FlexPod Data Center with Oracle Database 11g R2 RAC with 7-Mode

Deployment Guide for FlexPod with Oracle Database 11g R2 RAC with Oracle Direct NFS Client

Last Updated: November 22, 2013

Building Architectures to Solve Business Problems

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

http://www.cisco.com/go/designzone.

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

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

Oracle RAC on FlexPod

Executive Summary

Data powers essentially every operation in a modern enterprise, from keeping the supply chain operating efficiently to managing relationships with customers. Database administrators and their IT departments face many challenges that demand needs for a simplified deployment and operation model providing high performance, availability and lower total cost of ownership (TCO). Cisco Validated Designs (CVD) can help you deploy data center applications in a wide range of operating environments including mission critical workloads. This CVD describes how the Cisco Unified Computing System™ (Cisco UCS®) can be used in conjunction with NetApp® FAS storage systems to implement an Oracle Real Application Clusters (RAC) solution that is an Oracle Certified Configuration. FlexPod® components are integrated and standardized to help you eliminate the guesswork and achieve timely, repeatable, and consistent deployments. FlexPod has been optimized to run a variety of mixed workloads and while offering flexible yet robust design configurations. Customers are generally able to accelerate their transition with the FlexPod data center solution, which integrates disparate compute, storage, and network components into a single architecture that scales to fit a variety of virtualized and non-virtualized customer environments.

The key benefits of FlexPod deployments are:

- Single platform from industry leaders in networking, computing, and storage.
- Pretested, validated solution platform to reduce risk and increase efficiencies.
- Flexible IT architecture for today's needs, yet scales for future growth.
- Cooperative support model for efficient and streamlined resolution.

For more information on NetApp FlexPod architecture, visit

http://www.netapp.com/us/technology/flexpod/

Target Audience

This document is intended to assist solution architects, project managers, infrastructure managers, sales engineers, field engineers, and consultants in planning, designing, and deploying Oracle Database 11g R2 RAC hosted on FlexPod. This document assumes that the reader has an architectural understanding of the Cisco Unified Computing System, Oracle Database 11gR2 GRID Infrastructure, Oracle Real Application Clusters, NetApp storage systems, and related software.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Purpose of this Guide

This FlexPod CVD demonstrates how enterprises can apply best practices to deploy Oracle Database 11g R2 RAC using Cisco Unified Computing System, Cisco Nexus family switches, and NetApp FAS storage systems. This validation effort exercised typical Online transaction processing (OLTP) and Decision-support systems (DSS) workloads to ensure expected stability, performance and resiliency design as demanded by mission critical data center deployments.

Business Needs

Business applications are moving into integrated stacks consisting of compute, network, and storage. This FlexPod solution helps to reduce costs and complexity of a traditional Oracle Database 11g R2 RAC deployment. Following business needs for Oracle Database 11g R2 RAC deployments are addressed by this solution.

- Increasing DBA's productivity by ease of provisioning and simplified yet scalable architecture.
- A balanced configuration that yields predictable purchasing guidelines at the compute, network and storage tiers for a given workload.
- Reduced risk for a solution that is tested for end-to-end interoperability of compute, storage, and network.

Solution Overview

Oracle Database 11g R2 RAC on FlexPod with Oracle Direct NFS Client

This solution provides an end-to-end architecture with Cisco Unified Computing System, Oracle, and NetApp technologies and demonstrates the FlexPod configuration benefits for running Oracle Database 11g R2 RAC with Cisco VICs (Virtual Interface Card) and Oracle Direct NFS Client.

The following infrastructure and software components are used for this solution:

- Oracle Database 11g R2 RAC
- Cisco Unified Computing System
- Cisco Nexus 5548UP switches
- NetApp storage systems and supporting components
- NetApp OnCommand® System Manager 2.2
- Swingbench, a benchmark kit for OLTP and DSS workloads

The following is the solution architecture connectivity layout for this FlexPod deployment.

Figure 1 Solution Architecture

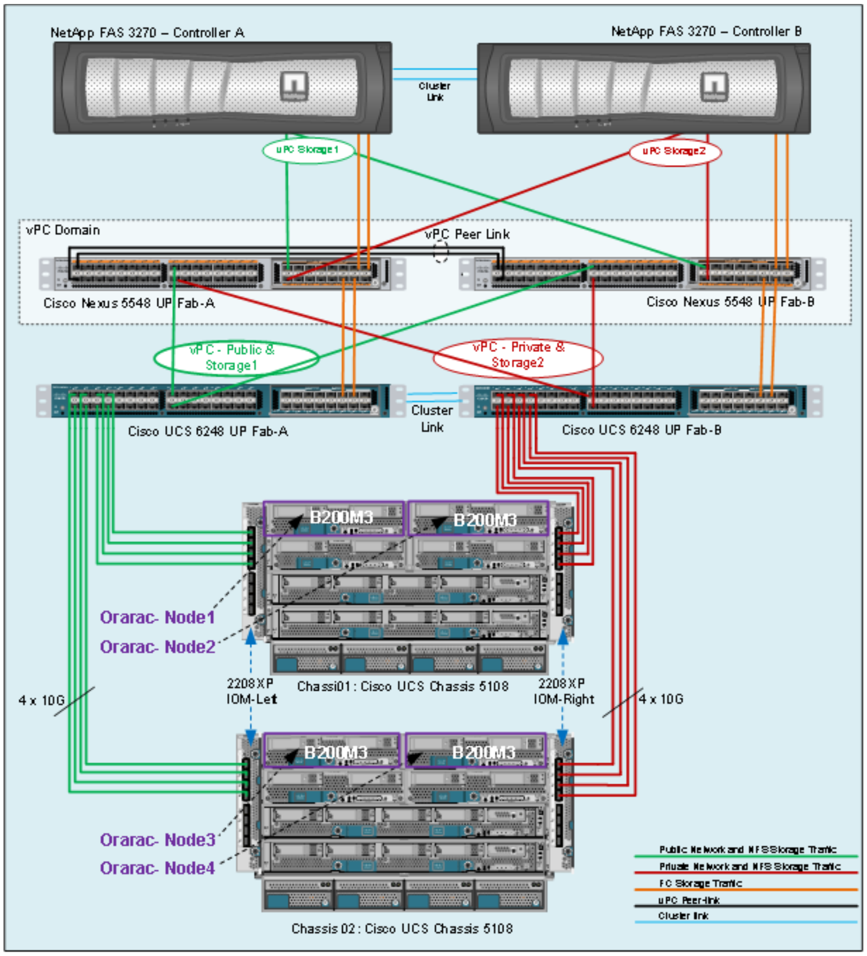

Technology Overview

This section describes the Cisco Unified Computing System (Figure 2).

Figure 2 Third Generation Fabric Computing

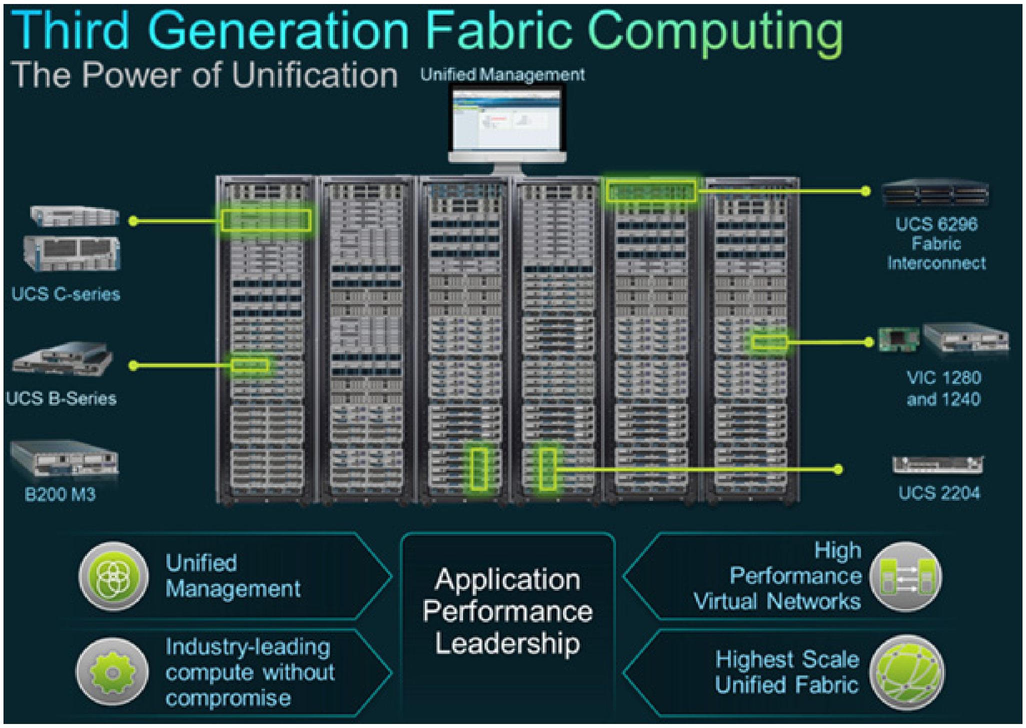
Third Generation Fabric Computing
The Power of Unification
Unified Management
UCS C-series
UCS B-Series
B200 M3
UCS 6296 Fabric Interconnect
VIC 1280 and 1240
UCS 2204
Unified Management
Industry-leading compute without compromise
Application Performance Leadership
High Performance Virtual Networks
Highest Scale Unified Fabric

The Cisco Unified Computing System is a third-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce TCO and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet (10GbE) unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain that is controlled and managed centrally.

Main Components of the Cisco Unified Computing System

The main components of the Cisco Unified Computing System are as follows:

- Compute—The system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon® E5-2600 Series Processors. Cisco UCS B-Series Blade Servers work with virtualized and non-virtualized applications to increase performance, energy efficiency, flexibility and productivity.
- Network—The system is integrated onto a low-latency, lossless, 80-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- Storage access—The system provides consolidated access to both storage area network (SAN) and network-attached storage (NAS) over the unified fabric. By unifying storage access, Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This provides customers with the options for setting storage access and investment protection. Additionally, server administrators can reassign storage-access policies for system connectivity to storage resources, thereby simplifying storage connectivity and management for increased productivity.

- Management—The system uniquely integrates all system components which enable the entire solution to be managed as a single entity by the Cisco UCS Manager. The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations.

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership (TCO), increased Return on Investment (ROI) and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system which unifies the technology in the data center. The system is managed, serviced and tested as a whole.
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand.
- Industry standards supported by a partner ecosystem of industry leaders.

Figure 3 Components of Cisco Unified Computing System

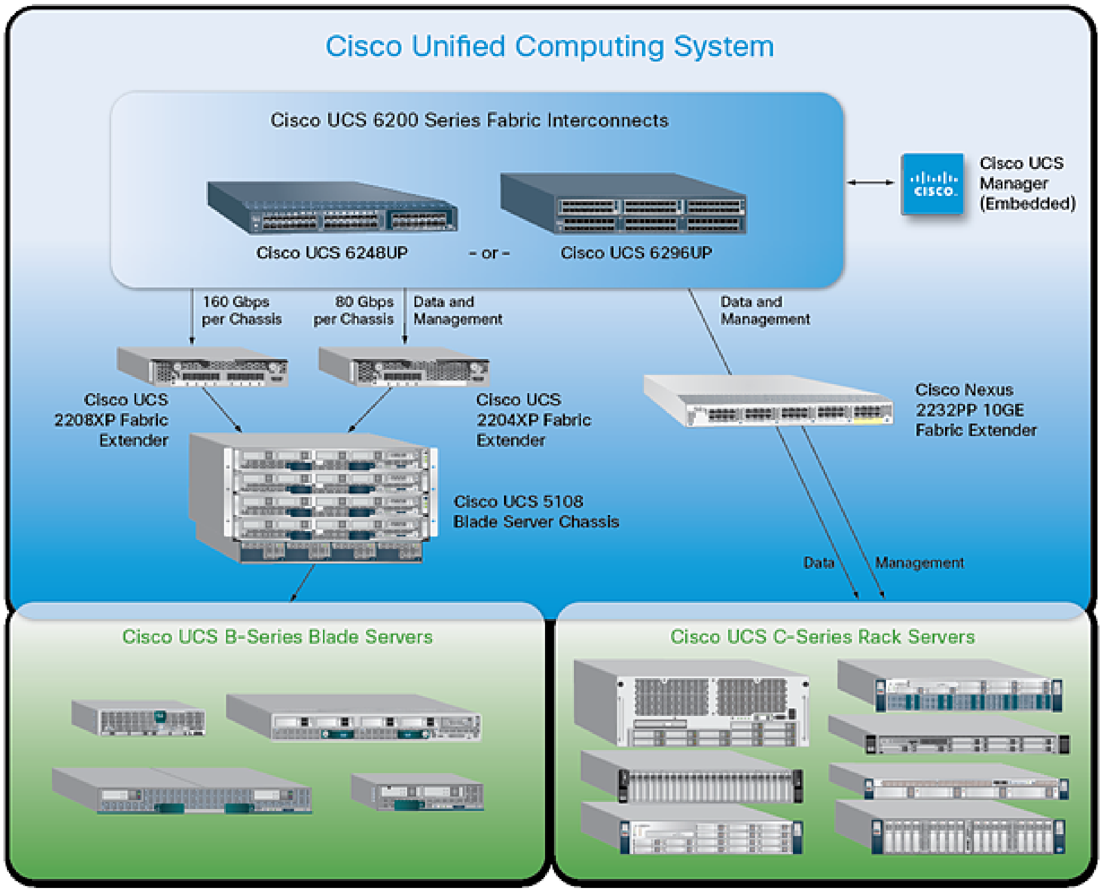
Cisco Unified Computing System
Cisco UCS 6200 Series Fabric Interconnects
Cisco UCS 6248UP
- or -
Cisco UCS 6296UP
Cisco UCS Manager (Embedded)
160 Gbps per Chassis
80 Gbps per Chassis
Data and Management
Data and Management
Cisco UCS 2208XP Fabric Extender
Cisco UCS 2204XP Fabric Extender
Cisco Nexus 2232PP 10GE Fabric Extender
Cisco UCS 5108 Blade Server Chassis
Data
Management
Cisco UCS B-Series Blade Servers
Cisco UCS C-Series Rack Servers

Cisco UCS Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis (Figure 4) is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis.

The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2208 XP Fabric Extenders.

A passive mid-plane provides up to 40 Gbps of I/O bandwidth per server slot and up to 80 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 80 Gigabit Ethernet standards.

Figure 4 Cisco Blade Server Chassis (front, back, and populated with blades)

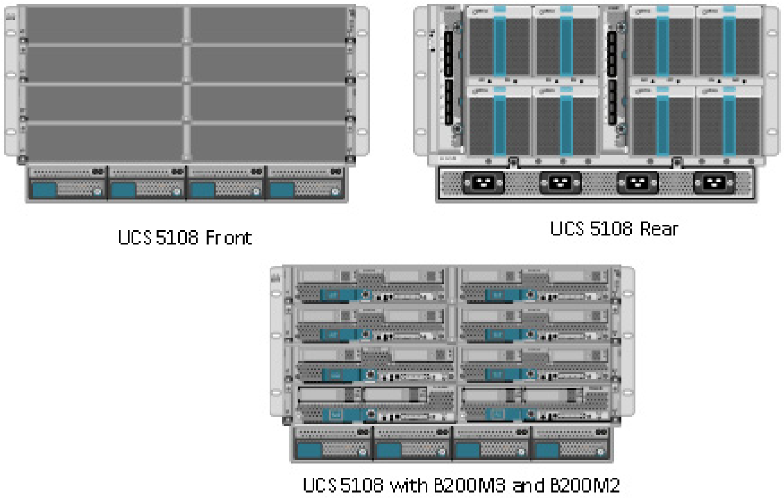

UCS 5108 Front

UCS 5108 Rear

UCS 5108 with B200M3 and B200M2

Cisco UCS B200 M3 Blade Server

The Cisco UCS B200 M3 Blade Server is a half-width, two-socket blade server. The system uses two Intel Xeon® E5-2600 Series Processors, up to 384 GB of DDR3 memory, two optional hot-swappable small form factor (SFF) serial attached SCSI (SAS) disk drives, and two VIC adapters that provides up to 80 Gbps of I/O throughput. The server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

Figure 5 Cisco UCS B200 M3 Blade Server

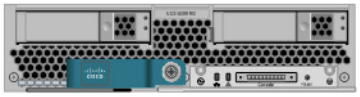

Cisco UCS Virtual Interface Card 1240

A Cisco innovation, the Cisco UCS VIC 1240 is a four-port 10 Gigabit Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

Cisco UCS 6248UP Fabric Interconnect

The Fabric interconnects provide a single point for connectivity and management for the entire system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by Cisco UCS Manager. The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in deterministic I/O latency regardless of a server or virtual machine's topological location in the system.

Cisco UCS 6200 Series Fabric Interconnects support the system's 80-Gbps unified fabric with low-latency, lossless, cut-through switching that supports IP, storage, and management traffic using a single set of cables. The fabric interconnects feature virtual interfaces that terminate both physical and virtual connections equivalently, establishing a virtualization-aware environment in which blade, rack servers, and virtual machines are interconnected using the same mechanisms. The Cisco UCS 6248UP is a 1-RU fabric interconnect that features up to 48 universal ports that can support 80 Gigabit Ethernet, Fibre Channel over Ethernet, or native Fibre Channel connectivity.

Figure 6 Cisco UCS 6248UP 20-Port Fabric (front and back)

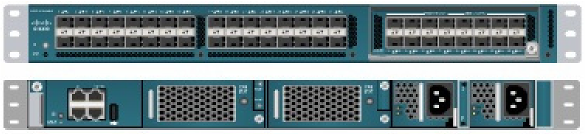

Cisco UCS Manager

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

Some of the key elements managed by Cisco UCS Manager include:

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

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

With model-based management, administrators manipulate a desired system configuration and associate a model's policy driven service profiles with hardware resources, and the system configures itself to match requirements. This automation accelerates provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations. This approach represents a radical simplification compared to traditional systems, reducing capital expenditures (CAPEX) and operating expenses (OPEX) while increasing business agility, simplifying and accelerating deployment, and improving performance.

Cisco UCS Service Profiles

Figure 7 show the traditional provisioning approach.

Figure 7 Compute, LAN, SAN Provisioning via Traditional Provisional Approach

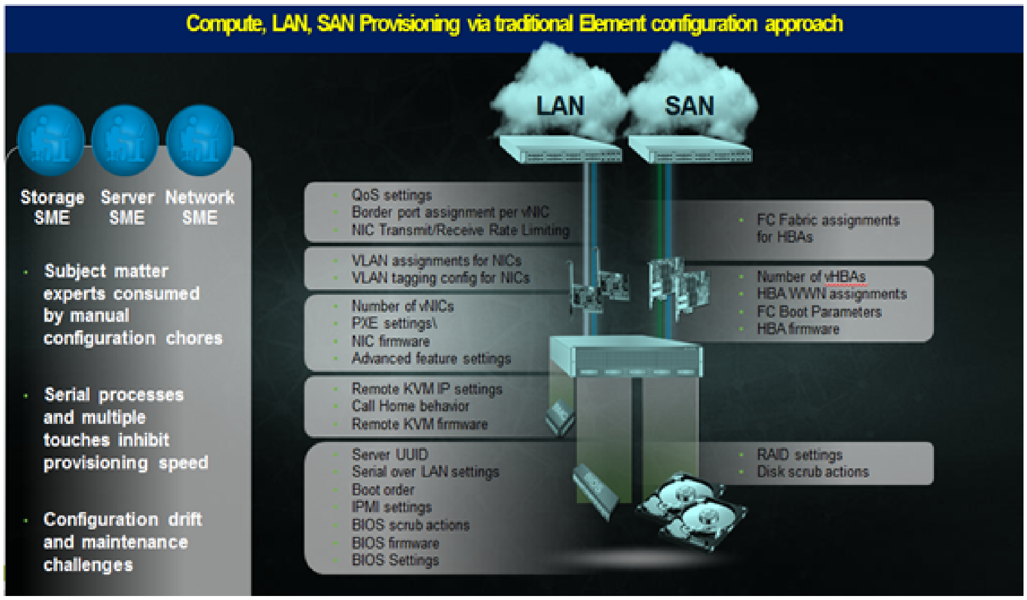

Compute, LAN, SAN Provisioning via traditional Element configuration approach

- Subject matter experts consumed by manual configuration chores
- Serial processes and multiple touches inhibit provisioning speed
- Configuration drift and maintenance challenges

- QoS settings
- Border port assignment per vNIC
- NIC Transmit/Receive Rate Limiting

- VLAN assignments for NICs
- VLAN tagging config for NICs

- Number of vNICs
- PXE settings\
- NIC firmware
- Advanced feature settings

- Remote KVM IP settings
- Call Home behavior
- Remote KVM firmware

- Server UUID
- Serial over LAN settings
- Boot order
- IPMI settings
- BIOS scrub actions
- BIOS firmware
- BIOS Settings

- FC Fabric assignments for HBAs

- Number of vHBAs
- HBA WWN assignments
- FC Boot Parameters
- HBA firmware

- RAID settings
- Disk scrub actions

A server's identity is made up of many properties such as UUID, boot order, IPMI settings, BIOS firmware, BIOS settings, etc. There are many areas that need to be configured to give the server its identity and make it unique from every other server within your data center. Some of these parameters are kept in the hardware of the server itself (like BIOS firmware version, BIOS settings, boot order, FC boot settings, etc.) while some settings are kept on your network and storage switches (like VLAN assignments, FC fabric assignments, QoS settings, ACLs, etc.).

Lengthy deployment cycles

- Every deployment requires coordination among server, storage, and network teams
- Need to ensure correct firmware & settings for hardware components
- Need appropriate LAN and SAN connectivity

Response time to business needs

- Tedious deployment process
- Manual, error prone processes, that are difficult to automate
- High OPEX costs, outages caused by human errors

Limited OS and application mobility

- Storage and network settings tied to physical ports and adapter identities
- Static infrastructure leads to over-provisioning, higher OPEX costs

Cisco Unified Computing System has uniquely addressed these challenges with the introduction of service profiles that enables integrated, policy based infrastructure management. UCS Service Profiles hold the DNA for nearly all configurable parameters required to set up a physical server. A set of user defined policies (rules) allow quick, consistent, repeatable, and secure deployments of UCS servers.

Figure 8 Cisco UCS Service Profiles—Integration and Management

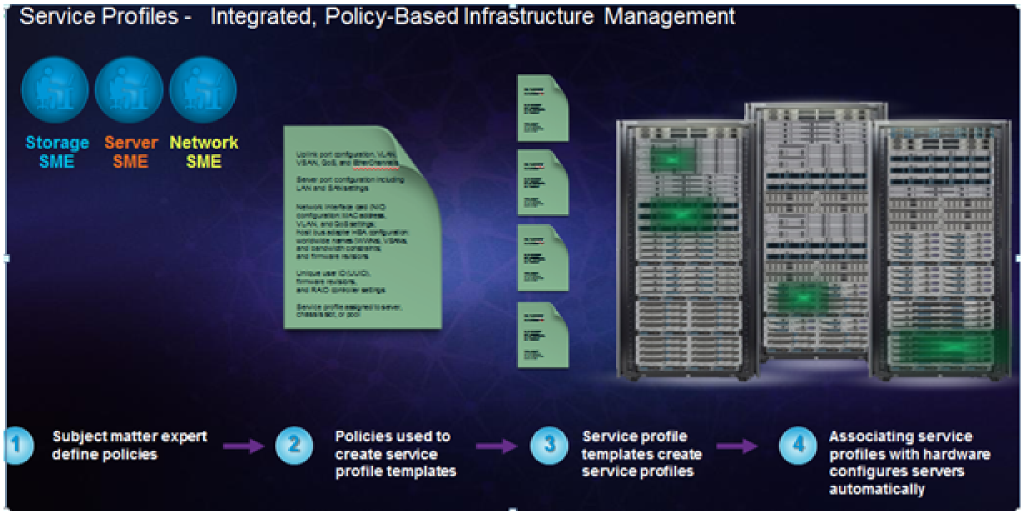
Service Profiles - Integrated, Policy-Based Infrastructure Management
Storage SME
Server SME
Network SME
1 Subject matter expert define policies
2 Policies used to create service profile templates
3 Service profile templates create service profiles
4 Associating service profiles with hardware configures servers automatically

Cisco UCS Service Profiles contain values for a server's property settings, including virtual network interface cards (vNICs), MAC addresses, boot policies, firmware policies, fabric connectivity, external management, and high availability information. By abstracting these settings from the physical server into a Cisco Service Profile, the Service Profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, Service Profiles can, at any time, be migrated from one physical server to another. This logical abstraction of the server personality separates the dependency of the hardware type or model and is a result of Cisco's unified fabric model (rather than overlaying software tools on top).

This innovation is still unique in the industry despite competitors claiming to offer similar functionality. In most cases, these vendors must rely on several different methods and interfaces to configure these server settings. Furthermore, Cisco is the only hardware provider to offer a truly unified management platform, with UCS Service Profiles and hardware abstraction capabilities extending to both blade and rack servers.

Some of key features and benefits of UCS service profiles are discussed below.

Service Profiles and Templates

A service profile contains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity. The Cisco UCS Manager provisions servers utilizing service profiles. The UCS Manager implements a role-based and policy-based management focused on service profiles and templates. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, enabling flexibility in the use of system resources.

Service profile templates are stored in the Cisco UCS 6200 Series Fabric Interconnects for reuse by server, network, and storage administrators. Service profile templates consist of server requirements and the associated LAN and SAN connectivity. 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.

The Cisco UCS Manager can deploy the service profile on any physical server at any time. When a service profile is deployed to a server, the Cisco UCS Manager automatically configures the server, adapters, Fabric Extenders, and Fabric Interconnects to match the configuration specified in the service profile. A service profile template parameterizes the UIDs that differentiate between server instances.

This automation of device configuration reduces the number of manual steps required to configure servers, Network Interface Cards (NICs), Host Bus Adapters (HBAs), and LAN and SAN switches.

Programmatically Deploying Server Resources

Cisco UCS Manager provides centralized management capabilities, creates a unified management domain, and serves as the central nervous system of the Cisco Unified Computing System. Cisco UCS Manager is embedded device management software that manages the system from end-to-end as a single logical entity through an intuitive GUI, CLI, or XML API. Cisco UCS Manager implements role- and policy-based management using service profiles and templates. This construct improves IT productivity and business agility. Now infrastructure can be provisioned in minutes instead of days, shifting IT's focus from maintenance to strategic initiatives.

Dynamic Provisioning

Cisco UCS resources are abstract in the sense that their identity, I/O configuration, MAC addresses and WWNs, firmware versions, BIOS boot order, and network attributes (including QoS settings, ACLs, pin groups, and threshold policies) all are programmable using a just-in-time deployment model.. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, 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.

Cisco Nexus 5548UP Switch

The Cisco Nexus 5548UP is a 1RU 1 Gigabit and 10 Gigabit Ethernet switch offering up to 960 gigabits per second throughput and scaling up to 48 ports. It offers 32 1/10 Gigabit Ethernet fixed enhanced Small Form-Factor Pluggable (SFP+) Ethernet/FCoE or 1/2/4/8-Gbps native FC unified ports and three expansion slots. These slots have a combination of Ethernet/FCoE and native FC ports.

Figure 9 Cisco Nexus 5548UP Switch

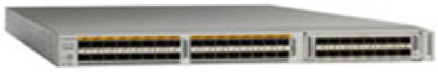

The Cisco Nexus 5548UP Switch delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages,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

Architectural Flexibility

- Unified ports that support traditional Ethernet, Fiber Channel (FC),and Fiber Channel over Ethernet (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 IEEE 802.1AE
- Offers converged Fabric extensibility, based on emerging standard IEEE 802.1BR, with Fabric Extender (FEX) Technology portfolio, including:
  - Cisco Nexus 2000 FEX

  - Adapter FEX
  - VM-FEX

Infrastructure Simplicity

- 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, FC, RoE, and IBoE
- Reduces management points with FEX Technology

Business Agility

- Meets 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 1 -rack-unit, 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 Fiber Channel, and Ethernet or FCoE
- Throughput of up to 960 Gbps

NetApp Storage Technologies and Benefits

NetApp storage platform can handle different type of files and data from various sources-including user files, e-mail, and databases. Data ONTAP is the fundamental NetApp software platform that runs on all NetApp storage systems. Data ONTAP is a highly optimized, scalable operating system that supports mixed NAS and SAN environments and a range of protocols, including Fiber Channel, iSCSI, FCoE, NFS, and CIFS. The platform includes the Write Anywhere File Layout (WAFL®) file system and storage virtualization capabilities. By leveraging the Data ONTAP platform, the NetApp Unified Storage Architecture offers the flexibility to manage, support, and scale to different business environments by using a common knowledge base and tools. This architecture enables users to collect, distribute, and manage data from all locations and applications at the same time. This allows the investment to scale by standardizing processes, cutting management time, and increasing availability. Figure 10 shows the various NetApp Unified Storage Architecture platforms.

Figure 10 NetApp Unified Storage Architecture Platforms

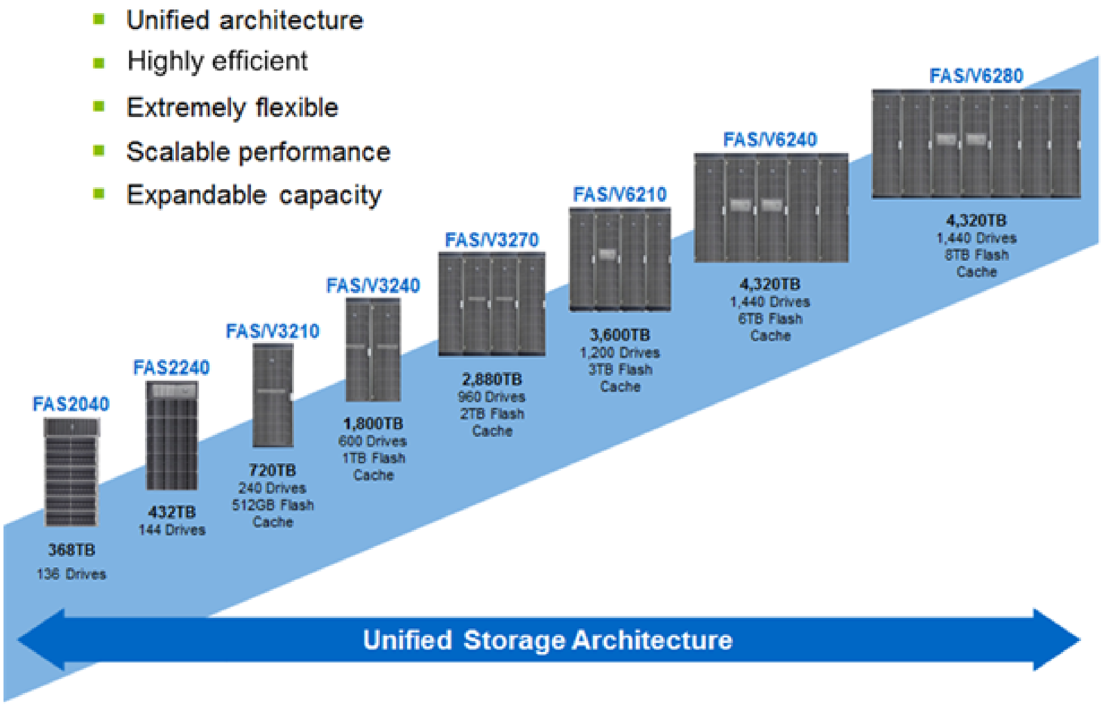

The NetApp storage hardware platform used in this solution is the FAS3270A. The FAS3200 series is an excellent platform for primary and secondary storage for an Oracle Database 11g R2 GRID Infrastructure deployment.

A number of NetApp tools and enhancements are available to augment the storage platform. These tools assist in deployment, backup, recovery, replication, management, and data protection. This solution makes use of a subset of these tools and enhancements.

Storage Architecture

The storage design for any solution is a critical element that is typically responsible for a large percentage of the solution's overall cost, performance, and agility.

The basic architecture of the storage system's software is shown in the figure below. A collection of tightly coupled processing modules handles CIFS, FCP, FCoE, HTTP, iSCSI, and NFS requests. A request starts in the network driver and moves up through network protocol layers and the file system, eventually generating disk I/O, if necessary. When the file system finishes the request, it sends a reply back to the network. The administrative layer at the top supports a command line interface (CLI) similar to UNIX® that monitors and controls the modules below. In addition to the modules shown, a simple real-time kernel provides basic services such as process creation, memory allocation, message passing, and interrupt handling.

The networking layer is derived from the same Berkeley code used by most UNIX systems, with modifications made to communicate efficiently with the storage appliance's file system. The storage appliance provides transport-independent seamless data access using block- and file-level protocols from the same platform. The storage appliance provides block-level data access over an FC SAN fabric using FCP and over an IP-based Ethernet network using iSCSI. File access protocols such as NFS, CIFS, HTTP, or FTP provide file-level access over an IP-based Ethernet network.

Figure 11 *NetApp Controller Storage Architecture*

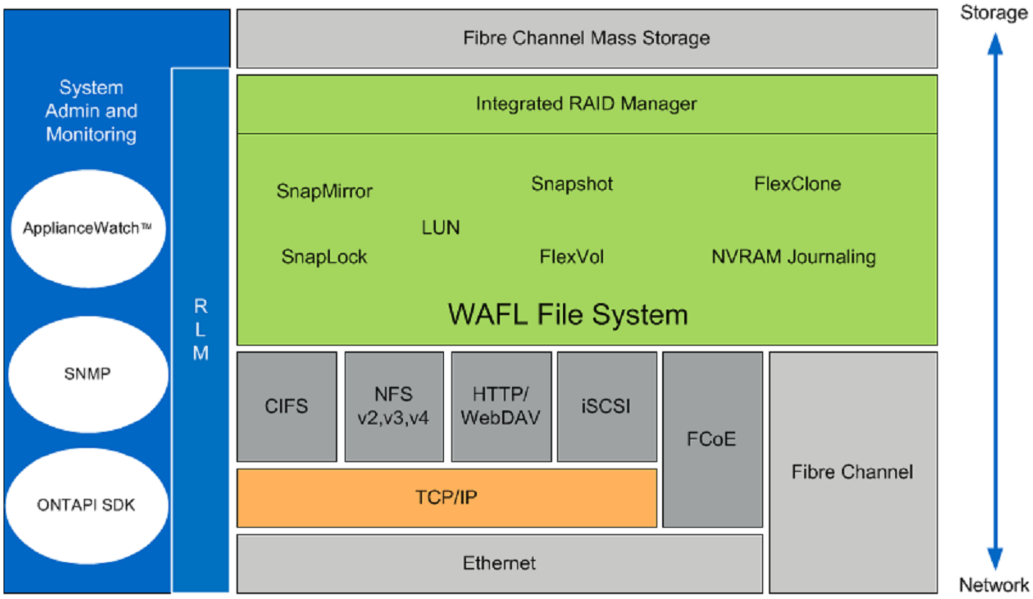

RAID-DP

RAID-DP® is NetApp's implementation of double-parity RAID 6, which is an extension of NetApp's original Data ONTAP WAFL® RAID 4 design. Unlike other RAID technologies, RAID-DP provides the ability to achieve a higher level of data protection without any performance impact, while consuming a minimal amount of storage. For more information on RAID-DP, see http://www.netapp.com/us/products/platform-os/raid-dp.html.

Snapshot

NetApp Snapshot technology provides zero-cost, near-instantaneous backup and point-in-time copies of the volume or LUN by preserving the Data ONTAP WAFL consistency points.

Creating Snapshot copies incurs minimal performance effect because data is never moved, as it is with other copy-out technologies. The cost for Snapshot copies is at the rate of block-level changes and not 100% for each backup, as it is with mirror copies. Using Snapshot can result in savings in storage cost for backup and restore purposes and opens up a number of efficient data management possibilities.

FlexVol

NetApp® FlexVol® storage-virtualization technology enables you to respond to changing storage needs fast, lower your overhead, avoid capital expenses, and reduce disruption and risk. FlexVol technology aggregates physical storage in virtual storage pools, so you can create and resize virtual volumes as your application needs change.

With FlexVol you can improve-even double-the utilization of your existing storage and save the expense of acquiring more disk space. In addition to increasing storage efficiency, you can improve I/O performance and reduce bottlenecks by distributing volumes across all available disk drives.

NetApp OnCommand System Manager 2.0

System Manager is a powerful management tool for NetApp storage that allows administrators to manage a single NetApp storage system as well as clusters, quickly and easily.

Some of the benefits of the System Manager Tool are:

- Easy to install
- Easy to manage from a Web browser
- Does not require storage expertise
- Increases storage productivity and response time
- Cost effective
- Leverages storage efficiency features such as thin provisioning and compression

Oracle Database 11g R2 RAC

Oracle Database 11g Release 2 provides the foundation for IT to successfully deliver more information with higher quality of service, reduce the risk of change within IT, and make more efficient use of IT budgets.

Oracle Database 11g R2 Enterprise Edition provides industry-leading performance, scalability, security, and reliability on a choice of clustered or single-servers with a wide range of options to meet user needs. Grid computing relieves users from concerns about where data resides and which computer processes their requests. Users request information or computation and have it delivered - as much as they want, whenever they want. For a DBA, the grid is about resource allocation, information sharing, and high availability. Oracle Database with Real Application Clusters provide the infrastructure for your database grid. Automatic Storage Management provides the infrastructure for a storage grid. Oracle Enterprise Manager Grid Control provides you with holistic management of your grid.

Oracle Database 11g Direct NFS Client

Direct NFS client is an Oracle developed, integrated, and optimized client that runs in user space rather than within the operating system kernel. This architecture provides for enhanced scalability and performance over traditional NFS v3 clients. Unlike traditional NFS implementations, Oracle supports asynchronous I/O across all operating system environments with Direct NFS client. In addition, performance and scalability are dramatically improved with its automatic link aggregation feature. This allows the client to scale across as many as four individual network pathways with the added benefit of improved resiliency when Network connectivity is occasionally compromised. It also allows Direct NFS clients to achieve near block level Performance. For more information on Direct NFS Client comparison to block protocols, see http://media.netapp.com/documents/tr-3700.pdf.

Design Topology

This section presents physical and logical high-level design considerations for Cisco UCS networking and computing on NetApp storage for Oracle Database 11g R2 RAC deployments.

Hardware and Software Used for this Solution

Table 1 shows the software and hardware used for Oracle Database 11g R2 GRID Infrastructure with RAC Option Deployment.

Table 1 Software and Hardware for Oracle Database 11g R2 GRID Infrastructure with RAC Option Deployment

Hardware	Quantity
Servers	
Cisco UCS 6248UP fabric interconnects	2 (configured as an active-active pair)
Cisco UCS 5108 B-Series blade chassis	2
Cisco UCS 2208 B-Series blade Fabric Extender modules	4 (2 per Blade UCS Blade Chassis)
Cisco UCS B200 M3 B-Series blade servers: 2-socket E5-2690 2.9GHz CPU & 128GB Memory & 1 x Cisco UCS 1240 virtual interface card VIC	4 nodes for RAC used for this solution
Network	
Cisco Nexus 5548UP switches	2
Storage	
NetApp FAS3270 storage controllers: 2 x 1TB Flash Cache per controller & 2 x Dual-port 10GbE PCIe adapters per controller	2 nodes, configured as an active-active pair
NetApp DS4243 disk shelves: 24 x 600GB SAS drives per shelf	4 shelf total, configured with 2 shelf per Controller

Software	Version
Redhat Enterprise Linux 5.8 64-bit	kernel-2.6.18-308.el5
Oracle Database 11gR2 GRID	11.2.0.3
Oracle Database 11gR2 Database	11.2.0.3
Cisco UCS Manager	2.1.(1a)
Cisco NXOS for Nexus 5548UP	5.0(3)N2(1)
NetApp Data ONTAP	8.1.2
NetApp OnCommand system Manager	2.1

Cisco UCS Networking and NetApp NFS Storage Topology

This section explains Cisco UCS networking and computing design considerations when deploying Oracle Database 11g R2 RAC in an NFS Storage Design. In this design, the NFS traffic is isolated from the regular management and application data network using the same Cisco UCS infrastructure by defining logical VLAN networks to provide better data security. Figure below, presents a detailed view of the physical topology, and some of the main components of Cisco UCS in an NFS network design.

Figure 12 Cisco UCS Networking and NFS Storage Network Topology

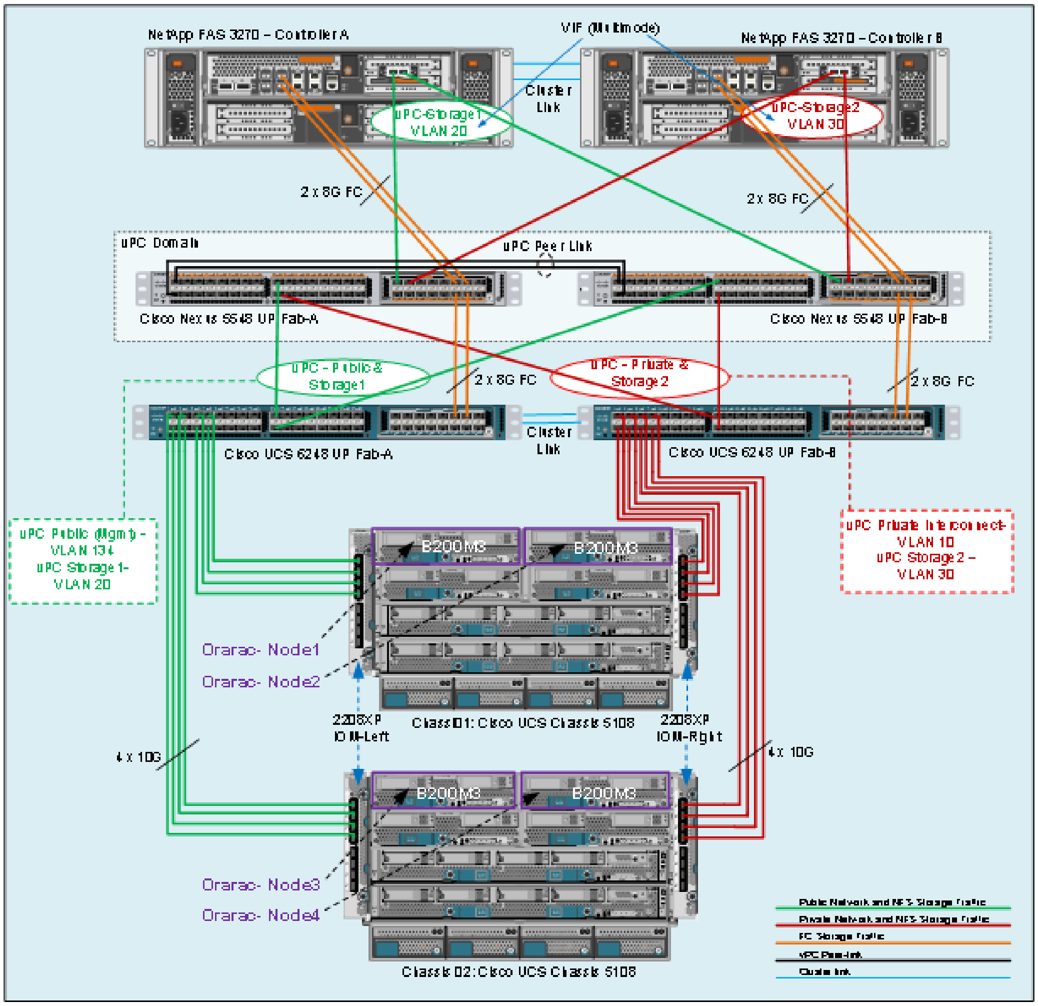

As shown above, a pair of Cisco UCS 6248UP fabric interconnects carries both storage and network traffic from the blades with the help of Cisco Nexus 5548UP switch. The 8GB Fiber channel traffic (shown in green lines) leaves the UCS Fabrics through Nexus5548 Switches to NetApp Array to boot the Operating system from SAN environment. This is a typical configuration that can be deployed in a customer's environment or customers can choose to boot other methods such as local disk boot as per business requirements. Boot from SAN is a recommended option to enable stateless computing for UCS servers. From the block diagram it looks like all 8GB FC links are blocking? As described it appears there is no valid path to the SAN.

Both the fabric interconnect and the Cisco Nexus 5548UP switch are clustered with the peer link between them to provide high availability. Two virtual Port Channels (vPCs) are configured to provide public network, private network and storage access paths for the blades to northbound switches. Each vPC has VLANs created for application network data, NFS storage data, and management data paths. For more information about vPC configuration on the Cisco Nexus 5548UP Switch, see http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html.

As illustrated in picture above, Eight (Four per chassis) links go to Fabric Interconnect "A" (ports 1 through 8). Similarly, Eight links go to Fabric Interconnect B. Fabric Interconnect-A links are used for Oracle Public network & NFS Storage Network traffic and Fabric Interconnect-B links are used for Oracle private interconnect traffic and NFS Storage network traffic.

Note For Oracle RAC configuration on UCS, we recommend to keep all private interconnects local on a single Fabric interconnect with NIC failover enabled. In such case, the private traffic will stay local to that fabric interconnect and will not be routed via northbound network switch. In other words, all inter blade (or RAC node private) communication will be resolved locally at the fabric interconnect and this significantly reduces latency for Oracle Cache Fusion traffic.

Cisco UCS Manager Configuration Overview

The following are the high-level steps involved for a Cisco UCS configuration:

1. Configuring Fabric Interconnects for Chassis and Blade Discovery
   a. Configuring Server Ports
2. Configuring LAN and SAN on UCS Manager
   a. Configure and Enable Ethernet LAN uplink Ports
   b. Configure and Enable FC SAN uplink Ports
   c. Configure VLAN
   d. Configure VSAN
3. Configuring UUID, MAC, WWWN and WWPN Pool
   a. UUID Pool Creation
   b. IP Pool and MAC Pool Creation
   c. WWNN Pool and WWPN Pool Creation
4. Configuring vNIC and vHBA Template
   a. Create vNIC templates
   b. Create Public vNIC template
   c. Create Private vNIC template
   d. Create Storage vNIC template
   e. Create HBA templates
5. Configuring Ethernet Uplink Port-Channels
6. Create Server Boot Policy for SAN Boot

Details for each step are discussed in subsequent sections below.

Configuring Fabric Interconnects for Blade Discovery

Cisco UCS 6248 UP Fabric Interconnects are configured for redundancy. It provides resiliency in case of failures. The first step is to establish connectivity between the blades and fabric interconnects.

Configuring Server ports

1. Click Equipment.
2. Click Fabric Interconnects.
3. Click Fabric Interconnect A.
4. Click Fixed Module.
5. Click Ethernet Ports and select the desired number of ports.
6. Right-click "Configure as Server Port" as show below.

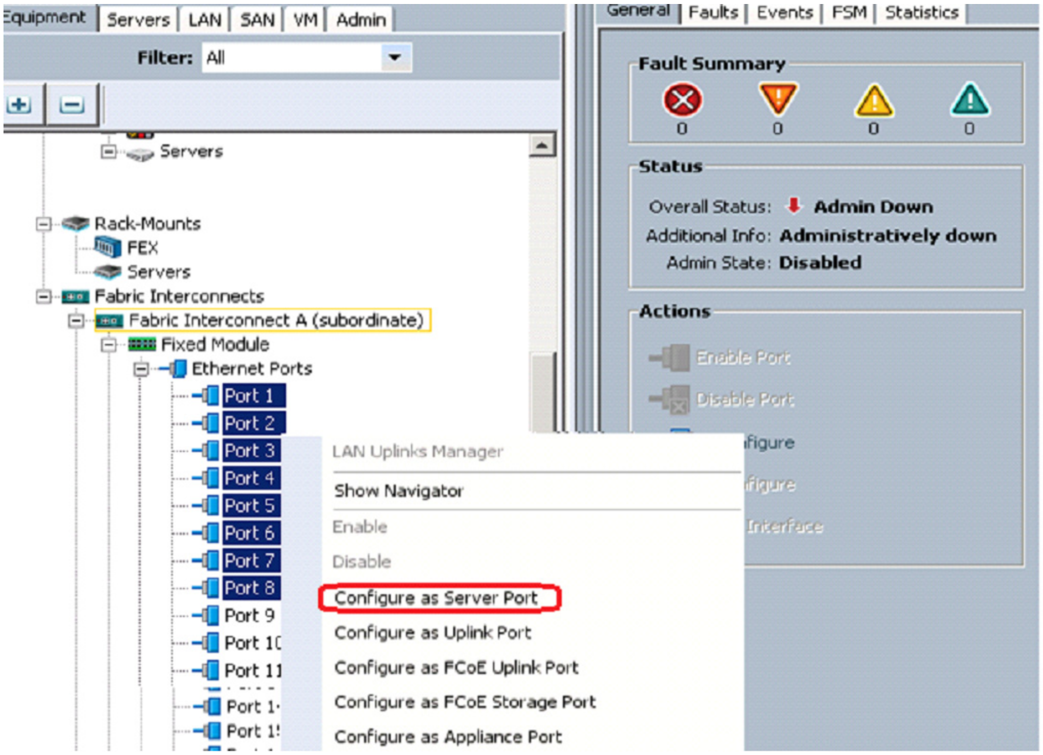

Filter: All

Overall Status: Admin Down
Additional Info: Administratively down
Admin State: Disabled

Configuring LAN and SAN on Cisco UCS Manager

Perform the LAN and SAN configuration steps in the Cisco UCS Manager as shown in the following section.

Configure and Enable Ethernet LAN uplink Ports

1. From the Equipment tab click Fabric Interconnects.
2. Click Fabric Interconnect A.
3. Click Fixed Module.

4. From the Ethernet Ports menu, select the desired number of ports and right-click Configure as Uplink Port.

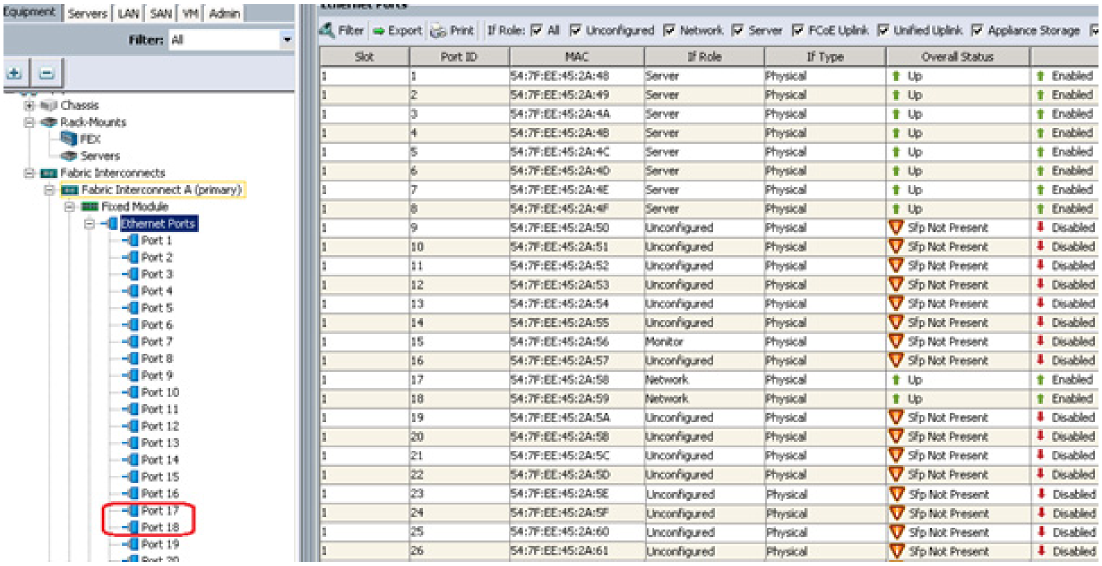

Slot	Port ID	MAC	If Role	If Type	Overall Status	
1	1	54:7F:EE:45:2A:48	Server	Physical	Up	Enabled
1	2	54:7F:EE:45:2A:49	Server	Physical	Up	Enabled
1	3	54:7F:EE:45:2A:4A	Server	Physical	Up	Enabled
1	4	54:7F:EE:45:2A:4B	Server	Physical	Up	Enabled
1	5	54:7F:EE:45:2A:4C	Server	Physical	Up	Enabled
1	6	54:7F:EE:45:2A:4D	Server	Physical	Up	Enabled
1	7	54:7F:EE:45:2A:4E	Server	Physical	Up	Enabled
1	8	54:7F:EE:45:2A:4F	Server	Physical	Up	Enabled
1	9	54:7F:EE:45:2A:50	Unconfigured	Physical	Sfp Not Present	Disabled
1	10	54:7F:EE:45:2A:51	Unconfigured	Physical	Sfp Not Present	Disabled
1	11	54:7F:EE:45:2A:52	Unconfigured	Physical	Sfp Not Present	Disabled
1	12	54:7F:EE:45:2A:53	Unconfigured	Physical	Sfp Not Present	Disabled
1	13	54:7F:EE:45:2A:54	Unconfigured	Physical	Sfp Not Present	Disabled
1	14	54:7F:EE:45:2A:55	Unconfigured	Physical	Sfp Not Present	Disabled
1	15	54:7F:EE:45:2A:56	Monitor	Physical	Sfp Not Present	Disabled
1	16	54:7F:EE:45:2A:57	Unconfigured	Physical	Sfp Not Present	Disabled
1	17	54:7F:EE:45:2A:58	Network	Physical	Up	Enabled
1	18	54:7F:EE:45:2A:59	Network	Physical	Up	Enabled
1	19	54:7F:EE:45:2A:5A	Unconfigured	Physical	Sfp Not Present	Disabled
1	20	54:7F:EE:45:2A:5B	Unconfigured	Physical	Sfp Not Present	Disabled
1	21	54:7F:EE:45:2A:5C	Unconfigured	Physical	Sfp Not Present	Disabled
1	22	54:7F:EE:45:2A:5D	Unconfigured	Physical	Sfp Not Present	Disabled
1	23	54:7F:EE:45:2A:5E	Unconfigured	Physical	Sfp Not Present	Disabled
1	24	54:7F:EE:45:2A:5F	Unconfigured	Physical	Sfp Not Present	Disabled
1	25	54:7F:EE:45:2A:60	Unconfigured	Physical	Sfp Not Present	Disabled
1	26	54:7F:EE:45:2A:61	Unconfigured	Physical	Sfp Not Present	Disabled

As shown in the screenshot above, we selected Port 17 and 18 on Fabric interconnect A and configured them as Ethernet Uplink ports. Repeat the same step on Fabric interconnect B to configure Port 17 and 18 as Ethernet uplink ports. We will use these ports to create Port-channels in later sections.

Configure and Enable FC SAN Uplink Ports

1. From the Equipment tab, click Fabric Interconnects.
2. Click Fabric Interconnect A.
3. Configure the Unified Ports menu.

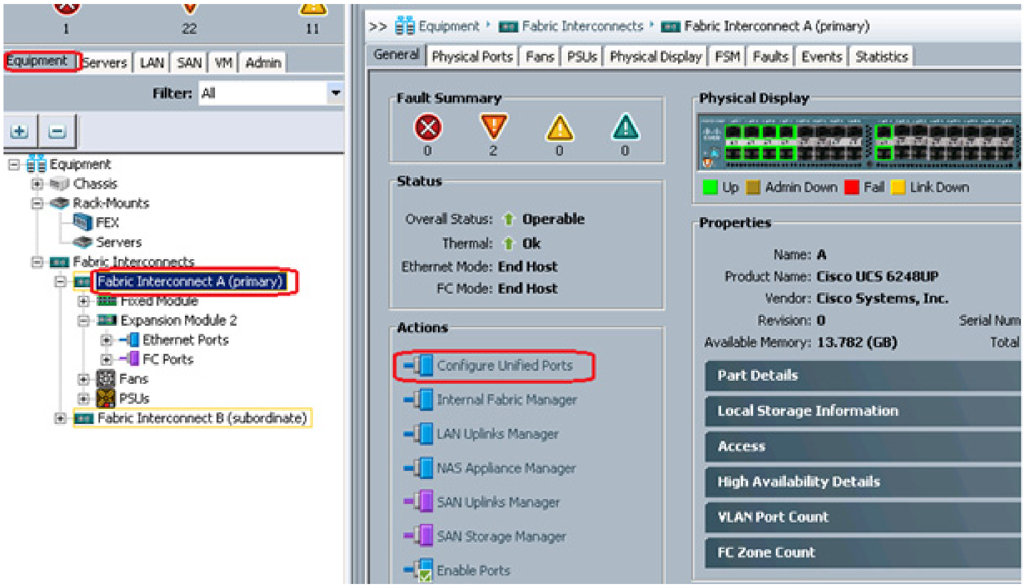

Status

Overall Status: Operable
Thermal: Ok
Ethernet Mode: End Host
FC Mode: End Host

Properties

Name: A
Product Name: Cisco UCS 6248UP
Vendor: Cisco Systems, Inc.
Revision: 0
Serial Num
Available Memory: 13.782 (GB)
Total

The Configure Expansion Module Ports displays.

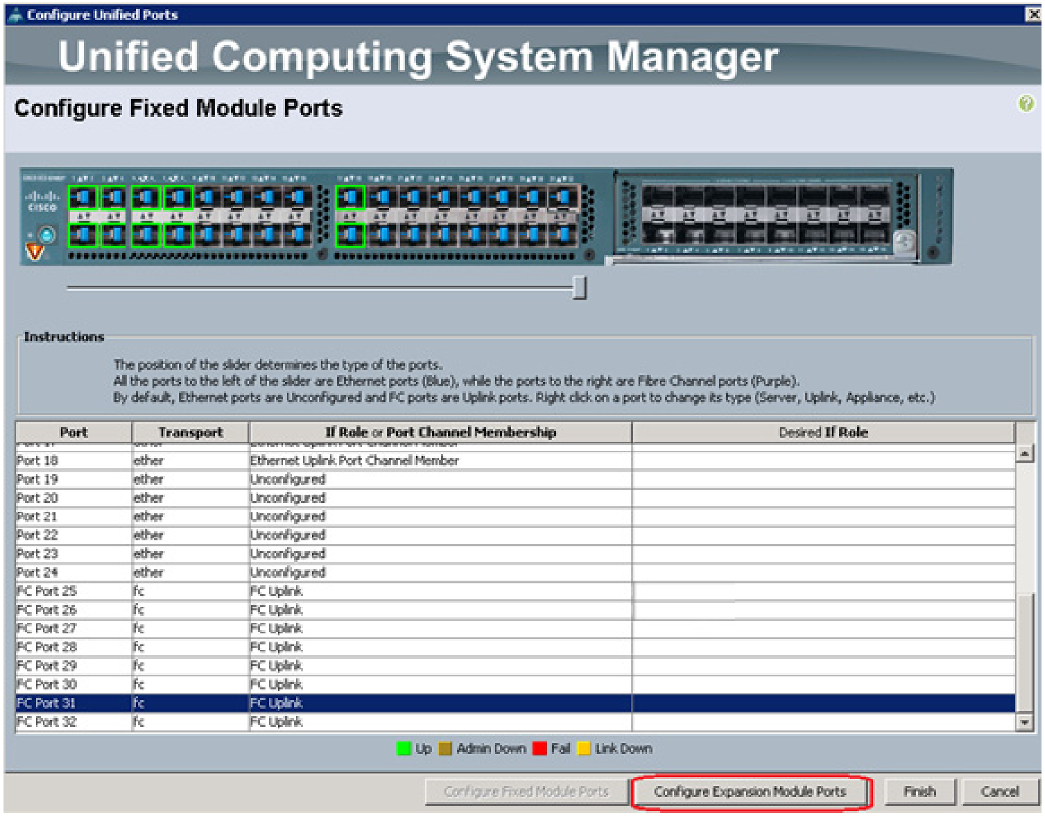
Configure Unified Ports
Unified Computing System Manager
Configure Fixed Module Ports
Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Ethernet ports (Blue), while the ports to the right are Fibre Channel ports (Purple).
By default, Ethernet ports are Unconfigured and FC ports are Uplink ports. Right click on a port to change its type (Server, Uplink, Appliance, etc.)

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 18	ether	Ethernet Uplink Port Channel Member	
Port 19	ether	Unconfigured	
Port 20	ether	Unconfigured	
Port 21	ether	Unconfigured	
Port 22	ether	Unconfigured	
Port 23	ether	Unconfigured	
Port 24	ether	Unconfigured	
FC Port 25	fc	FC Uplink	
FC Port 26	fc	FC Uplink	
FC Port 27	fc	FC Uplink	
FC Port 28	fc	FC Uplink	
FC Port 29	fc	FC Uplink	
FC Port 30	fc	FC Uplink	
FC Port 31	fc	FC Uplink	
FC Port 32	fc	FC Uplink	

🟩 Up 🟫 Admin Down 🟥 Fail 🟨 Link Down

4. Click "Configure Expansion Module Ports" button. The Configure Expansion Module Ports window displays.

XXXX

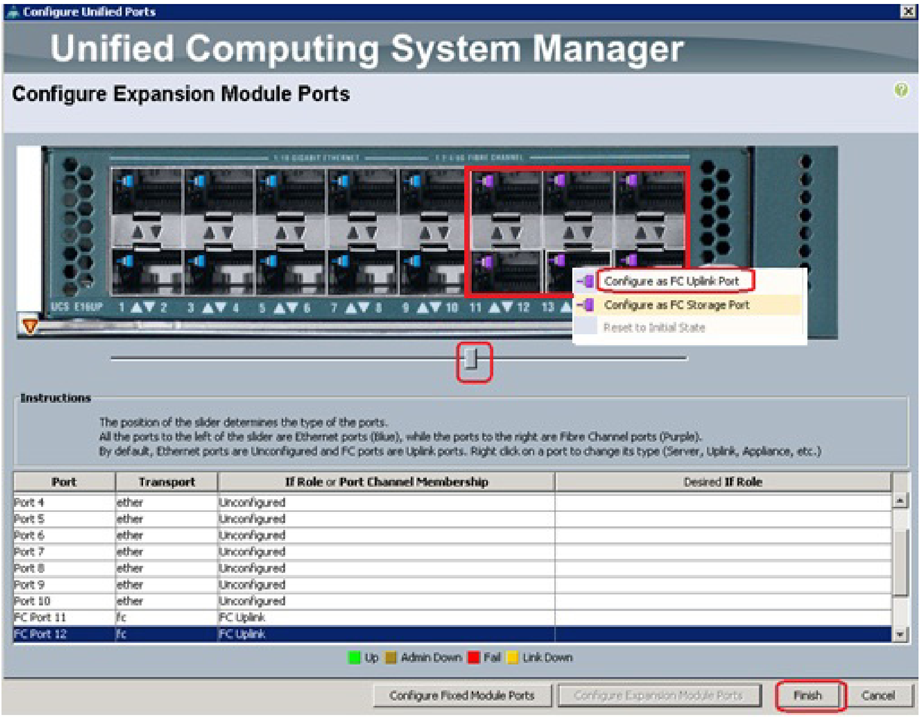
Configure Unified Ports
Unified Computing System Manager
Configure Expansion Module Ports
UCS E16UP
1 2 3 4 5 6 7 8 9 10 11 12 13
Configure as FC Uplink Port
Configure as FC Storage Port
Reset to Initial State

Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Ethernet ports (Blue), while the ports to the right are Fibre Channel ports (Purple).
By default, Ethernet ports are Unconfigured and FC ports are Uplink ports. Right click on a port to change its type (Server, Uplink, Appliance, etc.)

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 4	ether	Unconfigured	
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
FC Port 11	fc	FC Uplink	
FC Port 12	fc	FC Uplink	

🟩 Up 🟫 Admin Down 🟥 Fail 🟨 Link Down
Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

You can use the slider to select a set of ports. Right Click on the selected ports and click on "Configure as FC Uplink Port" and Click on Finish button to Save the configuration changes.

Repeat for same steps on Fabric Interconnect B switch to enable LAN uplink and FC uplink ports. Next step is to configure VLANs for this configuration. Before we get into vLAN configuration, here are a couple of Best practices for Oracle RAC configuraion.

Important Oracle RAC Best Practices and Recommendations for vLANs and vNIC Configuration

For Direct NFS clients running on Linux, best practices recommend always to use multipaths in separate subnets. If multiple paths are configured in the same subnet, the operating system invariably picks the first available path from the routing table. All traffic flows through this path and the load balancing and scaling do not work as expected. Please refer to Oracle metalink note 822481.1 for more details.

Note For this configuration, we created VLAN 20 and VLAN 30 for storage access.

Oracle Grid Infrastructure can activate a maximum of four private network adapters for availability and bandwidth requirements. In our testing, we observed that a single Cisco UCS 10GE private vNIC configured with failover did not require multiple vnics configuration from bandwidth and availability perspective. If you want to configure multiple vnics for your private interconnect ,we strongly recommend using a separate VLAN for each private vNIC. As a general best practice, it is a good idea to localize all private interconnect traffic to single fabric interconnect. For more information on Oracle HAIP, please refer to Oracle metalink note 1210883.1.

When you have decided on VLAN and vNICs, we are ready to configure vLANs for this setup.

Configure VLAN

1. In Cisco UCS manager, click LAN

2. Go to LAN Cloud > VLAN and right-click Create VLANs. In this solution, we need to create 4 VLANs: one for private (VLAN 10), one for public network (VLAN 134), and two more for storage traffic (VLAN 20 and 30). These four VLANS will be used in the vNIC templates discussed later.

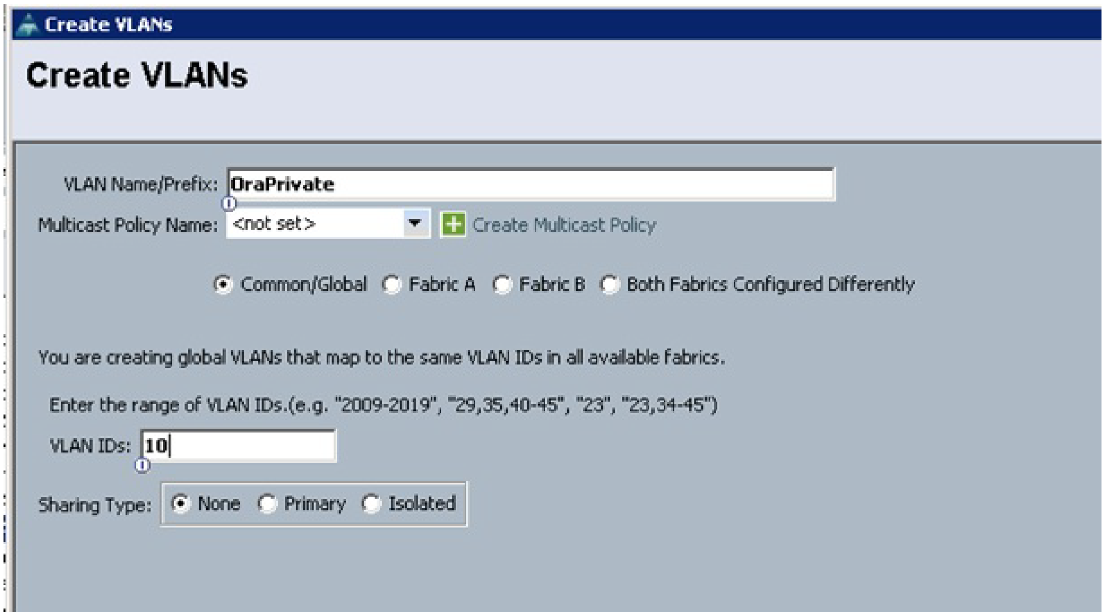

Create VLANs

VLAN Name/Prefix: OraPrivate

Multicast Policy Name: <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: 10

Sharing Type: ● None ○ Primary ○ Isolated

In the screenshot above, we have highlighted VLAN 10 creation for Private network. It is also very important that you create both VLANs as global across both fabric interconnects. This way, VLAN identity is maintained across the fabric interconnects in case of NIC failover.

Repeat the process for creating Public vlans & Storage vlans when using Oracle HAIP feature, you may have to configure additional vlans to be associated with additional vnics as well.

Here is the summary of VLANs once you complete VLAN creation.

- VLAN 10 for Oracle RAC private interconnect interfaces.
- VLAN 134 for public interfaces.
- VLAN 20 and VLAN 30 for storage access.

Note Even though private VLAN traffic stays local within UCS domain during normal operating conditions, it is necessary to configure entries for these private VLANS in northbound network switch. This will allow the switch to route interconnect traffic appropriately in case of partial link failures. These scenarios and traffic routing are discussed in details in later sections.

The figure below summarizes all the VLANs for Public and Private network and Storage access.

Figure 13 VLAN Summary

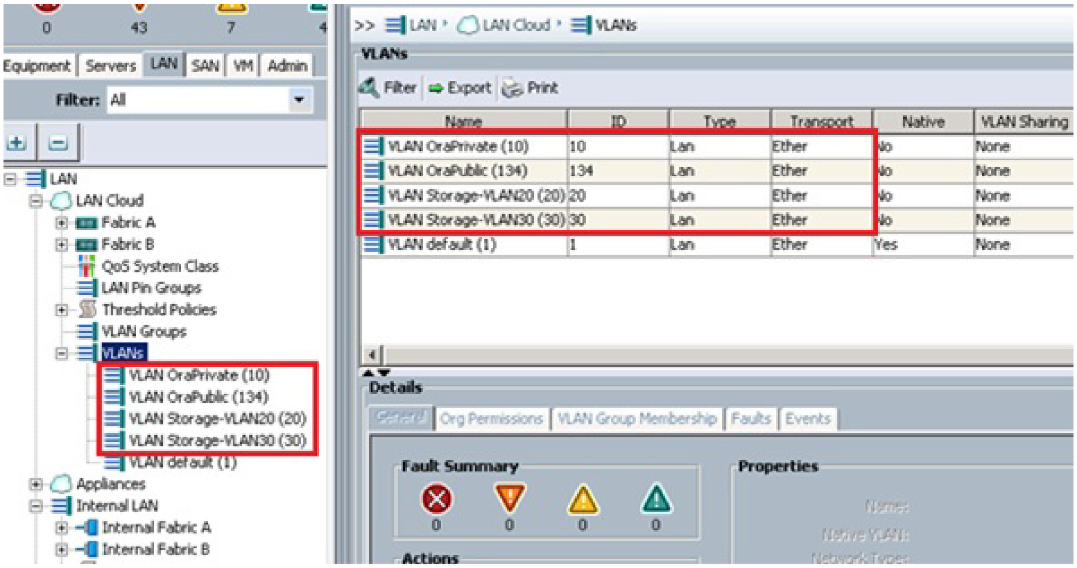

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN OraPrivate (10)	10	Lan	Ether	No	None
VLAN OraPublic (134)	134	Lan	Ether	No	None
VLAN Storage-VLAN20 (20)	20	Lan	Ether	No	None
VLAN Storage-VLAN30 (30)	30	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	Yes	None

Configure VSAN

In Cisco UCS manager, click SAN > SAN Cloud > VSANs and right-click to Create VSAN. In this study we created VSAN 25 for SAN Boot.

Figure 14 Configuring VSAN in Cisco UCS Manager

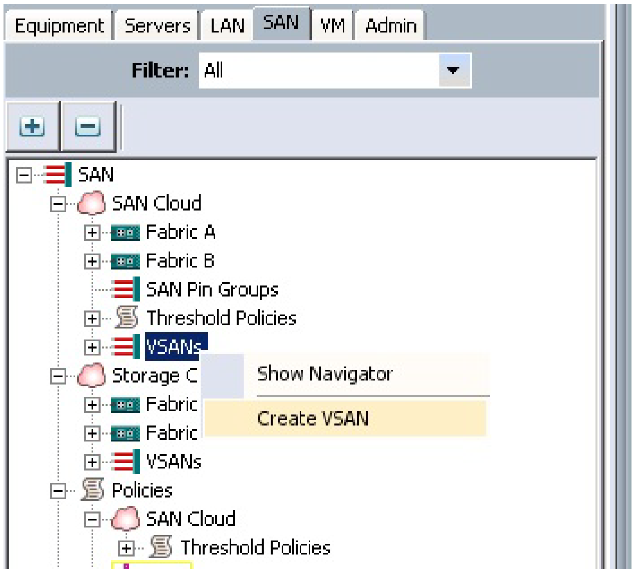

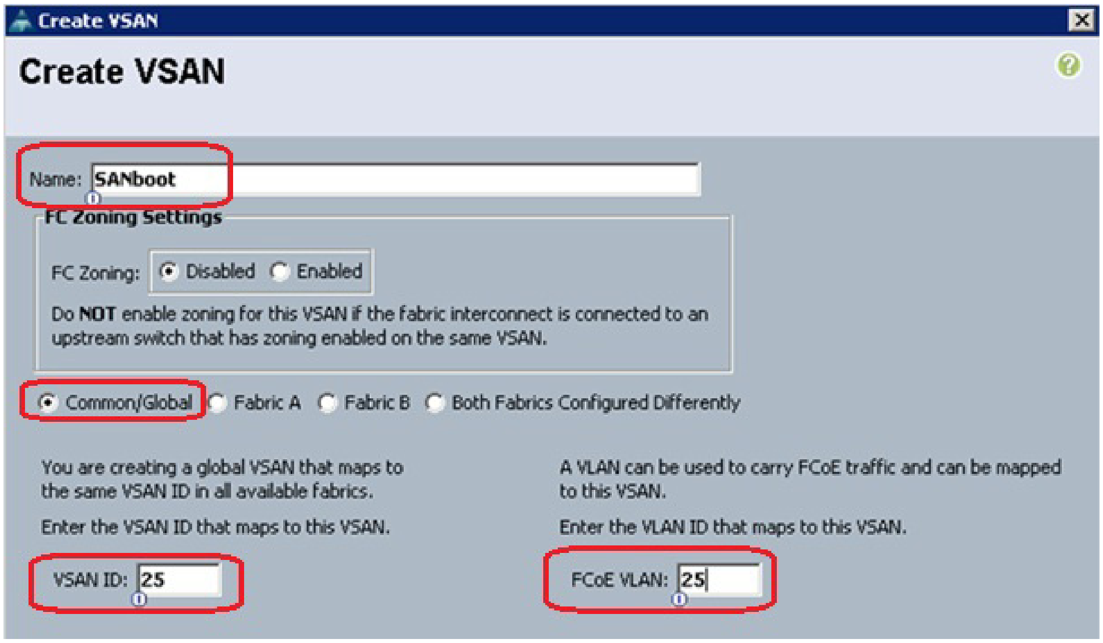

Create VSAN

Create VSAN

Name: SANboot

FC Zoning Settings

FC Zoning: ● Disabled ○ Enabled

Do **NOT** enable zoning for this VSAN if the fabric interconnect is connected to an upstream switch that has zoning enabled on the same VSAN.

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating a global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 25

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 25

In this study, we created VSAN Name as "SAN boot" and Selected "Common/Global" to create VSAN on both the Fabrics and specified VSAN ID as "25" and FCoE VLAN ID as "25".

Note If FCoE traffic for SAN Storage is not used, it is mandatory to specify VLAN ID.

Figure 15 VSAN Summary

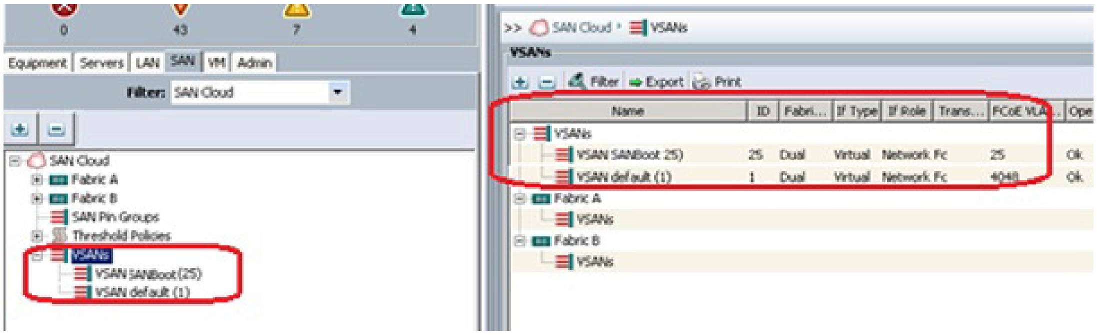

Name	ID	Fabri...	If Type	If Role	Trans...	FCoE VLA...	Ope
VSANs							
VSAN SANBoot 25)	25	Dual	Virtual	Network	Fc	25	Ok
VSAN default (1)	1	Dual	Virtual	Network	Fc	4048	Ok
Fabric A							
VSANs							
Fabric B							
VSANs							

Configure Pools

When VLANs and VSAN are created, we need to configure pools for UUID, MAC Addresses, Management IP and WWN.

UUID Pool Creation

In Cisco UCS Manager, click Servers > Pools > UUID Suffix Pools and right-click to "Create UUID Suffix Pool", create a new pool.

Figure 16 Create UUID Pools

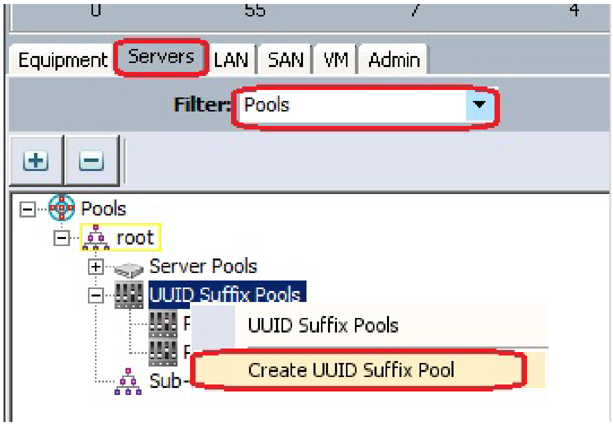

As shown in Figure 17, we created UUID Pool as OraFlex-UUID-Pools.

Figure 17 Post UUID Pool Creation

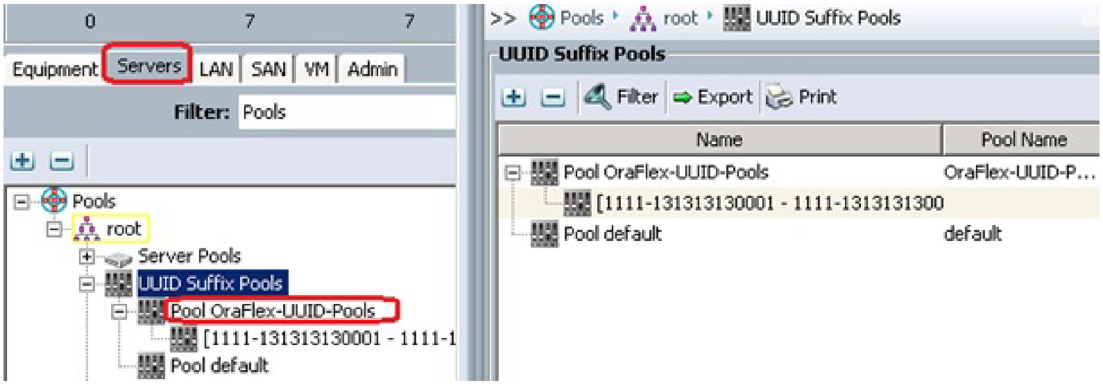

Name	Pool Name
Pool OraFlex-UUID-Pools	OraFlex-UUID-P...
[1111-131313130001 - 1111-1313131300	
Pool default	default

IP Pool and MAC Pool Creation

In Cisco UCS Manager, click LAN > Pools > IP Pools and right-click to "Create IP Pool Ext-mgmt". We created ext-mgmt IP pool as shown below. Next, click MAC Pools to "Create MAC Pools". We created OraFlex-MAC-Pools for the all vNIC MAC addresses.

Figure 18 Create IP Pool and MAC Pool

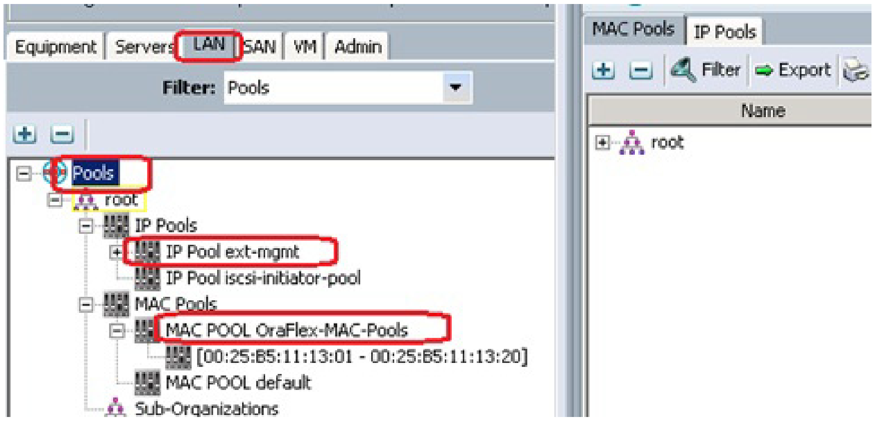

The IP pools will be used for console management, while MAC addresses for the vnics being carved out later.

WWNN Pool and WWPN pool Creation

In Cisco UCS Manager, click SAN > Pools > WWNN Pools and right-click to "Create WWNN Pools". Next, click WWPN Pools to "Create WWPN Pools". These WWNN and WWPN entries will be used for Boot from SAN configuration.

We created OraFlex-WWPN-Pools and OraFlex_WWNN-Pools as shown below.

Figure 19 Create WWNN and WWPN Pool

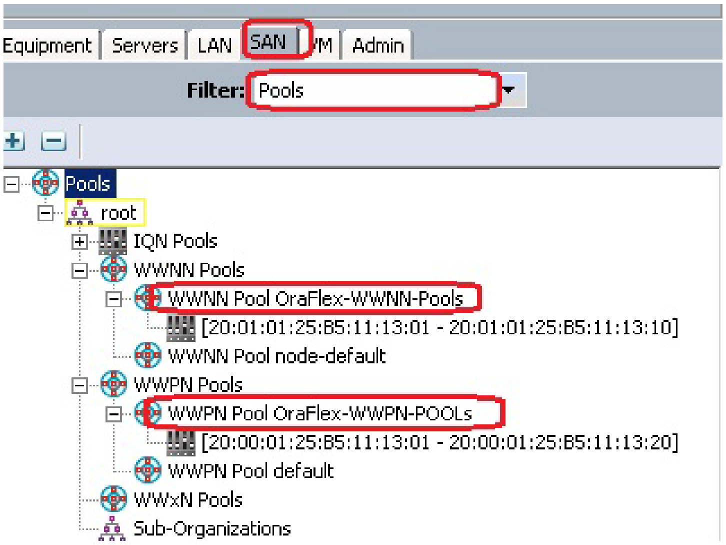
Equipment
Servers
LAN
SAN
VM
Admin
Filter: Pools
Pools
root
IQN Pools
WWNN Pools
WWNN Pool OraFlex-WWNN-Pools
[20:01:01:25:B5:11:13:01 - 20:01:01:25:B5:11:13:10]
WWNN Pool node-default
WWPN Pools
WWPN Pool OraFlex-WWPN-POOLs
[20:00:01:25:B5:11:13:01 - 20:00:01:25:B5:11:13:20]
WWPN Pool default
WWxN Pools
Sub-Organizations

Pool creation is complete. The next step is to create vNIC and vHBA templates.

Configure vNIC and vHBA Templates

Create a vNIC Template

In Cisco UCS Manager, click LAN > Policies > vNIC templates and right-click “Create vNIC Template.”

Figure 20 Create vNIC Template

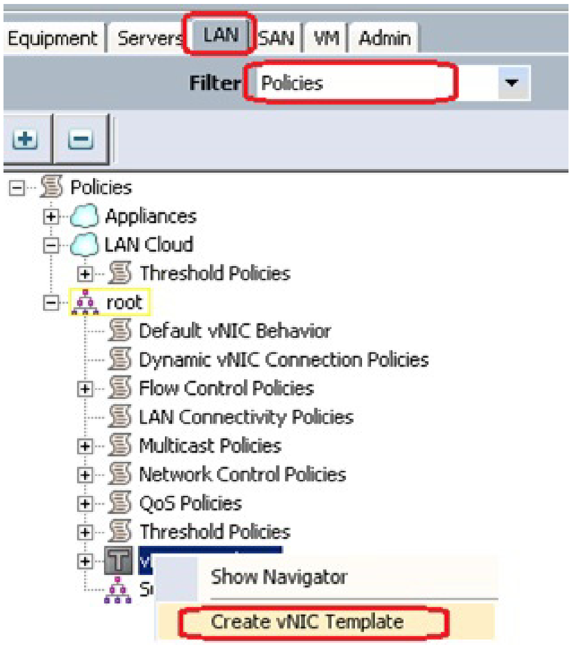

We created four vNIC templates for this Oracle RAC on FlexPod configuration: one for public network, one for private network and two for storage network. The private network is for Oracle RAC internal heartbeat and cache fusion traffic while Public network for external clients like middle tiers and ssh sessions to the Oracle hosts. The storage networks are for NFS data traffic access of Oracle DSS and OLTP database volumes.

Create Private vNIC Template

For Oracle private network vNIC template, we strongly recommend to set 9000 MTU. We also will pin this template to Fabric B with vNIC failover enabled. For private vNICs derived from this template, all communication among those private vNICs will stay local to Fabric Interconnect B. In case, of a failure on Fabric B, the appropriate vNICs will failover to Fabric A. As shown in the screenshot below, please make sure to use OraFlex-MAC-Pools as the MAC Pool for MAC addresses.

Figure 21 Create Private vNIC Template

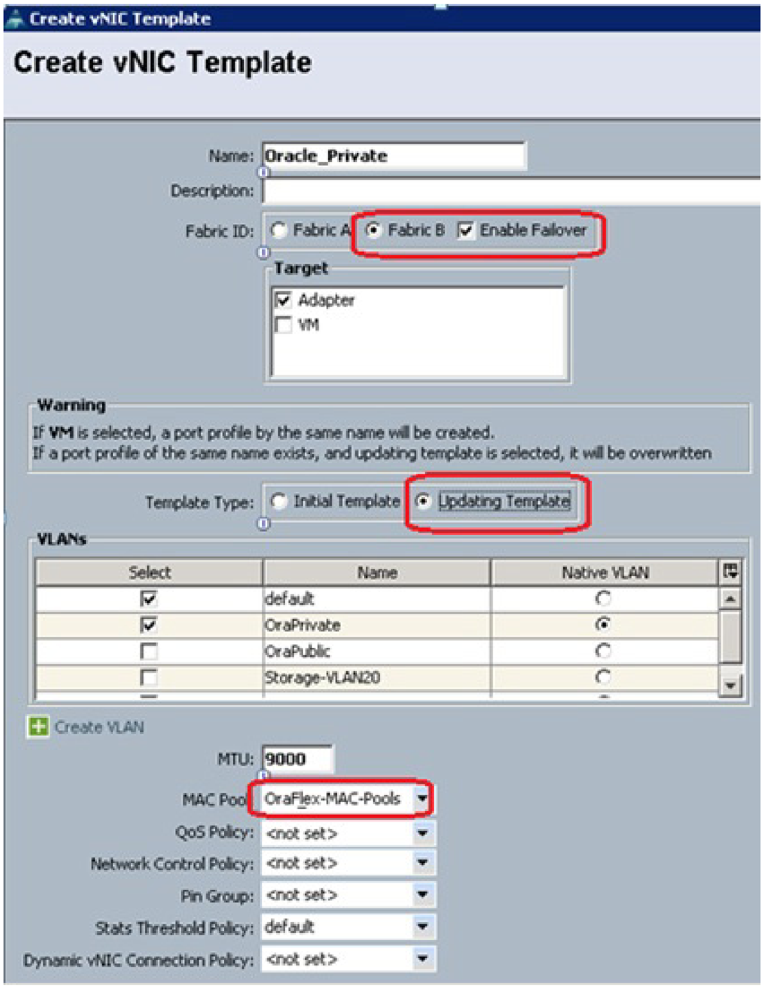

Create vNIC Template

Name: Oracle_Private

Description:

Fabric ID: ○ Fabric A ● Fabric B ☑ Enable Failover

Target

☑ Adapter

□ VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
☑	default	○
☑	OraPrivate	●
□	OraPublic	○
□	Storage-VLAN20	○

MTU: 9000

MAC Pool: OraFlex-MAC-Pools

QoS Policy: <not set>

Network Control Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

Dynamic vNIC Connection Policy: <not set>

Create Public vNIC Template

Next, create a vNIC template for public network. We used MTU=1500 and OraFlex-MAC-Pools as the MAC Pool for this template. We selected "enable failover" for public network.

Figure 22 Create Public vNIC Template

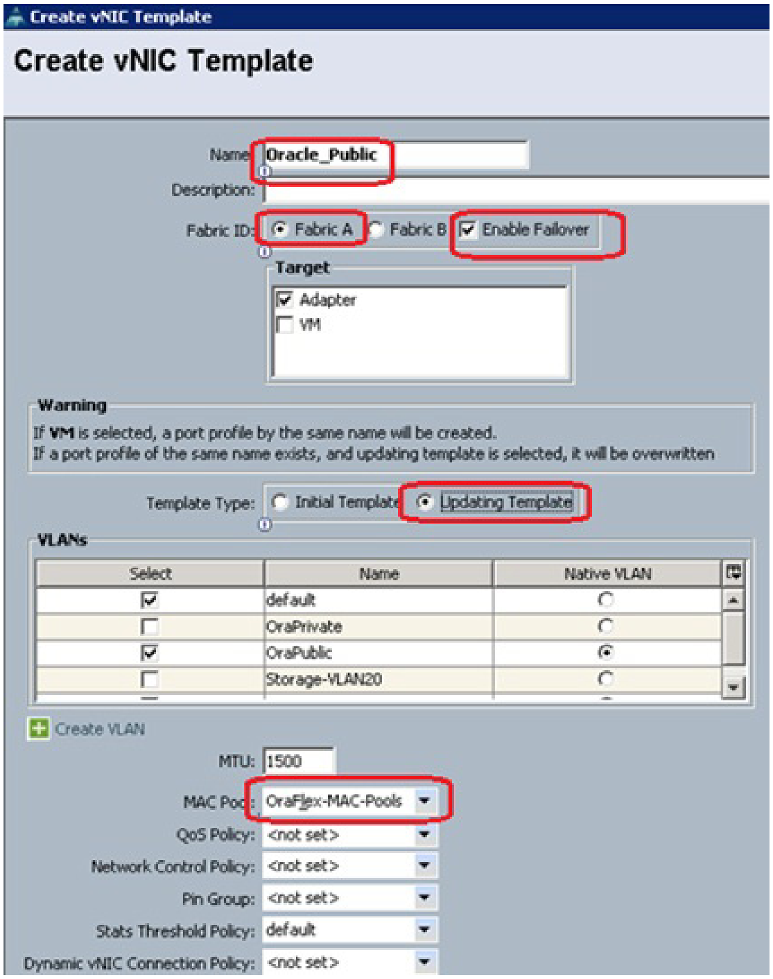

Create vNIC Template

Create vNIC Template

Name: Oracle_Public

Description:

Fabric ID: ● Fabric A ○ Fabric B ☑ Enable Failover

Target

☑ Adapter
□ VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
☑	default	○
□	OraPrivate	○
☑	OraPublic	●
□	Storage-VLAN20	○

Create VLAN

MTU: 1500

MAC Pool: OraFlex-MAC-Pools

QoS Policy: <not set>

Network Control Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

Dynamic vNIC Connection Policy: <not set>

Create Storage vNIC Template

For storage vNIC template, we used VLAN 20 and pinned it Fabric interconnect A. The vNICs derived from this template will drive database NFS traffic so we also used Jumbo frames (MTU=9000) as shown below.

Figure 23 Create Storage vNIC Template on Fabric A

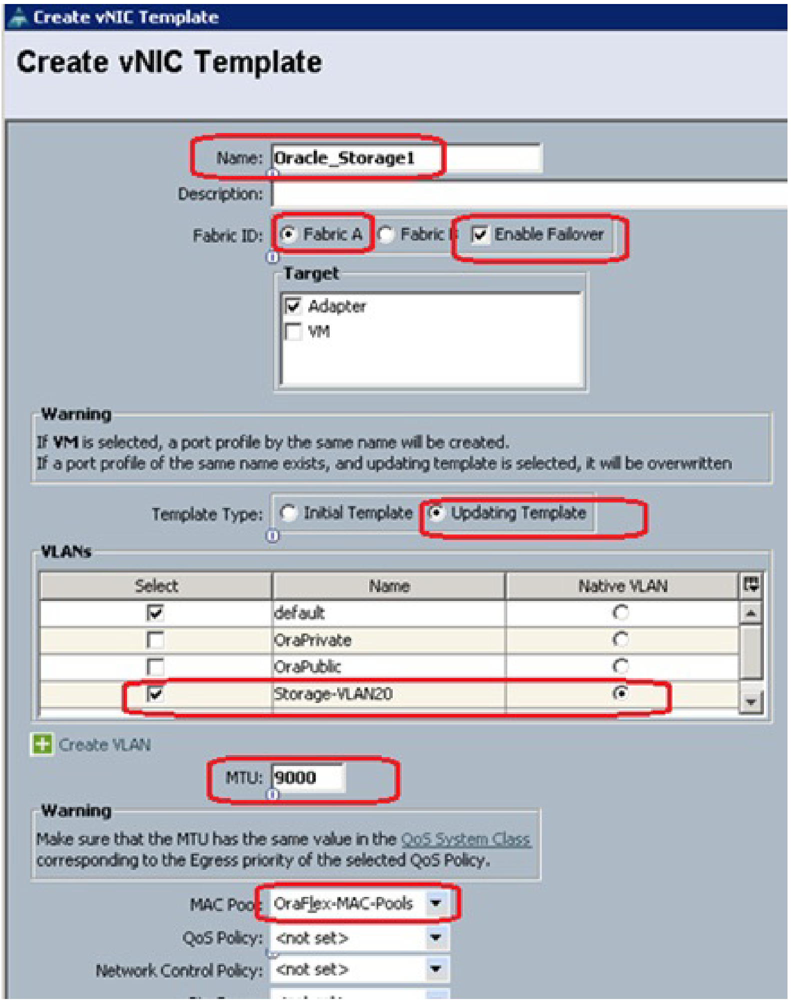

Create vNIC Template

Create vNIC Template

Name: Oracle_Storage1

Description:

Fabric ID: ● Fabric A ○ Fabric B ☑ Enable Failover

Target

☑ Adapter

□ VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
☑	default	○
□	OraPrivate	○
□	OraPublic	○
☑	Storage-VLAN20	●

Create VLAN

MTU: 9000

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

MAC Pool: OraFlex-MAC-Pools

QoS Policy: <not set>

Network Control Policy: <not set>

Repeat the same process to create a storage vNIC template for Fabric B.

Figure 24 Create Storage vNIC Template on Fabric B

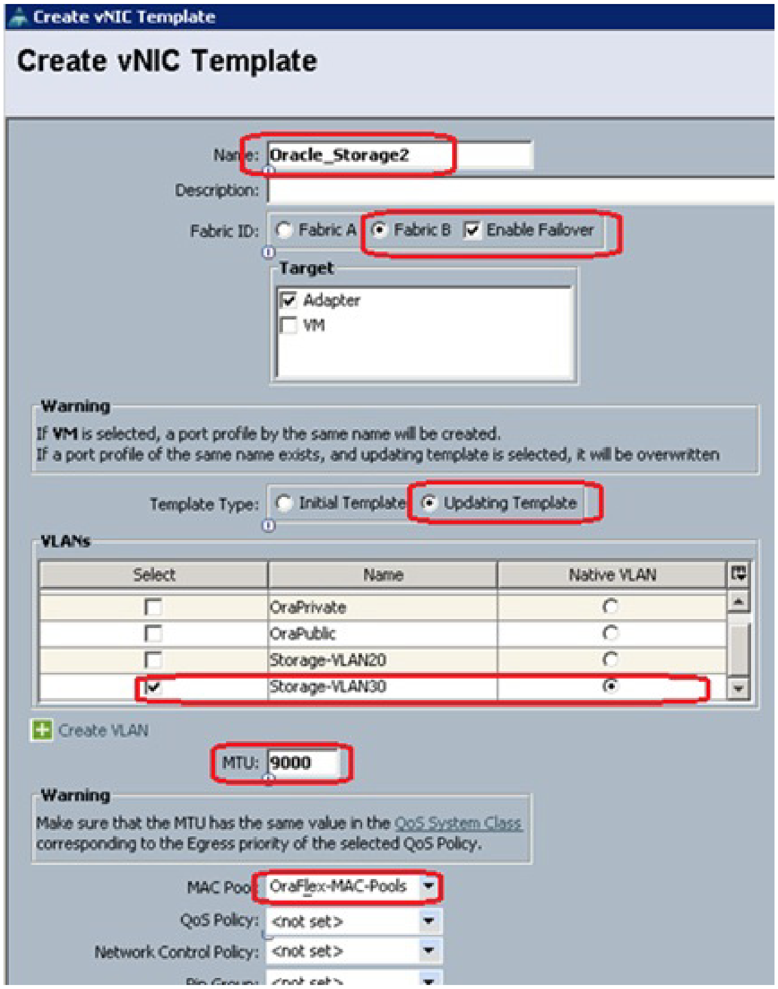

Create vNIC Template

Name: Oracle_Storage2

Description:

Fabric ID: ○ Fabric A ● Fabric B ☑ Enable Failover

Target

☑ Adapter

☐ VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
☐	OraPrivate	○
☐	OraPublic	○
☐	Storage-VLAN20	○
☑	Storage-VLAN30	●

Create VLAN

MTU: 9000

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

MAC Pool: OraFlex-MAC-Pools

QoS Policy: <not set>

Network Control Policy: <not set>

Pin Group: <not set>

The following is the vNIC template summary after all necessary templates for this FlexPod configuration are created.

Figure 25 vNIC Template Summary

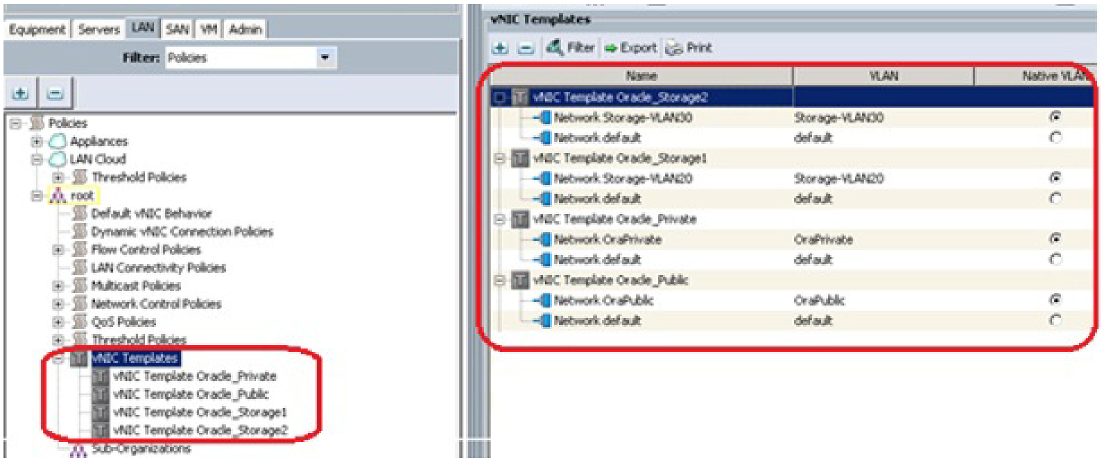

Name	VLAN	Native VLAN
vNIC Template Oracle_Storage2		
Network Storage-VLAN30	Storage-VLAN30	●
Network default	default	○
vNIC Template Oracle_Storage1		
Network Storage-VLAN20	Storage-VLAN20	●
Network default	default	○
vNIC Template Oracle_Private		
Network OraPrivate	OraPrivate	●
Network default	default	○
vNIC Template Oracle_Public		
Network OraPublic	OraPublic	●
Network default	default	○

Create HBA Templates

In Cisco UCS Manager, click SAN > Policies > vHBA templates and right-click "Create vHBA Template."

Figure 26 Create vHBA Templates

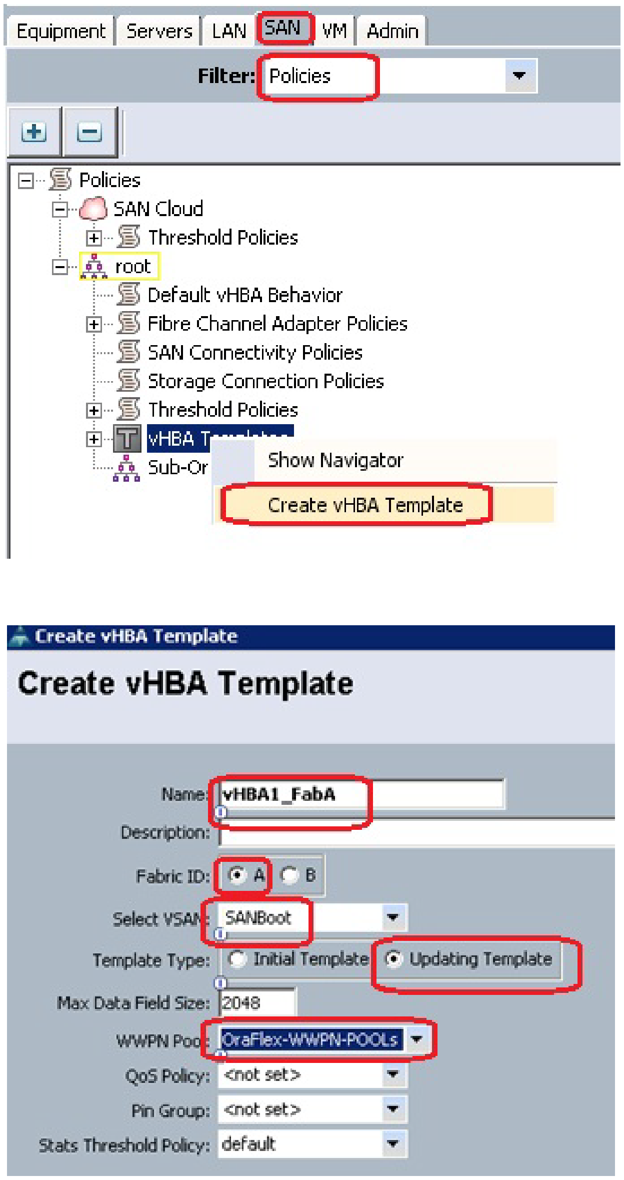

Create vHBA Template

Create vHBA Template

Name: vHBA1_FabA

Description:

Fabric ID: ● A ○ B

Select VSAN: SANBoot

Template Type: ○ Initial Template ● Updating Template

Max Data Field Size: 2048

WWPN Pool: OraFlex-WWPN-POOLs

QoS Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

We created two vHBA templates as vHBA1_FabA and vHBA2_FabB as shown below.

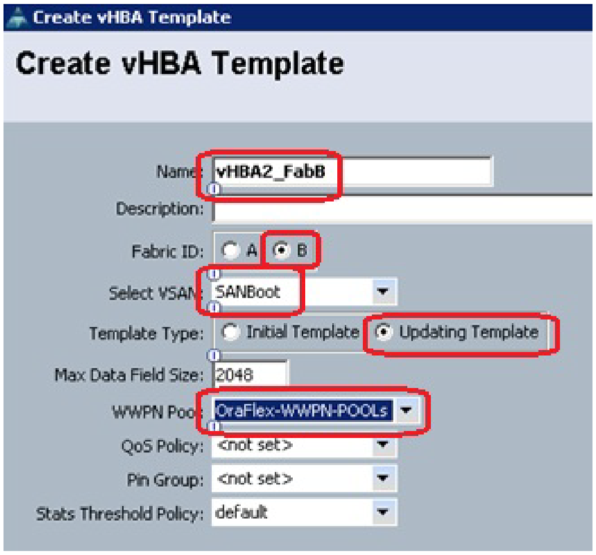

Create vHBA Template

Create vHBA Template

Name: vHBA2_FabB

Description:

Fabric ID: ○ A ● B

Select VSAN: SANBoot

Template Type: ○ Initial Template ● Updating Template

Max Data Field Size: 2048

WWPN Pool: OraFlex-WWPN-POOLs

QoS Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

Next, we will configure Ethernet uplink port channels.

Configure Ethernet Uplink Port Channels

To configure Port Channels, click LAN > LAN Cloud > Fabric A > Port Channels and right-click “Create Port-Channel.” Select the desired Ethernet Uplink ports configured earlier.

Repeat these steps to create Port Channels on Fabric B. In the current setup, we used ports 17 and 18 on Fabric A and configured as Port Channel 10. Similarly, ports 17 and 18 on Fabric B are configured to create Port Channel 11.

Figure 27 Configuring Port Channels

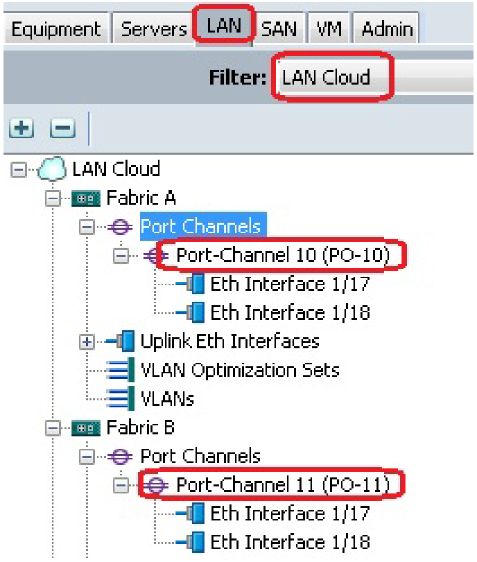

Ether Port Channel Details on Fabric A

Figure 28 Port Channel-10 Status and Properties

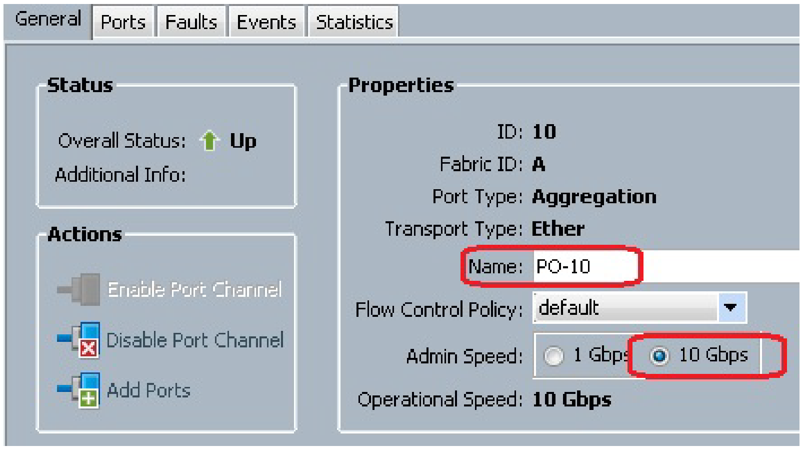

Status

Overall Status: Up

Additional Info:

Actions

Properties

ID: 10

Fabric ID: A

Port Type: Aggregation

Transport Type: Ether

Name: PO-10

Flow Control Policy: default

Admin Speed: ○ 1 Gbps ● 10 Gbps

Operational Speed: 10 Gbps

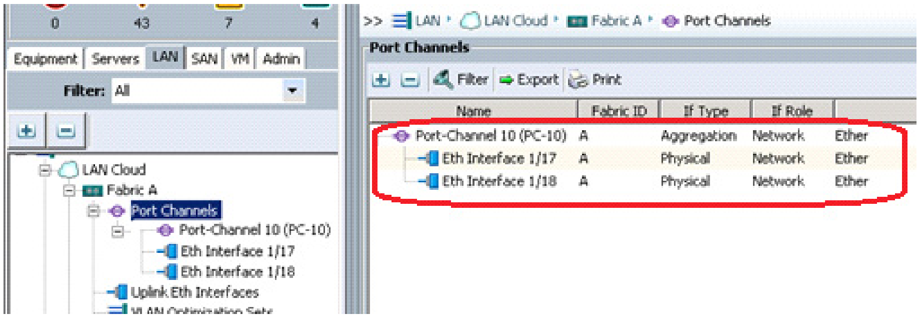

Filter: All

>> LAN ▸ LAN Cloud ▸ Fabric A ▸ Port Channels

Port Channels

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

Ether Port Channel Details on Fabric B

Figure 29 Port Channel-11 Status and Properties

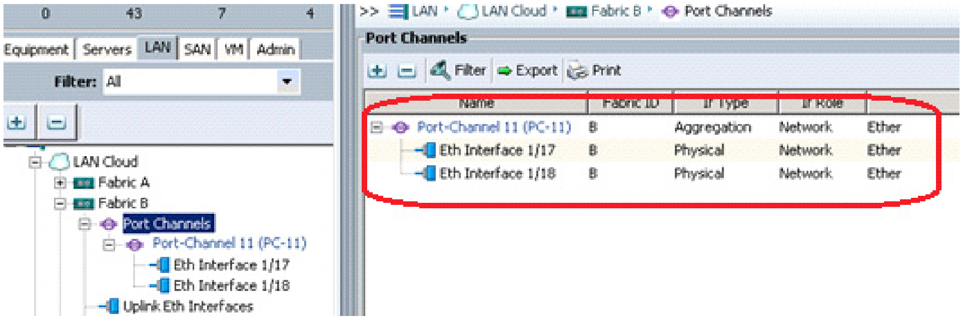

Filter: All

>> LAN ▸ LAN Cloud ▸ Fabric B ▸ Port Channels

Port Channels

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

Create Server Boot Policy for SAN Boot

Navigate to Cisco UCS Manager > Servers > Policies > Boot Policies > right-click Create Boot Policy. From the Local devices "Add CD-ROM" and click vHBA's to "add SAN Boot" as shown below.

Figure 30 Cisco UCS Server Boot Policy

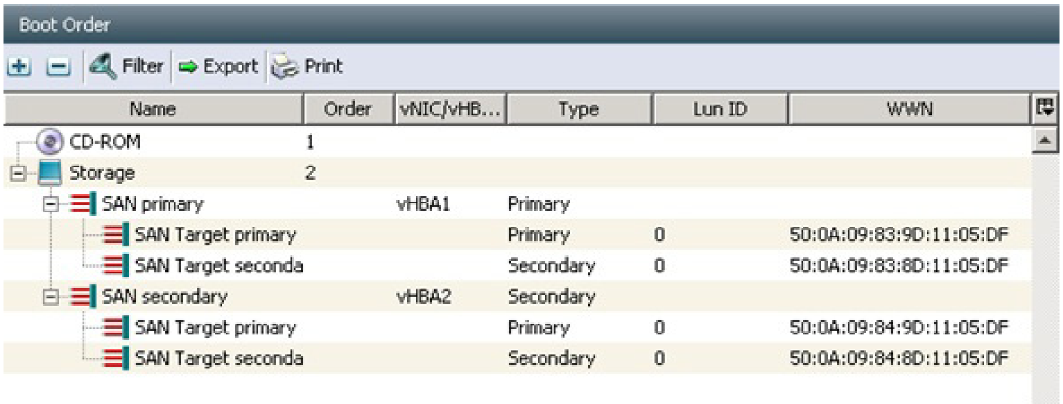

Name	Order	vNIC/vHB...	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		vHBA1	Primary		
SAN Target primary			Primary	0	50:0A:09:83:9D:11:05:DF
SAN Target seconda			Secondary	0	50:0A:09:83:8D:11:05:DF
SAN secondary		vHBA2	Secondary		
SAN Target primary			Primary	0	50:0A:09:84:9D:11:05:DF
SAN Target seconda			Secondary	0	50:0A:09:84:8D:11:05:DF

Note WWN for Storage has to be identified from NetApp storage.

When the above preparation steps are complete we are ready to create a service template from which the service profiles can be easily created.

Service Profile Creation and Association to Cisco UCS Blades

Service profile templates enable policy based server management that helps ensure consistent server resource provisioning suitable to meet predefined workload needs.

Create Service Profile Template

In Cisco UCS Manager, click Servers > Service Profile Templates > root and right-click root to “Create Service Profile Template.”

Figure 31 Create Service Profile Template

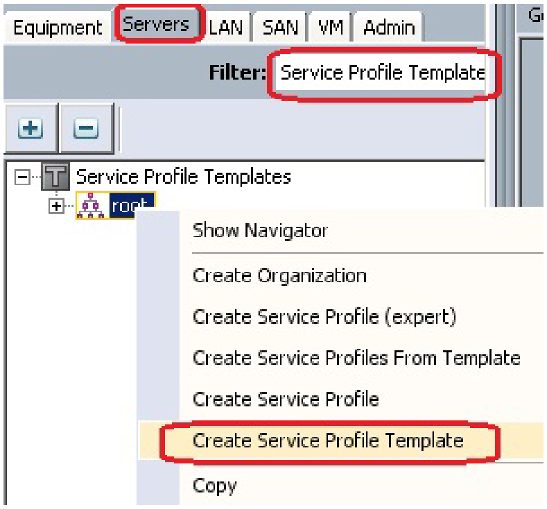

Equipment
Servers
LAN
SAN
VM
Admin
Service Profile Templates
root
Show Navigator
Create Organization
Create Service Profile (expert)
Create Service Profiles From Template
Create Service Profile
Create Service Profile Template
Copy

Filter: Service Profile Template

Enter the template name and select the UUID Pool that was created earlier and move on to the next screen.

Figure 32 Identify Service Profile Template

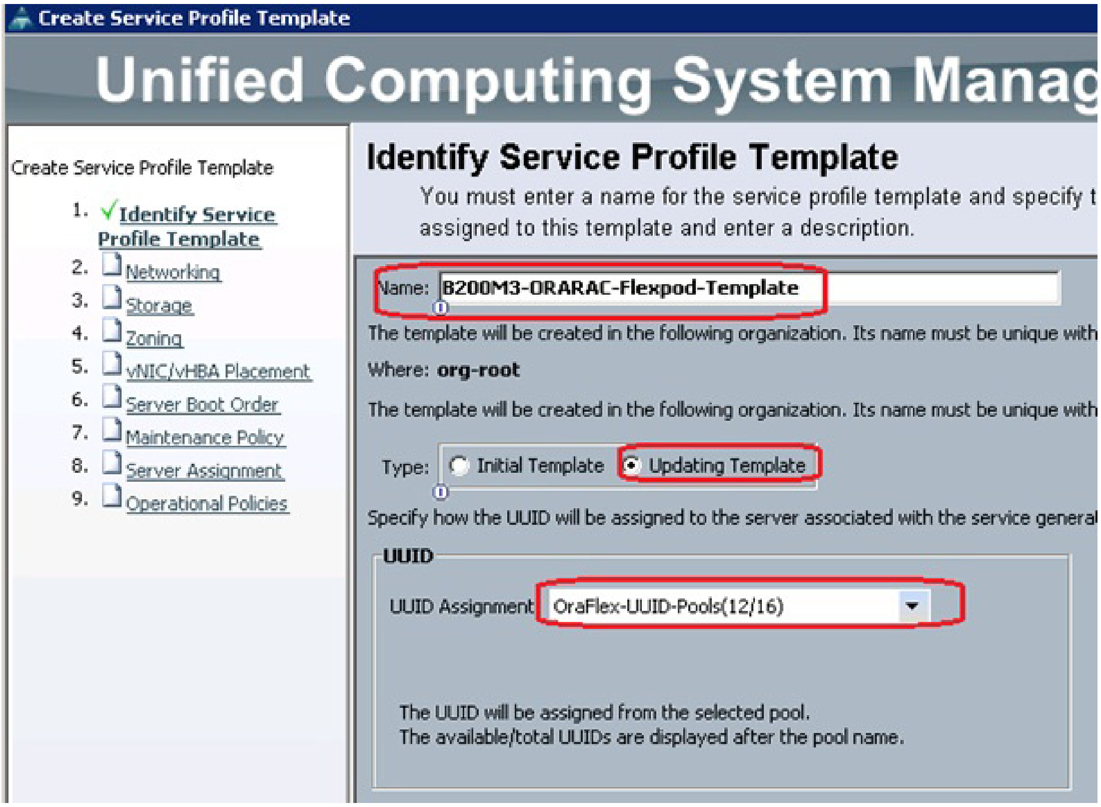

Identify Service Profile Template

You must enter a name for the service profile template and specify t assigned to this template and enter a description.

Name: B200M3-ORARAC-Flexpod-Template

The template will be created in the following organization. Its name must be unique with

Where: **org-root**

The template will be created in the following organization. Its name must be unique with

Type: ○ Initial Template ● Updating Template

Specify how the UUID will be assigned to the server associated with the service generat

UUID

UUID Assignment: OraFlex-UUID-Pools(12/16)

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Figure 33 Networking—LAN Configurtation

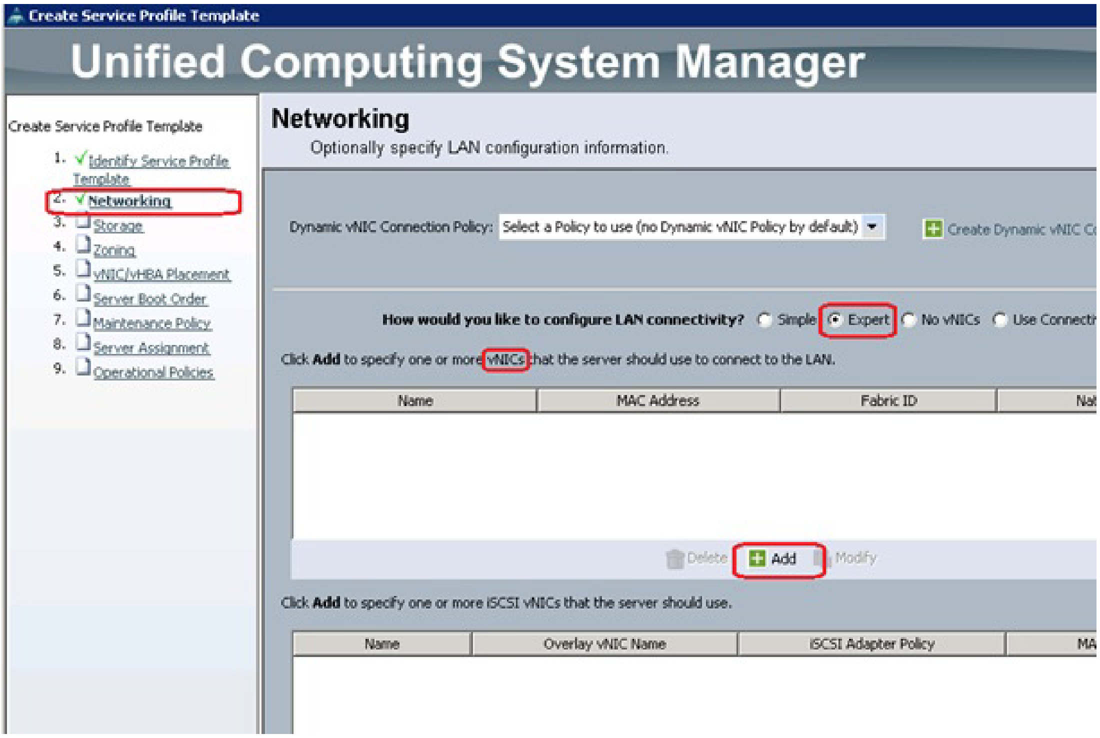

Create Service Profile Template

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: Select a Policy to use (no Dynamic vNIC Policy by default)

How would you like to configure LAN connectivity? ○ Simple ● Expert ○ No vNICs ○ Use Connecti

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

Name	MAC Address	Fabric ID	Na

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

Name	Overlay vNIC Name	iSCSI Adapter Policy	MA

In the Networking page create vNICs; one on each fabric and associate them with the VLAN policies created earlier. Select expert mode, and click on add on the section that specifies add one or more vNICs that the server should use to connect to the LAN.

In the create vNIC page, select "Use vNIC template" and adapter policy as Linux. In the page below vNIC1 was selected for Oracle public network.

Figure 34 Creation of vNICs using vNIC Template

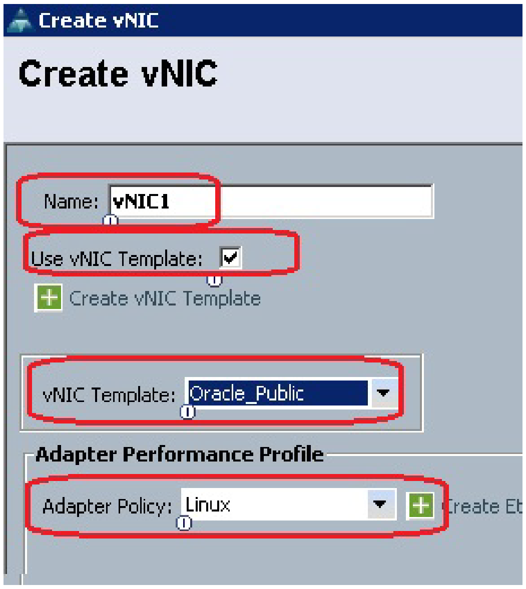

Create vNIC

Create vNIC

Name: vNIC1

Use vNIC Template: ☑

vNIC Template: Oracle_Public

Adapter Performance Profile

Adapter Policy: Linux

Similarly, create vNIC2 for private, vNIC3 for Storage side A, and VNIC4 for Storage Side B with appropriate vNIC template mapping for each vNIC.

Figure 35 Post vNIC Creation Using vNIC Template

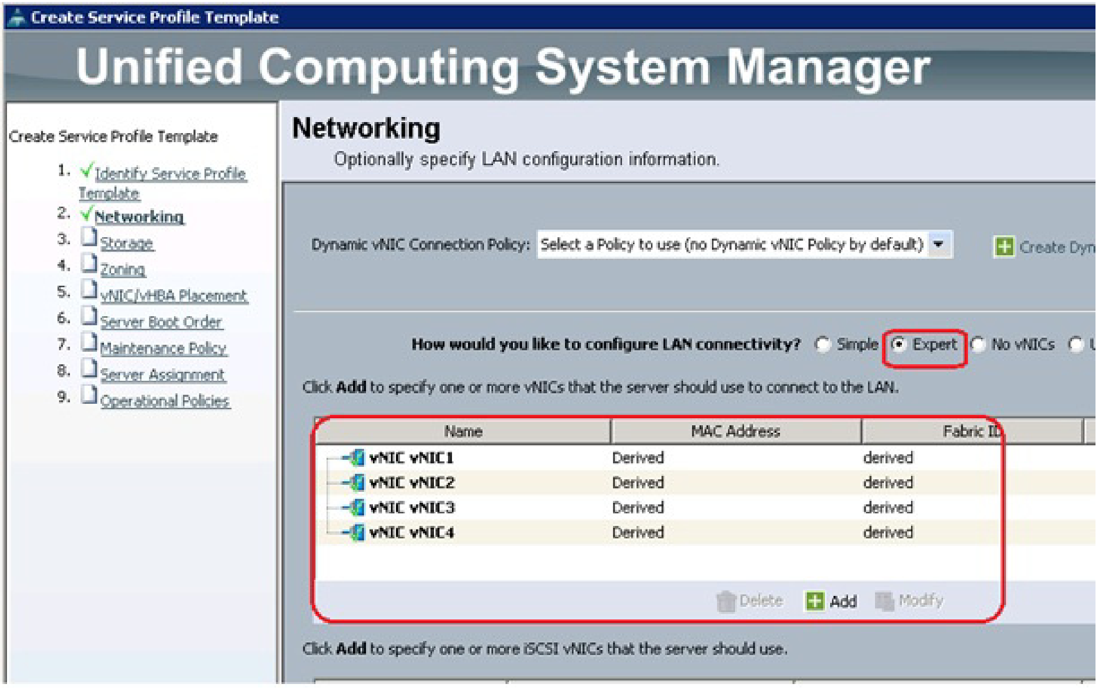
Create Service Profile Template
Unified Computing System Manager
Create Service Profile Template
1. ✓ Identify Service Profile Template
2. ✓ Networking
3. □ Storage
4. □ Zoning
5. □ vNIC/vHBA Placement
6. □ Server Boot Order
7. □ Maintenance Policy
8. □ Server Assignment
9. □ Operational Policies
Create Dyn
Delete
Add
Modify

Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default)

How would you like to configure LAN connectivity? ○ Simple ● Expert ○ No vNICs ○ U

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

Name	MAC Address	Fabric ID
vNIC vNIC1	Derived	derived
vNIC vNIC2	Derived	derived
vNIC vNIC3	Derived	derived
vNIC vNIC4	Derived	derived

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

When vNICs are created, you will need to create vHBA's.

In the storage page, select expert mode, choose the WWNN pool created earlier and click the Add button to create vHBA's.

Figure 36 Storage—SAN Configuration

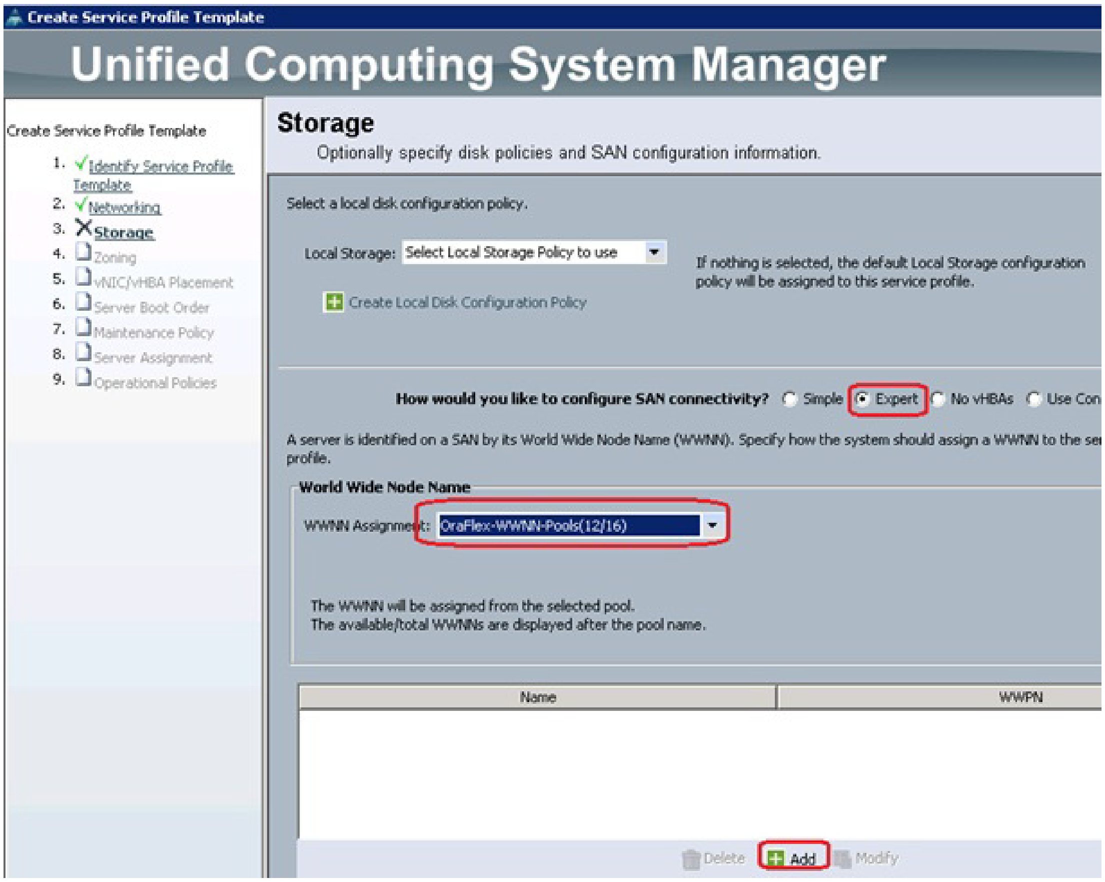

Create Service Profile Template

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

Storage

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage: Select Local Storage Policy to use

If nothing is selected, the default Local Storage configuration policy will be assigned to this service profile.

Create Local Disk Configuration Policy

How would you like to configure SAN connectivity? ○ Simple ● Expert ○ No vHBAs ○ Use Con

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the se profile.

World Wide Node Name

WWNN Assignment: OraFlex-WWNN-Pools(12/16)

The WWNN will be assigned from the selected pool.
The available/total WWNNs are displayed after the pool name.

Name	WWPN

Delete Add Modify

We created two vHBA's:

- vHBA1 using template vHBA1_FabA
- vHBA2 using template vHBA2_FabB

Figure 37 Create vHBAs Using vHBA Template

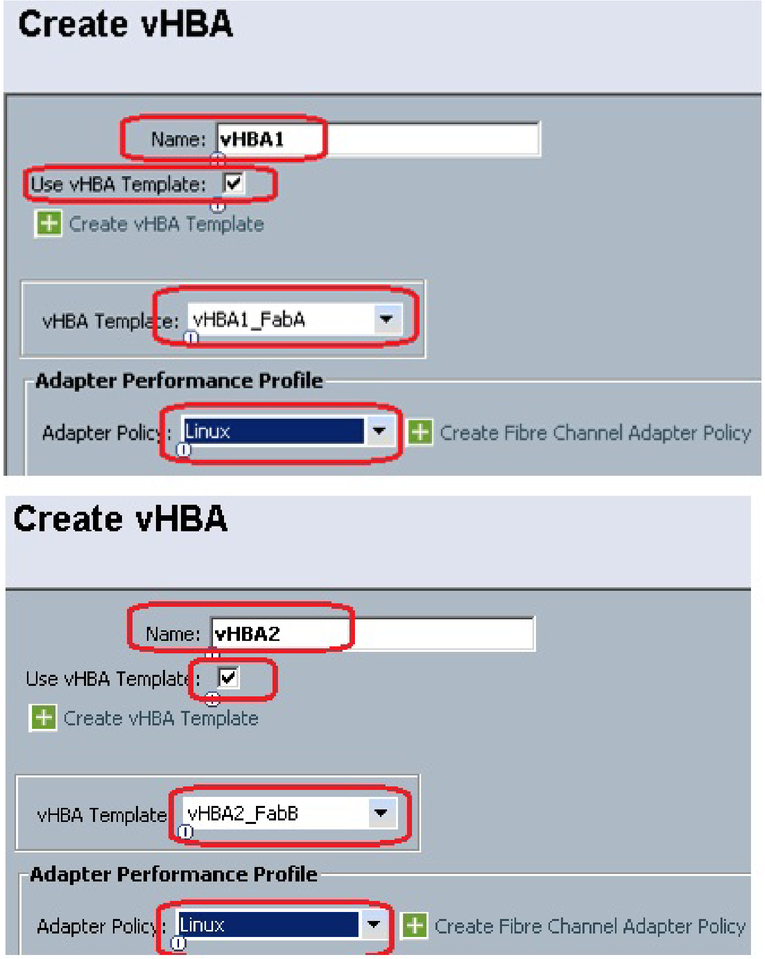

Create vHBA

Name: vHBA1

Use vHBA Template: ☑

vHBA Template: vHBA1_FabA

Adapter Performance Profile

Adapter Policy: Linux

Create vHBA

Name: vHBA2

Use vHBA Template: ☑

vHBA Template: vHBA2_FabB

Adapter Performance Profile

Adapter Policy: Linux

Create vHBA Template
Create Fibre Channel Adapter Policy
Create vHBA Template
Create Fibre Channel Adapter Policy

Figure 38 Worldwide Node Name Assignment

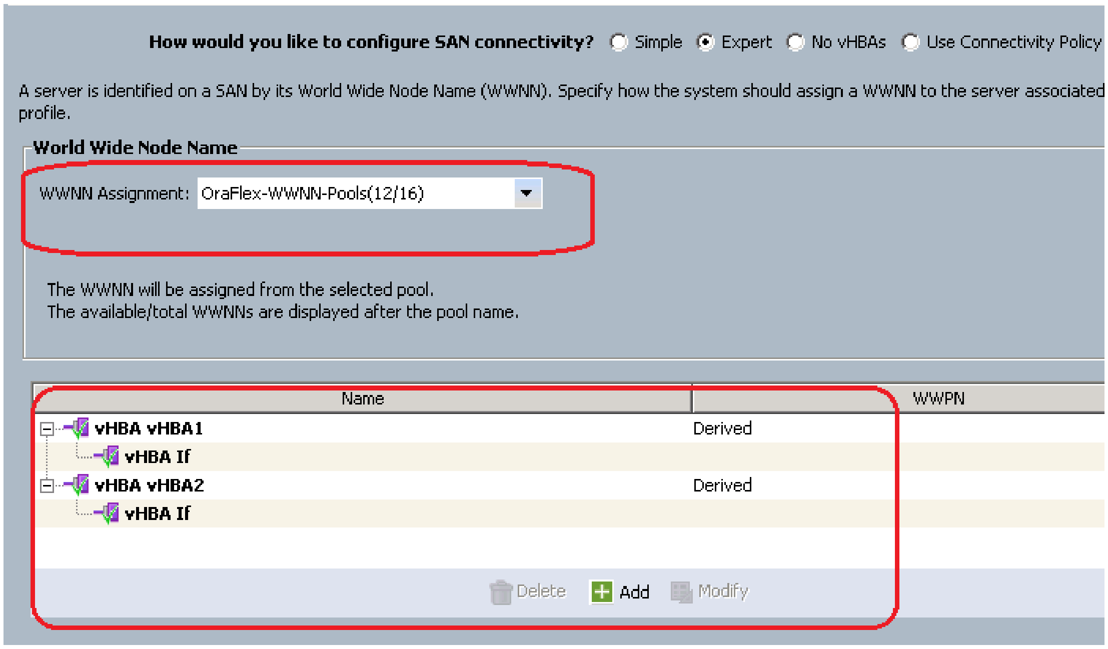

How would you like to configure SAN connectivity? ○ Simple ● Expert ○ No vHBAs ○ Use Connectivity Policy

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated profile.

World Wide Node Name

WWNN Assignment: OraFlex-WWNN-Pools(12/16)

The WWNN will be assigned from the selected pool.
The available/total WWNNs are displayed after the pool name.

Name	WWPN
vHBA vHBA1	Derived
vHBA If	
vHBA vHBA2	Derived
vHBA If	

For this FlexPod configuration, we used Cisco Nexus 5548UP for zoning so we will skip the zoning section and use default vNIC/vHBA placement.

Figure 39 vNIC/vHBA Placement

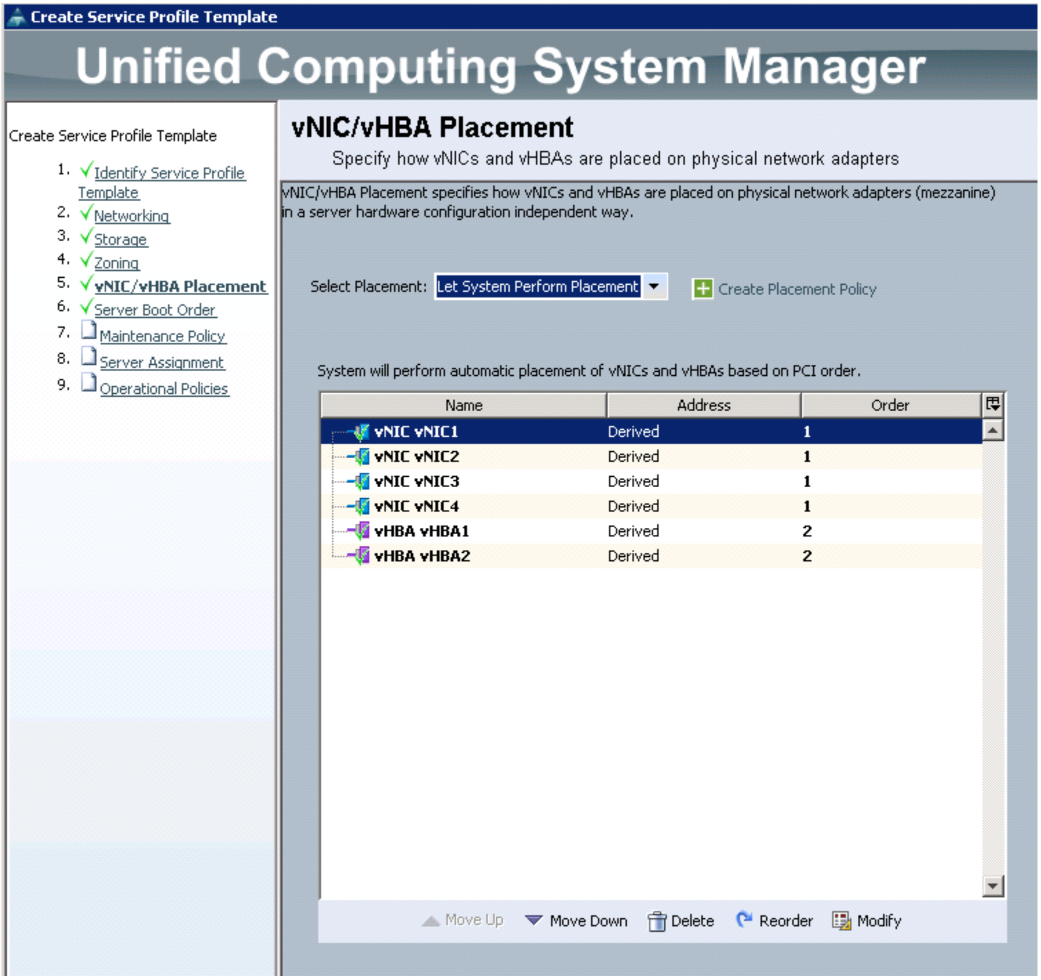

Create Service Profile Template

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

vNIC/vHBA Placement

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: Let System Perform Placement

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order
vNIC vNIC1	Derived	1
vNIC vNIC2	Derived	1
vNIC vNIC3	Derived	1
vNIC vNIC4	Derived	1
vHBA vHBA1	Derived	2
vHBA vHBA2	Derived	2

Server Boot Policy

In the Server Boot Order page, choose the Boot Policy created for SAN boot and click Next.

Figure 40 Configure Server Boot Policy During Service Profile Template Creation

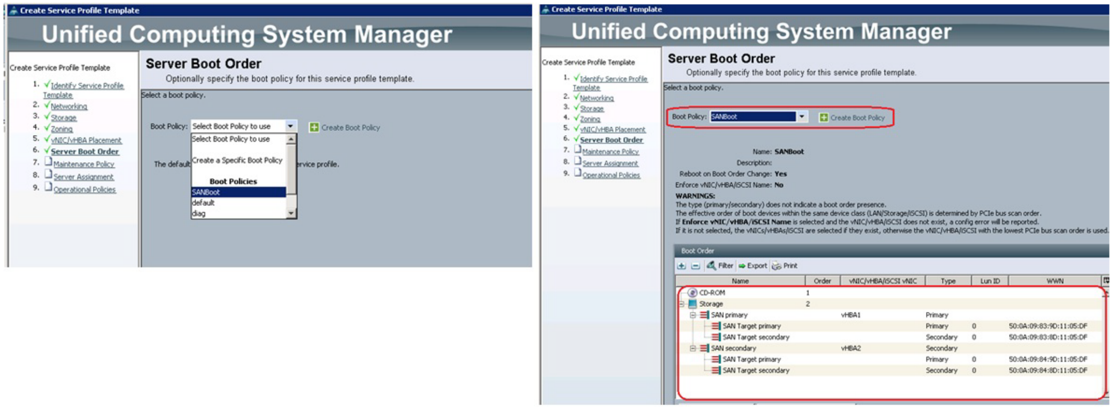

Create Service Profile Template

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

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Select Boot Policy to use

The default ... ervice profile.

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

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: SANBoot

Name: **SANBoot**
Description:
Reboot on Boot Order Change: **Yes**
Enforce vNIC/vHBA/iSCSI Name: **No**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		vHBA1	Primary		
SAN Target primary			Primary	0	50:0A:09:83:9D:11:05:DF
SAN Target secondary			Secondary	0	50:0A:09:83:8D:11:05:DF
SAN secondary		vHBA2	Secondary		
SAN Target primary			Primary	0	50:0A:09:84:9D:11:05:DF
SAN Target secondary			Secondary	0	50:0A:09:84:8D:11:05:DF

Figure 41 Server Boot Order Assignment

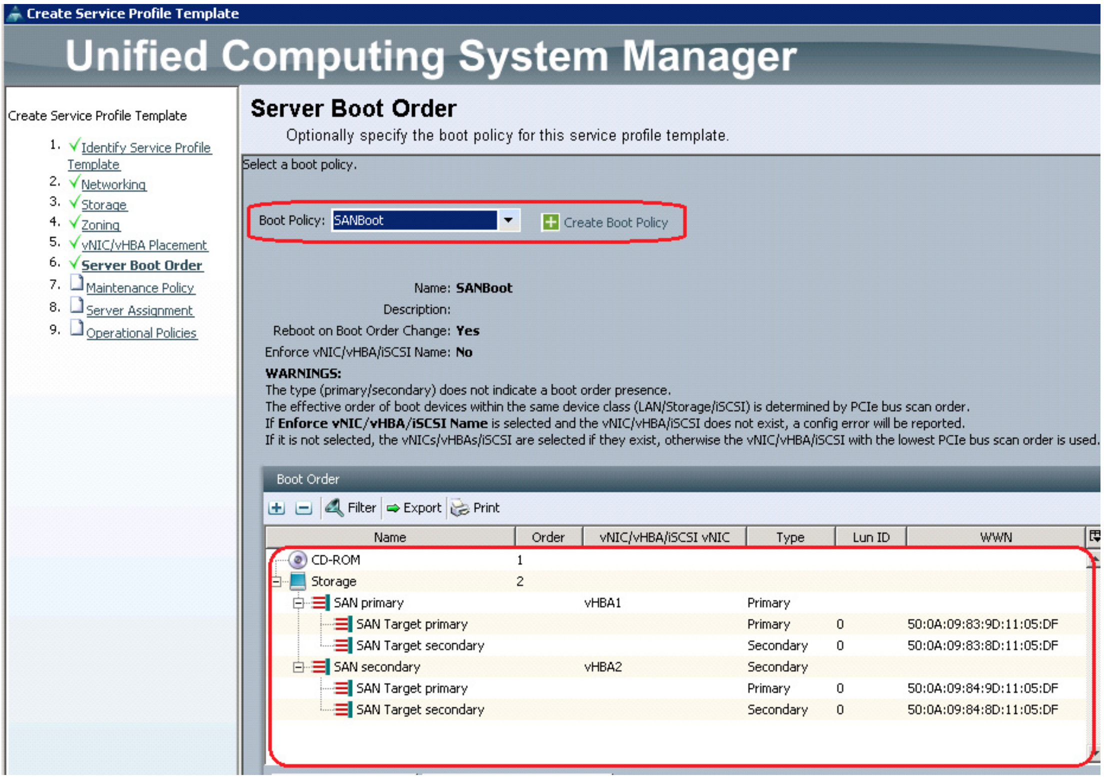

Create Service Profile Template
Unified Computing System Manager

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

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: SANBoot

Name: **SANBoot**
Description:
Reboot on Boot Order Change: **Yes**
Enforce vNIC/vHBA/iSCSI Name: **No**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		vHBA1	Primary		
SAN Target primary			Primary	0	50:0A:09:83:9D:11:05:DF
SAN Target secondary			Secondary	0	50:0A:09:83:8D:11:05:DF
SAN secondary		vHBA2	Secondary		
SAN Target primary			Primary	0	50:0A:09:84:9D:11:05:DF
SAN Target secondary			Secondary	0	50:0A:09:84:8D:11:05:DF

The maintenance and server assignment policies were left to default in our configuration. However, they may vary from site to site depending on your work loads, best practices and policies.

Create Service Profiles from Service Profile Templates

In Cisco UCS Manager, click Servers > Service Profile Templates and right-click Create Service Profiles from Template.

Figure 42 Create Service Profile from Service Profile Template

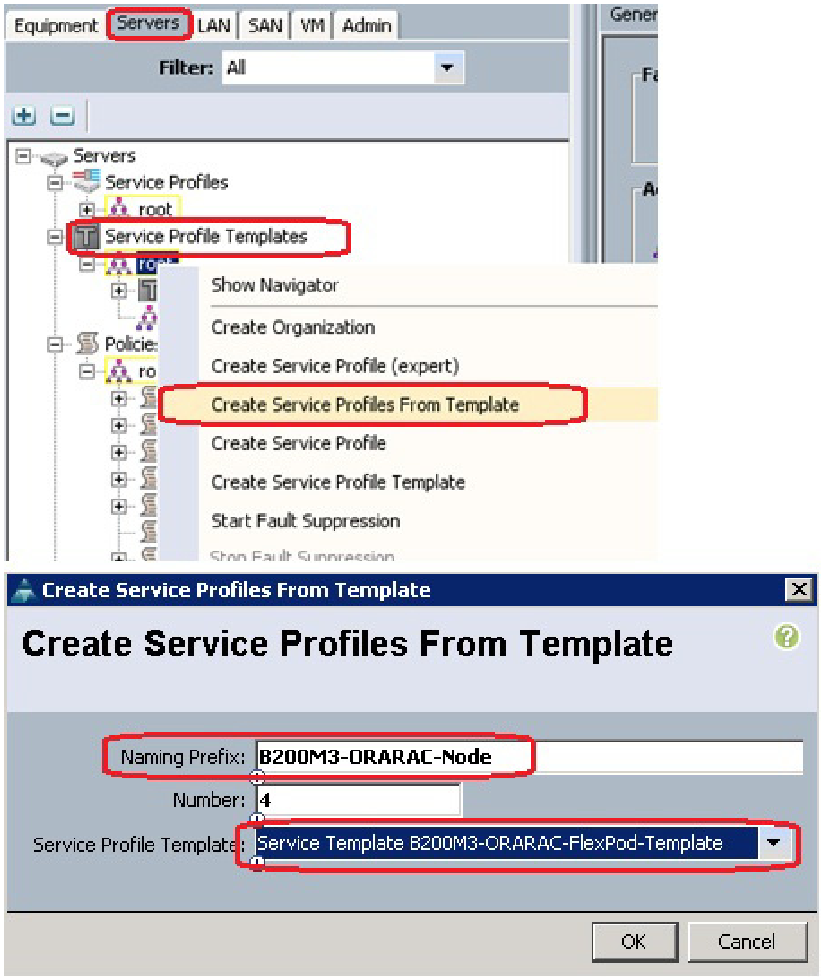

Create Service Profiles From Template

Naming Prefix: B200M3-ORARAC-Node

Number: 4

Service Profile Template: Service Template B200M3-ORARAC-FlexPod-Template

We created four service profiles with the name prefix "B200M3-ORARAC-Node" as listed below:

- B200M3-ORARAC-Node1
- B200M3-ORARAC-Node2
- B200M3-ORARAC-Node3
- B200M3-ORARAC-Node4

Figure 43 Post Service Profiles Creation

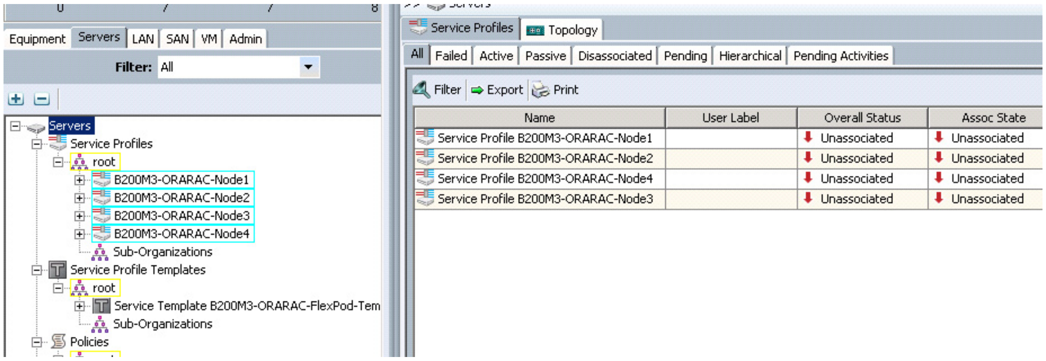

Name	User Label	Overall Status	Assoc State
Service Profile B200M3-ORARAC-Node1		Unassociated	Unassociated
Service Profile B200M3-ORARAC-Node2		Unassociated	Unassociated
Service Profile B200M3-ORARAC-Node4		Unassociated	Unassociated
Service Profile B200M3-ORARAC-Node3		Unassociated	Unassociated

Associating a Service Profile to Servers

As service profiles are created we are ready to associate them to the servers.

1. Under the Servers tab, select the desired service profile and select Change Service Profile Association.

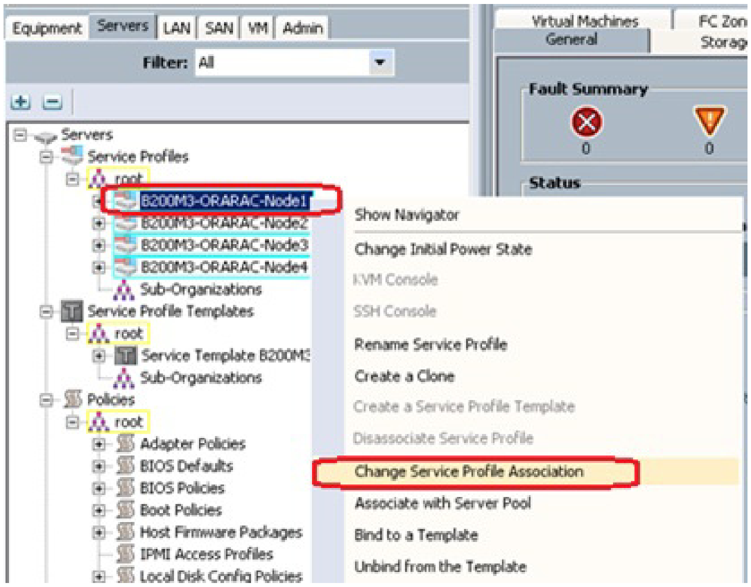

2. Click Change Service Profile Association under the General tab, select the existing server that you would like to assign and click OK.

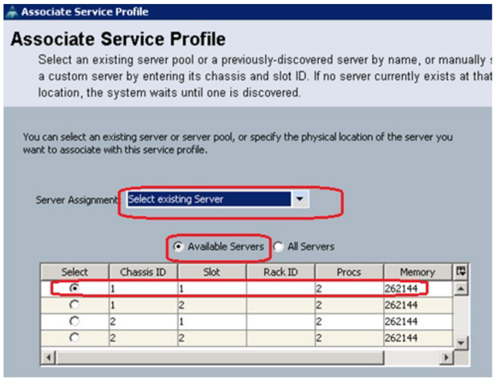

Associate Service Profile

Associate Service Profile

Select an existing server pool or a previously-discovered server by name, or manually s
a custom server by entering its chassis and slot ID. If no server currently exists at that
location, the system waits until one is discovered.

You can select an existing server or server pool, or specify the physical location of the server you want to associate with this service profile.

Server Assignment: Select existing Server

● Available Servers ○ All Servers

Select	Chassis ID	Slot	Rack ID	Procs	Memory
●	1	1		2	262144
○	1	2		2	262144
○	2	1		2	262144
○	2	2		2	262144

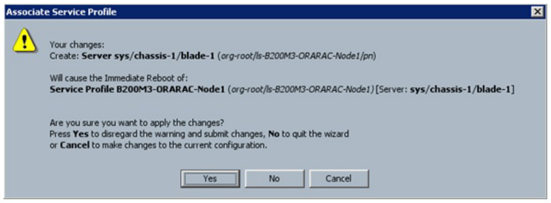

Associate Service Profile

Your changes:
Create: **Server sys/chassis-1/blade-1** (*org-root/ls-B200M3-ORARAC-Node1/pn*)

Will cause the Immediate Reboot of:
Service Profile B200M3-ORARAC-Node1 (*org-root/ls-B200M3-ORARAC-Node1*) [Server: **sys/chassis-1/blade-1**]

Are you sure you want to apply the changes?
Press **Yes** to disregard the warning and submit changes, **No** to quit the wizard
or **Cancel** to make changes to the current configuration.

Repeat the same steps to associate the remaining three service profiles for the respective blade servers as shown below:

- B200M3-ORARAC-Node1 - Chassis1 - Slot1
- B200M3-ORARAC-Node2 - Chassis1 - Slot2
- B200M3-ORARAC-Node3 - Chassis2 - Slot1
- B200M3-ORARAC-Node4 - Chassis2 - Slot2

When you start associating the service profiles, the overall status will be shown as "Config"

Figure 44 Overall Status During Service Profile Association

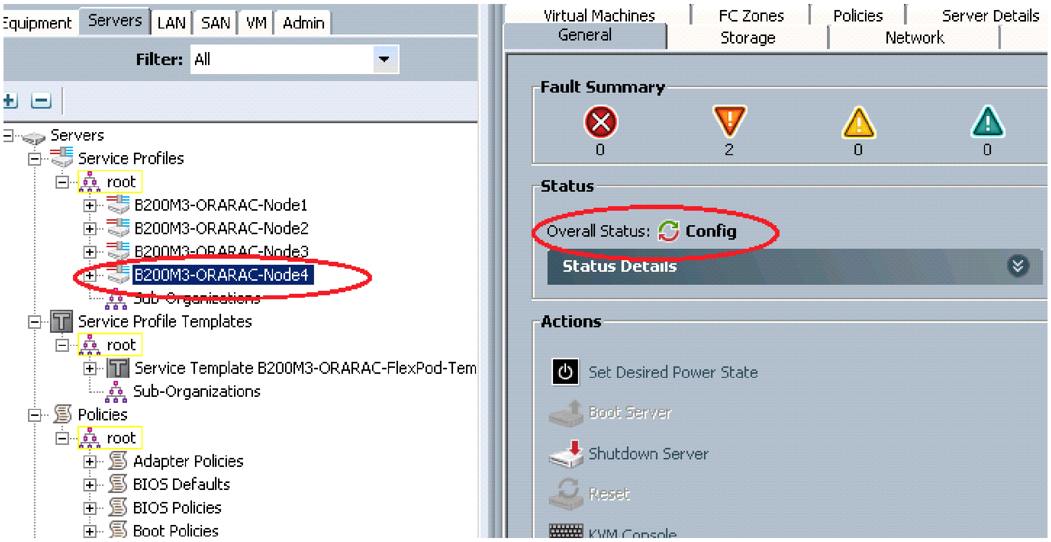

Filter: All

Overall Status: Config

Note Make sure the FSM (Final State Machine) Association progress status completes by 100 percent.

Figure 45 FSM (Final State Machine) Status During Service Profile Association

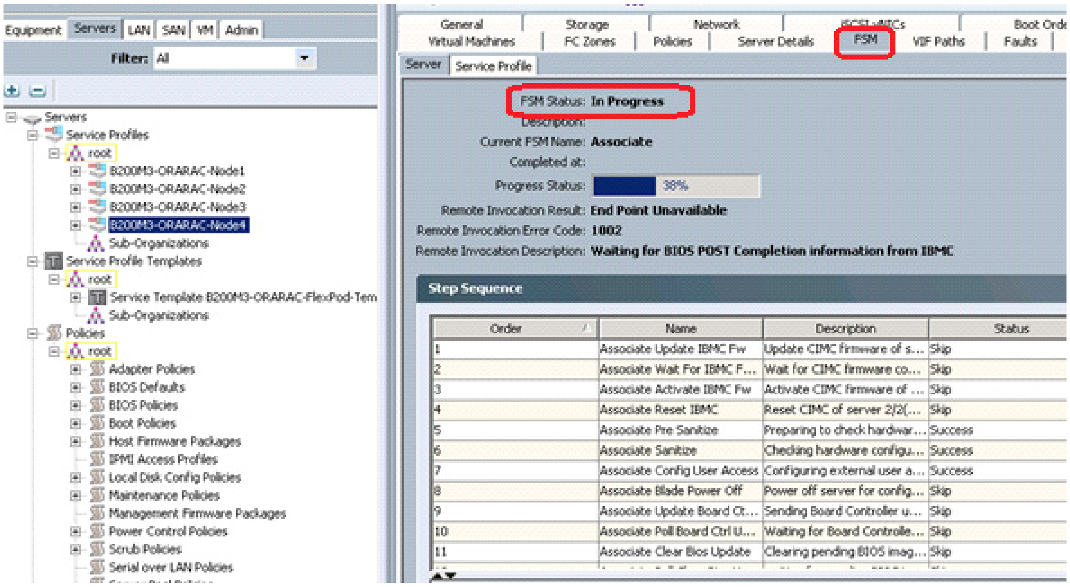

FSM Status: **In Progress**
Description:
Current FSM Name: **Associate**
Completed at:
Progress Status: 38%
Remote Invocation Result: **End Point Unavailable**
Remote Invocation Error Code: **1002**
Remote Invocation Description: **Waiting for BIOS POST Completion information from IBMC**

Order	Name	Description	Status
1	Associate Update IBMC Fw	Update CIMC firmware of s...	Skip
2	Associate Wait For IBMC F...	Wait for CIMC firmware co...	Skip
3	Associate Activate IBMC Fw	Activate CIMC firmware of ...	Skip
4	Associate Reset IBMC	Reset CIMC of server 2/2(...	Skip
5	Associate Pre Sanitize	Preparing to check hardwar...	Success
6	Associate Sanitize	Checking hardware configu...	Success
7	Associate Config User Access	Configuring external user a...	Success
8	Associate Blade Power Off	Power off server for config...	Skip
9	Associate Update Board Ct...	Sending Board Controller u...	Skip
10	Associate Poll Board Ctrl U...	Waiting for Board Controlle...	Skip
11	Associate Clear Bios Update	Clearing pending BIOS imag...	Skip

Make sure all the service profiles are associated as shown below.

Figure 46 Post Service Profile Association

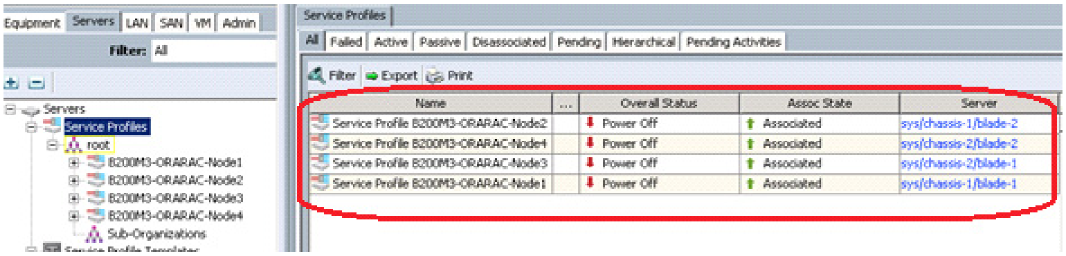

Name	...	Overall Status	Assoc State	Server
Service Profile B200M3-ORARAC-Node2		Power Off	Associated	sys/chassis-1/blade-2
Service Profile B200M3-ORARAC-Node4		Power Off	Associated	sys/chassis-2/blade-2
Service Profile B200M3-ORARAC-Node3		Power Off	Associated	sys/chassis-2/blade-1
Service Profile B200M3-ORARAC-Node1		Power Off	Associated	sys/chassis-1/blade-1

Cisco Nexus 5548UP Configuration

The following are the general steps involved in Nexus 5548UP configuration.

1. Configure NPIV and FCoE features
2. Fiber Channel Fabric Zoning for SAN Boot
3. Setting up VLAN and VSAN Configuration
4. Configuring Virtual Port Channel for Oracle Data Network and Storage Network
    a. Configure Virtual Port Channel on Nexus 5548
    b. Configure Virtual Port Channel for Data Network
    c. Configure Virtual Port Channel on NFS Storage Network
5. Setting up Jumbo Frames on N5548UP

In the following sections describe the steps in detail.

Configure NPIV and FCoE Features

If the NPIV feature is not enabled or if it is a new setup, enable the NPIV feature on the Cisco Nexus N5548 switches. The NPIV feature is required to register FLOGI (Fabric Login) of the Cisco UCS blade WWPN (Host Initiator).

Steps to enable NPIV and FCoE features on Cisco Nexus 5548 switch A:

```
N5548-Fab-A# config terminal
N5548-Fab-A(config)# feature npiv
N5548-Fab-A(config)# feature fcoe
```

Repeat the steps to enable the NPIV feature on Switch B.

Configure Fiber Channel Fabric Zoning for SAN Boot

Make sure you have (8 GB) SFP+ modules connected to the Nexus 5548UP ports. The port mode is set to AUTO as well as the speed is set to AUTO. Rate mode is "dedicated" and when everything is configured correctly.

Note A Nexus 5548 series switch supports multiple VSAN configurations and lot of additional features. We are highlighting only a single VSAN and only relevant features that are required for this study. It is beyond the scope of this document to highlight other features and their use.

Table2 lists the plan for zones and their associated members that are used in the testing and discussed in this document. In the current setup, we used a total of 4 paths: 2 paths from each Fabrics and Cisco Nexus 5548's to the storage.

Table 2 Zones for Oracle RAC Node Setup

Zone Name	Host (HBA) - WWPN	NetApp Storage visible from Zone (PWWN)
B200M3-CH1-BL1-ORARAC1-vHBA1	20:00:01:25:b5:11:13:02	ControllerA – Port OC 50:0a:09:83:9d:11:05:df & ControllerB – Port OC 50:0a:09:83:8d:11:05:df
B200M3-CH1-BL2-ORARAC2-vHBA1	20:00:01:25:b5:11:13:04	
B200M3-CH2-BL1-ORARAC3-vHBA1	20:00:01:25:b5:11:13:06	
B200M3-CH2-BL2-ORARAC4-vHBA1	20:00:01:25:b5:11:13:08	
B200M3-CH1-BL1-ORARAC1-vHBA2	20:00:01:25:b5:11:13:01	ControllerA – Port OD 50:0a:09:84:9d:11:05:df & ControllerB – Port OD 50:0a:09:84:8d:11:05:df
B200M3-CH1-BL2-ORARAC2-vHBA2	20:00:01:25:b5:11:13:03	
B200M3-CH2-BL1-ORARAC3-vHBA2	20:00:01:25:b5:11:13:05	
B200M3-CH2-BL2-ORARAC4-vHBA2	20:00:01:25:b5:11:13:07	

To find out the host adapter WWPN's for each of the HBA's, launch Cisco UCS Manager. From the Cisco UCS Manager, select the desired server from Equipment > chassis > servers > vHBAs menu.

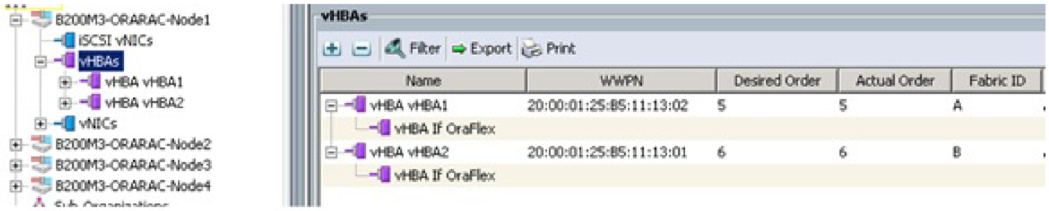

B200M3-ORARAC-Node1
iSCSI vNICs
vHBAs
vHBA vHBA1
vHBA vHBA2
vNICs
B200M3-ORARAC-Node2
B200M3-ORARAC-Node3
B200M3-ORARAC-Node4
Sub Organizations

Name	WWPN	Desired Order	Actual Order	Fabric ID
vHBA vHBA1	20:00:01:25:B5:11:13:02	5	5	A
vHBA If OraFlex				
vHBA vHBA2	20:00:01:25:B5:11:13:01	6	6	B
vHBA If OraFlex				

The WWPN numbers for both HBAs for server 1 are illustrated above. In the current setup, we used a total of 4 paths: 2 paths from each Fabrics and Cisco Nexus 5548's to the storage. The details are shown below.

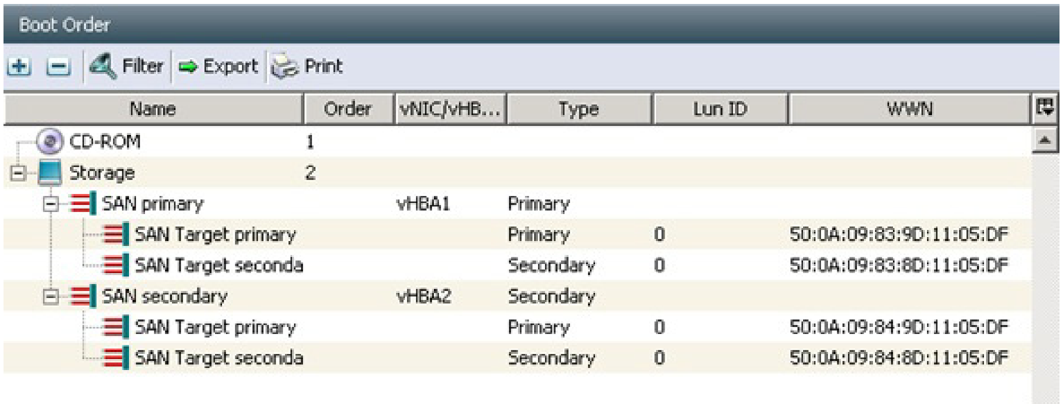

Name	Order	vNIC/vHB...	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		vHBA1	Primary		
SAN Target primary			Primary	0	50:0A:09:83:9D:11:05:DF
SAN Target seconda			Secondary	0	50:0A:09:83:8D:11:05:DF
SAN secondary		vHBA2	Secondary		
SAN Target primary			Primary	0	50:0A:09:84:9D:11:05:DF
SAN Target seconda			Secondary	0	50:0A:09:84:8D:11:05:DF

Setting up VLAN and VSAN Configuration

VLAN Configuration for Cisco Nexus 5548 Fabric A

```
VLAN 134 Configuration for Public traffic

N5548-Fab-A# config terminal
N5548-Fab-A(config)# VLAN 134
N5548-Fab-A(config-VLAN)# name Oracle_RAC_Public_Traffic
N5548-Fab-A(config-VLAN)# no shutdown
N5548-Fab-A(config-VLAN)# exit

VLAN 134 Configuration for Private traffic

N5548-Fab-A(config)# VLAN 10
N5548-Fab-A(config-VLAN)# name Oracle_RAC_Private_Traffic
N5548-Fab-A(config-VLAN)# no ip igmp snooping
N5548-Fab-A(config-VLAN)# no shutdown
N5548-Fab-A(config-VLAN)# exit

VLAN 20 Configuration for storage traffic

N5548-Fab-A(config)# VLAN 20
```

```
N5548-Fab-A(config-VLAN)# name Storage_Traffic for ControllerA
N5548-Fab-A(config-VLAN)# no shutdown
N5548-Fab-A(config-VLAN)# exit

VLAN 30 Configuration for storage traffic

N5548-Fab-A(config)# VLAN 30
N5548-Fab-A(config-VLAN)# name Storage_Traffic for ControllerB
N5548-Fab-A(config-VLAN)# no shutdown
N5548-Fab-A(config-VLAN)# exit


N5548-Fab-B(config)# Copy running-config startup-config
```

VSAN Configuration for Cisco Nexus 5548 Fabric A

```
N5548-Fab-A(config)# vsan database
N5548-Fab-A(config-vsan-database)# vsan 25
N5548-Fab-A(config-vsan-database)# vsan 25 interface fc 1/29-32
N5548-Fab-A(config)# exit
N5548-Fab-B(config)# Copy running-config startup-config
```

Repeat the VLAN and VSAN configuration steps on Cisco Nexus 5548 Fabric B.

Configuring the Virtual Port Channel for Oracle Data Network and Storage Network

Configure Virtual Port Channel on Cisco Nexus 5548

Cisco Nexus 5548UP vPC configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers is as shown in Table 3. To provide Layer 2 and Layer 3 switching, a pair of Cisco Nexus 5548UP Switches with upstream switching are deployed, providing high availability in the event of failure to Cisco UCS to handle management, application, and Network storage data traffic. In the Cisco Nexus 5548UP switch topology, a single vPC feature is enabled to provide high availability, faster convergence in the event of a failure, and greater throughput.

Table 3 vPC Mapping for the Cisco Nexus 5548UP Switch

vPC Domain	vPC Name	vPC ID
1	Peer-Link	1
1	vPC-Public	17
1	vPC-Private	18
1	vPC-Storage1	20
1	vPC-Storage2	30

In the vPC design table, a single vPC domain, Domain ID 1 is created across Cisco Nexus 5548UP member switches to define vPCs to carry specific VLAN network traffic. In this topology, we defined 5 vPCs with IDs. vPC ID 1 is defined to Peer link communication between 2 x Nexus switches in Fabric A & B. vPC IDs 17 and 18 are defined for traffic from Cisco UCS fabric interconnects, and vPC IDs 20 and 30 are defined for NFS Storage traffic to NetApp Array. These vPCs are managed within the Cisco Nexus 5548UP, which connects Cisco UCS fabric interconnects and the NetApp storage system.

The following are the steps to configure vPC.

1. Login into N5548-A as admin.

```
N5548-Fab-A# config terminal
N5548-Fab-A(config)#feature vpc
N5548-Fab-A(config)#vpc domain 1
N5548-Fab-A(config-vpc-domain)# peer-keepalive destination  <Mgmt. IP Address of
peer-N5548-B>
N5548-Fab-A(config-vpc-domain)# exit

N5548-Fab-A(config)# interface port-channel  1
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# vpc peer-link
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 1,10,134,20,30
N5548-Fab-A(config-if)# spanning-tree port type network
N5548-Fab-A(config-if)# exit
N5548-Fab-A(config)# interface Ethernet1/1
N5548-Fab-A(config-if)# description Peer link connected to N5548B-Eth1/1
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 1,10,20,30,134
N5548-Fab-A(config-if)# channel-group 1 mode active
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
N5548-Fab-A(config)# interface Ethernet1/2
N5548-Fab-A(config-if)# description Peer link connected to N5548B-Eth1/2
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 1,10,20,30,134
N5548-Fab-A(config-if)# channel-group 1 mode active
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
N5548-Fab-A(config)# copy running-config startup-config
```

2. Login into N5548-B as admin.

```
N5548-Fab-B(config)# conf term
N5548-Fab-B(config)# feature vpc
N5548-Fab-B(config)# vpc domain 1
N5548-Fab-B(config-vpc-domain# peer-keepalive destination  <Mgmt. IP Address of
peer-N5548-A>
N5548-Fab-B(config-vpc-domain)# exit

N5548-Fab-B(config)# interface port-channel  1
N5548-Fab-B(config-if)# description Port-Channel for vPC Peer-link
N5548-Fab-B(config-if)# switchport mode trunk
N5548-Fab-B(config-if)# vpc peer-link
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 1,10,134,20,30
N5548-Fab-B(config-if)# spanning-tree port type network
N5548-Fab-B(config-if)# exit
N5548-Fab-B(config)# interface Ethernet1/1
N5548-Fab-B(config-if)# description Peer link connected to N5548A-Eth1/1
N5548-Fab-B(config-if)# switchport mode trunk
```

```
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 1,10,20,30,134
N5548-Fab-B(config-if)# channel-group 1 mode active
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
N5548-Fab-B(config)# interface Ethernet1/2
N5548-Fab-B(config-if)# description Peer link connected to N5548A-Eth1/2
N5548-Fab-B(config-if)# switchport mode trunk
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 1,10,20,30,134
N5548-Fab-B(config-if)# channel-group 1 mode active
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
N5548-Fab-B(config)# copy running-config startup-config
```

Configure the Virtual Port Channel for the Data Network

Table 4 shows vPC configuration details for Cisco UCS 6248UP Fabric Interconnects A and B with required vPC IDs, VLAN IDs, and Ethernet uplink ports.

Table 4 vPC Port Channel for the Cisco Nexus 5548UP Switch

vPC Name	vPC ID on Nexus 5548UP	Fabric Interconnect Uplink Ports	Nexus 5548 Uplink Ports	VLAN ID
vPC-Public 1	17	Fab-A Eth 1/17 & Fab-A Eth 1/18	Fab-A Eth 1/17 & Fab-B Eth 1/17	134 (management), 10 (Private Interconnect), Public Access, Virtual IP, SCAN IP 20 (NFS Storage)
vPC-Private 2	18	Fab-B Eth 1/17 & Fab-B Eth 1/18	Fab-A Eth 1/18 & Fab-B Eth 1/18	134 (management), 10 (Private Interconnect), Public Access, Virtual IP, SCAN IP 30 (NFS Storage)

On Cisco UCS Fabric Interconnect A, Ethernet uplink ports 17 and 18 are connected to Cisco Nexus 5548UP Fabric A (port 1/17) and Cisco Nexus 5548UP Fabric B (port 1/17), which are part of vPC ID 17 and have access to Oracle DB Public and Private VLAN IDs 134 and 10 & also have access to Storage Network VLAN IDs 20 & 30. The same configuration is replicated for vPC ID 18 on Fabric interconnect B, with ports 17 and 18 connected to port 8 of Cisco Nexus 5548UP Fabric A (port 1/18) and Cisco Nexus 5548UP Fabric B (port 1/18) and it has same access like Nexus 5548 Fabric A switch.

The following are the configuration details for the Cisco Nexus 5548UP switches.

Port Channel Configuration on the Cisco Nexus 5548 Fabric-A

```
N5548-Fab-A(config)# interface Port-channel 17
N5548-Fab-A(config-if)# description Port-Channel for Fabric InterconnectA
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 134,10,20,30
N5548-Fab-A(config-if)# vPC 17
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
```

```
N5548-Fab-A(config)# interface Port-channel 18
N5548-Fab-A(config-if)# description Port-Channel for Fabric InterconnectB
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 134,10,20,30
N5548-Fab-A(config-if)# vPC 18
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit

N5548-Fab-A# config Terminal
N5548-Fab-A(config)# interface eth 1/17
N5548-Fab-A(config-if)# description Connection from Fabric Interconnect-A
Eth1/17
N5548-Fab-A(config-if)# channel-group 17 mode active
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
N5548-Fab-A(config)# interface eth 1/18
N5548-Fab-A(config-if)# description Connection from Fabric Interconnect-B
Eth1/17
N5548-Fab-A(config-if)# channel-group 18 mode active
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
```

Port Channel Configuration on the Cisco Nexus 5548 Fabric-B

```
N5548-Fab-B(config)# interface Port-channel 17
N5548-Fab-B(config-if)# description Port-Channel for Fabric InterconnectA
N5548-Fab-B(config-if)# switchport mode trunk
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 134,10,20,30
N5548-Fab-B(config-if)# vPC 17
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
N5548-Fab-B(config)# interface Port-channel 18
N5548-Fab-B(config-if)# description Port-Channel for Fabric InterconnectB
N5548-Fab-B(config-if)# switchport mode trunk
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 134,10,20,30
N5548-Fab-B(config-if)# vPC 18
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit

N5548-Fab-B# config Terminal
N5548-Fab-B(config)# interface eth 1/17
N5548-Fab-B(config-if)# description Connection from Fabric Interconnect-A
Eth1/18
N5548-Fab-B(config-if)# channel-group 17 mode active
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
N5548-Fab-B(config)# interface eth 1/18
N5548-Fab-B(config-if)# description Connection from Fabric Interconnect-B
Eth1/18
N5548-Fab-B(config-if)# channel-group 18 mode active
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
```

This establishes the required network connectivity on the Cisco UCS server side and now you will need to complete the steps to connect to NetApp storage.

Configure Virtual Port Channel for NFS Storage Network

On the Cisco Nexus 5548UP Switch, a separate vPC is created to access NetApp shared storage for NFS data access. The vPC is created with the vPC name and corresponding vPC ID and required VLAN IDs, as shown in Table 5.

Table 5 vPC Configuration Details for the Cisco Nexus 5548UP for NetApp Storage Access

vPC Name	vPC ID	10GigE Eth Ports (Controllers A and B)	Nexus 5548 Ethernet Ports	VLAN ID
vPC-Storage 1	20	e1b and e1c (Controller A)	Fab-A Eth 2/1 & Fab-B Eth 2/1	20
vPC-Storage 2	30	e1b and e1c (Controller B)	Fab-A Eth 2/3 & Fab-B Eth 2/3	30

On NetApp Storage Controller A, Ethernet 10-Gbps port e1b is connected to Cisco Nexus 5548-Fab-A (Eth 2/1), and Ethernet port e1c is connected to Cisco Nexus 5548-Fab-B (Eth 2/1), which are part of vPC-Storage1 with vPC ID 20 that allows traffic from VLAN ID 20. On NetApp Storage Controller B, Ethernet 10-Gbps port e1b is connected to Cisco Nexus 5548-Fab-A (Eth 2/3), and Ethernet port e1c is connected to Cisco Nexus 5548-Fab-B (Eth 2/3), which are part of vPC-Storage2 with vPC ID 30 that allows traffic from VLAN ID 30.

Port Channel Configuration on the Cisco Nexus 5548 Fabric-A

```
N5548-Fab-A(config)# interface Port-channel20
N5548-Fab-A(config-if)# description PortChannel for multimode VIF from
ControllerA-10G
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# switchport trunk native VLAN 20
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 20,30
N5548-Fab-A(config-if)# vPC 20
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit

N5548-Fab-A(config)# interface Port-channel30
N5548-Fab-A(config-if)# description PortChannel for multimode VIF from
ControllerB-10G
N5548-Fab-A(config-if)# switchport mode trunk
N5548-Fab-A(config-if)# switchport trunk native VLAN 30
N5548-Fab-A(config-if)# switchport trunk allowed VLAN 20,30
N5548-Fab-A(config-if)# vPC 30
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
```

Interface Configuration

```
N5548-Fab-A(config)# interface Ethernet 2/1
N5548-Fab-A(config-if)# description Connection to NetApp Controller-A-Port-e1a
N5548-Fab-A(config-if)# channel-group 20 mode active
N5548-Fab-A(config-if)# spanning-tree portfast
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
```

```
N5548-Fab-A(config)# interface Ethernet 2/3
N5548-Fab-A(config-if)# description Connection to NetApp Controller-B-Port-e1a
N5548-Fab-A(config-if)# channel-group 30 mode active
N5548-Fab-A(config-if)# spanning-tree portfast
N5548-Fab-A(config-if)# no shutdown
N5548-Fab-A(config-if)# exit
```

Port Channel Configuration on the Cisco Nexus 5548 Fabric-B

```
N5548-Fab-B(config)# interface Port-channel20
N5548-Fab-B(config-if)# description PortChannel for multimode VIF from
ControllerA-10G
N5548-Fab-B(config-if)# switchport mode trunk
N5548-Fab-B(config-if)# switchport trunk native VLAN 20
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 20,30
N5548-Fab-B(config-if)# vPC 20
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit

N5548-Fab-B(config)# interface Port-channel30
N5548-Fab-B(config-if)# description PortChannel for multimode VIF from
ControllerB-10G
N5548-Fab-B(config-if)# switchport mode trunk
N5548-Fab-B(config-if)# switchport trunk native VLAN 30
N5548-Fab-B(config-if)# switchport trunk allowed VLAN 20,30
N5548-Fab-B(config-if)# vPC 30
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
```

Interface Configuration

```
N5548-Fab-B(config)# interface Ethernet 2/1
N5548-Fab-B(config-if)# description Connection to NetApp Controller-A-Port-e1b
N5548-Fab-B(config-if)# channel-group 20 mode active
N5548-Fab-B(config-if)# spanning-tree portfast
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit

N5548-Fab-B(config)# interface Ethernet 2/3
N5548-Fab-B(config-if)# description Connection to NetApp Controller-B-Port-e1b
N5548-Fab-B(config-if)# channel-group 30 mode active
N5548-Fab-B(config-if)# spanning-tree portfast
N5548-Fab-B(config-if)# no shutdown
N5548-Fab-B(config-if)# exit
```

When configuring the Cisco Nexus 5548UP with vPCs, be sure that the status for all vPCs is "Up" for connected Ethernet ports by running the commands shown in Figure 14 from the CLI on the Cisco Nexus 5548UP Switch.

Figure 47 Port Channel Status on the Cisco Nexus 5548UP

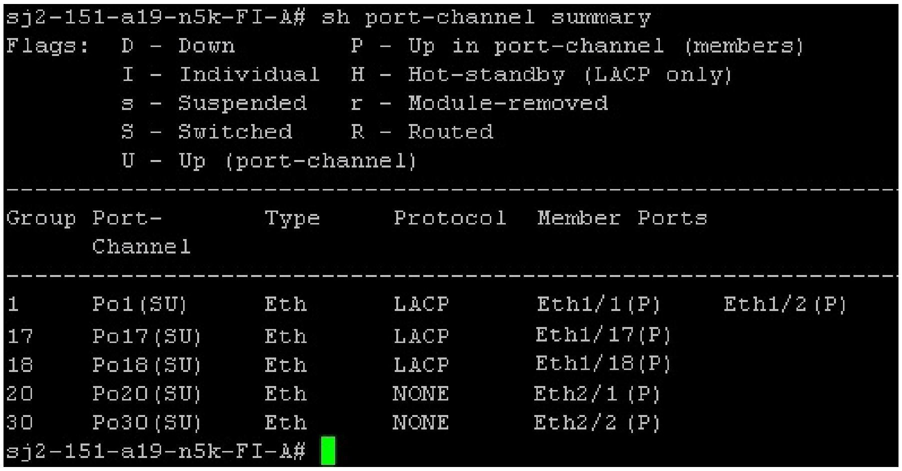

```
sj2-151-a19-n5k-FI-A# sh port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
17    Po17(SU)    Eth      LACP      Eth1/17(P)
18    Po18(SU)    Eth      LACP      Eth1/18(P)
20    Po20(SU)    Eth      NONE      Eth2/1(P)
30    Po30(SU)    Eth      NONE      Eth2/2(P)
sj2-151-a19-n5k-FI-A#
```

Figure 48 Port Channel Summary on the Cisco Nexus 5548UP

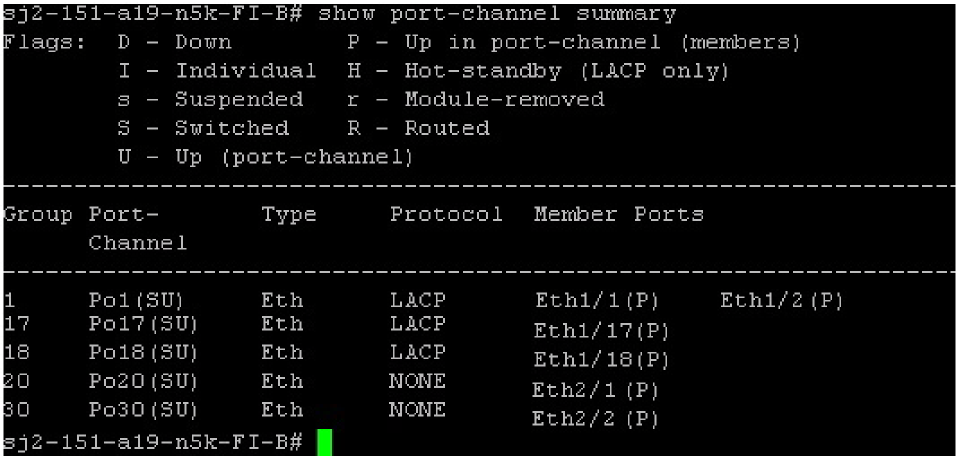

```
sj2-151-a19-n5k-FI-B# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
17    Po17(SU)    Eth      LACP      Eth1/17(P)
18    Po18(SU)    Eth      LACP      Eth1/18(P)
20    Po20(SU)    Eth      NONE      Eth2/1(P)
30    Po30(SU)    Eth      NONE      Eth2/2(P)
sj2-151-a19-n5k-FI-B#
```

Show vPC status should display the following information for a successful configuration.

Figure 49 Virtual Port Channel Status on the Cisco Nexus 5548UP Fabric A Switch

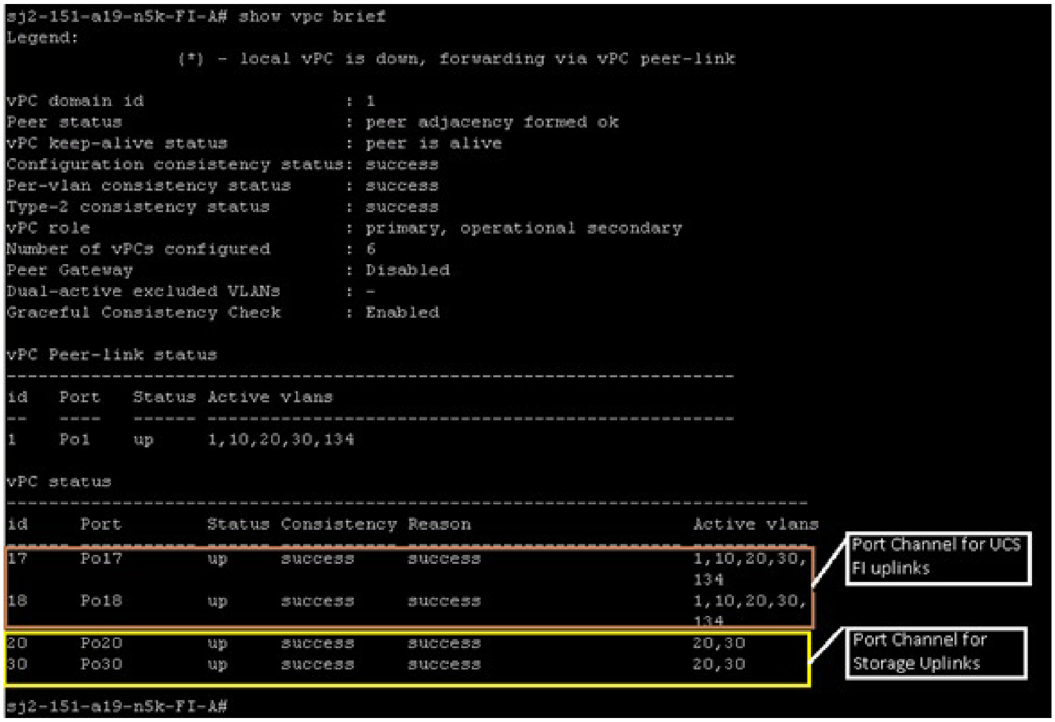

```
sj2-151-a19-n5k-FI-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 6
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     1,10,20,30,134

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                      Active vlans
------ ----------- ------ ----------- --------------------------- ------------
17     Po17        up     success     success                     1,10,20,30,
                                                                  134
18     Po18        up     success     success                     1,10,20,30,
                                                                  134
20     Po20        up     success     success                     20,30
30     Po30        up     success     success                     20,30
sj2-151-a19-n5k-FI-A#
```

Port Channel for UCS FI uplinks
Port Channel for Storage Uplinks

Figure 50 Port Channel Status on the Cisco Nexus 5548UUP on Fabric B Switch

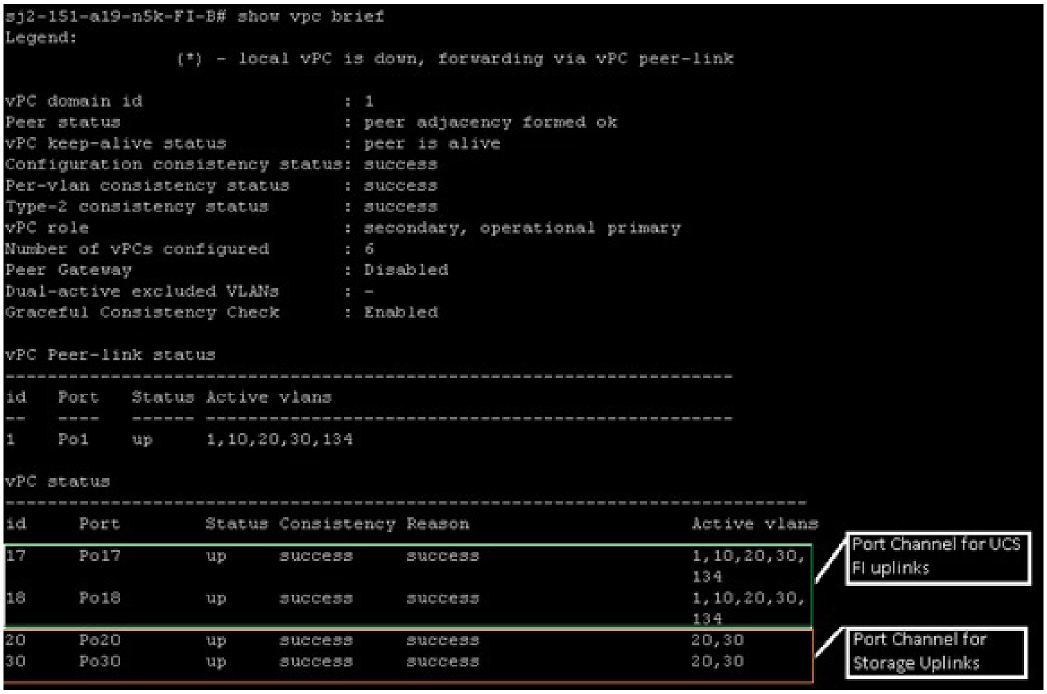

```
sj2-151-a19-n5k-FI-B# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary, operational primary
Number of vPCs configured         : 6
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     1,10,20,30,134

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                      Active vlans
17     Po17        up     success     success                     1,10,20,30,
                                                                  134
18     Po18        up     success     success                     1,10,20,30,
                                                                  134
20     Po20        up     success     success                     20,30
30     Po30        up     success     success                     20,30
```

Port Channel for UCS FI uplinks
Port Channel for Storage Uplinks

Setting Up Jumbo Frames on the Cisco Nexus 5548UP

Jumbo frames with an MTU=9000 have to be setup on both Nexus5548UP switches. Please note that Oracle private interconnect traffic does not leave the Cisco UCS domain (Fabric Interconnect) in normal operating conditions. However, if there is a partial link or IOM failure, the private interconnect traffic has to be routed to the immediate northbound switch (Cisco Nexus5548 UP in our case). Since we are using Jumbo Frames as a best practice for Oracle private interconnect, we need to have jumbo frames configured on the Cisco Nexus 5548UP switches.

The following section describes how to enable jumbo frames on the Cisco Nexus 5548UP Fabric A Switch.

```
N5548-Fab-A# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
N5548-Fab-A(config)# class-map type network-qos class-platinum
N5548-Fab-A(config-cmap-nq)# exit
N5548-Fab-A(config)# policy-map type network-qos jumbo
N5548-Fab-A(config-pmap-nq)# class type network-qos class-default
N5548-Fab-A(config-pmap-nq-c)# mtu 9216
N5548-Fab-A(config-pmap-nq-c)# multicast-optimize
N5548-Fab-A(config-pmap-nq-c)# exit
N5548-Fab-A(config-pmap-nq)# system qos
N5548-Fab-A(config-sys-qos)# service-policy type network-qos jumbo
N5548-Fab-A(config-sys-qos)# exit
N5548-Fab-A(config)# copy running-config startup-config
[########################################] 100%
N5548-Fab-A(config)#
```

Repeat these steps to configure Jumbo Frames on Nexus 5548UP Fabric B Switch. This completes the Cisco Nexus 5548 switch configuration. The next step is to configure the NetApp storage.

NetApp Storage Configuration Overview

This section discusses the NetApp storage layout design considerations when deploying an Oracle Database 11g R2 RAC on FlexPod.

Figure 51 depicts a high-level storage design overview of a NetApp FAS3270 HA storage system

Figure 51 Design Overview of NetApp High-Availability Storage

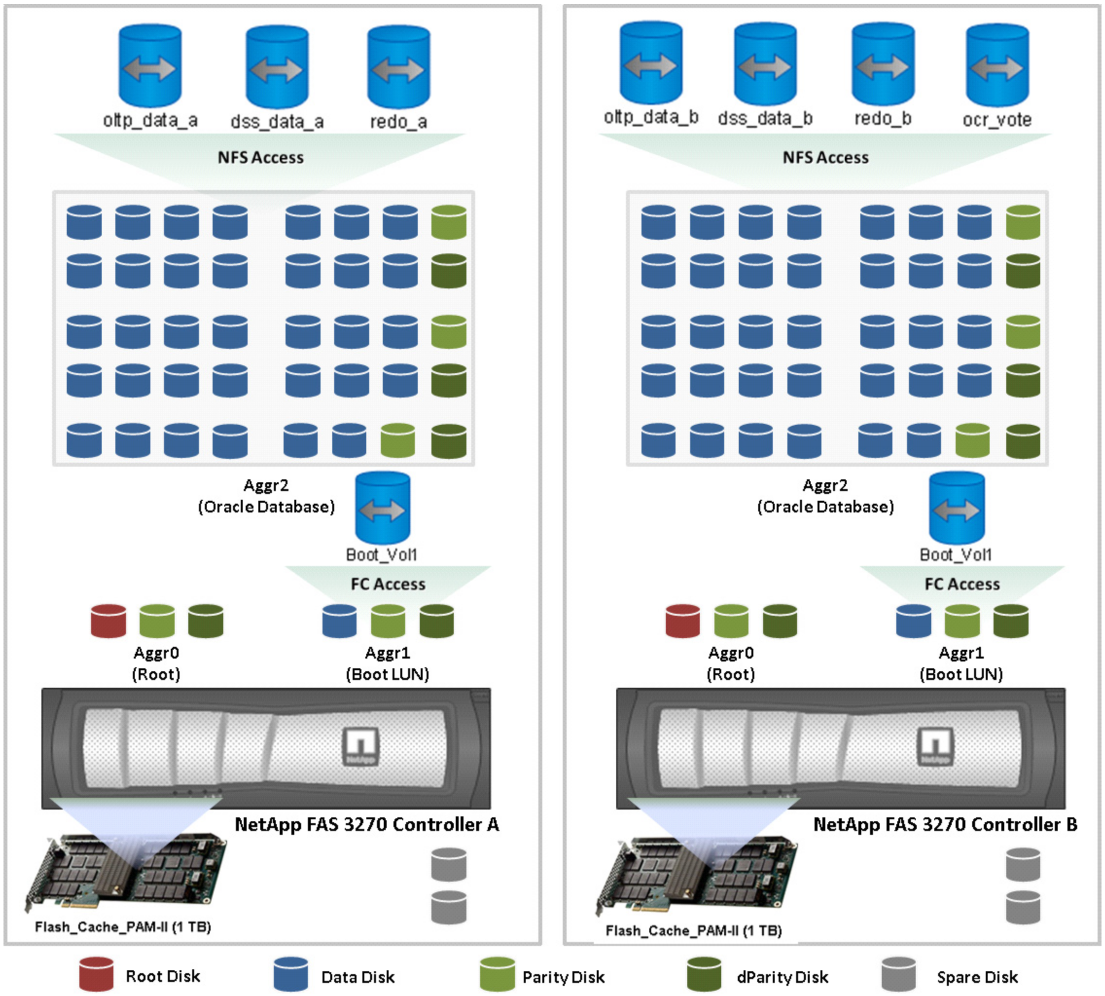

Table 6 shows the NetApp storage layout with volumes and LUNs created for various purposes.

Table 6 NetApp Storage Layout with Volumes and LUNs

NetApp Storage Layout		
Aggregation and NetApp Controller	**NetApp FlexVol or LUN**	**Comments (LUNs are only used for Boot)**
Aggr1 on Controller A	Boot_Vol1	FC SAN boot LUN for Oracle DB host for node 1 and node 2 of the failover cluster with Cisco UCS B200M3 blade server
Aggr2 on Controller A	oltp_data_a	oltp datafiles(odd number files), spfiles, and control files.
Aggr2 on Controller A	dss_data_a	dss datafiles(odd number files), spfiles, and control files.
Aggr2 on Controller A	redo_a	redo log files and control files.

Aggr1 on Controller B	Boot_Vol2	FC SAN boot LUN for Oracle DB host for node 3 and node 4 of the failover cluster with Cisco UCS B200M3 blade server
Aggr2 on Controller B	oltp_data_b	oltp datafiles (even number files), spfiles, and the copy of control files.
Aggr2 on Controller B	dss_data_b	dss datafiles (even number files), spfiles, and the copy of control files.
Aggr2 on Controller B	redo_b	redo log files and copy of control files.
Aggr2 on Controller B	OCR_VOTE_VOL	OCR and voting disk using NFS

Use the following commands to configure the NetApp storage systems to implement the storage layout design described here.

Storage Configuration for SAN Boot

Create and Configure Aggregate, Volumes and Boot LUNs

NetApp FAS3270HA Controller A

1. Create Aggregate and Volumes for Boot LUN is detailed in the SAN boot NetApp Storage Configuration section. (Shown the same steps as below)
2. Create Aggr1 with a RAID group size of 4, 4 disks, and RAID_DP redundancy for hosting NetApp FlexVol volumes and LUNs.
3. FAS3270HA-Controller A> aggr create aggr1 -t raid_dp -r 4 4
4. Create NetApp FlexVol volumes on Aggr1 for hosting FC Boot LUNs. These volumes are exposed to Cisco UCS blades for Booting Oracle Linux over SAN
5. FAS3270HA-Controller A> vol create Boot_VOL1 aggr1 900g
6. Create Boot LUNs on NetApp FlexVol volumes for booting Oracle Linux 5.8 over SAN. The example is of creating Boot LUN for Orarac_Node1.
7. FAS3270HA-Controller A> lun create -s 200g -t BootLUN-Node1
8. /vol/Boot_VOL1
9. Repeat step 3, to create Boot LUNs for the hosts Orarac-Node2.

NetApp FAS3270HA Controller B

1. Create Aggr1 with a RAID group size of 4, 4 disks, and RAID_DP redundancy for hosting NetApp FlexVol volumes and LUNs.
2. FAS3270HA-Controller B> aggr create aggr1 -t raid_dp -r 4 4
3. Create NetApp FlexVol volumes on Aggr1 for hosting FC Boot LUNs. These volumes are exposed to Cisco UCS blades for booting Oracle Linux over SAN
4. FAS3270HA-Controller B> vol create Boot_VOL1 aggr1 900g
5. Create Boot LUNs on NetApp FlexVol volumes for booting Oracle Linux 5.8 over SAN. The example is of creating Boot LUN for Orarac_Node3.

6. FAS3270HA-Controller A> lun create -s 200g -t BootLUN-Node3 /vol/Boot_VOL1
7. Repeat step 3, to create Boot LUNs for the hosts Orarac-Node4.

Create and Configure Initiator Group (igroup) and LUN Mapping

NetApp FAS3270HA Controller A

1. Create Initiator group (Igroup) and map the LUNs to the specific host OraRac-node1.

```
FAS3270HA-Controller A> igroup create -i -t linux Orarac_Node1-  group1
20:00:01:25:b5:11:13:02
FAS3270HA-Controller A> lun map /vol/Boot_Vol1/BootLUN-Node1 Orarac_Node1-group1
0
```

2. Repeat step 1, to create Initiator group and map Boot LUNs to the hosts Orarac-Node2.

NetApp FAS3270HA Controller B

1. Create Initiator group (Igroup) and map the LUNs to the specific host OraRac-node3.

```
FAS3270HA-Controller A> igroup create -i -t linux Orarac_Node3-
group1 20:00:01:25:b5:11:13:06
FAS3270HA-Controller A> lun map /vol/Boot_Vol1/BootLUN-Node3
Orarac_Node3-group1 0
```

2. Repeat step 1, to create Initiator group and map Boot LUNs to the hosts Orarac-Node4.

Storage Configuration for NFS Storage Network:

Create and Configure Aggregate and Volumes

NetApp FAS3270HA Controller A

1. Create Aggr2 with a RAID group size of 10, 40 disks, and RAID_DP redundancy for hosting NetApp FlexVol volumes and LUNs.

```
FAS3270HA-Controller A> aggr create aggr2 -t raid_dp -r 10 40
```

2. Create NetApp FlexVol volumes on Aggr2 for oltp & dss data files. These volumes are exposed to Oracle RAC nodes.

```
FAS3270HA-Controller A> vol create oltp_data_a aggr2 6144g
FAS3270HA-Controller A> vol create dss_data_a aggr2 3096g
FAS3270HA-Controller A> vol create redo_a aggr2 500g
```

NetApp FAS3270HA Controller B

1. Create Aggr2 with a RAID group size of 10, 40 disks, and RAID_DP redundancy for hosting NetApp FlexVol volumes and LUNs.

```
FAS3270HA-Controller B> aggr create aggr2 -t raid_dp -r 10 40
```

2. Create NetApp FlexVol volumes on Aggr2 for oltp & dss data files. These volumes are exposed to Oracle RAC nodes.

```
FAS3270HA-Controller B> vol create oltp_data_b aggr2 6144g
FAS3270HA-Controller B> vol create dss_data_b aggr2 3096g
FAS3270HA-Controller B> vol create redo_b aggr2 500g
FAS3270HA-Controller B> vol create ocr_vote aggr2 25g
```

NFS exports all the flexible volumes (data volumes, redo log volumes, and OCR and voting disk volumes) from both Controller A and Controller B, providing read/write access to the root user of all hosts created in the previous steps.

Create and Configure VIF Interface (Multimode)

Make sure that the NetApp multimode virtual interface (VIF) feature is enabled on NetApp storage systems on 10 Gigabit Ethernet ports (e1a and e1b) for NFS Storage access. We used the same VIF to access all flexible volumes created to store Oracle Database files that are using the NFS protocol. Your best practices may vary depending upon setup.

VIF Configuration on Controller A

```
ControllerA>ifgrp create multimode VIF19 -b ip e1a e1b
ControllerA>ifconfig VIF19 10.10.20.5 netmask 255.255.255.0 mtusize 9000 partner
VIF20
ControllerA>ifconfig VIF19 up
```

VIF Configuration on Controller B

```
ControllerB>ifgrp create multimode VIF20 -b ip e1a e1b
ControllerB>ifconfig VIF19 10.10.30.5 netmask 255.255.255.0 mtusize 9000 partner
VIF19
ControllerB>ifconfig VIF20 up
```

Note Make the changes persistent

```
ControllerA:: /etc/rc
 Hostname CONTROLLERA
 vif create multimode VIF19 -b ip e1b e1a
 ifconfig net `hostname`-net mediatype auto netmask 255.255.255.0 partner VIF20
 route add default 10.29.150.1 1
 routed on
 options dns.domainname example.com
 options dns.enable on
 options nis.enable off
 savecore

ControllerB:: /etc/rc
 hostname CONTROLLERB
 vif create multimode VIF20 -b ip e1b e1a
 ifconfig net `hostname`-net mediatype auto netmask 255.255.255.0 partner VIF19
 route add default 10.29.150.1 1
 routed on
 options dns.domainname example.com
 options dns.enable on
 options nis.enable off
 savecore
```

Check the NetApp configuration

```
Controller A:> Vif status VIF 19
Controller B:> Vif status VIF 20
```

Make sure that the MTU is set to 9000 and that jumbo frames are enabled on the Cisco UCS static and dynamic vNICs and on the upstream Cisco Nexus 5548UP switches.

Figure 52 shows the virtual interface "VIF19" & VIF "20" created with the MTU size set to 9000 and the trunk mode set to multiple, using two 10 Gigabit Ethernet ports (e1a and e1b) on NetApp storage Controller A and Controller B respectively.

Figure 52 Virtual Interface (VIF19) Creation on Controller A

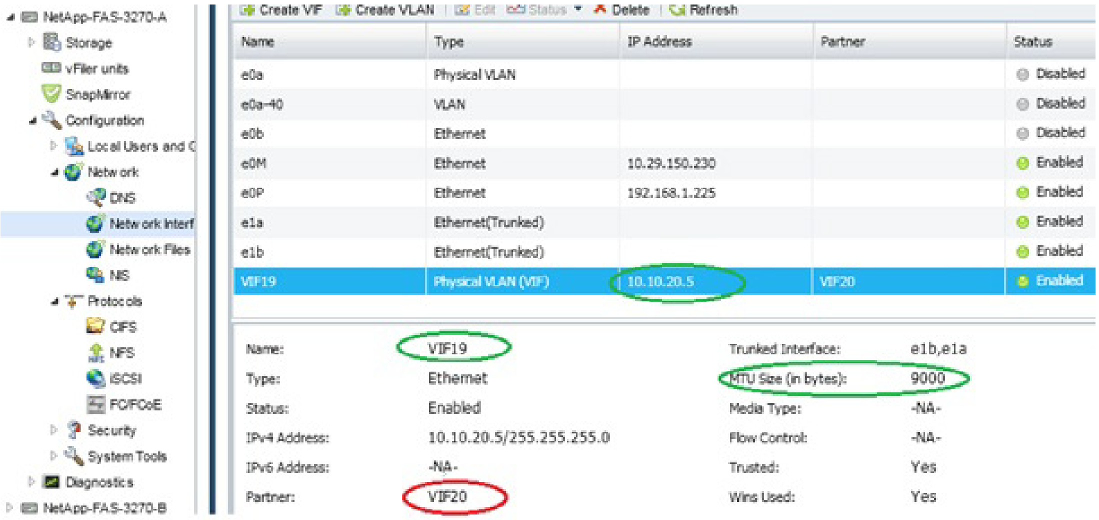

Name	Type	IP Address	Partner	Status
e0a	Physical VLAN			Disabled
e0a-40	VLAN			Disabled
e0b	Ethernet			Disabled
e0M	Ethernet	10.29.150.230		Enabled
e0P	Ethernet	192.168.1.225		Enabled
e1a	Ethernet(Trunked)			Enabled
e1b	Ethernet(Trunked)			Enabled
VIF19	Physical VLAN (VIF)	10.10.20.5	VIF20	Enabled

Name:	VIF19	Trunked Interface:	e1b,e1a
Type:	Ethernet	MTU Size (in bytes):	9000
Status:	Enabled	Media Type:	-NA-
IPv4 Address:	10.10.20.5/255.255.255.0	Flow Control:	-NA-
IPv6 Address:	-NA-	Trusted:	Yes
Partner:	VIF20	Wins Used:	Yes

Figure 53 Virtual Interface (VIF20) Creation on Controller B

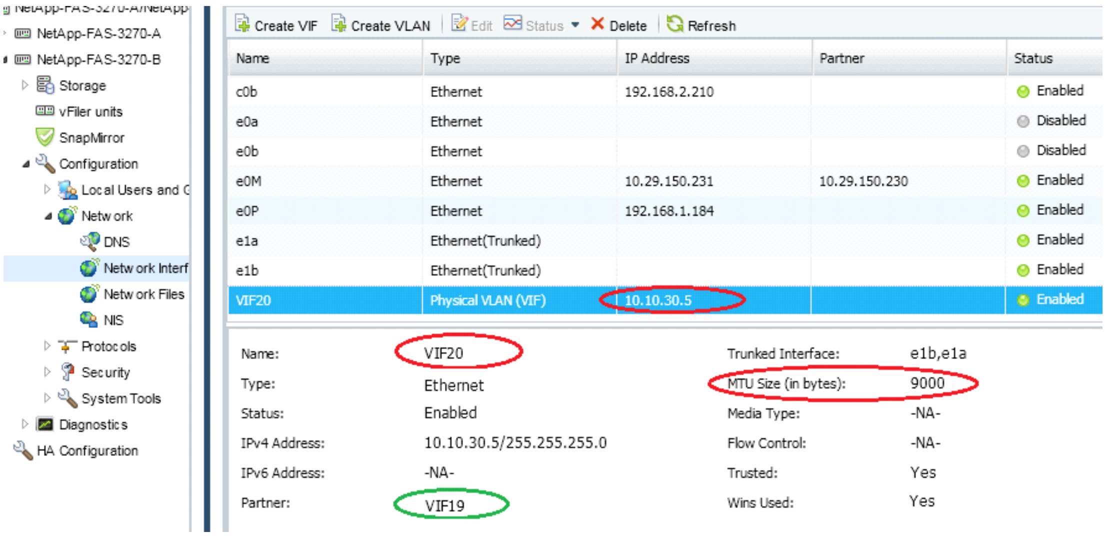

Name	Type	IP Address	Partner	Status
c0b	Ethernet	192.168.2.210		Enabled
e0a	Ethernet			Disabled
e0b	Ethernet			Disabled
e0M	Ethernet	10.29.150.231	10.29.150.230	Enabled
e0P	Ethernet	192.168.1.184		Enabled
e1a	Ethernet(Trunked)			Enabled
e1b	Ethernet(Trunked)			Enabled
VIF20	Physical VLAN (VIF)	10.10.30.5		Enabled

Name:	VIF20	Trunked Interface:	e1b,e1a
Type:	Ethernet	MTU Size (in bytes):	9000
Status:	Enabled	Media Type:	-NA-
IPv4 Address:	10.10.30.5/255.255.255.0	Flow Control:	-NA-
IPv6 Address:	-NA-	Trusted:	Yes
Partner:	VIF19	Wins Used:	Yes

Figure 54 Virtual Network Interface (VIF20) Properties

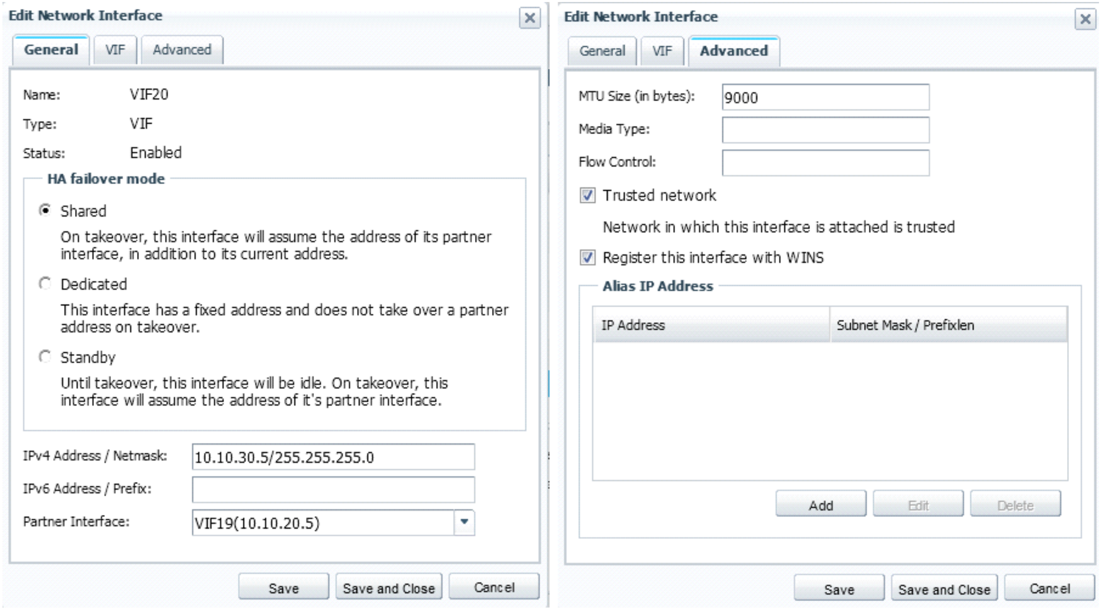

Name: VIF20
Type: VIF
Status: Enabled

HA failover mode

● Shared
On takeover, this interface will assume the address of its partner interface, in addition to its current address.

○ Dedicated
This interface has a fixed address and does not take over a partner address on takeover.

○ Standby
Until takeover, this interface will be idle. On takeover, this interface will assume the address of it's partner interface.

IPv4 Address / Netmask: 10.10.30.5/255.255.255.0
IPv6 Address / Prefix:
Partner Interface: VIF19(10.10.20.5)

MTU Size (in bytes): 9000
Media Type:
Flow Control:

☑ Trusted network
Network in which this interface is attached is trusted

☑ Register this interface with WINS

Alias IP Address

IP Address	Subnet Mask / Prefixlen

This completes the storage configuration. The next step is to review the boot from SAN details.

Cisco UCS Blade Servers and Stateless Computing with SAN Boot

Boot from SAN Benefits

Booting from SAN is another key feature which helps in moving towards stateless computing in which there is no static binding between a physical server and the OS / applications it is tasked to run. The OS is installed on a SAN LUN and boot from SAN policy is applied to the service profile template or the service profile. If the service profile were to be moved to another server, the pwwn of the HBAs and the Boot from SAN (BFS) policy also moves along with it. The new server now takes the same exact character of the old server, providing the true unique stateless nature of the UCS Blade Server.

The key benefits of booting from the network:

- Reduce Server Footprints: Boot from SAN alleviates the necessity for each server to have its own direct-attached disk, eliminating internal disks as a potential point of failure. Thin diskless servers also take up less facility space, require less power, and are generally less expensive because they have fewer hardware components.
- Disaster and Server Failure Recovery: All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. If a disaster destroys functionality of the servers at the primary site, the remote site can take over with minimal downtime.

  Recovery from server failures is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of its image. As a result, boot from SAN can greatly reduce the time required for server recovery.

- High Availability: A typical data center is highly redundant in nature - redundant paths, redundant disks and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for mechanical failure of a local disk.
- Rapid Redeployment: Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage a cost effective endeavor.

With Boot from SAN, the image resides on a SAN LUN and the server communicates with the SAN through a host bus adapter (HBA). The HBAs BIOS contain the instructions that enable the server to find the boot disk. All FC-capable Converged Network Adapter (CNA) cards supported on Cisco UCS B-series blade servers support Boot from SAN.

After power on self-test (POST), the server hardware component fetches the boot device that is designated as the boot device in the hardware BOIS settings. When the hardware detects the boot device, it follows the regular boot process.

Summary for Boot from SAN Configuration

The following boot from SAN configuration steps have been completed:

- SAN Zoning configuration on the Nexus 5548UP switches
- NetApp Storage Array Configuration for Boot LUN
- Cisco UCS configuration of Boot from SAN policy in the service profile

The next step is to install the OS.

Note The steps to complete an OS installation in a SAN Boot configuration are not detailed in this document.

OS Installation Steps and Recommendations

For this solution, we configured a 4-node Oracle Database 11g R2 RAC cluster using Cisco B200 M3 servers. The servers used boot from SAN to enable stateless computing in case a need arises to replace/swap the server using UCS unique service profile capabilities. While OS boot is using Fiber channel, the databases and grid infrastructure components were configured to use NFS protocol on the NetApp storage. Oracle Linux 5.8 64-bit Operating System is installed on each server.

Table 7 summarizes the hardware and software configuration details.

Table 7 Host Configuration

Component	Details	Description
Server	4xB200 M3	2 Sockets with 8 cores with HT enabled
Memory	256 GB	Physical memory
NIC1	Public Access	Management and Public Access, MTU Size 1500
NIC2	Private Interconnect	Private Interconnect configured for HAIP, MTU Size 9000
NIC3	NFS Storage Access	Database access through NFS Storage to Filer A, MTU size 9000

NIC4	NFS Storage Access	Database access through NFS Storage to Filer B, MTU size 9000
SWAP Space	64 GB	Swap Space

Oracle Database 11g R2 GRID Infrastructure with RAC Option Deployment

This section describes the high -evel steps for Oracle Database 11g R2 RAC install. Prior to GRID and database install, verify all the prerequisites are completed. As per best practice you can install Oracle validated RPM that will ensure all prerequisites are met before Oracle grid install. We will not cover step-by-step install for Oracle GRID in this document but will provide partial summary of details that might be relevant.

Use the following Oracle document for pre-installation tasks, such as setting up the kernel parameters, RPM packages, user creation, etc.

http://download.oracle.com/docs/cd/E11882_01/install.112/e10812/prelinux.htm#BABHJHCJ

Oracle Grid Infrastructure ships with the Cluster Verification Utility (CVU) that can be run to validate pre and post installation configurations. For more details on CVU please see Oracle® Clusterware Administration and Deployment Guide 11g Release 2 (11.2) Appendix A.

We created following local directory structure and ownerships on each RAC nodes and

```
mkdir -p /u01/app/11.2.0/grid
mkdir -p /u01/app/oracle
mkdir /data_1
mkdir /data_2
mkdir /dss_data_a
mkdir /dss_data_b
mkdir /redo_1
mkdir /redo_2
mkdir /ocrvote
chown -R oracle:oinstall  /u01/app/oracle  /data_1   /data_2   /dss_data_a
/dss_data_b /redo_1 /redo_2
chmod -R 775   /u01/app/oracle /data_1   /data_2   dss_data_a /dss_data_b
/redo_1 /redo_2
chown -R grid:oinstall  /u01/app   /ocrvote
chmod -R 775   /u01/app  /ocrvote
```

In this test case, we used local directory for GRID Installation and Database binary Installation. Alternatively, these binaries can be installed in a shared directory on NFS volumes though you lose the advantage of GRID Infrastructure and Database rolling upgrades.

Table 8 summarizes NFS Volume mappings with mount points for each RAC node.

Table 8 Local Mount Points and NetApp NFS Volumes

Location	NetApp NFS Volumes	Owner	Purpose
/u01/app/11.2.0/grid	NA	grid	Oracle GRID binary installation
/u01/app/oracle	NA	oracle	Oracle Database binary installation
/data_1	/vol/oltp_data_a	oracle	OLTP Data and control files
/data_2	/vol/oltp_data_b	oracle	OLTP Data and control files

/redo_1	/vol/redo_a	oracle	Redo log files for OLTP and DSS workloads
/redo_2	/vol/redo_b	oracle	Redo log files for OLTP and DSS workloads
/dss_data_a	/vol/dss_data_a	oracle	DSS data and control files
/dss_data_b	/vol/dss_data_b	oracle	DSS data and control files
/ocrvote	/vol/ocrvote	grid	OCR and voting disks

1. Edit /etc/fstab file in each RAC node and add mount points for all database and GRID NFS volumes with the appropriate mount options. Please note that these mount points need to be created first.

Note Oracle Direct NFS (dNFS) configuration steps will need to be performed at a later stage after database creation.

The following is a sample output from mount command on Node 1:

```
[root@orarac1 ~]# mount
…
…
10.10.30.5:/vol/ocrvote on /ocrvote type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.30.5)

10.10.20.5:/vol/redo_a on /redo_1 type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.20.5)

10.10.30.5:/vol/redo_b on /redo_2 type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.30.5)

10.10.20.5:/vol/oltp_data_a on /data_1 type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.20.5)

10.10.30.5:/vol/oltp_data_b on /data_2 type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.30.5)

10.10.20.5:/vol/dss_data_a on /dss_data_a type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.20.5)

10.10.30.5:/vol/dss_data_b on /dss_data_b type nfs
(rw,bg,hard,nointr,rsize=65536,wsize=65536,tcp,actimeo=0,nfsvers=3,timeo=600,add
r=10.10.30.5)
```

To determine the proper mount options for different file systems of Oracle 11g R2, see https://kb.netapp.com/support/index?page=content&id=3010189&actp=search&viewlocale=en_US&searchid

Note An rsize and wsize of 65536 is supported by NFS v3 and used in this configuration to improve performance.

2. Configure the private and public NICs with the appropriate IP addresses.

Multicast Requirements for Oracle Grid Infrastructure

With Oracle Grid Infrastructure release 2 (11.2), on each cluster member node, the Oracle mDNS daemon uses multicasting on all interfaces to communicate with other nodes in the cluster. Multicasting needs to be enabled for Oracle RAC

Across the broadcast domain as defined for the private interconnect

On the IP address subnet ranges 224.0.0.251/24 and/or 230.0.1.0/24. Multicast configuration support is required only one of this subnet ranges.

Subnet ranges of 230.0.1.0/24 are required only for Oracle releases prior to 11.2.0.3. To address this issue, Oracle has released Patch: 9974223 on top of 11.2.0.2. This patch must be applied before executing root.sh on each node in the cluster. This patch makes use of the 224.0.0.251 multicast network address in addition to the 230.0.1.0 (port 42424) multicast address. For Oracle releases 11.2.0.3 onwards oracle uses 224.0.0.251/24 and no patch is needed. Hence it is recommended to install 11.2.0.3 which is a full install. Running Cluster Verification Utility will spill out any requirements around multicasting. For details on multicasting with Oracle on 11gr2 please refer to metalink note 1212703.1

3. Identify the virtual IP addresses and SCAN IPs and have them setup in DNS per Oracle's recommendation. Alternatively, you can update the /etc/hosts file with all the details (private, public, SCAN and virtual IP) if you do not have DNS services available

4. Create files for OCR and voting devices under /ocrvote local directories as follows.

   Login as "grid" user from any one node and create the following raw files

   dd if=/dev/zero of=/ocrvote/ocr/ocr1 bs=1m count=1024

   dd if=/dev/zero of=/ocrvote/ocr/ocr2 bs=1m count=1024

   dd if=/dev/zero of=/ocrvote/ocr/ocr3 bs=1m count=1024

   dd if=/dev/zero of=/ocrvote/vote/vote1 bs=1m count=1024

   dd if=/dev/zero of=/ocrvote/vote/vote2 bs=1m count=1024

   dd if=/dev/zero of=/ocrvote/vote/vote3 bs=1m count=1024

5. Configure ssh option (with no password) for the Oracle user and grid user. For more information about ssh configuration, refer to the Oracle installation documentation.

Note Oracle Universal Installer also offers automatic SSH connectivity configuration and testing.

6. Configure "/etc/sysctl.conf" and update shared memory and semaphore parameters required for Oracle GRID Installation. Also configure "/etc/security/limits.conf" file by adding user limits for oracle and grid users.

Note Do not perform these steps if Oracle Validated RPM is installed.

7. Configure hugepages.

   Hugepages is a method to have larger page size that is useful for working with very large memory. For Oracle Databases, using HugePages reduces the operating system maintenance of page states, and increases Translation Lookaside Buffer (TLB) hit ratio.

Advantages of HugePages

- HugePages are not swappable so there is no page-in/page-out mechanism overhead.

- Hugepage uses fewer pages to cover the physical address space, so the size of "book keeping" (mapping from the virtual to the physical address) decreases, so it requiring fewer entries in the TLB and so TLB hit ratio improves.
- Hugepages reduces page table overhead.
- Eliminated page table lookup overhead: Since the pages are not subject to replacement, page table lookups are not required.
- Faster overall memory performance: On virtual memory systems each memory operation is actually two abstract memory operations. Since there are fewer pages to work on, the possible bottleneck on page table access is clearly avoided.

For our configuration, we used hugepages for both OLTP and DSS workloads. Please refer to Oracle metalink document 361323.1 for hugepages configuration details.

When hugepages are configured, You are now ready to install Oracle Database 11g R2 GRID Infrastructure with RAC option and the database.

Installing Oracle Database 11g R2 RAC

It is not within the scope of this document to include the specifics of an Oracle RAC installation; you should refer to the Oracle installation documentation for specific installation instructions for your Environment.

To install Oracle, follow these steps:

Step 1 Download the Oracle Database 11g Release 2 Grid Infrastructure (11.2.0.3.0) and Oracle Database 11g

Release 2 (11.2.0.3.0) for Linux x86-64.

Step 2 For this configuration, we used NFS shared volumes for OCR and voting disks for Oracle Grid Infrastructure install. Alternatively, you can install and use Oracle ASM for OCR and voting disks only. For more details, see Grid Infrastructure Installation Guide for Linux :

http://www.oracle.com/pls/db112/to_toc?pathname=install.112/e10812/toc.htm

Figure 55 Setting Up Oracle Grid Infrastructure for a Cluster

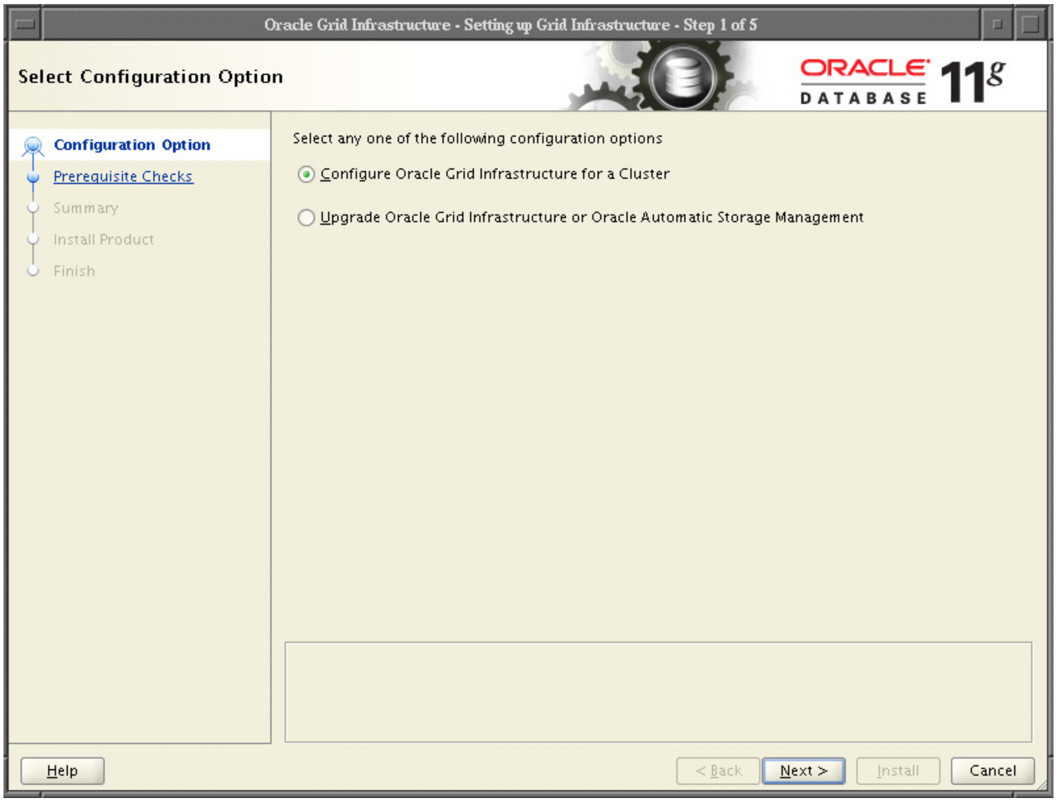

Oracle Grid Infrastructure - Setting up Grid Infrastructure - Step 1 of 5

Select Configuration Option

Select any one of the following configuration options

● Configure Oracle Grid Infrastructure for a Cluster

○ Upgrade Oracle Grid Infrastructure or Oracle Automatic Storage Management

Figure 56 Grid Plug and Play Information

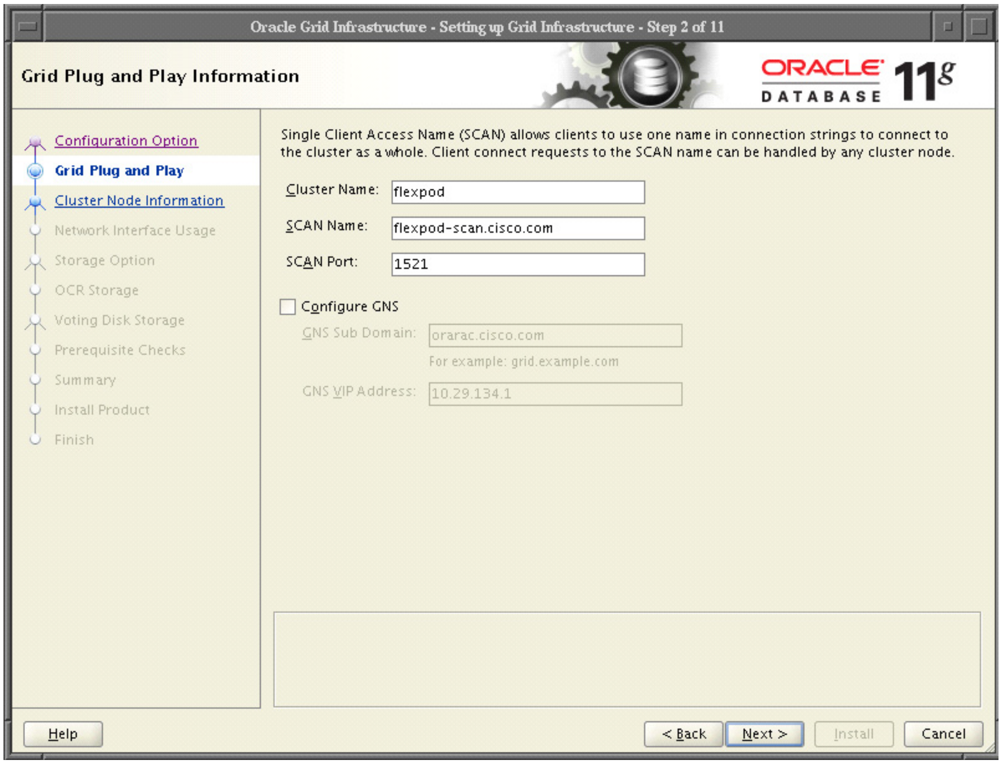

Single Client Access Name (SCAN) allows clients to use one name in connection strings to connect to the cluster as a whole. Client connect requests to the SCAN name can be handled by any cluster node.

Cluster Name: flexpod

SCAN Name: flexpod-scan.cisco.com

SCAN Port: 1521

□ Configure GNS

GNS Sub Domain: orarac.cisco.com

For example: grid.example.com

GNS VIP Address: 10.29.134.1

Figure 57 Configure Cluster Nodes Information

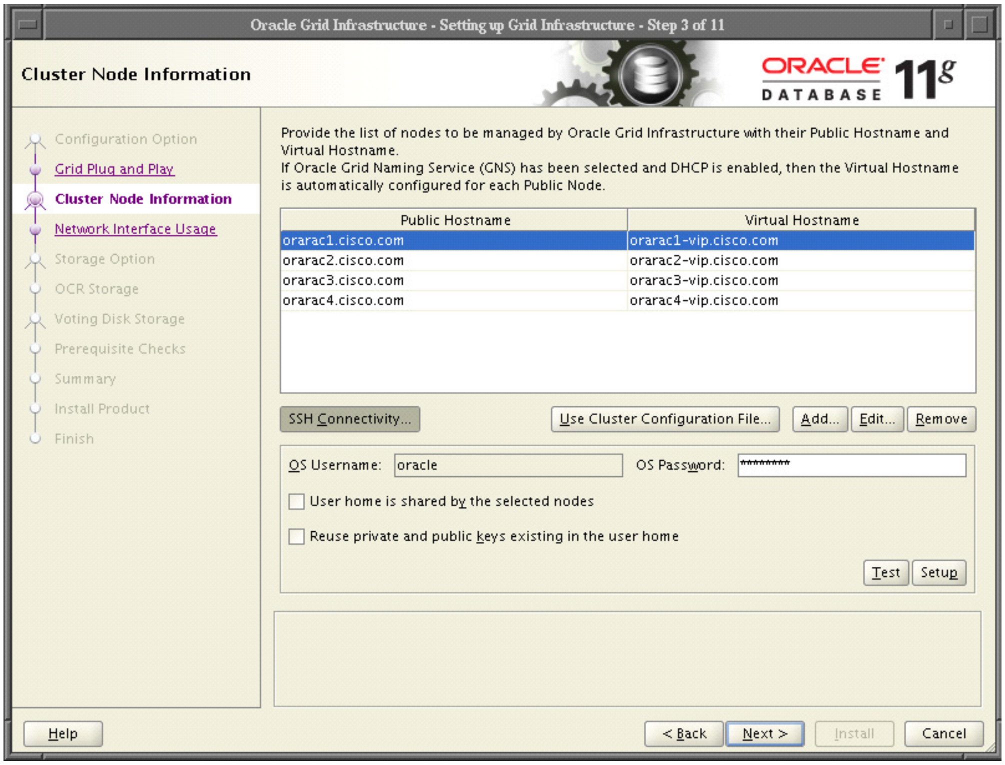

Provide the list of nodes to be managed by Oracle Grid Infrastructure with their Public Hostname and Virtual Hostname.
If Oracle Grid Naming Service (GNS) has been selected and DHCP is enabled, then the Virtual Hostname is automatically configured for each Public Node.

Public Hostname	Virtual Hostname
orarac1.cisco.com	orarac1-vip.cisco.com
orarac2.cisco.com	orarac2-vip.cisco.com
orarac3.cisco.com	orarac3-vip.cisco.com
orarac4.cisco.com	orarac4-vip.cisco.com

OS Username: oracle   OS Password: *********

□ User home is shared by the selected nodes

□ Reuse private and public keys existing in the user home

Note Make sure to select "Shared File System" if you are using NFS volumes for OCR and voting files.

Figure 58 Choose Storage Option Information

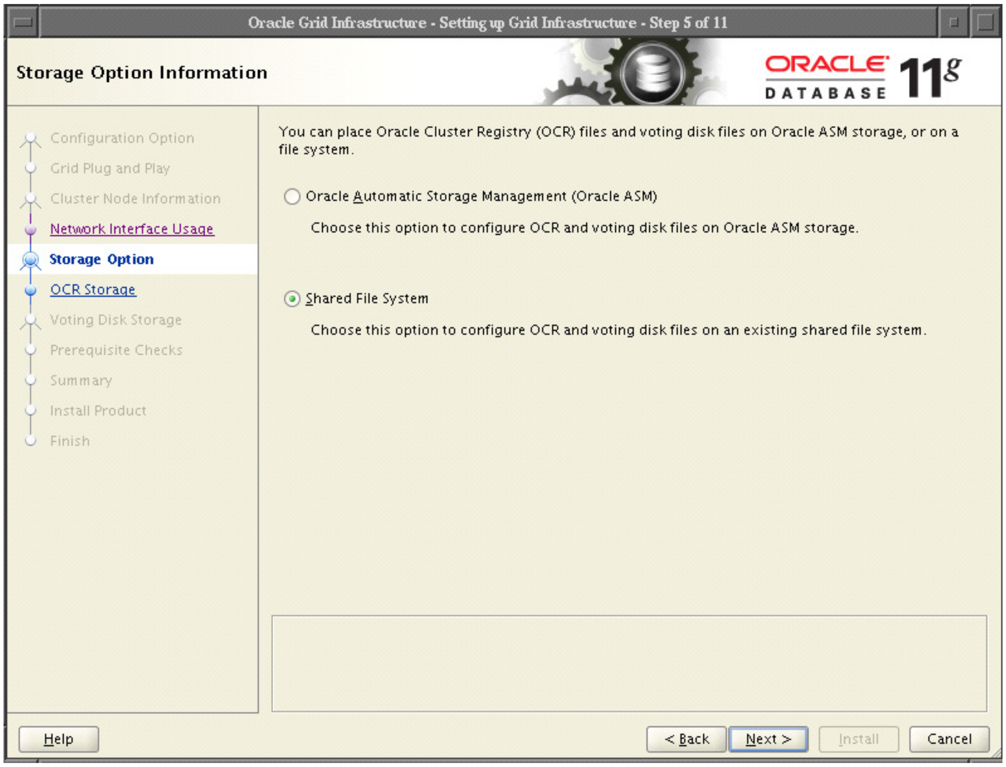

You can place Oracle Cluster Registry (OCR) files and voting disk files on Oracle ASM storage, or on a file system.

○ Oracle Automatic Storage Management (Oracle ASM)

Choose this option to configure OCR and voting disk files on Oracle ASM storage.

● Shared File System

Choose this option to configure OCR and voting disk files on an existing shared file system.

Figure 59 Configure OCR Storage Option

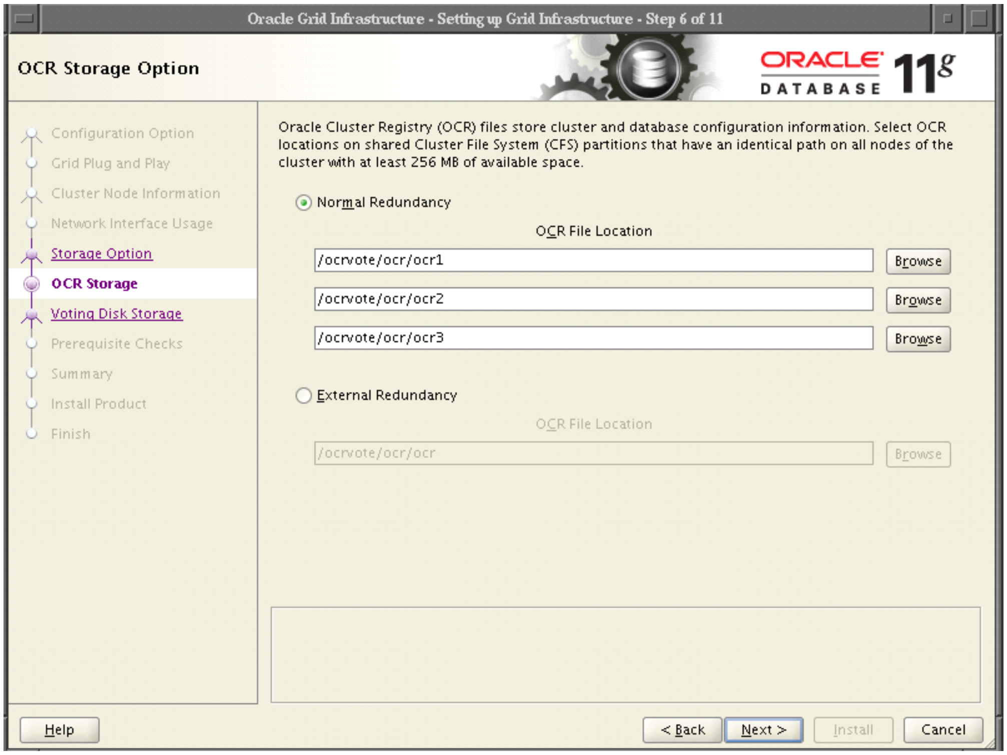

Oracle Cluster Registry (OCR) files store cluster and database configuration information. Select OCR locations on shared Cluster File System (CFS) partitions that have an identical path on all nodes of the cluster with at least 256 MB of available space.

● Normal Redundancy

OCR File Location

/ocrvote/ocr/ocr1 Browse

/ocrvote/ocr/ocr2 Browse

/ocrvote/ocr/ocr3 Browse

○ External Redundancy

OCR File Location

/ocrvote/ocr/ocr Browse

Figure 60 Configure Voting Disk Storage Option

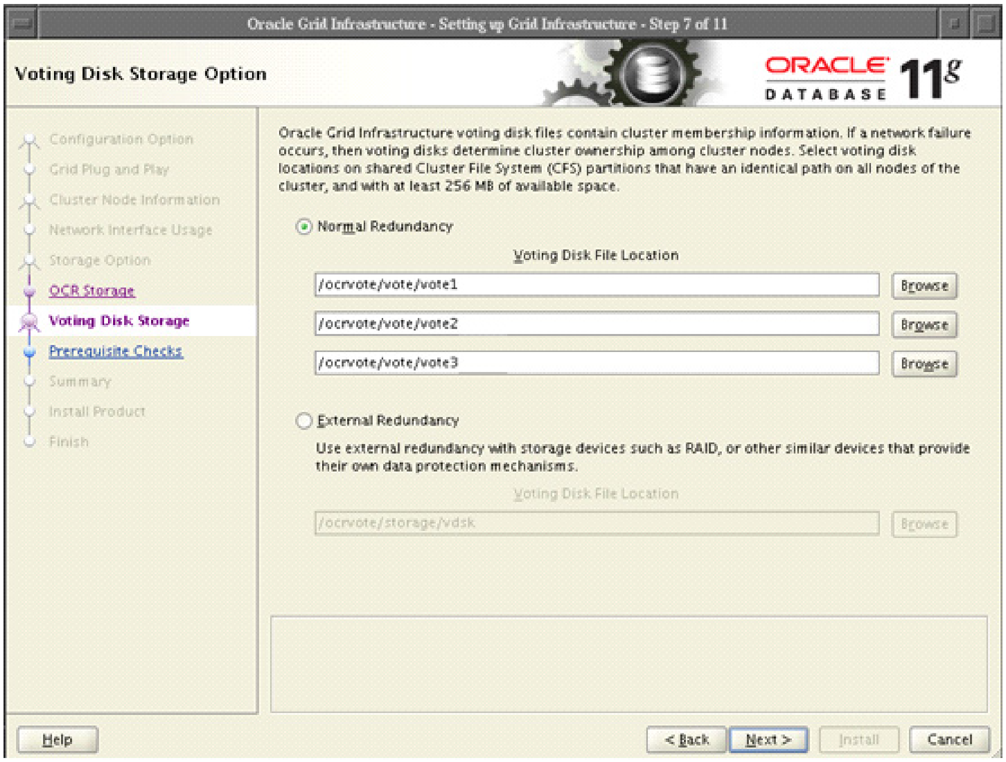

Oracle Grid Infrastructure - Setting up Grid Infrastructure - Step 7 of 11

Voting Disk Storage Option

Oracle Grid Infrastructure voting disk files contain cluster membership information. If a network failure occurs, then voting disks determine cluster ownership among cluster nodes. Select voting disk locations on shared Cluster File System (CFS) partitions that have an identical path on all nodes of the cluster, and with at least 256 MB of available space.

● Normal Redundancy

Voting Disk File Location

/ocrvote/vote/vote1

/ocrvote/vote/vote2

/ocrvote/vote/vote3

○ External Redundancy

Use external redundancy with storage devices such as RAID, or other similar devices that provide their own data protection mechanisms.

Voting Disk File Location

/ocrvote/storage/vdsk

Figure 61 Perform Prerequisites Check

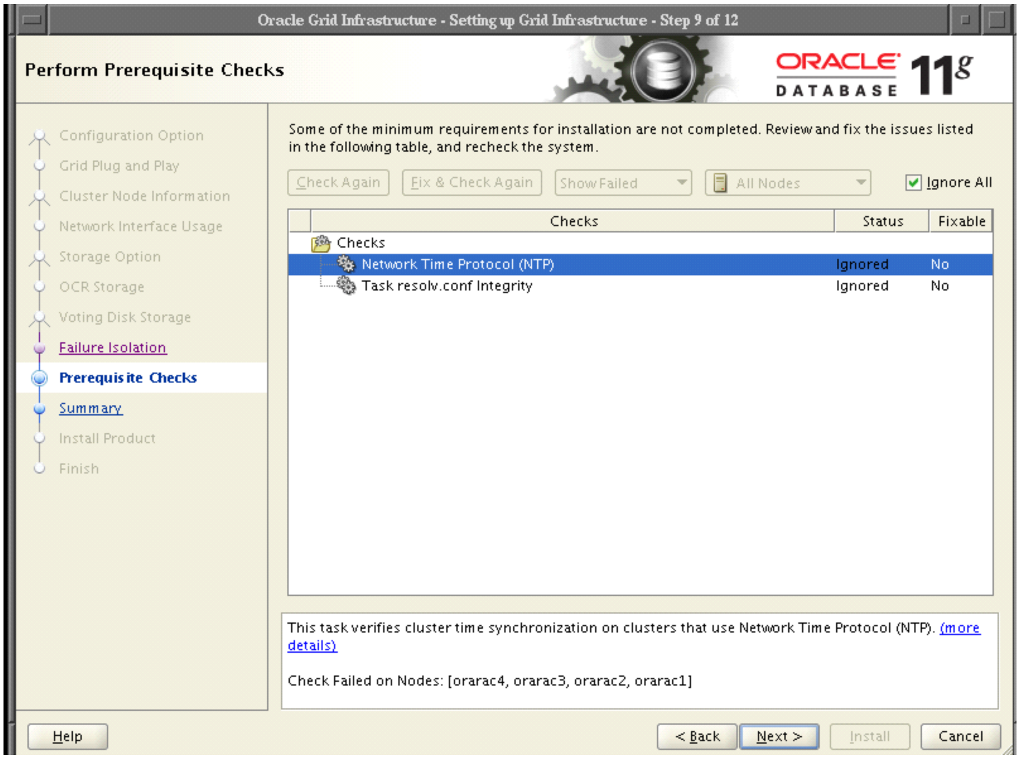

Some of the minimum requirements for installation are not completed. Review and fix the issues listed in the following table, and recheck the system.

☑ Ignore All

Checks	Status	Fixable
Checks		
Network Time Protocol (NTP)	Ignored	No
Task resolv.conf Integrity	Ignored	No

This task verifies cluster time synchronization on clusters that use Network Time Protocol (NTP). (more details)

Check Failed on Nodes: [orarac4, orarac3, orarac2, orarac1]

Figure 62 Grid Infrastructure Configuration Summary

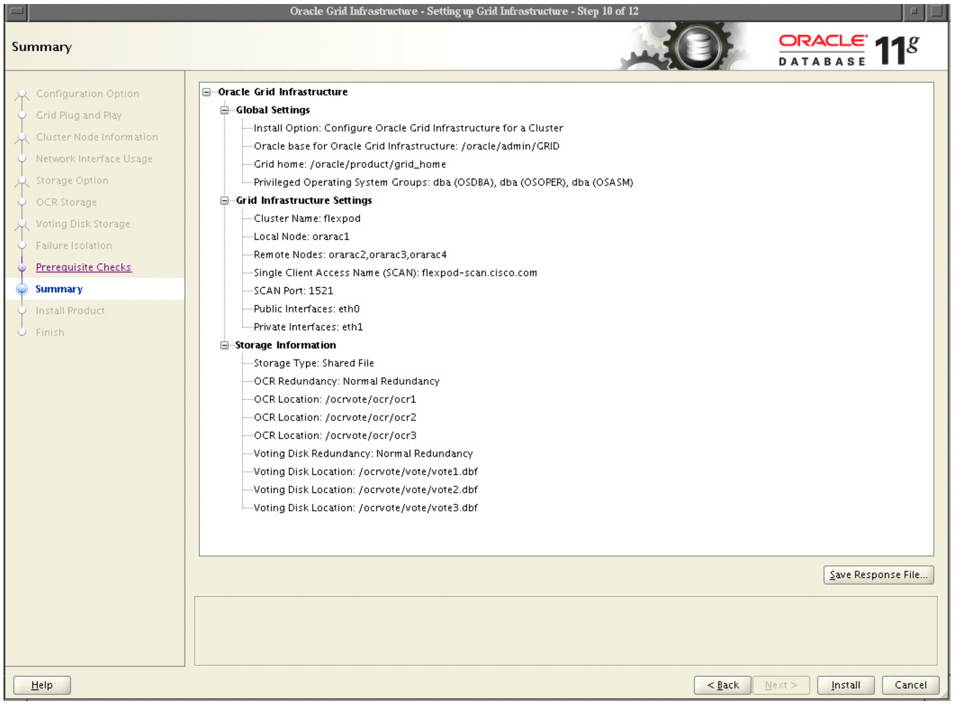

```
Oracle Grid Infrastructure
  Global Settings
    Install Option: Configure Oracle Grid Infrastructure for a Cluster
    Oracle base for Oracle Grid Infrastructure: /oracle/admin/GRID
    Grid home: /oracle/product/grid_home
    Privileged Operating System Groups: dba (OSDBA), dba (OSOPER), dba (OSASM)
  Grid Infrastructure Settings
    Cluster Name: flexpod
    Local Node: orarac1
    Remote Nodes: orarac2,orarac3,orarac4
    Single Client Access Name (SCAN): flexpod-scan.cisco.com
    SCAN Port: 1521
    Public Interfaces: eth0
    Private Interfaces: eth1
  Storage Information
    Storage Type: Shared File
    OCR Redundancy: Normal Redundancy
    OCR Location: /ocrvote/ocr/ocr1
    OCR Location: /ocrvote/ocr/ocr2
    OCR Location: /ocrvote/ocr/ocr3
    Voting Disk Redundancy: Normal Redundancy
    Voting Disk Location: /ocrvote/vote/vote1.dbf
    Voting Disk Location: /ocrvote/vote/vote2.dbf
    Voting Disk Location: /ocrvote/vote/vote3.dbf
```

Click Install and finish the remaining steps such as executing root.sh on all nodes.

When the install is complete, run few health checks on the system to validate that Oracle RACinstalled fine. Few of the Oracle RAC Health check commands are mentioned below.

```
crsctl check cluster -all
srvctl status nodeapps
crsctl status res -t
ocrcheck
crsctl query css votedisk
olsnodes -n -s
```

Step 3 When Oracle Grid install is complete, install Oracle Database 11g Release 2 Database "Software Only"; do not create the database after Oracle GRID Installation as oracle user. See Real Application Clusters Installation Guide for Linux and UNIX for detailed installation instructions:

http://www.oracle.com/pls/db112/to_toc?pathname=install.112/e10813/toc.htm

Step 4 Run the dbca tool as oracle user to create OLTP and DSS databases. Make sure to place the datafiles, redo logs and control files in proper directory paths as created in above steps. We will discuss additional details about OLTP and DSS schema creation in workload section.

Step 5 Configure Direct NFS client.

For improved NFS performance, Oracle recommends using the Direct NFS Client shipped with Oracle

11g. The direct NFS client looks for NFS details in the following locations:

- $ORACLE_HOME/dbs/oranfstab
- /etc/oranfstab
- /etc/mtab

In RAC configuration with Direct NFS ,the oranfstab must be configured on all nodes. Here is oranfstab configuration from RAC node 1.

```
[oracle@orarac1 dbs]$ vi oranfstab
server:10.10.20.5
path:10.10.20.5
local:10.10.20.101
local:10.10.30.101
server:10.10.30.5
path:10.10.30.5
local:10.10.30.101
local:10.10.20.101
export:/ocrvote mount:/vol/ocrvote
export:/redo_1 mount:/vol/redo_a
export:/redo_2 mount:/vol/redo_b
export:/data_1 mount:/vol/oltp_data_a
export:/data_2 mount:/vol/oltp_data_b
export:/dss_data_a mount:/vol/dss_data_a
export:/dss_data_b mount:/vol/dss_data_b
```

Since the NFS mount point details were defined in the "/etc/fstab", and therefore the "/etc/mtab" file, there is no need to configure any extra connection details. When setting up your NFS mounts, reference the Oracle documentation for guidance on what types of data can/cannot be accessed via Direct NFS Client. For the client to work we need to switch the libodm11.so library for the libnfsodm11.so library, as shown below.

```
srvctl stop database -d rac1db
cd $ORACLE_HOME/lib
```

```
mv libodm11.so libodm11.so_stub
ln -s libnfsodm11.so libodm11.so
srvctl start database -d rac1db
```

Note For 11.2, DNFS can also be enabled via "make -f ins_rdbms.mk dnfs_on" command.

With the configuration complete, you can see the direct NFS client usage via the following views:

- v$dnfs_servers
- v$dnfs_files
- v$dnfs_channels
- v$dnfs_stats

The following is an example from OLTP database configuration:

```
SQL> select SVRNAME, DIRNAME from  v$DNFS_SERVERS;

SVRNAME         DIRNAME
--------------- --------------------
10.10.30.5      /vol/oltp_data_b
10.10.20.5      /vol/oltp_data_a
10.10.20.5      /vol/redo_a
10.10.30.5      /vol/redo_b
```

Note The Direct NFS Client supports direct I/O and asynchronous I/O by default.

Workload and Database Configuration

We used Swingbench for workload testing. Swingbench is a simple to use, free, Java based tool to generate database workload and perform stress testing using different benchmarks in Oracle database environments. Swingbench provides four separate benchmarks, namely, Order Entry, Sales History, Calling Circle, and Stress Test. For the tests described in this paper, Swingbench Order Entry benchmark was used for OLTP workload testing and the Sales History benchmark was used for the DSS workload testing. The Order Entry benchmark is based on SOE schema and is TPC-C like by types of transactions. The workload uses a very balanced read/write ratio around 60/40 and can be designed to run continuously and test the performance of a typical Order Entry workload against a small set of tables, producing contention for database resources. The Sales History benchmark is based on the SH schema and is TPC-H like. The workload is query (read) centric and is designed to test the performance of queries against large tables.

As discussed in the previous section, two independent databases were created earlier for Oracle Swingbench OLTP and DSS workloads. Next step is to pre-create the order entry and sales history schema for OLTP and DSS workload. Swingbench Order Entry (OLTP) workload uses SOE tablespace and Sales History workload uses SH tablespaces. We pre-created these schemas in order to associate multiple datafiles with tablespaces and evenly distribute them across two storage controllers. For our setup, we created 160 datafiles for SOE tablespace with odd number files for storage controller A and even number of files for storage controller B. We used 80 datafiles for Sales history workload and evenly distributed them across both storage controllers. Once the schemas for these workloads are created, we populated both databases with Swingbench datagenerator as shown below.

OLTP Database

The OLTP database was populated with the following data:

```
[oracle@orarac1 ~]$ sqlplus soe/soe
SQL*Plus: Release 11.2.0.3.0 Production on Wed Mar 27 12:02:01 2013
Copyright (c) 1982, 2011, Oracle.  All rights reserved.
Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.3.0 - 64bit Production
With the Partitioning, Real Application Clusters, Oracle Label Security, OLAP,
Data Mining, Oracle Database Vault and Real Application Testing options

SQL> select table_name, num_rows from user_tables;

TABLE_NAME                       NUM_ROWS
------------------------------ ----------
CUSTOMERS                      4299999998
ORDER_ITEMS                   18125234562
ORDERS                         6175512689
LOGON                          1750000000
ORDERENTRY_METADATA                     4
PRODUCT_DESCRIPTIONS                 1000
PRODUCT_INFORMATION                  1000
INVENTORIES                        949031
WAREHOUSES                           1000
```

DSS (Sales History) Database

The DSS database was populated with the following data:

```
[oracle@orarac1 ~]$ sqlplus sh/sh
SQL*Plus: Release 11.2.0.3.0 Production on Wed Mar 27 12:11:45 2013
Copyright (c) 1982, 2011, Oracle.  All rights reserved.
Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.3.0 - 64bit Production
With the Partitioning, Real Application Clusters, Oracle Label Security, OLAP,
Data Mining, Oracle Database Vault and Real Application Testing options

SQL> select table_name, num_rows from user_tables;

TABLE_NAME                       NUM_ROWS
------------------------------ ----------
COUNTRIES                              23
PROMOTIONS                            503
CUSTOMERS                      2177687296
PRODUCTS                               72
SUPPLEMENTARY_DEMOGRAPHICS     2177687296
CHANNELS                                5
SALES                         10888436544
TIMES                                6209
```

Typically encountered in the real world deployments, we tested scalability and stress related scenarios that ran on current 4-node Oracle RAC cluster configuration.

1. OLTP user scalability and OLTP cluster scalability representing small and random transactions
2. DSS workload representing larger transactions

3. Mixed workload featuring OLTP and DSS workloads running simultaneously for 24 hours

Test Performance Data

When the databases were created, we started out with OLTP database calibration,about the number of concurrent users and database configuration. For Order Entry workload, we used 96GB SGA and ensured that hugepages were in use. Each OLTP scalability test was run for at least 12 hours and we ensured that results are consistent for the duration of the full run.

OLTP Workload

For OLTP workloads, the common measurement metrics are Transactions Per Minute (TPM), users scalability with IOPs and CPU utilization. Here are the scalability charts for Order Entry workload.

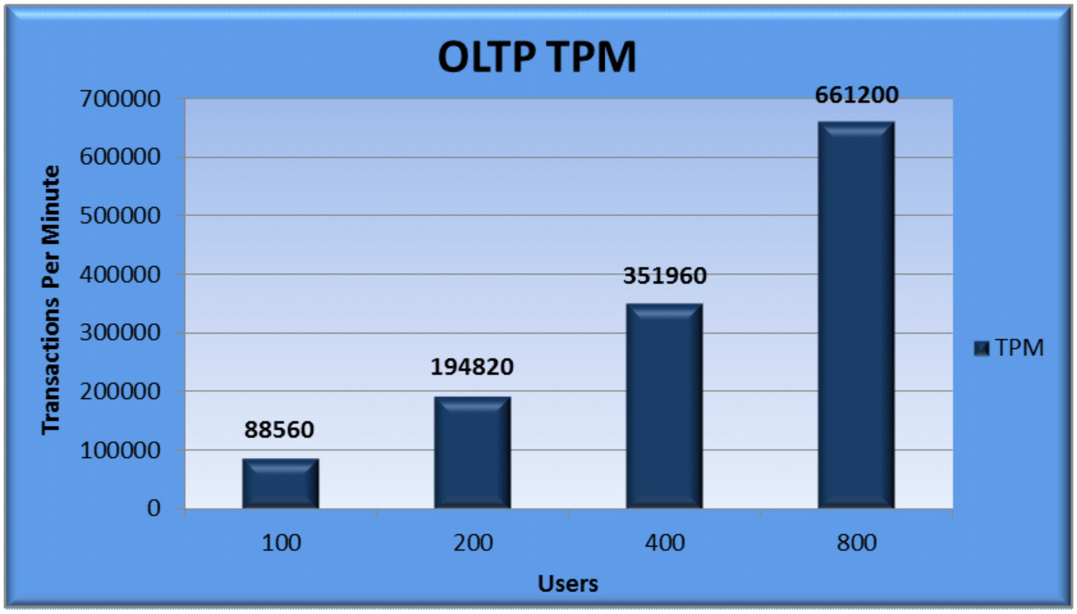

For OLTP TPM tests, we ran tests with 100, 200, 400 and 800 users across 4-node cluster. During the tests, we validated that Oracle SCAN listener fairly and evenly load balanced users across all 4 nodes of the cluster. We also observed appropriate scalability in TPMs as number or users across clusters increased. Next graph shows increased IO and scalability as number of users across the nodes in the cluster increased.

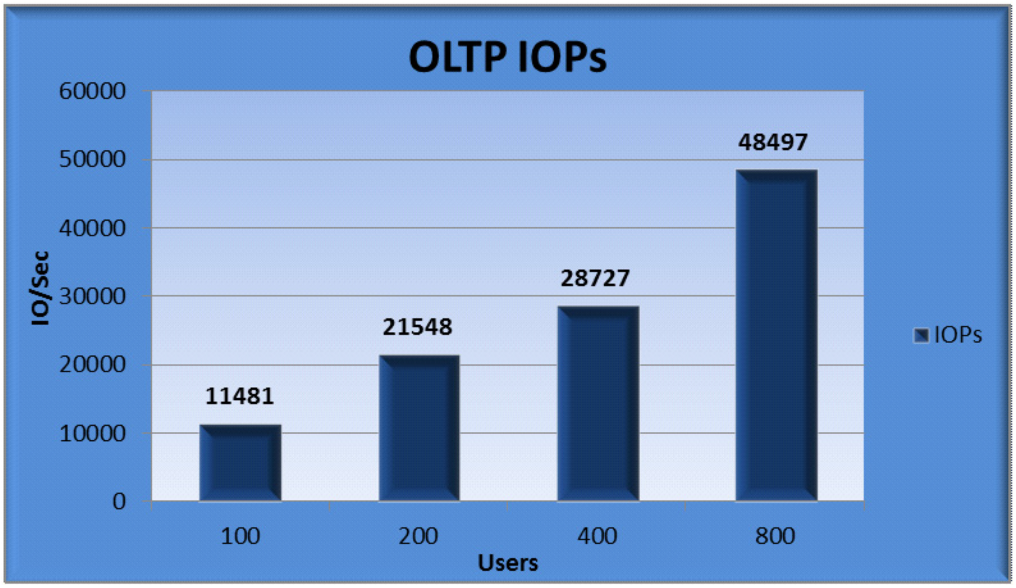

As indicated in the graph above, we observed about 48,497 IO/Sec across 4-node cluster. The Oracle AWR report below also summarizes Physical Reads/Sec and Physical Writes/Sec per instance. During OLTP tests, we observed some resource utilization variations due to random nature of the workload as depicted by 400 user IOPs. We ran each test multiple times to ensure consistent numbers that are presented in this solution.

Table 9 *Summary of Oracle AWR Report*

System	Statistics	-	Per	Second	DB/Inst:	FLEXPOD /flexpod1	Snaps:	130-132		
I#	**Logical Reads/s**	**Physical Reads/s**	**Physical Writes/s**	**Redo Size(k)/s**	**Block Changes /s**	**User Calls/s**	**Execs/s**	**Parses/s**	**Logons/s**	**Txns/s**
1	303,182.03	7,404.50	5,585.90	9,856.40	67,784.90	3,058.30	29,233.80	3.3	0.12	4,599.60
2	280,835.48	7,000.10	4,652.50	8,906.70	60,686.40	2,747.00	26,262.70	3.3	0.12	4,131.30
3	270,783.24	6,930.00	4,665.50	8,862.30	60,335.80	2,732.20	26,116.70	3.2	0.12	4,109.10
4	300,397.02	7,308.00	4,950.00	9,592.60	65,615.30	2,914.20	27,862.00	3.7	0.13	4,382.90
~~	~~~~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~~~	~~~~~~~~~~	~~~~~~~~~~~~
Sum	1,155,197.77	28,642.60	19,853.90	37,218.00	254,422.4	11,451.70	109,475.2	13.5	0.5	17,222.90
Avg	288799.44	7,160.70	4,963.50	9,304.50	63,605.60	2,862.90	27,368.80	3.4	0.12	4,305.70
Std	15,592.42	231.1	437.1	497.1	3,684.10	154.2	1,473.40	0.3	0	232

The table below shows interconnect traffic for the 4-node Oracle RAC cluster during 800 user run. The average interconnect traffic was 350 MB/Sec for the duration of the run.

Table 10 Interconnect Traffic Bandwidth Usage

Interconnect Traffic	Sent (MB/s) Total	Received (MB/s) Total
Instance 1	90.5	91.1
Instance 2	84.2	84
Instance 3	87.4	86.8
Instance 4	88	87.3
Total MB/Sec	350.1	349.2

The chart below indicates cluster CPU utilization as the number of users scale from 25 users/node to 200 users/node.

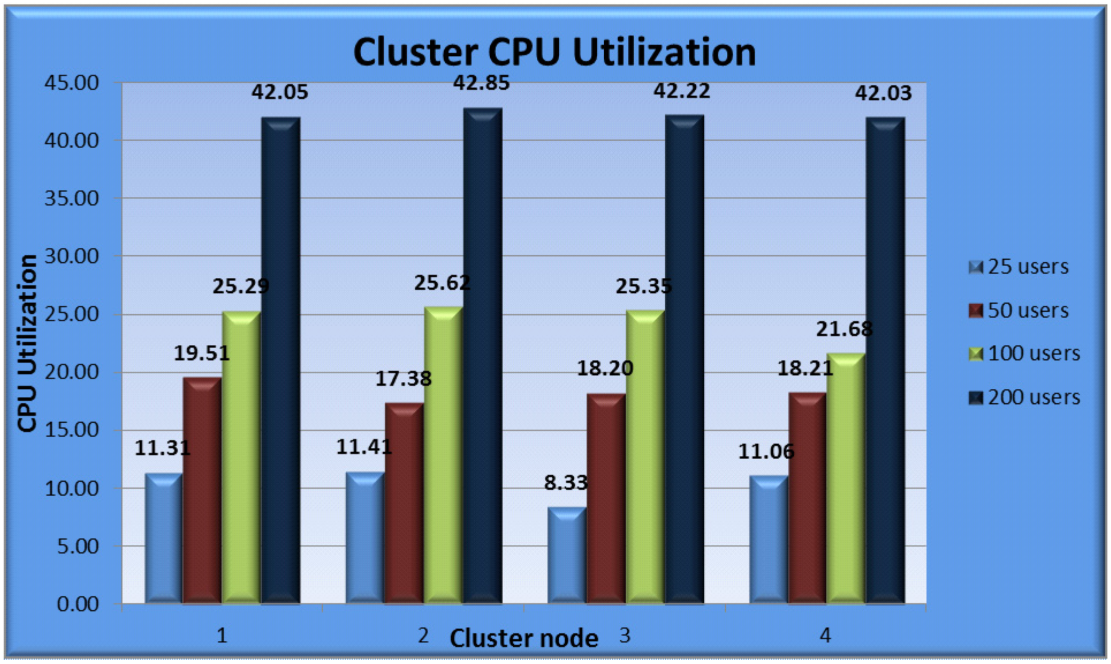

DSS Workload

DSS workloads are generally sequential in nature, read intensive and exercise large IO size. DSS workloads run a small number of users that typically exercise extremely complex queries that run for hours. For our tests, we ran Swingbench Sales history workload with 20 users. The charts below show DSS workload results.

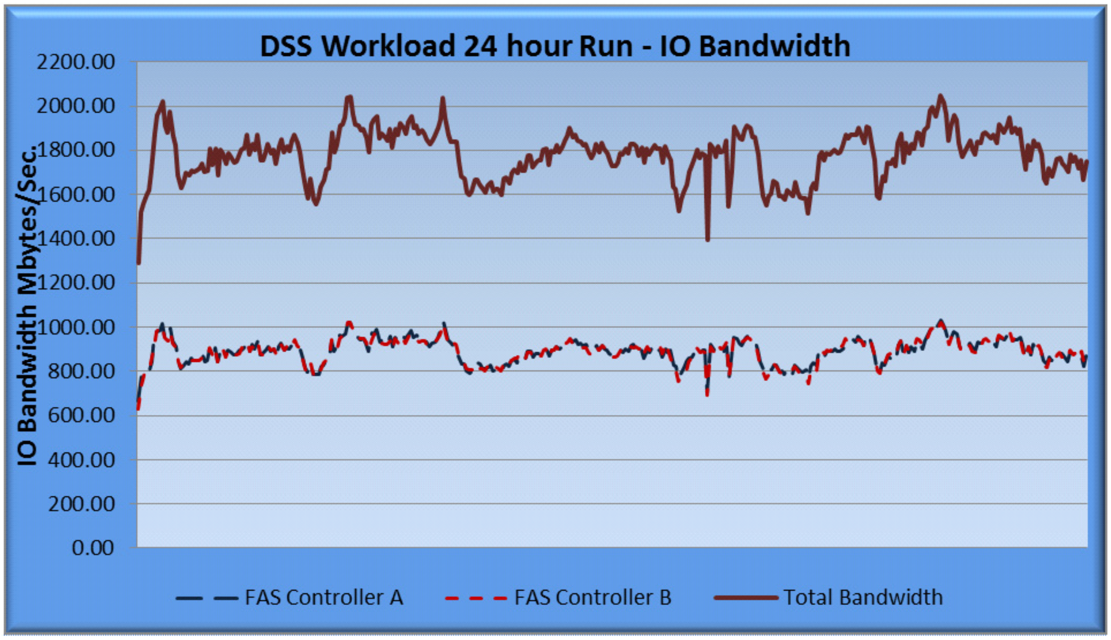

For 24 hour DSS workload test, we observed total IO bandwidth ranging between 1.8 GBytes/Sec and 2.0 GBytes/Sec. As indicated on the charts, the IO was also evenly distributed across both NetApp FAS storage controllers and we did not observe any significant dips in performance and IO bandwidth for a sustained period of time.

Mixed Workload

The next test is to run both OLTP and DSS workloads simultaneously. This test will ensure that configuration in this test is able to sustain small random queries presented via OLTP along with large and sequential transactions submitted via DSS workload. We ran the tests for 24 hours. Here are the results.

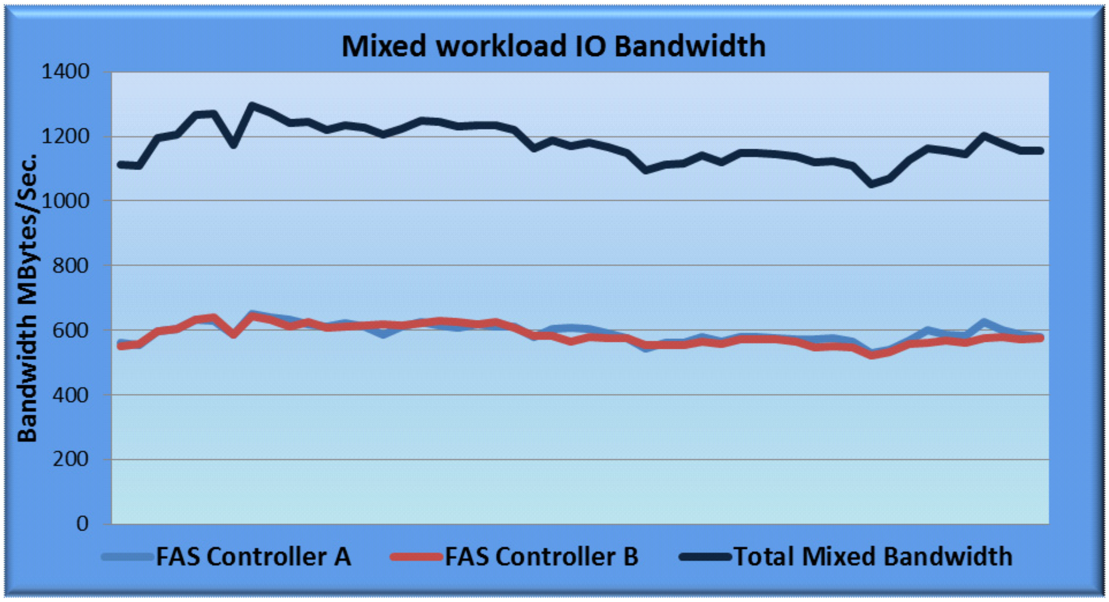

For mixed workloads running for 24 hours, we observed approximately 1.2 GBytes/Sec. IO bandwidth. The OLTP transactions also averaged between 300K and 340K transactions per minute.

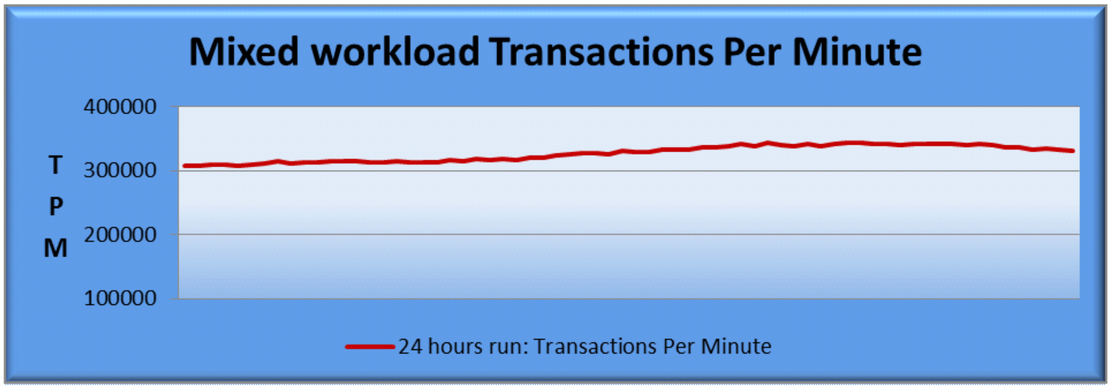

Destructive and Hardware Failover Tests

The goal of these tests is to ensure that reference architecture withstands commonly occurring failures either due to unexpected crashes, hardware failures or human errors. We conducted many hardware, software (process kills) and OS specific failures that simulate real world scenarios under stress conditions. In the destructive testing, we also demonstrate unique failover capabilities of Cisco VIC 1240 adapter. We have highlighted some of those test cases below.

Table 11 Hardware Failover Scenarios

Scenario	Test	Status
Test 1 – Chassis 1-IOM2 Link Failure test	Run the system on full mixed work load. Disconnect the private links from first chassis and reconnect the links after 5 minutes.	Network Traffic from IOM2 will failover without any disruption to IOM1.
Test 2 – UCS 6248 Fabric-B Failure test	Run the system on full load as above. Reboot Fabric B, let it join the cluster back and then Reboot Fabric A.	Fabric failovers did not cause any disruption to Private and Storage traffic.
Test 3 – Nexus 5548 Fabric-A Failure test	Run the system on full mixed work load. Reboot the Nexus5548 Fabric-A Switch, wait for 5 minutes, connect it back and repeat for Nexus5548 Fabric-B Switch.	No disruption to the Public and Storage Traffic

Figure 63 Network Traffic Before Failover Test

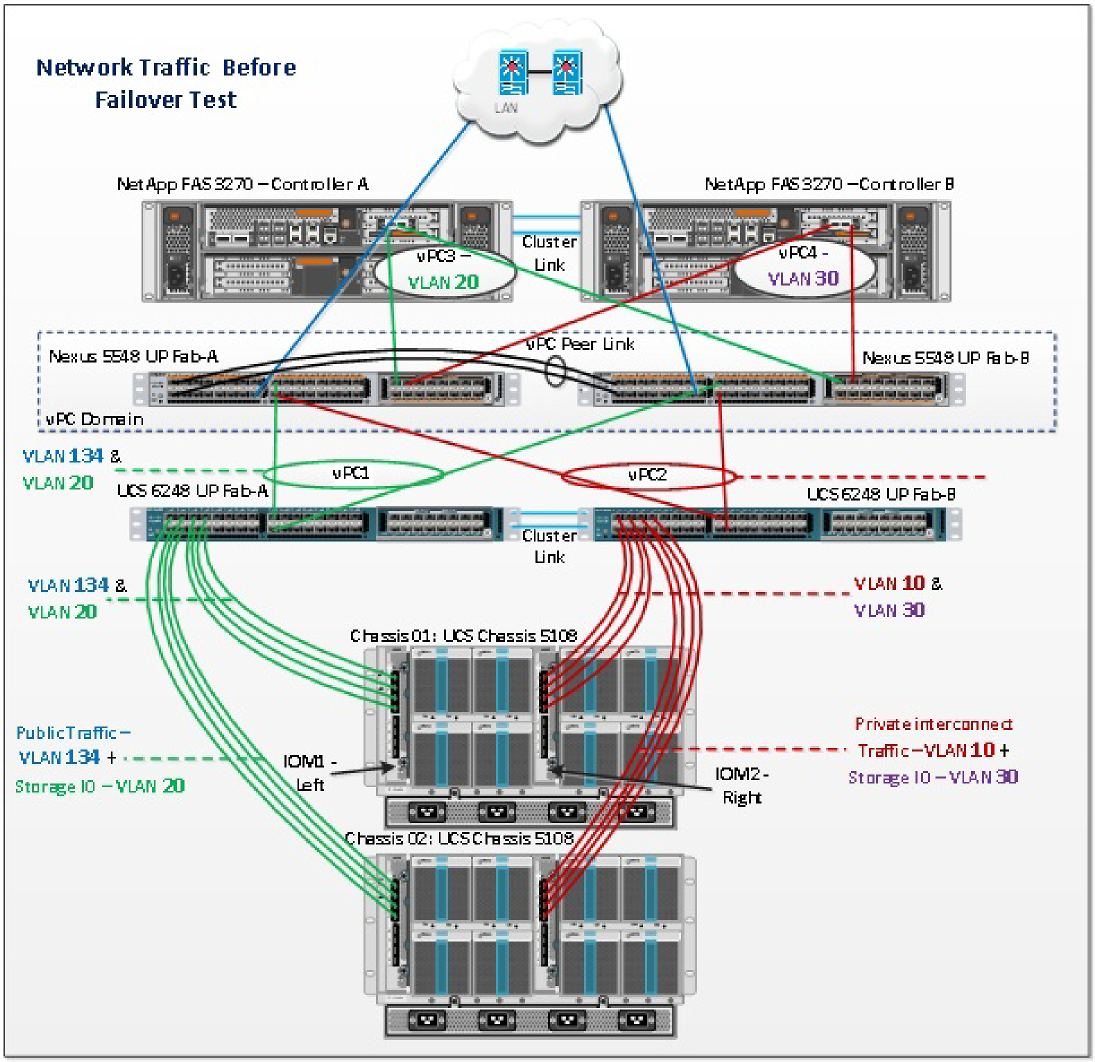
Network Traffic Before Failover Test
LAN
NetApp FAS3270 – Controller A
NetApp FAS3270 – Controller B
Cluster Link
vPC3 – VLAN 20
vPC4 - VLAN 30
vPC Peer Link
Nexus 5548 UP Fab-A
Nexus 5548 UP Fab-B
vPC Domain
VLAN 134 & VLAN 20
vPC1
vPC2
UCS 6248 UP Fab-A
UCS 6248 UP Fab-B
Cluster Link
VLAN 134 & VLAN 20
VLAN 10 & VLAN 30
Chassis 01: UCS Chassis 5108
Public Traffic – VLAN 134 + Storage IO – VLAN 20
IOM1 - Left
IOM2 - Right
Private interconnect Traffic – VLAN 10 + Storage IO – VLAN 30
Chassis 02: UCS Chassis 5108

The screenshot below shows a MAC address table and VLAN information for Cisco UCS Fabric interconnects and Nexus switches before failover test. To get this information, login to Cisco UCS Fabric interconnect and execute connect nxos A.

Figure 64 MAC Address Table Information on Fabric Interconnect A

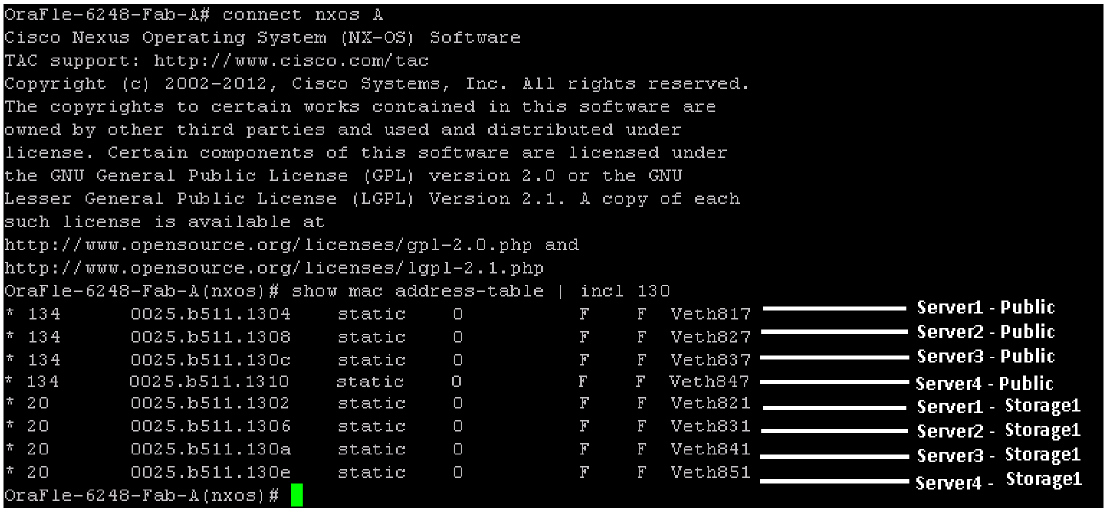

```
OraFle-6248-Fab-A# connect nxos A
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
OraFle-6248-Fab-A(nxos)# show mac address-table | incl 130
* 134      0025.b511.1304    static    0         F    F  Veth817
* 134      0025.b511.1308    static    0         F    F  Veth827
* 134      0025.b511.130c    static    0         F    F  Veth837
* 134      0025.b511.1310    static    0         F    F  Veth847
* 20       0025.b511.1302    static    0         F    F  Veth821
* 20       0025.b511.1306    static    0         F    F  Veth831
* 20       0025.b511.130a    static    0         F    F  Veth841
* 20       0025.b511.130e    static    0         F    F  Veth851
OraFle-6248-Fab-A(nxos)#
```

Log in to Cisco UCS Fabric Interconnect B and type "connect nxos B" as shown below.

Figure 65 MAC Address Table Information on Fabric Interconnect B

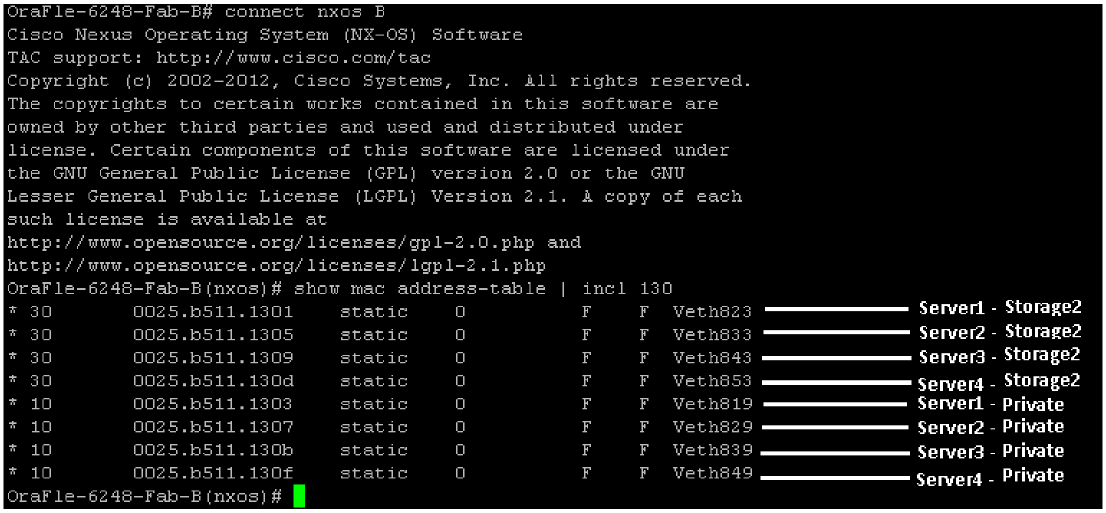

```
OraFle-6248-Fab-B# connect nxos B
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
OraFle-6248-Fab-B(nxos)# show mac address-table | incl 130
* 30       0025.b511.1301    static    0         F    F  Veth823
* 30       0025.b511.1305    static    0         F    F  Veth833
* 30       0025.b511.1309    static    0         F    F  Veth843
* 30       0025.b511.130d    static    0         F    F  Veth853
* 10       0025.b511.1303    static    0         F    F  Veth819
* 10       0025.b511.1307    static    0         F    F  Veth829
* 10       0025.b511.130b    static    0         F    F  Veth839
* 10       0025.b511.130f    static    0         F    F  Veth849
OraFle-6248-Fab-B(nxos)#
```

We configured a Virtual Port channel on Cisco Nexus 5548 Switches, both the N5548 switches register the same Mac address via vPC peer link.

Below is a VLAN and Mac Address table from Cisco UCS Nexus5548 Fabric A before the failover test.

Figure 66 MAC Address Table Information on Nexus 5548 Switch Fabric A

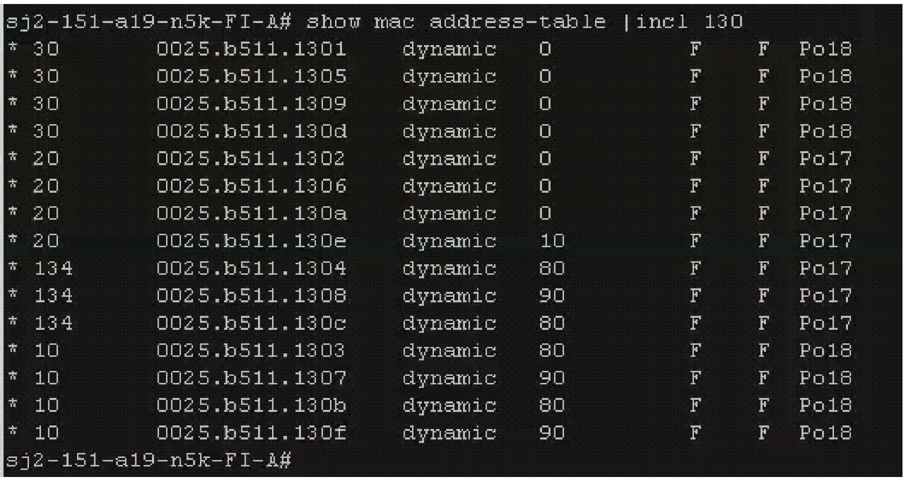

```
sj2-151-a19-n5k-FI-A# show mac address-table |incl 130
* 30        0025.b511.1301    dynamic   0          F     F  Po18
* 30        0025.b511.1305    dynamic   0          F     F  Po18
* 30        0025.b511.1309    dynamic   0          F     F  Po18
* 30        0025.b511.130d    dynamic   0          F     F  Po18
* 20        0025.b511.1302    dynamic   0          F     F  Po17
* 20        0025.b511.1306    dynamic   0          F     F  Po17
* 20        0025.b511.130a    dynamic   0          F     F  Po17
* 20        0025.b511.130e    dynamic   10         F     F  Po17
* 134       0025.b511.1304    dynamic   80         F     F  Po17
* 134       0025.b511.1308    dynamic   90         F     F  Po17
* 134       0025.b511.130c    dynamic   80         F     F  Po17
* 10        0025.b511.1303    dynamic   80         F     F  Po18
* 10        0025.b511.1307    dynamic   90         F     F  Po18
* 10        0025.b511.130b    dynamic   80         F     F  Po18
* 10        0025.b511.130f    dynamic   90         F     F  Po18
sj2-151-a19-n5k-FI-A#
```

Below is a VLAN and Mac Address table from Cisco UCS Nexus5548 Fabric B.

Figure 67 MAC Address Table Information on Nexus 5548 Switch Fabric B

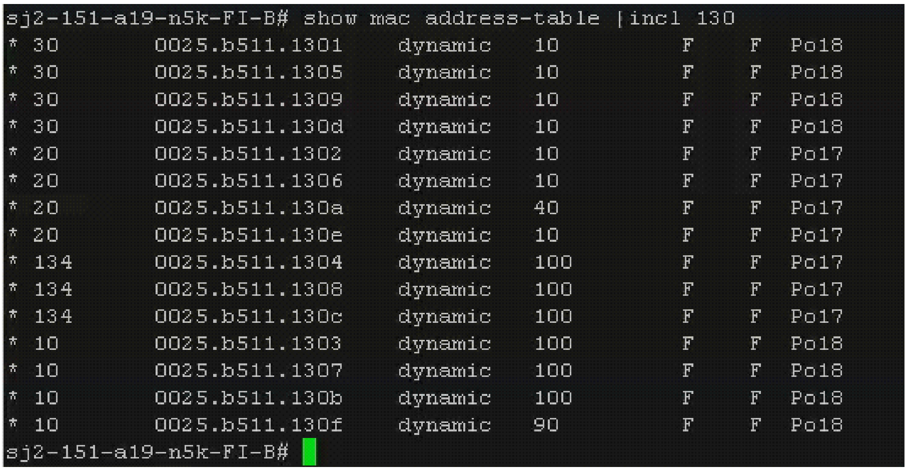

```
sj2-151-a19-n5k-FI-B# show mac address-table |incl 130
* 30        0025.b511.1301    dynamic   10         F     F  Po18
* 30        0025.b511.1305    dynamic   10         F     F  Po18
* 30        0025.b511.1309    dynamic   10         F     F  Po18
* 30        0025.b511.130d    dynamic   10         F     F  Po18
* 20        0025.b511.1302    dynamic   10         F     F  Po17
* 20        0025.b511.1306    dynamic   10         F     F  Po17
* 20        0025.b511.130a    dynamic   40         F     F  Po17
* 20        0025.b511.130e    dynamic   10         F     F  Po17
* 134       0025.b511.1304    dynamic   100        F     F  Po17
* 134       0025.b511.1308    dynamic   100        F     F  Po17
* 134       0025.b511.130c    dynamic   100        F     F  Po17
* 10        0025.b511.1303    dynamic   100        F     F  Po18
* 10        0025.b511.1307    dynamic   100        F     F  Po18
* 10        0025.b511.130b    dynamic   100        F     F  Po18
* 10        0025.b511.130f    dynamic   90         F     F  Po18
sj2-151-a19-n5k-FI-B#
```

Figure 68 Test 1 - Cisco UCS Chassis - IOM2 Link Failure Test

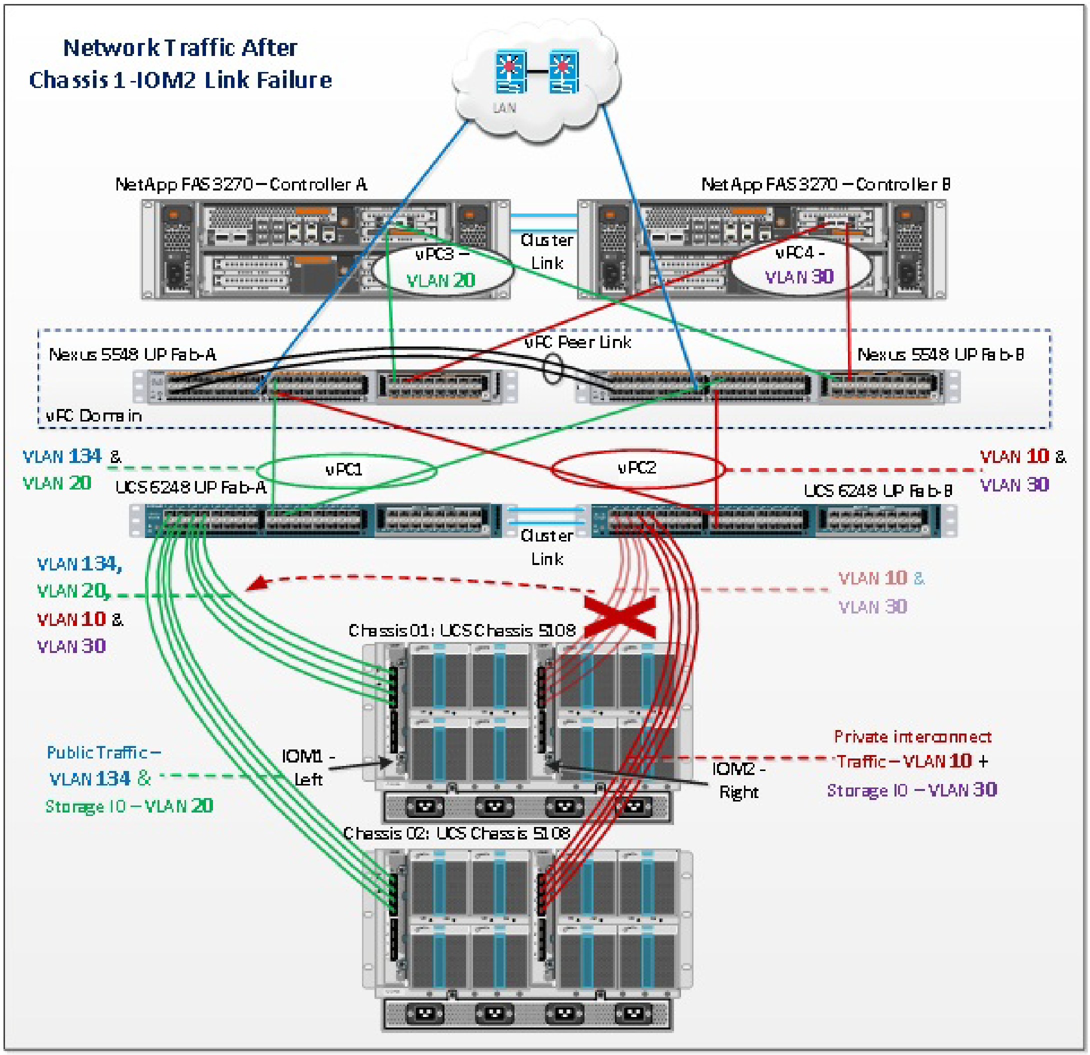
Network Traffic After
Chassis 1 -IOM2 Link Failure
LAN
NetApp FAS3270 – Controller A
NetApp FAS3270 – Controller B
Cluster Link
vPC3 – VLAN 20
vPC4 - VLAN 30
vPC Peer Link
Nexus 5548 UP Fab-A
Nexus 5548 UP Fab-B
vPC Domain
VLAN 134 & VLAN 20
vPC1
vPC2
VLAN 10 & VLAN 30
UCS 6248 UP Fab-A
UCS 6248 UP Fab-B
Cluster Link
VLAN 134, VLAN 20, VLAN 10 & VLAN 30
VLAN 10 & VLAN 30
Chassis 01: UCS Chassis 5108
Public Traffic – VLAN 134 & Storage IO – VLAN 20
IOM1 - Left
IOM2 - Right
Private interconnect Traffic – VLAN 10 + Storage IO – VLAN 30
Chassis 02: UCS Chassis 5108

As shown above, during Chassis01 - IOM2 (Fabric-B) link failure, the respective blades (Server1 and Server2) on chassis01 will failover the mac address and its VLAN to IOM1 on Fabric A side.

Below is a failed over MAC address and VLAN information on Cisco UCS Fabric Interconnect A.

Figure 69 MAC Address Table Information on Fabric Interconnect A

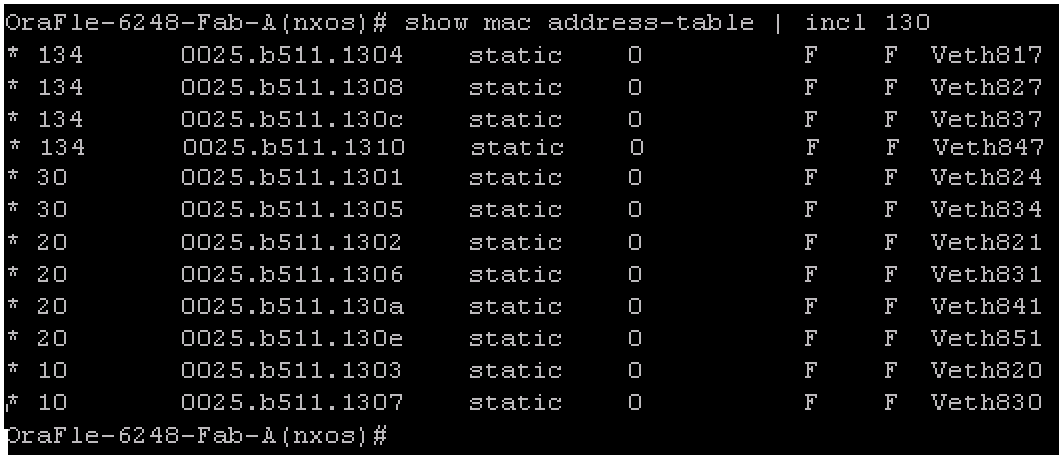

```
OraFle-6248-Fab-A(nxos)# show mac address-table | incl 130
* 134      0025.b511.1304    static    0          F    F   Veth817
* 134      0025.b511.1308    static    0          F    F   Veth827
* 134      0025.b511.130c    static    0          F    F   Veth837
* 134      0025.b511.1310    static    0          F    F   Veth847
* 30       0025.b511.1301    static    0          F    F   Veth824
* 30       0025.b511.1305    static    0          F    F   Veth834
* 20       0025.b511.1302    static    0          F    F   Veth821
* 20       0025.b511.1306    static    0          F    F   Veth831
* 20       0025.b511.130a    static    0          F    F   Veth841
* 20       0025.b511.130e    static    0          F    F   Veth851
* 10       0025.b511.1303    static    0          F    F   Veth820
* 10       0025.b511.1307    static    0          F    F   Veth830
OraFle-6248-Fab-A(nxos)#
```

Below is a MAC address and VLAN information on Cisco UCS Fabric Interconnect B.

Figure 70 MAC Address Table Information on Fabric Interconnect B

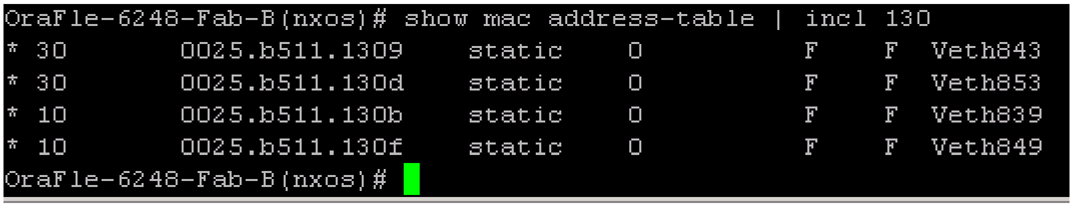

```
OraFle-6248-Fab-B(nxos)# show mac address-table | incl 130
* 30       0025.b511.1309    static    0          F    F   Veth843
* 30       0025.b511.130d    static    0          F    F   Veth853
* 10       0025.b511.130b    static    0          F    F   Veth839
* 10       0025.b511.130f    static    0          F    F   Veth849
OraFle-6248-Fab-B(nxos)#
```

Figure 71 Test 2 - Cisco UCS 6248 FI Fabric B Failure

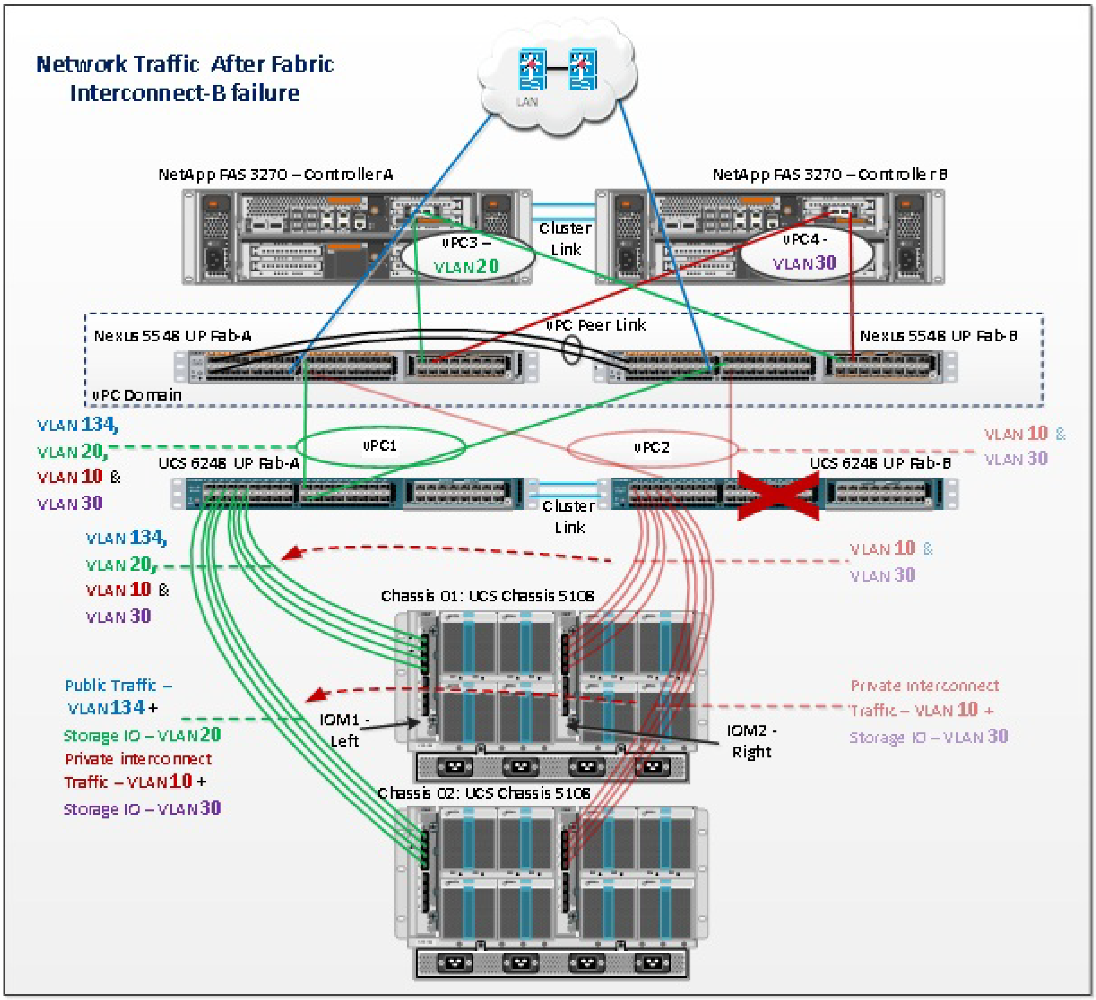
Network Traffic After Fabric Interconnect-B failure
LAN
NetApp FAS 3270 – Controller A
NetApp FAS 3270 – Controller B
Cluster Link
vPC3 – VLAN 20
vPC4 - VLAN 30
Nexus 5548 UP Fab-A
vPC Peer Link
Nexus 5548 UP Fab-B
vPC Domain
VLAN 134, VLAN 20, VLAN 10 & VLAN 30
vPC1
vPC2
UCS 6248 UP Fab-A
UCS 6248 UP Fab-B
VLAN 10 & VLAN 30
Cluster Link
VLAN 134, VLAN 20, VLAN 10 & VLAN 30
VLAN 10 & VLAN 30
Chassis 01: UCS Chassis 5108
Public Traffic – VLAN 134 +
Storage IO – VLAN 20
Private interconnect Traffic – VLAN 10 +
Storage IO – VLAN 30
IOM1 - Left
IOM2 - Right
Private interconnect Traffic – VLAN 10 +
Storage IO – VLAN 30
Chassis 02: UCS Chassis 5108

As shown above, during Fabric Interconnect B Switch failure, the respective blades (Server1 & Server2) on chassis01 (Server 3 and Server4) on chassis02 will failover the MAC addresses and its VLAN to Fabric Interconnect B.

Below is the MAC address and VLAN information on Cisco UCS Fabric Interconnect A.

Figure 72 MAC Address Table Information on Nexus 5548 Switch Fabric A

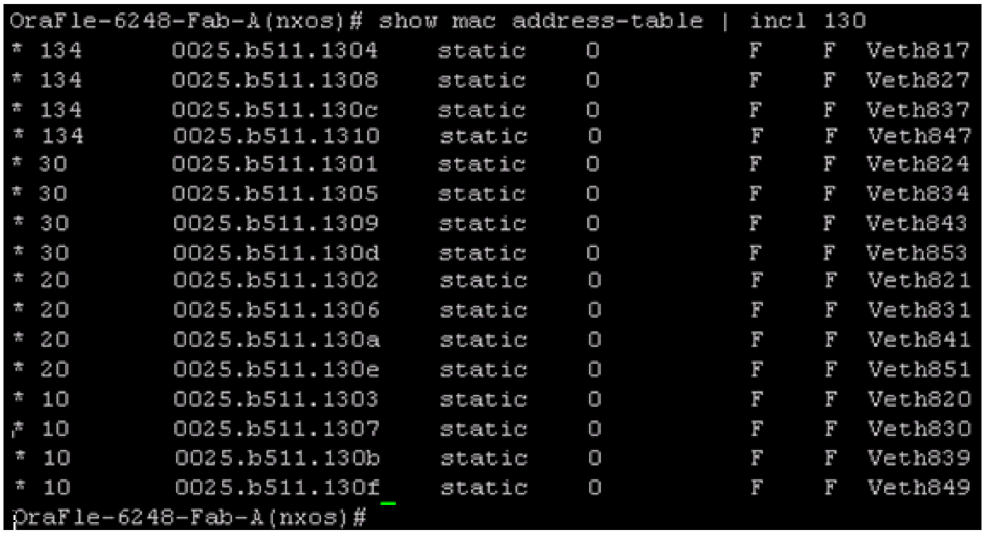

```
OraFle-6248-Fab-A(nxos)# show mac address-table | incl 130
* 134      0025.b511.1304    static    0         F    F  Veth817
* 134      0025.b511.1308    static    0         F    F  Veth827
* 134      0025.b511.130c    static    0         F    F  Veth837
* 134      0025.b511.1310    static    0         F    F  Veth847
* 30       0025.b511.1301    static    0         F    F  Veth824
* 30       0025.b511.1305    static    0         F    F  Veth834
* 30       0025.b511.1309    static    0         F    F  Veth843
* 30       0025.b511.130d    static    0         F    F  Veth853
* 20       0025.b511.1302    static    0         F    F  Veth821
* 20       0025.b511.1306    static    0         F    F  Veth831
* 20       0025.b511.130a    static    0         F    F  Veth841
* 20       0025.b511.130e    static    0         F    F  Veth851
* 10       0025.b511.1303    static    0         F    F  Veth820
* 10       0025.b511.1307    static    0         F    F  Veth830
* 10       0025.b511.130b    static    0         F    F  Veth839
* 10       0025.b511.130f    static    0         F    F  Veth849
OraFle-6248-Fab-A(nxos)#
```

Figure 73 Test 3 - Cisco Nexus 5548 Fabric A Failure

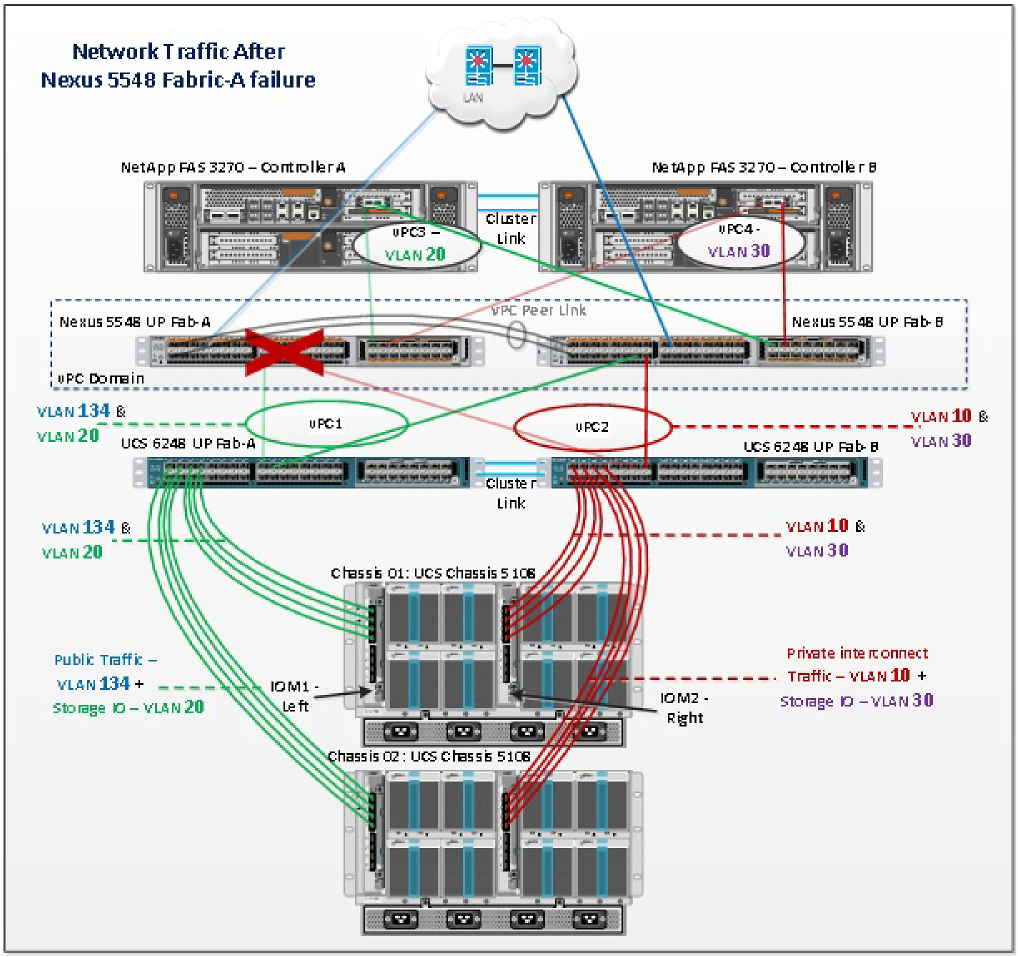

As shown above, during the Cisco Nexus 5548 Switch A failure, the Oracle servers will send the Public and Storage traffic to N5548 Fabric B switch. Below is the VLAN and MAC address table from Cisco UCS Nexus5548 B switch.

Figure 74 MAC Address Table Information on Nexus 5548 Switch Fabric B

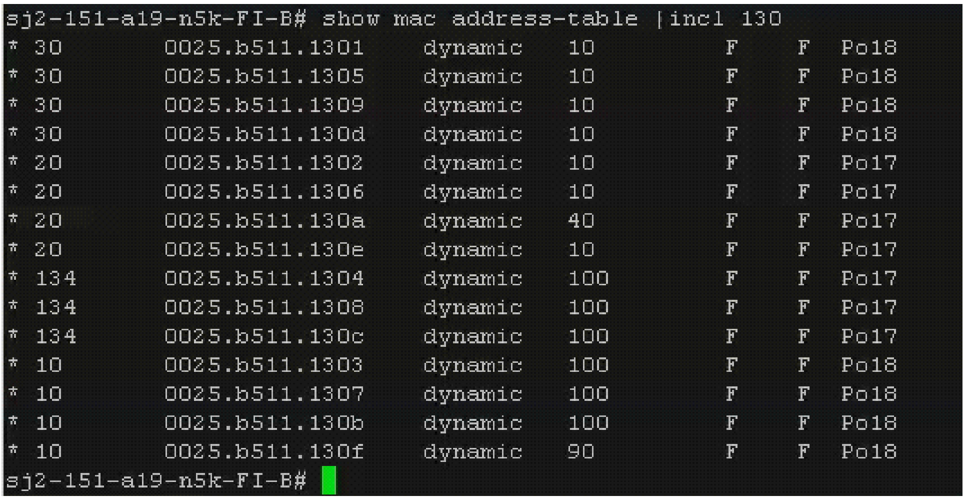

```
sj2-151-a19-n5k-FI-B# show mac address-table |incl 130
* 30       0025.b511.1301    dynamic   10         F     F   Po18
* 30       0025.b511.1305    dynamic   10         F     F   Po18
* 30       0025.b511.1309    dynamic   10         F     F   Po18
* 30       0025.b511.130d    dynamic   10         F     F   Po18
* 20       0025.b511.1302    dynamic   10         F     F   Po17
* 20       0025.b511.1306    dynamic   10         F     F   Po17
* 20       0025.b511.130a    dynamic   40         F     F   Po17
* 20       0025.b511.130e    dynamic   10         F     F   Po17
* 134      0025.b511.1304    dynamic   100        F     F   Po17
* 134      0025.b511.1308    dynamic   100        F     F   Po17
* 134      0025.b511.130c    dynamic   100        F     F   Po17
* 10       0025.b511.1303    dynamic   100        F     F   Po18
* 10       0025.b511.1307    dynamic   100        F     F   Po18
* 10       0025.b511.130b    dynamic   100        F     F   Po18
* 10       0025.b511.130f    dynamic   90         F     F   Po18
sj2-151-a19-n5k-FI-B#
```

Conclusion

FlexPod is built on leading computing, networking, storage, and infrastructure software components. With the FlexPod based solution, customers can leverage a secure, integrated, and optimized stack that includes compute, network and storage resources that are sized, configured and deployed as a fully tested unit running industry standard applications such as Oracle Database 11g RAC over D-NFS (Direct NFS).

Here's what makes the combination of Cisco UCS with NetApp storage so powerful for Oracle environments:

- Cisco UCS stateless computing architecture provided by the Service Profile capability of Cisco UCS allows for fast, non-disruptive workload changes to be executed simply and seamlessly across the integrated UCS infrastructure and Cisco x86 servers.
- Cisco UCS combined with a highly scalable NAS platform from NetApp provides the ideal combination for Oracle's unique, scalable, and highly available NFS technology.
- All of this is made possible by Cisco's Unified Fabric with its focus on secure IP networks as the standard interconnect for the server and data management solutions.

As a result, customers can achieve dramatic cost savings when leveraging Ethernet based products plus deploy any application on a scalable Shared IT infrastructure built on Cisco and NetApp technologies. Finally, FlexPod™, jointly developed by NetApp and Cisco, is a flexible infrastructure platform composed of pre-sized storage, networking, and server components. It's designed to ease your IT transformation and operational challenges with maximum efficiency and minimal risk.

FlexPod differs from other solutions by providing:

- Integrated, validated technologies from industry leaders and top-tier software partners.
- A single platform, built from unified compute, fabric, and storage technologies, that lets you scale to large-scale data centers without architectural changes.
- Centralized, simplified management of infrastructure resources, including end-to-end automation.
- A choice of validated FlexPod management solutions from trusted partners who work through our open APIs

- A flexible cooperative support model that resolves issues rapidly and spans across new and legacy products.

Appendix A—Nexus 5548 UP Fabric Zoning Configuration

The following is an example which shows Nexus 5548 Fabric Zoning Configuration for all the Oracle RAC Servers.

Login to Cisco Nexus 5548 through.ssh and issue the following:

Nexus 5548 Fabric A Zoning Configuration

```
N5548-Fab-A# config terminal
N5548-Fab-A(config)# zoneset name ORARAC-FLEX-FAB-A vsan 25
N5548-Fab-A(config-zoneset)# zone name B200M3-CH1-BL1-ORARAC1-vHBA1
N5548-Fab-A(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:02 - (node1, 
Host Initiator1)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:9d:11:05:df  - 
(Controller_A, Port_0C)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:8d:11:05:df  - 
(Controller_B, Port_0C)
N5548-Fab-A(config-zoneset-zone)# exit
N5548-Fab-A(config-zoneset)# zone name B200M3-CH1-BL2-ORARAC2-vHBA1
N5548-Fab-A(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:04 -  (node2, 
Host Initiator1)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:9d:11:05:df  - 
(Controller_A, Port_0C)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:8d:11:05:df  - 
(Controller_B, Port_0C)
N5548-Fab-A(config-zoneset-zone)# exit
N5548-Fab-A(config-zoneset)# zone name B200M3-CH2-BL1-ORARAC3-vHBA1
N5548-Fab-A(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:06 - (node3, 
Host Initiator1)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:9d:11:05:df  - 
(Controller_A, Port_0C)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:8d:11:05:df  - 
(Controller_B, Port_0C)
N5548-Fab-A(config-zoneset-zone)# exit
N5548-Fab-A(config-zoneset)# zone name B200M3-CH2-BL2-ORARAC4-vHBA1
N5548-Fab-A(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:08 - (node4, 
Host Initiator1)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:9d:11:05:df  - 
(Controller_A, Port_0C)
N5548-Fab-A(config-zoneset-zone)# member pwwn 50:0a:09:83:8d:11:05:df  - 
(Controller_B, Port_0C)
N5548-Fab-A(config-zoneset-zone)# exit
N5548-Fab-A(config-zoneset)# exit
N5548-Fab-A(config)# zoneset activate name ORARAC-FLEX-FAB-A vsan 25
N5548-Fab-A(config)# Copy running-config startup-config
```

Nexus 5548 Fabric B Zoning Configuration

```
N5548-Fab-B# Config terminal
```

```
N5548-Fab-B(config)# zoneset name ORARAC-FLEX-FAB-B vsan 25
N5548-Fab-B(config-zoneset)# zone name B200M3-CH1-BL1-ORARAC1-vHBA2
N5548-Fab-B(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:01 - (node1, 
Host Initiator2)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:9d:11:05:df  - 
(Controller_A, Port_0D)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:8d:11:05:df  - 
(Controller_B, Port_0D)
N5548-Fab-B(config-zoneset-zone)# exit

N5548-Fab-B(config-zoneset)# zone name B200M3-CH1-BL2-ORARAC2-vHBA2
N5548-Fab-B(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:03 - (node2, 
Host Initiator2)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:9d:11:05:df  - 
(Controller_A, Port_0D)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:8d:11:05:df  - 
(Controller_B, Port_0D)
N5548-Fab-B(config-zoneset-zone)# exit
N5548-Fab-B(config-zoneset)# zone name B200M3-CH2-BL1-ORARAC3-vHBA2
N5548-Fab-B(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:05 - (node3, 
Host Initiator2)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:9d:11:05:df  - 
(Controller_A, Port_0D)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:8d:11:05:df  - 
(Controller_B, Port_0D)
N5548-Fab-B(config-zoneset-zone)# exit
N5548-Fab-B(config-zoneset)# zone name B200M3-CH2-BL2-ORARAC4-vHBA2
N5548-Fab-B(config-zoneset-zone)# member pwwn 20:00:01:25:b5:11:13:07 - (node4, 
Host Initiator2)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:9d:11:05:df  - 
(Controller_A, Port_0D)
N5548-Fab-B(config-zoneset-zone)# member pwwn 50:0a:09:84:8d:11:05:df  - 
(Controller_B, Port_0D)
N5548-Fab-B(config-zoneset-zone)# exit
N5548-Fab-B(config-zoneset)# exit
N5548-Fab-B(config)# zoneset activate name ORARAC-FLEX-FAB-B vsan 25
N5548-Fab-B(config)# Copy running-config startup-config
```

Appendix B—Cisco Nexus 5548UP Switch Running Configuration

Show running configuration for Nexus 5548 Fabric A Switch: (Partial section)

```
  version 5.0(3)N2(1)
feature fcoe
feature npiv
feature fport-channel-trunk
feature telnet
cfs ipv4 distribute
cfs eth distribute
feature interface-VLAN
feature hsrp
feature lacp
feature vpc
feature lldp
```

```
logging level aaa 5
logging level cdp 6
logging level vpc 6
logging level lldp 5
logging level flogi 5
logging level radius 5
logging level monitor 6
logging level session-mgr 6
logging level spanning-tree 6
logging level interface-VLAN 5
username admin password 5 $1$SS5YPZF2$FT4bt4bHDqZDBMAIeueAV1  role network-admin
no password strength-check
ip domain-lookup
hostname N5548-Fab-A
class-map type qos class-fcoe
class-map type qos match-all class-platinum
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos system_qos_policy
  class class-platinum
    set qos-group 2
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-platinum
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
    multicast-optimize
policy-map type network-qos system_nq_policy
  class type network-qos class-platinum
    mtu 9216
    pause no-drop
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-default
    mtu 9216
    multicast-optimize
system qos
  service-policy type network-qos jumbo
slot 1
slot 2
port 11-16 type fc
snmp-server user admin network-admin auth md5 0x12662623c9fe652fc2205684d4feb333
 priv 0x12662623c9fe652fc2205684d4feb333 localizedkey
snmp-server enable traps entity fru
vrf context management
  ip route 0.0.0.0/0 10.29.134.1
VLAN 1
VLAN 10
```

```
  name Oracle_RAC_Private_Traffic
  no ip igmp snooping
VLAN 20
  name Storage-VLAN
VLAN 30
  name Storage-VLAN30
VLAN 134
  name Oracle_RAC_Public_Traffic"
vpc domain 1
  peer-keepalive destination 10.29.134.15
vsan database
  vsan 15 name "Fabric_A"
  vsan 25

fcdomain fcid database
  vsan 1 wwn 20:4d:54:7f:ee:45:27:c0 fcid 0xc60000 dynamic
  vsan 1 wwn 20:4e:54:7f:ee:45:27:c0 fcid 0xc60001 dynamic
  vsan 1 wwn 20:4f:54:7f:ee:45:27:c0 fcid 0xc60002 dynamic
  vsan 1 wwn 20:50:54:7f:ee:45:27:c0 fcid 0xc60003 dynamic
  vsan 1 wwn 20:1f:54:7f:ee:73:74:80 fcid 0xc60004 dynamic
  vsan 1 wwn 20:20:54:7f:ee:73:74:80 fcid 0xc60005 dynamic
  vsan 1 wwn 50:0a:09:84:9d:11:05:df fcid 0xc60006 dynamic
  vsan 1 wwn 50:0a:09:84:8d:11:05:df fcid 0xc60007 dynamic
  vsan 1 wwn 50:0a:09:83:9d:11:05:df fcid 0xc60008 dynamic
  vsan 1 wwn 50:0a:09:83:8d:11:05:df fcid 0xc60009 dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:06 fcid 0xc6000a dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:04 fcid 0xc6000b dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:02 fcid 0xc6000c dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:08 fcid 0xc6000d dynamic
  vsan 25 wwn 20:1f:54:7f:ee:73:74:80 fcid 0x560000 dynamic
  vsan 25 wwn 20:20:54:7f:ee:73:74:80 fcid 0x560001 dynamic
  vsan 25 wwn 50:0a:09:83:9d:11:05:df fcid 0x560002 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:02 fcid 0x560003 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:04 fcid 0x560004 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:06 fcid 0x560005 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:08 fcid 0x560006 dynamic
  vsan 25 wwn 50:0a:09:83:8d:11:05:df fcid 0x560007 dynamic


interface Vlan1

interface Vlan20
  no shutdown
  description Storage-VLAN
  ip address 10.10.20.2/24
  hsrp version 2
  hsrp 20
    preempt
    priority 110
    ip 10.10.20.1

interface Vlan30
  no shutdown
  ip address 10.10.30.2/24
  hsrp version 2
  hsrp 30
    preempt
    priority 111
    ip 10.10.30.1

interface port-channel1
description Port-Channel for vPC Peer-link
  switchport mode trunk
  vpc peer-link
  switchport trunk allowed vlan 1,10,20,30,134
  spanning-tree port type network

interface port-channel17
  description Port-Channel for Fabric InterconnectA
  switchport mode trunk
  vpc 17
  switchport trunk allowed vlan 1,10,20,30,134
  spanning-tree port type edge trunk

interface port-channel18
  description Port-Channel for Fabric InterconnectB
  switchport mode trunk
  vpc 18
  switchport trunk allowed vlan 1,10,20,30,134
  spanning-tree port type edge trunk

interface port-channel20
  description PortChannel for multimode VIF from ControllerA-10G
  switchport mode trunk
  vpc 20
  switchport trunk native vlan 20
  switchport trunk allowed vlan 20,30
  spanning-tree port type edge trunk

interface port-channel30
  description PortChannel for multimode VIF from ControllerB-10G
  switchport mode trunk
  vpc 30
  switchport trunk native vlan 30
  switchport trunk allowed vlan 20,30
  spanning-tree port type edge trunk

vsan database
  vsan 25 interface fc2/1
  vsan 25 interface fc2/2
  vsan 25 interface fc2/3
  vsan 25 interface fc2/4

interface fc2/1
  no shutdown

interface fc2/2
  no shutdown

interface fc2/3
  no shutdown

interface fc2/4
  no shutdown

interface fc2/5
```

```
  no shutdown

interface fc2/6
  no shutdown

interface Ethernet1/1
  description Peer link connected to N5548B-Eth1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 1 mode active

interface Ethernet1/2
  description Peer link connected to N5548B-Eth1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 1 mode active

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16
  description "Public Traffic to 3750"
  switchport mode trunk
  switchport trunk native vlan 134
  switchport trunk allowed vlan 1,10-11,134

interface Ethernet1/17
  description Connection from Fabric Interconnect-A Eth1/17
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 17 mode active

interface Ethernet1/18
  description Connection from Fabric Interconnect-B Eth1/17
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 18 mode active

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet2/1
  description Connection to NetApp Controller-A-Port-e1a
  switchport mode trunk
  switchport trunk native vlan 20
  switchport trunk allowed vlan 20,30
  channel-group 20


interface Ethernet2/2
  description Connection to NetApp Controller-B-Port-e1a
  switchport mode trunk
  switchport trunk native vlan 30
  switchport trunk allowed vlan 20,30
  channel-group 30

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6
 interface Ethernet2/7

interface Ethernet2/8

interface Ethernet2/9
interface Ethernet2/10


interface mgmt0
  ip address 10.29.134.16/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.0.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.0.3.N2.1.bin
no ip igmp snooping
interface fc2/1
interface fc2/2
interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
logging logfile messages 6
!Full Zone Database Section for vsan 25
zone name B200M3-CH1-BL1-ORARAC1-vHBA1 vsan 25
    member pwwn 20:00:01:25:b5:11:13:02
    member pwwn 50:0a:09:83:9d:11:05:df
    member pwwn 50:0a:09:83:8d:11:05:df

zone name B200M3-CH1-BL2-ORARAC2-vHBA1 vsan 25
    member pwwn 20:00:01:25:b5:11:13:04
    member pwwn 50:0a:09:83:9d:11:05:df
    member pwwn 50:0a:09:83:8d:11:05:df

zone name B200M3-CH2-BL1-ORARAC3-vHBA1 vsan 25
    member pwwn 50:0a:09:83:8d:11:05:df
     member pwwn 50:0a:09:83:9d:11:05:df
    member pwwn 20:00:01:25:b5:11:13:06

zone name B200M3-CH2-BL2-ORARAC4-vHBA1 vsan 25
    member pwwn 20:00:01:25:b5:11:13:08
    member pwwn 50:0a:09:83:9d:11:05:df
    member pwwn 50:0a:09:83:8d:11:05:df

zoneset name ORARAC-FLEX-FAB-A vsan 25
    member B200M3-CH1-BL1-ORARAC1-vHBA1
    member B200M3-CH1-BL2-ORARAC2-vHBA1
    member B200M3-CH2-BL1-ORARAC3-vHBA1
    member B200M3-CH2-BL2-ORARAC4-vHBA1

zoneset activate name ORARAC-FLEX-FAB-A vsan 25
```

Show running configuration for Nexus 5548 Fabric B Switch:

```
N5548-Fab-B(config)# sh running-config
!Command: show running-config
!Time: Fri Mar  8 20:48:15 2013

version 5.0(3)N2(1)
feature fcoe
feature npiv
feature fport-channel-trunk
feature telnet
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
logging level aaa 5
logging level cdp 6
logging level vpc 6
logging level lldp 5
logging level flogi 5
logging level radius 5
logging level monitor 6
logging level session-mgr 6
logging level spanning-tree 6
logging level interface-vlan 5
username admin password 5 $1$ykW.XCAE$dvZtFvO5bBkGxyjQ3gOBh/  role network-admin
no password strength-check
ip domain-lookup
hostname N5548-Fab-B
class-map type qos class-fcoe
class-map type qos match-all class-platinum
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos system_qos_policy
  class class-fcoe
    set qos-group 1
policy-map type queuing system_q_in_policy
  class type queuing class-fcoe
    bandwidth percent 46
   class type queuing class-default
    bandwidth percent 9
policy-map type queuing system_q_out_policy
  class type queuing class-fcoe
    bandwidth percent 46
  class type queuing class-default
    bandwidth percent 9
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-platinum
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
    multicast-optimize
policy-map type network-qos system_nq_policy
  class type network-qos class-platinum
    mtu 9000
    pause no-drop
 class type network-qos class-fcoe
```

```
    pause no-drop
    mtu 2158
  class type network-qos class-default
    mtu 9000
    multicast-optimize
system qos
  service-policy type network-qos jumbo
slot 1
slot 2
port 11-16 type fc
snmp-server user admin network-admin auth md5 0x9c61fe8ad238d5db5ea985e10b5fa661
 priv 0x9c61fe8ad238d5db5ea985e10b5fa661 localizedkey
snmp-server enable traps entity fru
vrf context management
  ip route 0.0.0.0/0 10.29.134.1
vlan 1
vlan 10
  name Oracle_RAC_Private_Traffic
  no ip igmp snooping
vlan 20
  name Storage-VLAN
 vlan 30
  name Storage-VLAN30
vlan 134
  name Oracle_RAC_Public_Traffic
vpc domain 1
  peer-keepalive destination 10.29.134.16
vsan database
  vsan 25
device-alias database
  device-alias name A2P0 pwwn 50:06:01:60:47:20:2c:af
  device-alias name A2P2 pwwn 50:06:01:62:47:20:2c:af
  device-alias name A3P0 pwwn 50:06:01:64:47:20:2c:af
  device-alias name A3P2 pwwn 50:06:01:66:47:20:2c:af
  device-alias name B2P0 pwwn 50:06:01:68:47:20:2c:af
  device-alias name B2P2 pwwn 50:06:01:6a:47:20:2c:af
  device-alias name B3P0 pwwn 50:06:01:6c:47:20:2c:af
  device-alias name B3P2 pwwn 50:06:01:6e:47:20:2c:af

device-alias commit

fcdomain fcid database
   vsan 1 wwn 50:06:01:60:47:20:2c:af fcid 0xea00ef dynamic
!            [A2P0]
  vsan 1 wwn 20:4d:54:7f:ee:76:cd:80 fcid 0xea0000 dynamic
  vsan 1 wwn 20:4e:54:7f:ee:76:cd:80 fcid 0xea0001 dynamic
  vsan 1 wwn 20:4f:54:7f:ee:76:cd:80 fcid 0xea0002 dynamic
  vsan 1 wwn 20:50:54:7f:ee:76:cd:80 fcid 0xea0003 dynamic
  vsan 1 wwn 50:06:01:62:47:20:2c:af fcid 0xea01ef dynamic
!            [A2P2]
  vsan 1 wwn 50:06:01:68:47:20:2c:af fcid 0xea02ef dynamic
!            [B2P0]
  vsan 1 wwn 50:06:01:6a:47:20:2c:af fcid 0xea03ef dynamic
!            [B2P2]
  vsan 1 wwn 50:06:01:64:47:20:2c:af fcid 0xea04ef dynamic
!            [A3P0]
  vsan 1 wwn 50:06:01:66:47:20:2c:af fcid 0xea05ef dynamic
!            [A3P2]
```

```
  vsan 1 wwn 50:06:01:6c:47:20:2c:af fcid 0xea06ef dynamic
 !            [B3P0]
  vsan 1 wwn 50:06:01:6e:47:20:2c:af fcid 0xea07ef dynamic
!            [B3P2]
  vsan 1 wwn 20:1f:54:7f:ee:73:74:80 fcid 0xea0004 dynamic
  vsan 1 wwn 20:20:54:7f:ee:73:74:80 fcid 0xea0005 dynamic
  vsan 1 wwn 50:0a:09:83:9d:11:05:df fcid 0xea0006 dynamic
  vsan 1 wwn 50:0a:09:84:9d:11:05:df fcid 0xea0007 dynamic
  vsan 1 wwn 20:1f:54:7f:ee:77:5b:c0 fcid 0xea0008 dynamic
  vsan 1 wwn 20:20:54:7f:ee:77:5b:c0 fcid 0xea0009 dynamic
  vsan 1 wwn 50:0a:09:83:8d:11:05:df fcid 0xea000a dynamic
  vsan 1 wwn 50:0a:09:84:8d:11:05:df fcid 0xea000b dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:05 fcid 0xea000c dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:03 fcid 0xea000d dynamic
   vsan 1 wwn 20:00:01:25:b5:11:13:01 fcid 0xea000e dynamic
  vsan 1 wwn 20:00:01:25:b5:11:13:07 fcid 0xea000f dynamic
  vsan 25 wwn 20:1f:54:7f:ee:77:5b:c0 fcid 0x540000 dynamic
  vsan 25 wwn 20:20:54:7f:ee:77:5b:c0 fcid 0x540001 dynamic
  vsan 25 wwn 50:0a:09:84:9d:11:05:df fcid 0x540002 dynamic
  vsan 25 wwn 50:0a:09:84:8d:11:05:df fcid 0x540003 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:01 fcid 0x540004 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:03 fcid 0x540005 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:05 fcid 0x540006 dynamic
  vsan 25 wwn 20:00:01:25:b5:11:13:07 fcid 0x540007 dynamic
interface Vlan1

interface Vlan20
  description Storage1-VLAN-ControllerA
  ip address 10.10.20.3/24
  hsrp version 2
  hsrp 20
    preempt
    ip 10.10.20.1

interface Vlan30
  description Storage2-VLAN-ControllerB
  no shutdown
  ip address 10.10.30.3/24
  hsrp version 2
  hsrp 30
    preempt
    priority 111
    ip 10.10.30.1

interface port-channel1
description Port-Channel for vPC Peer-link
  switchport mode trunk
  vpc peer-link
  switchport trunk allowed vlan 1,10,20,30,134
  spanning-tree port type network

interface port-channel17
  switchport mode trunk
  vpc 17   description Port-Channel for Fabric InterconnectA
  switchport trunk allowed vlan 1,10,20,30,134
  spanning-tree port type edge trunk

interface port-channel18
```

```
  switchport mode trunk
  vpc 18
  description Port-Channel for Fabric InterconnectB
  switchport trunk allowed vlan 1,10,20,30,134
  spanning-tree port type edge trunk

interface port-channel20
  description PortChannel for multimode VIF from ControllerA-10G
  switchport mode trunk
  vpc 20
  switchport trunk native vlan 20
  switchport trunk allowed vlan 20,30
  spanning-tree port type edge trunk

interface port-channel30
  description PortChannel for multimode VIF from ControllerB-10G
  switchport mode trunk
  vpc 30
  switchport trunk native vlan 30
  switchport trunk allowed vlan 20,30
  spanning-tree port type edge trunk

vsan database
  vsan 25 interface fc2/1
  vsan 25 interface fc2/2
  vsan 25 interface fc2/3
  vsan 25 interface fc2/4


interface fc2/1
  no shutdown

interface fc2/2
  no shutdown

interface fc2/3
  no shutdown

interface fc2/4
  no shutdown

interface fc2/5
  no shutdown
 interface fc2/6
  no shutdown

interface Ethernet1/1
  description Peer link connected to N5548A-Eth1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 1 mode active

interface Ethernet1/2
  description Peer link connected to N5548A-Eth1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 1 mode active
```

```
interface Ethernet1/3

interface Ethernet1/4
 interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

 interface Ethernet1/16
  description "Public Traffic to 3750"
  switchport mode trunk
  switchport access vlan 134
  switchport trunk native vlan 134
  switchport trunk allowed vlan 1,10-11,134

interface Ethernet1/17
  description Connection from Fabric Interconnect-A Eth1/18
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 17 mode active

interface Ethernet1/18
  description Connection from Fabric Interconnect-B Eth1/18
  switchport mode trunk
  switchport trunk allowed vlan 1,10,20,30,134
  channel-group 18 mode active

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet2/1
  description Connection to NetApp Controller-A-Port-e1b
  switchport mode trunk
  switchport trunk native vlan 20
  switchport trunk allowed vlan 20,30
  channel-group 20

interface Ethernet2/2

interface Ethernet2/3
  description Connection to NetApp Controller-B-Port-e1b
  switchport mode trunk
  switchport trunk native vlan 30
  switchport trunk allowed vlan 20,30
  channel-group 30

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6

interface Ethernet2/7

interface Ethernet2/8

interface Ethernet2/9

interface Ethernet2/10

interface mgmt0
  ip address 10.29.134.15/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.0.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.0.3.N2.1.bin
no ip igmp snooping
interface fc2/1
interface fc2/2
interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
logging logfile messages 6
```

```
!Full Zone Database Section for vsan 25
zone name B200M3-CH1-BL1-ORARAC1-vHBA2 vsan 25
    member pwwn 20:00:01:25:b5:11:13:01
    member pwwn 50:0a:09:84:9d:11:05:df
    member pwwn 50:0a:09:84:8d:11:05:df

zone name B200M3-CH1-BL2-ORARAC2-vHBA2 vsan 25
    member pwwn 20:00:01:25:b5:11:13:03
    member pwwn 50:0a:09:84:8d:11:05:df
    member pwwn 50:0a:09:84:9d:11:05:df

zone name B200M3-CH2-BL1-ORARAC3-vHBA2 vsan 25
    member pwwn 20:00:01:25:b5:11:13:05
     member pwwn 50:0a:09:84:8d:11:05:df
    member pwwn 50:0a:09:84:9d:11:05:df

zone name B200M3-CH2-BL2-ORARAC4-vHBA2 vsan 25
    member pwwn 20:00:01:25:b5:11:13:07
    member pwwn 50:0a:09:84:8d:11:05:df
    member pwwn 50:0a:09:84:9d:11:05:df

zoneset name ORARAC-FLEX-FAB-B vsan 25
    member B200M3-CH1-BL1-ORARAC1-vHBA2
    member B200M3-CH1-BL2-ORARAC2-vHBA2
    member B200M3-CH2-BL1-ORARAC3-vHBA2
    member B200M3-CH2-BL2-ORARAC4-vHBA2

zoneset activate name ORARAC-FLEX-FAB-B vsan 25
```

References

Cisco UCS:

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

NetApp Data Storage Systems:

http://www.netapp.com/us/products/storage-systems/

Cisco Nexus:

http://www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html

Cisco Nexus 5000 Series NX-OS Software Configuration Guide:

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide.html

NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection:

http://www.netapp.com/us/library/technical-reports/tr-3298.html