D)	-
Eth1/3	1	eth	fabric	up	none	10G(D)	-
Eth1/4	1	eth	fabric	up	none	10G(D)	-
Eth1/5	1	eth	fabric	up	none	10G(D)	-
Eth1/6	1	eth	fabric	up	none	10G(D)	-
Eth1/7	1	eth	fabric	up	none	10G(D)	-
Eth1/8	1	eth	fabric	up	none	10G(D)	-
...( partial list )							

## show service-profile assoc detail

Service Profile Name: oraAppsDB1  
Association: Associated  
Server: 1/1  
Selected Server: sys/chassis-1/blade-1  
Server Pool:

Service Profile Name: oraAppsEM1  
Association: Associated  
Server: 1/5  
Selected Server: sys/chassis-1/blade-5

Service Profile Name: oraAppsMT1  
 Association: Associated  
 Server: 2/5  
 Selected Server: sys/chassis-2/blade-5  
 Server Pool:

Service Profile Name: oraAppsMT2  
 Association: Associated  
 Server: 2/6  
 Selected Server: sys/chassis-2/blade-6  
 Server Pool:

Service Profile Name: oraAppsMT3  
 Association: Associated  
 Server: 2/7  
 Selected Server: sys/chassis-2/blade-7  
 Server Pool:  
 ...(partial list )

## Appendix B: Nexus 5548UP Zoning Details

### Switch A

zoneset	name	OraAppsZoneset	vsan	101	
	zone	name	OraAppsDB_1		
	*			20:00:00:25:b5:bb:00:0c	OraAppsDB_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsDB_3		
	*			20:00:00:25:b5:bb:00:0e	OraAppsDB_vHBA3
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsMT1_1		
	*			20:00:00:25:b5:bb:00:0a	OraAppsMT1_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsMT2_1		
	*			20:00:00:25:b5:bb:00:07	OraAppsMT2_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsMT3_1		
	*			20:00:00:25:b5:bb:00:05	OraAppsMT3_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsEM_1		
	*			20:00:00:25:b5:bb:00:03	OraAppsEM_vHBA1

zoneset	name	OraAppsZoneset	vsan	101	
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2

#### Switch B

zoneset	name	OraAppsZoneset	vsan	101	
	zone	name	OraAppsDB_1		
	*			20:00:00:25:b5:bb:00:0c	OraAppsDB_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsDB_3		
	*			20:00:00:25:b5:bb:00:0e	OraAppsDB_vHBA3
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsMT1_1		
	*			20:00:00:25:b5:bb:00:0a	OraAppsMT1_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsMT2_1		
	*			20:00:00:25:b5:bb:00:07	OraAppsMT2_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsMT3_1		
	*			20:00:00:25:b5:bb:00:05	OraAppsMT3_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2
	zone	name	OraAppsEM_1		
	*			20:00:00:25:b5:bb:00:03	OraAppsEM_vHBA1
	*			50:06:01:60:3d:e0:21:f6	OraAppsSPA2
	*			50:06:01:68:3d:e0:21:f6	OraAppsSPB2

#### Appendix C: Cisco ACE Configuration

```
ebizcvd-4710-1/Admin# show run
Generating configuration....
```

```
logging enable
logging timestamp
logging trap 5
```

```
boot system image:c4710ace-t1k9-mz.A5_2_0.bin
```

```
login timeout 0
```

```
hostname ebizcvd-4710-1
```

```
interface gigabitEthernet 1/1
  switchport access vlan 135
  no shutdown
interface gigabitEthernet 1/2
  switchport access vlan 135
  no shutdown
interface gigabitEthernet 1/3
  shutdown
interface gigabitEthernet 1/4
  shutdown

access-list everyone line 10 extended permit ip any any

probe http http
  interval 5
  passdetect interval 2
  passdetect count 1
  expect status 200 200

rserver host oraappsmt1
  ip address 10.29.135.61
  inservice
rserver host oraappsmt2
  ip address 10.29.135.63
  inservice

serverfarm host web-farm
  predictor leastconns
  probe http-probe
  rserver oraappsmt1
    inservice
  rserver oraappsmt2
    inservice

sticky http-cookie ace-id app-cookie
  cookie insert browser-expire
  serverfarm web-farm

class-map match-all vip-web
  2 match virtual-address 10.29.135.90 tcp eq 8000

policy-map type management first-match remote-mgmt
  class class-default
    permit

policy-map type loadbalance first-match web-policy
  class class-default
```

```
sticky-serverfarm app-cookie

policy-map multi-match client-vips
class vip-web
  loadbalance vip inservice
  loadbalance policy web-policy
  loadbalance vip icmp-reply active

interface vlan 135
ip address 10.29.135.80 255.255.255.0
access-group input everyone
service-policy input remote-mgmt
service-policy input client-vips
no shutdown
interface vlan 135
ip address 10.29.135.80 255.255.255.0
access-group input everyone
service-policy input remote-mgmt
service-policy input client-vips
no shutdown

ip route 0.0.0.0 0.0.0.0 10.29.135.1

username admin password 5 $1$a/1PvTdQ$5yE.pNd7YdKyObJXKJEDo/ role Admin domain
default-domain
username www password 5 $1$TJ.Ed4o3$z8kCowI18gINHBVcQQbU10 role Admin domain
default-domain
ssh key rsa 4096 force
ssh key rsa1 768 force
```

**Note:** This is only a partial list and is just for reference.

#### Appendix D: ASM parameters

Some of the parameters in ASM were increased for taking batch loads and also to support multiple asmcp commands to move the database files from the file system to ASM.

```
asm_diskgroups='DATA','REDOCTL','FRA'
asm_power_limit=1
memory_target=1023M
large_pool_size=12M
sessions=600
processes=400
```

#### Appendix E: Database Instance Parameters

```
sga_target=60G
sessions=6000
db_writer_processes=4
```



---

db\_cache\_size=25G

## Appendix F: Oracle Apps Context File Parameters

```
<webentryurlprotocol oa_var="s_webentryurlprotocol">http</webentryurlprotocol>
<webentryhost oa_var="s_webentryhost">ebizonucs</webentryhost>
<webentrydomain oa_var="s_webentrydomain">cisco.com</webentrydomain>
<activewebport oa_var="s_active_webport" oa_type="DUP_PORT" base="8000" step="1"
range="-1" label="Active Web Port">8000</activewebport>
<login_page
oa_var="s_login_page">http://ebizonucs.cisco.com:8000/OA\_HTML/AppsLogin</login_page>
<externURL oa_var="s_external_url">http://ebizonucs.cisco.com:8000</externURL>
```



**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](http://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](http://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)