Oracle 12c on Cisco UCS and EMC VNX2

Deploying Oracle RAC 12cR1 (12.1.0.1) on Oracle Linux 6 using Cisco Unified Computing System 2.1 and EMC VNX8000

Last Updated: August 18, 2014

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Author

RamaKrishna Nishtala, Technical Leader, UCS Product Management and DC Solutions, Cisco Systems, Inc.

Ramakrishna Nishtala, has 20 years of experience with Oracle databases, IT infrastructure, virtualization and cloud computing. He has worked on optimizing, performance tuning and RISC migration of Oracle databases, ebusiness Suite and Fusion architectures. In his current role, Rama focuses on solutions and best practices for implementing Oracle solutions on Cisco Unified Computing System. The author is thankful to his Cisco colleagues and its partners, both Oracle and EMC for their help on this CVD on Oracle 12c.

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.

Oracle 12c on Cisco UCS and EMC VNX2

About this Document

This Cisco Validated Design describes how the Cisco Unified Computing System™ can be used in conjunction with EMC® VNX® storage systems to implement an Oracle Real Application Clusters (RAC) solution that is an Oracle Certified Configuration. The Cisco Unified Computing System provides the compute, network, and storage access components of the cluster, deployed as a single cohesive system. The result is an implementation that addresses many of the challenges that database administrators and their IT departments face today, including needs for a simplified deployment and operation model, high performance for Oracle 12c RAC software, and lower total cost of ownership (TCO). This document introduces the Cisco Unified Computing System and provides instructions for implementing it; it concludes with an analysis of the cluster's performance and reliability characteristics.

Audience

The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Introduction

Data powers essentially every operation in a modern enterprise, from keeping the supply chain operating efficiently to managing relationships with customers. Oracle 12c brings an innovative approach to the challenges of rapidly increasing amounts of data and demand for high performance with its containers. Oracle 12c RAC also uses a horizontal scaling (or scale-out) model that allows organizations to take advantage of the fact that the price of one-to-four-socket x86-architecture servers continues to drop while their processing power increases unabated. The clustered approach allows each server to contribute its processing power to the overall cluster's capacity, enabling a new approach to managing the cluster's performance and capacity.

Leadership from Cisco

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

Historically, enterprise database management systems have run on costly symmetric multiprocessing servers that use a vertical scaling (or scale-up) model. However, as the cost of one-to-four-socket x86-architecture servers continues to drop while their processing power increases, a new model has emerged. Oracle RAC uses a horizontal scaling, or scale-out, model, in which the active-active cluster uses multiple servers, each contributing its processing power to the cluster, increasing performance, scalability, and availability. The cluster balances the workload across the servers in the cluster, and the cluster can provide continuous availability in the event of a failure.

Oracle Certified Configuration

All components in an Oracle RAC implementation must work together flawlessly, and Cisco has worked closely with EMC and Oracle to create, test, and certify a configuration of Oracle RAC on the Cisco Unified Computing System. Cisco's Oracle Certified Configurations provide an implementation of Oracle Database with Real Application Clusters technology consistent with industry best practices. For back-end SAN storage, the certification environment included an EMC VNX storage system with a mix of SAS drives and state-of-the-art Flash drives (FDs) to further speed performance.

Benefits of the Configuration

The Oracle Certified Configuration of Oracle RAC on the Cisco Unified Computing System offers a number of important benefits.

Simplified Deployment and Operation

Because the entire cluster runs on a single cohesive system, database administrators no longer need to painstakingly configure each element in the hardware stack independently. The system's compute, network, and storage-access resources are essentially stateless, provisioned dynamically by Cisco® UCS Manager. This role-based and policy-based embedded management system handles every aspect of system configuration, from a server's firmware and identity settings to the network connections that connect storage traffic to the destination storage system. This capability dramatically simplifies the process of scaling an Oracle RAC configuration or re-hosting an existing node on an upgrade server. Cisco UCS Manager uses the concept of service profiles and service profile templates to consistently and accurately configure resources. The system automatically configures and deploys servers in minutes, rather than the hours or days required by traditional systems composed of discrete, separately managed components. Indeed, Cisco UCS Manager can simplify server deployment to the point where it can automatically discover, provision, and deploy a new blade server when it is inserted into a chassis.

The system is based on a 10-Gbps unified network fabric that radically simplifies cabling at the rack level by consolidating both IP and Fiber Channel traffic onto the same rack-level 10-Gbps converged network. This "wire-once" model allows in-rack network cabling to be configured once, with network features and configurations all implemented by changes in software rather than by error-prone changes in physical cabling. This Cisco Validated Configuration not only supports physically separate public and private networks; it provides redundancy with automatic failover.

High-Performance Platform for Oracle RAC

The Cisco UCS B-Series Blade Servers used in this certified configuration feature Intel Xeon E5- 4650 series processors that deliver intelligent performance, automated energy efficiency, and flexible virtualization. Intel Turbo Boost Technology automatically boosts processing power through increased frequency and use of hyper threading to deliver high performance when workloads demand and thermal conditions permit.

The Cisco Unified Computing System's 10-Gbps unified fabric delivers standards-based Ethernet and Fiber Channel over Ethernet (FCoE) capabilities that simplify and secure rack-level cabling while speeding network traffic compared to traditional Gigabit Ethernet networks. The balanced resources of the Cisco Unified Computing System allow the system to easily process an intensive online transaction processing (OLTP) and decision-support system (DSS) workload with no resource saturation.

Safer Deployments with Certified and Validated Configurations

Cisco and Oracle are working together to promote interoperability of Oracle's next-generation database and application solutions with the Cisco Unified Computing System, helping make the Cisco Unified Computing System a simple and reliable platform on which to run Oracle software.

Implementation Instructions

This Cisco Validated Design introduces the Cisco Unified Computing System and discusses the ways it addresses many of the challenges that database administrators and their IT departments face today. The document provides an overview of the certified Oracle 12c RAC configuration along with instructions for setting up the Cisco Unified Computing System and the EMC VNX2 storage system, including database table setup and the use of flash drives. The document reports on Cisco's performance measurements for the cluster and a reliability analysis that demonstrates how the system continues operation even when commonly encountered hardware faults occur.

Introducing the Cisco Unified Computing System

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

Comprehensive Management

The system uses an embedded, end-to-end management system that uses a high-availability active-standby configuration. Cisco UCS Manager uses role and policy-based management that allows IT departments to continue to use subject-matter experts to define server, network, and storage access policy. After a server and its identity, firmware, configuration, and connectivity are defined, the server, or a number of servers like it, can be deployed in minutes, rather than the hours or days that it typically takes to move a server from the loading dock to production use. This capability relieves database administrators from tedious, manual assembly of individual components and makes scaling an Oracle RAC configuration a straightforward process.

Radical Simplification

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

High Performance

The system's blade servers are based on the Intel Xeon 4650 series processors. These processors adapt performance to application demands, increasing the clock rate on specific processor cores as workload and thermal conditions permit. The system is integrated within a 10 Gigabit Ethernet-based unified fabric that delivers the throughput and low-latency characteristics needed to support the demands of the cluster's public network, storage traffic, and high-volume cluster messaging traffic.

Overview of Cisco Unified Computing System

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

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

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

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

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

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

Fabric Interconnects

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

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

The Cisco UCS 6248 Fabric Interconnect is a one-rack-unit (1RU), 10 Gigabit Ethernet, IEEE Data Center Bridging (DCB), and FCoE interconnect built to provide 960 Gbps throughput with very low latency. It has 48 high density ports in 1RU including one expansion module with 16 unified ports. Like its predecessors, it can be seamlessly managed with Cisco UCS Manager.

Fabric Extenders

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

The Cisco UCS 2208XP Fabric Extender has eight 10 Gigabit Ethernet, FCoE-capable, enhanced small Form-Factor Pluggable (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis. Each fabric extender on either sides of the chassis are connected through 8 x 10 Gb links to the fabric interconnects and offer:

- Connection of the Cisco UCS blade chassis to the Fabric Interconnect
- Eight 10 Gigabit Ethernet, FCoE-capable SFP+ ports
- Built-in chassis management function to manage the chassis environment (the power supply and fans as well as the blades) along with the Fabric Interconnect, eliminating the need for separate chassis management modules
- Full management by Cisco UCS Manager through the Fabric Interconnect
- Support for up to two Fabric Extenders, enabling increased capacity as well as redundancy
- Up to 160 Gbps of bandwidth per chassis

Blade Chassis

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

Cisco UCS Manager

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

Cisco UCS VIC 1280 Adapters

Cisco VIC 1280 is the second generation of Mezzanine Adapters from Cisco. VIC 1280 supports up to 256 PCI-e devices and up to 80 Gbps of throughput. Compared with its earlier generation of Palo Adapters it had doubled the capacity in throughput and PCI-e devices and is complaint with many OS and storage Vendors.

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 B420 M3 High-Performance Blade Servers

The Cisco UCS B420 M3 High-Performance Blade Servers are full-slot, 4-socket, high-performance blade servers offering the performance and reliability of the Intel Xeon processor E5-4650 product family and up to 1 TB of memory. The Cisco UCS B420 M3 supports four Small Form Factor (SFF) SAS and SSD drives and two converged network adapter (CNA) mezzanine slots up to 160 Gbps of I/O throughput. The Cisco UCS B420 M3 blade server extends Cisco UCS by offering increased levels of performance, scalability, and reliability for mission-critical workloads.

The Cisco UCS components used in the certified configuration are shown in Figure 1.

Figure 1 Cisco Unified Computing System Components

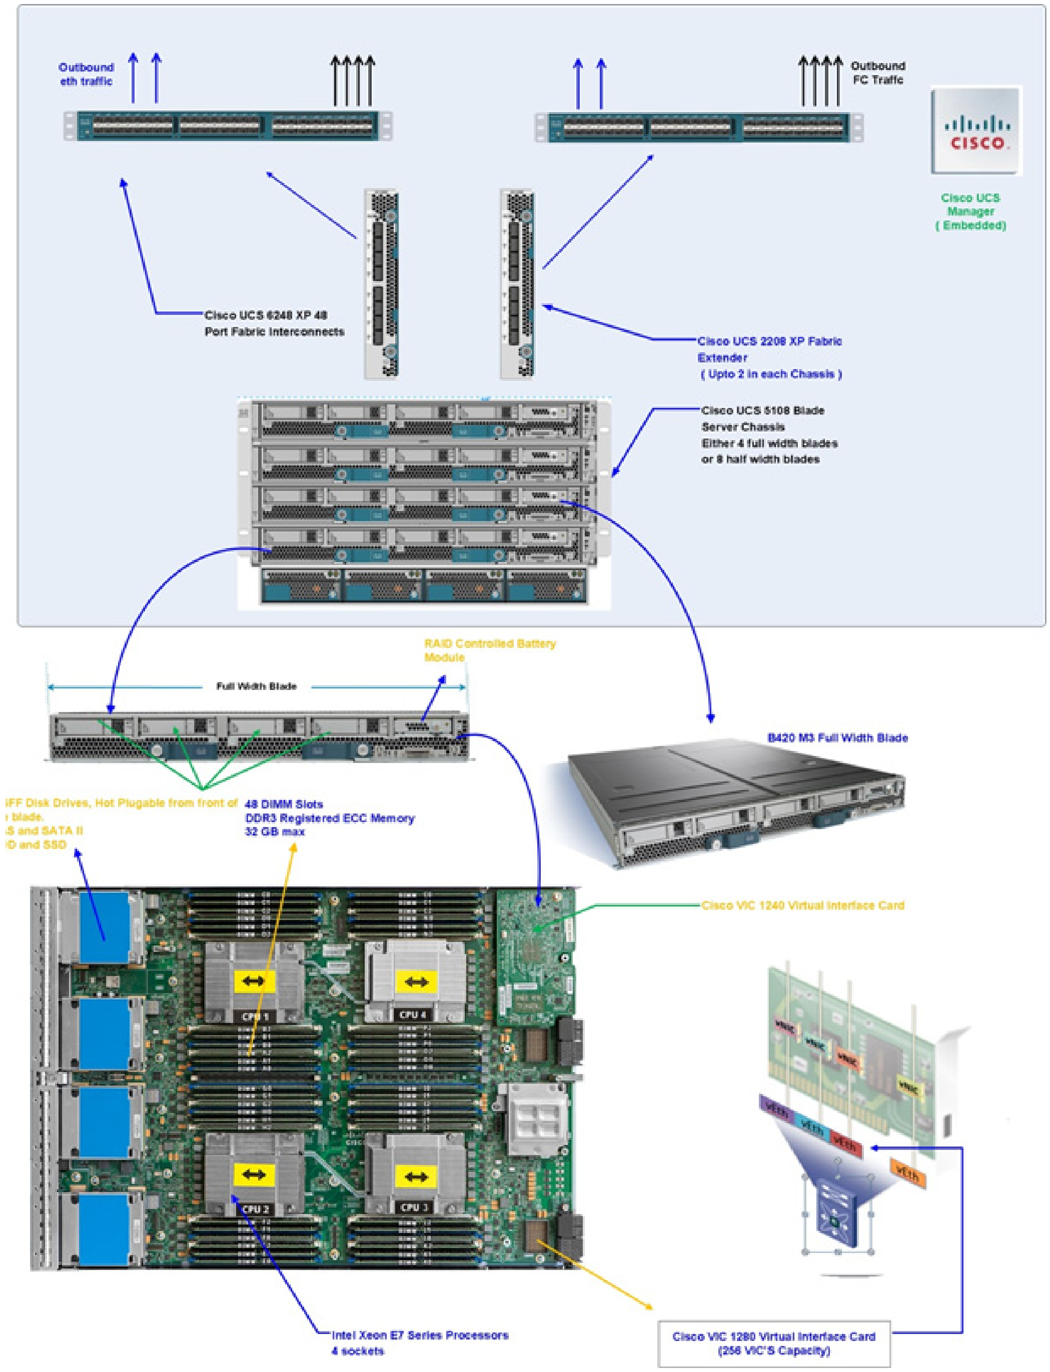

Service Profiles: Cisco Unified Computing System Foundation Technology

Cisco UCS resources are abstract in the sense that their identity, I/O configuration, MAC addresses and worldwide names (WWNs), firmware versions, BIOS boot order, and network attributes (including quality of service (QoS) settings, pin groups, and threshold policies) are all programmable using a just-in-time deployment model. The manager stores this identity, connectivity, and configuration information in service profiles that reside on the Cisco UCS 6200 Series Fabric Interconnects. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. 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.

Service Profile - Description, Overview and Elements

Service Profile - Description

Conceptually, a service profile is an extension of the virtual machine abstraction applied to physical servers. The definition has been expanded to include elements of the environment that span the entire data center, encapsulating the server identity (LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN physical port, and quality-of-service [QoS] policies) in logical "service profiles" that can be dynamically created and associated with any physical server in the system within minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation. It enables migration of identities between servers in the environment without requiring any physical configuration changes and facilitates rapid bare metal provisioning of replacements for failed servers. Service profiles also include operational policy information, such as information about firmware versions.

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

Service Profile - Overview

A service profile typically includes four types of information:

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

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

Service Profile Elements

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

Figure 2 below figure represents how Server, Network, and Storage Policies are encapsulated in a service profile.

Figure 2 Service Profile inclusions

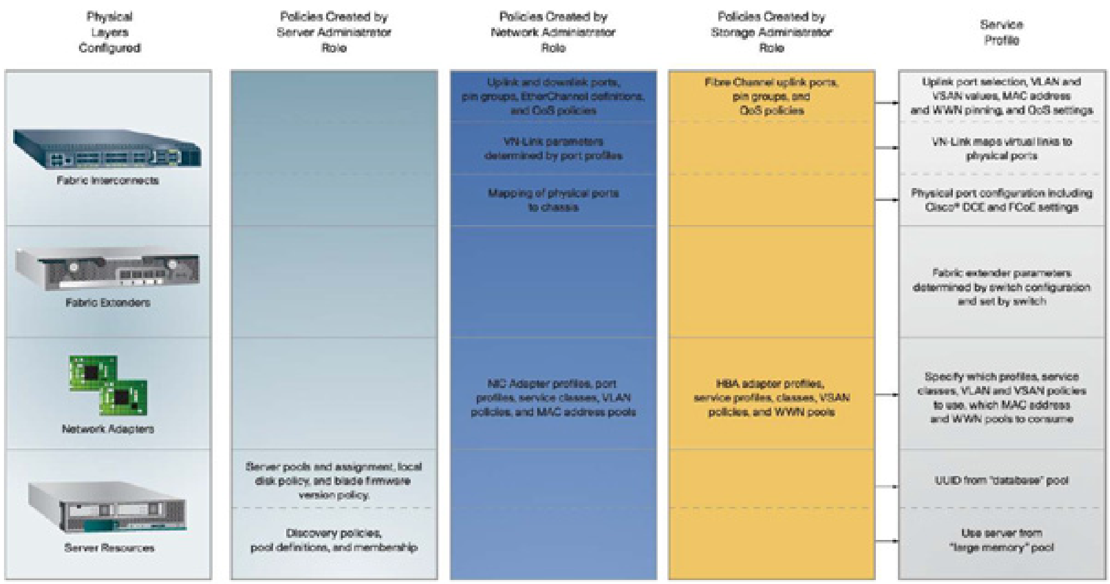

Understanding the Service Profile Template

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

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

The Cisco Unified Computing System used for the certified configuration is based on Cisco B-Series Blade Servers; however, the breadth of Cisco's server and network product line suggests that similar product combinations will meet the same requirements.

The system used to create the Oracle Certified Configuration is built from the hierarchy of components illustrated in Figure 1.

- The Cisco UCS 6248 XP 48-Port Fabric Interconnect provides low-latency, lossless, 10-Gbps unified fabric connectivity for the cluster. The fabric interconnect provides connectivity to blade server chassis and the enterprise IP network. Through a 16-port, 8-Gbps Fiber Channel expansion card, the fabric interconnect provides native Fiber Channel access to the EMC VNX storage system. Two fabric interconnects are configured in the cluster, providing physical separation between the public and private networks and also providing the capability to securely host both networks in the event of a failure.
- The Cisco UCS 2208XP Fabric Extender brings the unified fabric into each blade server chassis. The fabric extender is configured and managed by the fabric interconnects, eliminating the complexity of blade-server-resident switches. Two fabric extenders are configured in each of the cluster's two blade server chassis.
- The Cisco UCS 5108 Blade Server Chassis houses the fabric extenders, up to four power supplies, and up to four full width blade servers. As part of the system's radical simplification, the blade server chassis is also managed by the fabric interconnects, eliminating another point of management. Two chassis were configured for the Oracle 12c RAC described in this document.

  The blade chassis supports up to eight half-width blades or up to four full-width blades. The certified configuration used four (two in each chassis) Cisco UCS B420 M3 full width Blade Servers, each equipped with four 8-core Intel Xeon E5-4650 series processors. Each blade server was configured with 256 GB of memory.
- The blade server form factor supports a range of mezzanine-format Cisco UCS network adapters, including a 80 Gigabit Ethernet network adapter designed for efficiency and performance, the Cisco UCS VIC 1240 and VIC 1280 Virtual Interface Cards designed to deliver outstanding performance and full compatibility with existing Ethernet and Fiber Channel environments. These adapters present both an Ethernet network interface card (NIC) and a Fiber Channel host bus adapter (HBA) to the host operating system. They make the existence of the unified fabric transparent to the operating system, passing traffic from both the NIC and the HBA onto the unified fabric. This certified configuration used Cisco UCS VIC 1240 and VIC 1280 Virtual Interface Network Adapters (2 adapters per blade) that provide 120 Gbps of performance per blade server.

Cisco Nexus 5548UP Switch

Figure 3 Cisco Nexus 5548UP Switch

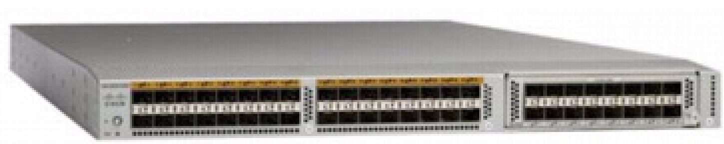

The Cisco Nexus 5548UP switch delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages,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

EMC VNX Unified Storage System

EMC VNX series unified storage systems deliver uncompromising scalability and flexibility, while providing market-leading simplicity and efficiency to minimize total cost of ownership.

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

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

A robust platform designed to deliver five 9s availability, the VNX series enable organizations to dynamically grow, share, and cost-effectively manage multi-protocol file systems and multi-protocol block storage access. The VNX series has been expressly designed to take advantage of the latest innovation in Flash drive technology, maximizing the storage system's performance and efficiency while minimizing cost per GB.

This VNX generation has active/active storage processors for higher performance when accessing LUNs, compared to the previous generation's active/passive scheme, which meant LUN access through the passive storage processor was slower than through the active one.

The VNX controlling software is now called its Operating Environment (OE) and it can use multi-core CPUs better, with work being spread across the available cores and hence increasing performance.

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

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

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

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

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

  For additional information about the EMC VNX Series refer to: http://www.emc.com/storage/vnx/vnx-series.htm

For details regarding EMC VNX Series Software Suites and the resulting value in performance, protection, and TCO that can be derived, refer to:
http://www.emc.com/collateral/software/data-sheet/h8509-vnx-software-suites-ds.pdf

Figure 4 EMC VNX Storage Systems

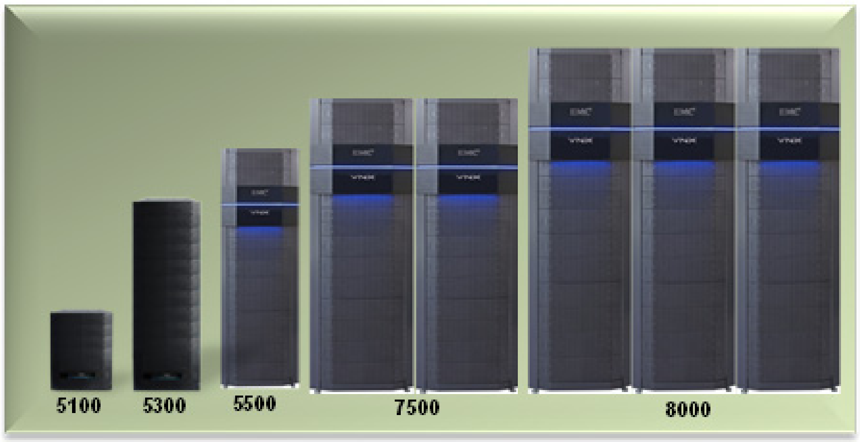

For additional information about the features available in the VNX product line pertaining to your Oracle deployment environment, refer to
http://www.emc.com/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf

Cisco Certified Configuration Inventory and Solution Overview

The inventory of the components used in the certification stack is detailed in Table 1.

The configuration presented in this Cisco Validated Design is based on the Oracle Database 12c Release 1 with Real Application Clusters technology certification environment specified for an Oracle RAC and EMC VNX storage system is as shown in Figure 5.

Inventory of the Certified Configuration

Table 1 Inventory of the Certified Configuration

Physical Cisco Unified Computing System Server Configuration	
Description	Quantity
Cisco UCS 5108 Blade Server Chassis, with 4 power supply units, 8 fans and 2 fabric extenders	2
Cisco UCS B420-M3 full width blades	8
4 Socket - Intel Xeon E5-4650 2.70 GHz processors, 8 cores per socket on 8 blades	256
16 GB DDR3-1600 MHz DIMM (16 per server, totaling 256 GB per blade server)	128
Cisco UCS VIC 1240 Virtual Interface Card, 256 PCI devices, Dual 4 x 10G (1 per server)	8
Cisco UCS VIC 1280 Virtual Interface Card, 256 PCI devices, Dual 4 x 10G (1 per server)	8
Hard Disk Drives (SAN Boot Install)	0
Cisco UCS - 6248XP 48 port Fabric Interconnect	2
16 port 8 Gbps Fiber Channel expansion module	2

LAN and SAN Components	
Description	Quantity
LAN and SAN Components	
Cisco Nexus 5548 UP Switch	2
VLAN Configuration	VLAN ID
Public VLAN	134
Private Network VLAN (Private traffic VLAN must be configured on the Nexus switches to ensure traffic flow in partial link failure as discussed later	10
VSAN Configuration	VSAN ID
Oracle database VSAN	15

Storage Configuration	
Description	Quantity
EMC VNX 8000 Storage System	1
600 GB 15k RPM, SAS disk drivers	300
200 GB flash drives	85
100 GB flash drives for cache	15
Bus	8
Enclosures	16

Operating System and RPM Components (installed on all Oracle nodes)	OS and RPM's
Operating System (64 bit)	Oracle Linux 6.3 x86_64(2.6.39-200.24.1.el6uek.x86_64)
Required RPM's by EMC (to be installed on all Cluster nodes to support EMC Power Path and Host agent)	EMCpower.LINUX-5.7.1.02.00-004.ol6_uek2_r2.x86_64 HostAgent-Linux-64-x86-en_US-1.0.0.1.0474-1.x86_64

Figure 5 Oracle Database 12cR1 with Real Application Clusters technology on Cisco Unified Computing System and EMC VNX Storage

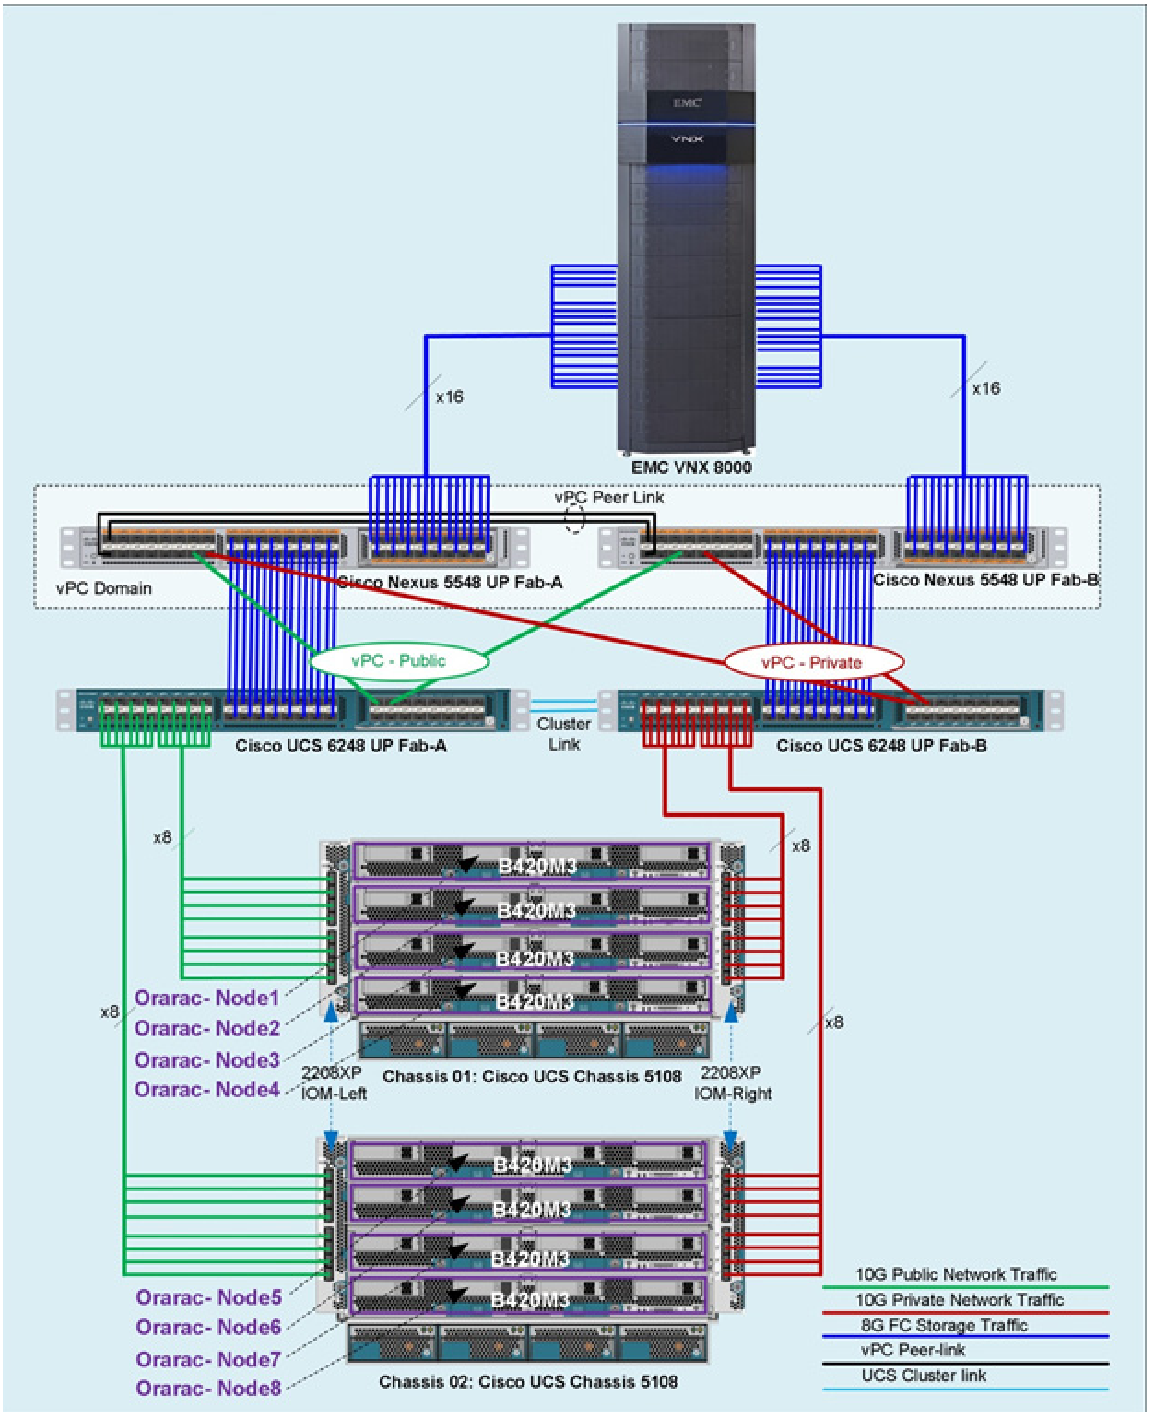

In Figure 5, the green lines indicate the public network connecting to Fabric Interconnect A, and the red lines indicate the private interconnects connecting to Fabric Interconnect B. For Oracle RAC environments, it is a best practice to keep all private interconnect (intra-blade) traffic to one Fabric interconnect. The public and private VLANs spanning the fabric interconnects help ensure the connectivity in case of link failure. Note that the FCoE communication takes place between the Cisco

Unified Computing System chassis and fabric interconnects. The Fiber channel traffic leaves the UCS Fabrics through their own N5k Switches to EMC (blue lines). This is a typical configuration that can be deployed in a customer's environment. The best practices and setup recommendations are described in subsequent sections of this document.

Configuring Cisco Unified Computing System for the 8 node Oracle RAC

For detailed information about configuring the Cisco UCS system refer to:
http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html

Configuring Fabric Interconnects

Two Cisco UCS 6248 UP Fabric Interconnects are configured for redundancy. This provides resiliency in case of a failure.

The first step is to establish connectivity between the blades and fabric interconnects. As shown in Figure 6, sixteen public (eight per chassis) links go to Fabric Interconnect "A" (ports 1 through 16). Similarly, sixteen private links go to Fabric Interconnect B. It is recommended to keep all private interconnects on a single Fabric interconnect. In such case, the private traffic will stay local to that fabric interconnect and will not go to northbound network switch. In other words, all inter blade (or RAC node private) communication will be resolved locally at the fabric interconnect.

Configuring Server Ports

Figure 6 Configuring Server Ports

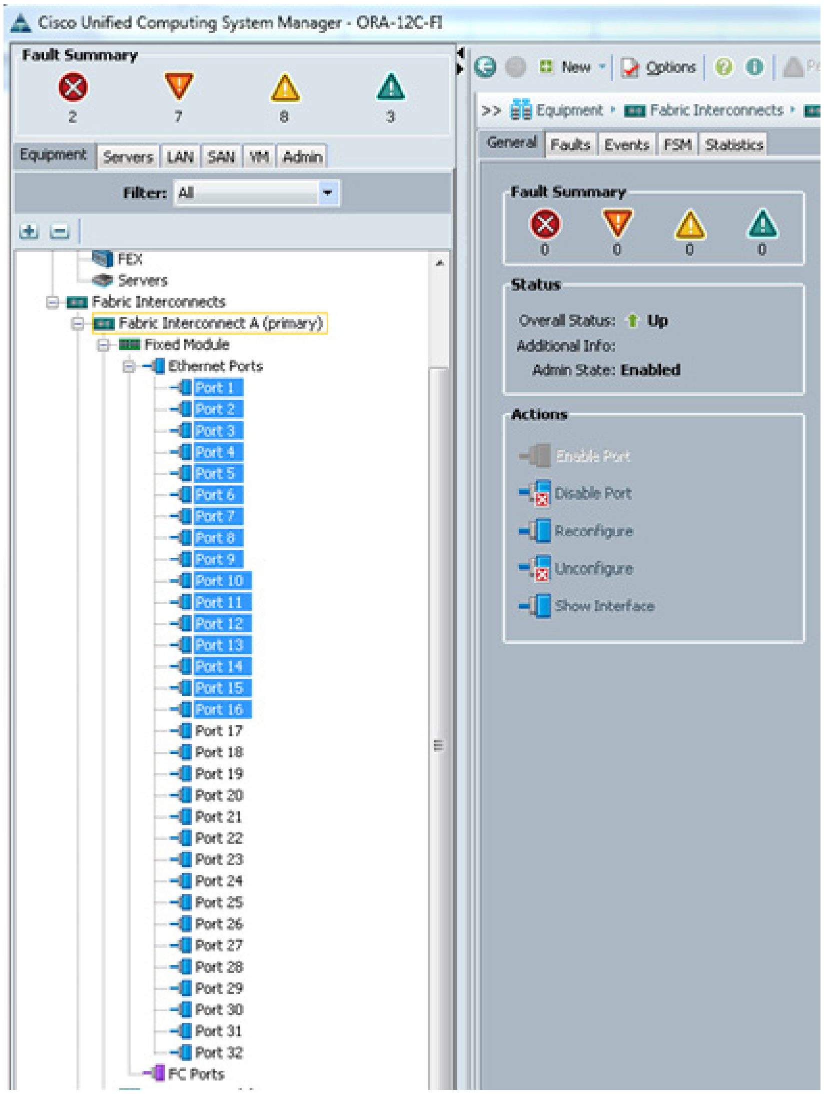

Configuring SAN and LAN with Cisco UCS Manager

Configuring SAN

On the SAN tab, create and configure the VSANs to be used for database as shown in Figure 7. On the test bed, we used vSAN 15 for the database.

Figure 7 Configuring SAN in Cisco UCS Manager

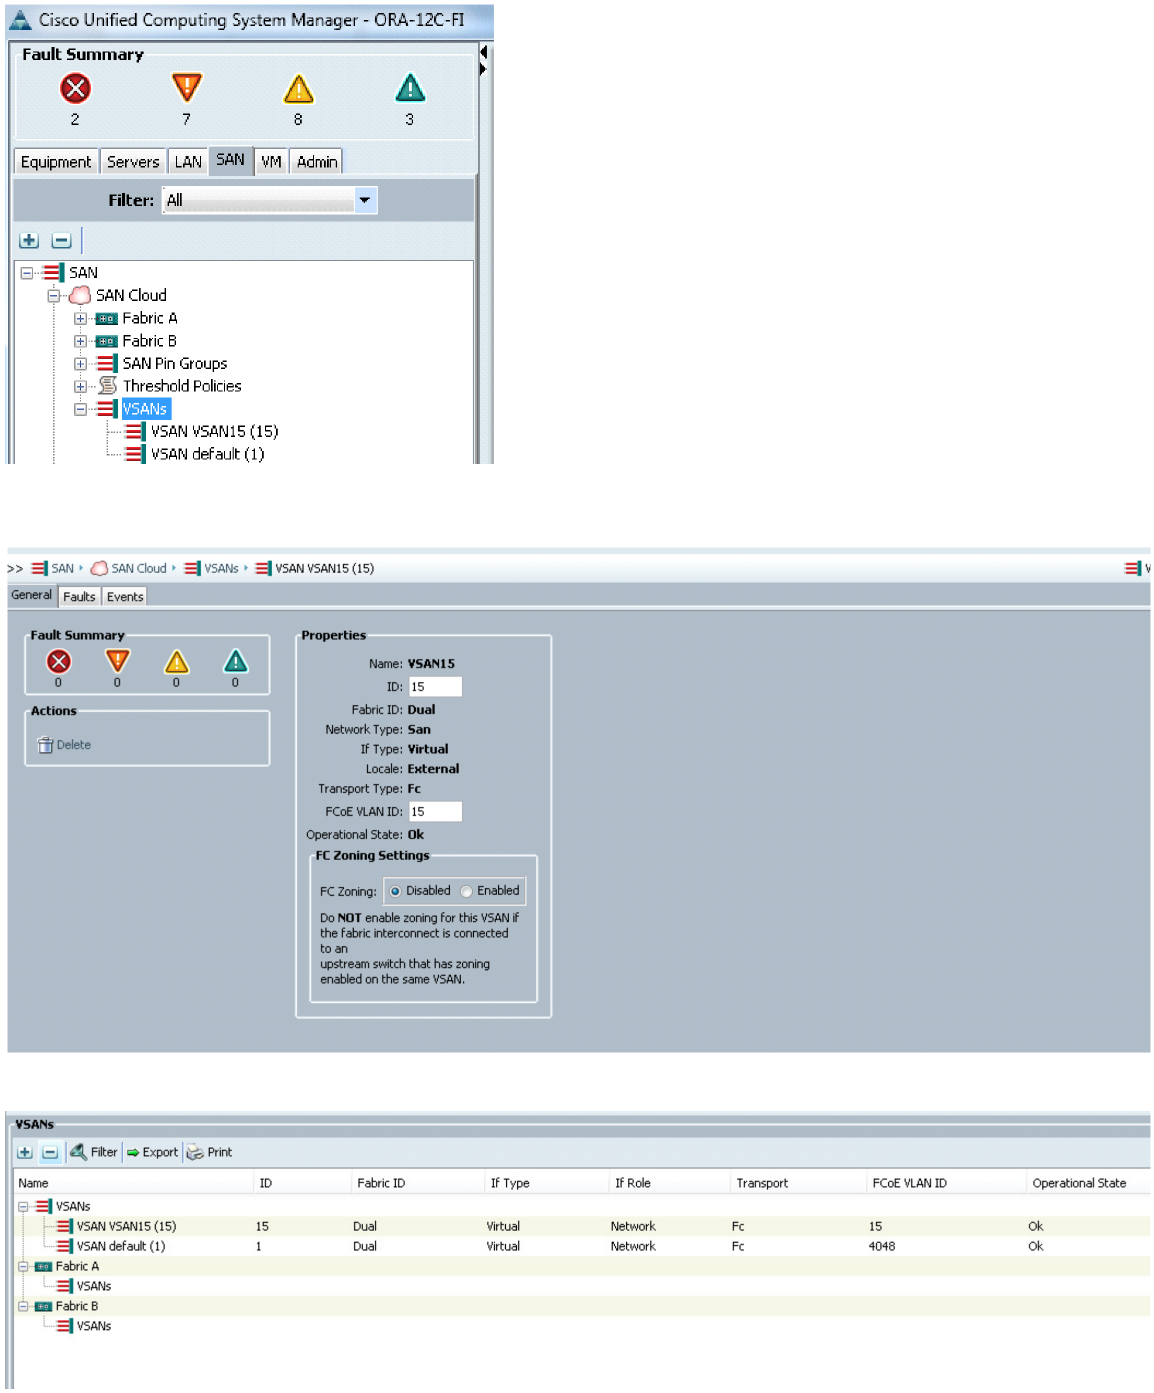

Configuring LAN

On the LAN tab as shown in Figure 8, create VLANs that will be used later for virtual NICs (vNICS) configured for private and public traffic for Oracle RAC. You can also set up MAC address pools for assignment to vNICS. For this setup, we used VLAN 134 for public interfaces and VLAN 10 for Oracle

RAC private interconnect interface. 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 failover.

Figure 8 Configuring LAN with Cisco UCS Manager

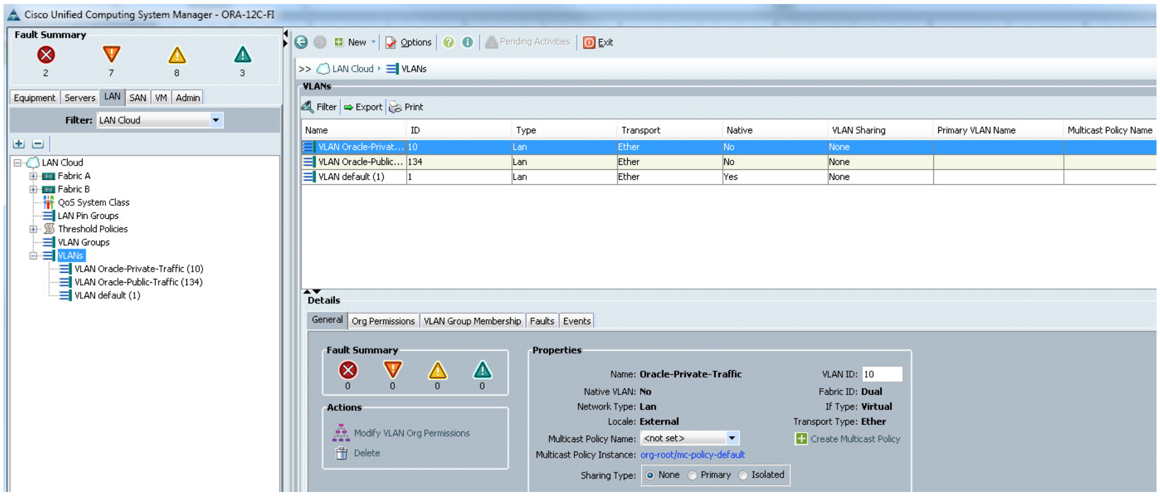

Even though private VLAN traffic stays local within Cisco UCS domain, it is necessary to configure entries for these private VLANS in northbound network switch. This will allow the switch to route interconnect traffic appropriately in case of partial link or IOM failures.

Configuring Jumbo Frames

Enable Jumbo Frames for Oracle Private Interconnect traffic.

Figure 9 Configuring Jumbo Frames

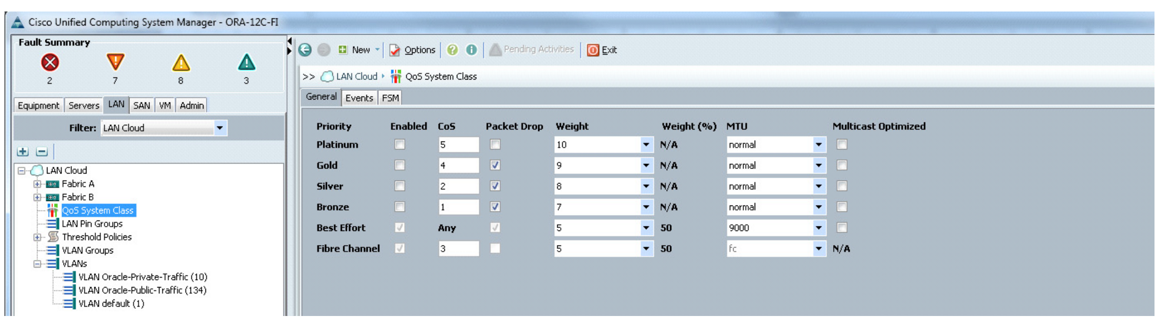

After these initial setups we can setup UCS service profile templates for the hardware configuration.

Configuring Ethernet Port-Channels

To configure Port-Channels, login to Cisco UCS Manager and from the LAN tab, filter on LAN cloud as shown in Figure 10.

Select Fabric A, right-click on port-channels and create port-channel. In the current Oracle RAC setup ports 17 and 18 on Fabric A were selected to be configured as port channel 10.

Similarly ports 17 and 18 on Fabric B were selected as port channel 11.

Figure 10 Configuring Port Channels

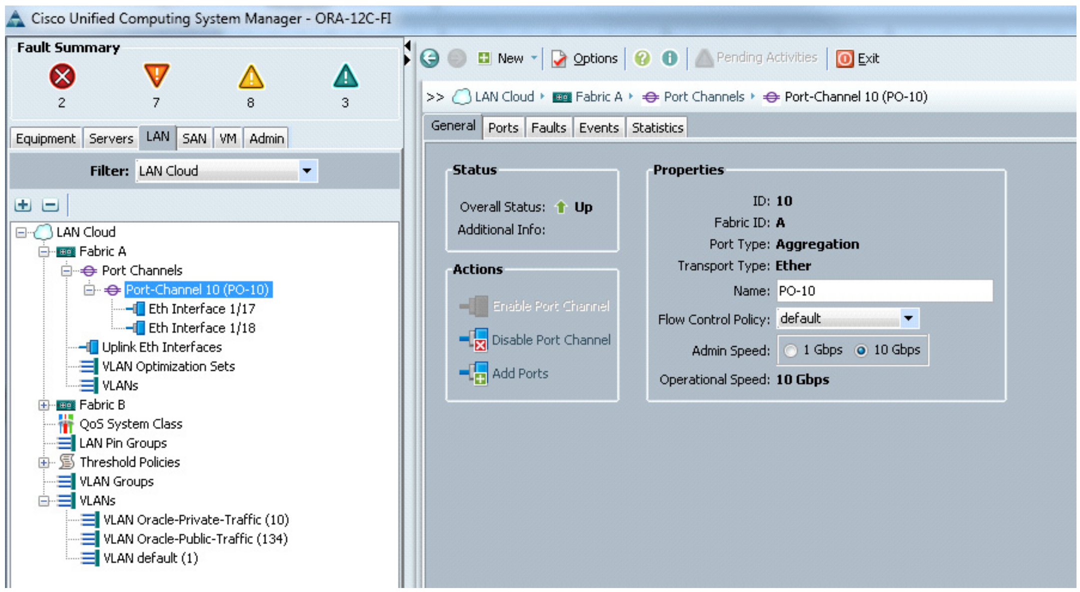

Port Channels on Fabric A

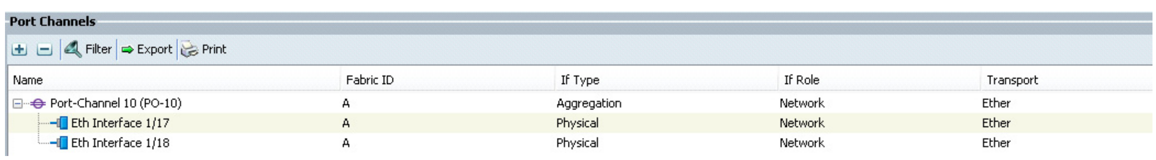

Port Channels

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

Port Channels on Fabric B

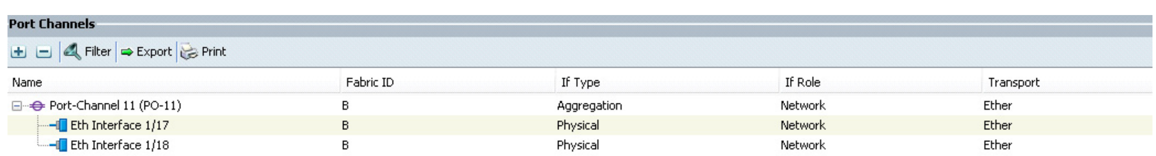

Port Channels

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

The next step is to set up VPC in n5k.

Preparatory Steps to Create Service Templates

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

UUID Pool

Click the servers tab, Filter on pools. Expand UUID suffix pools and create a new pool as shown in Figure 11.

Figure 11 Configuring UUID Pool in Cisco UCS

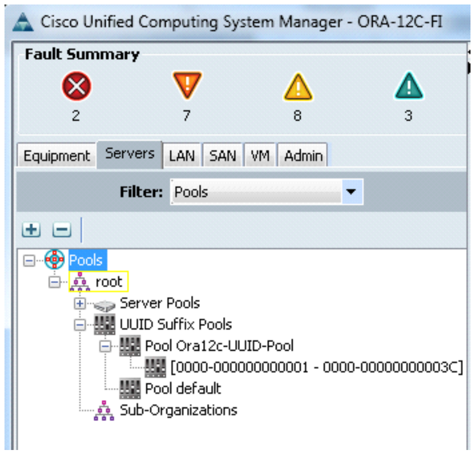

Create a default pool as shown below:

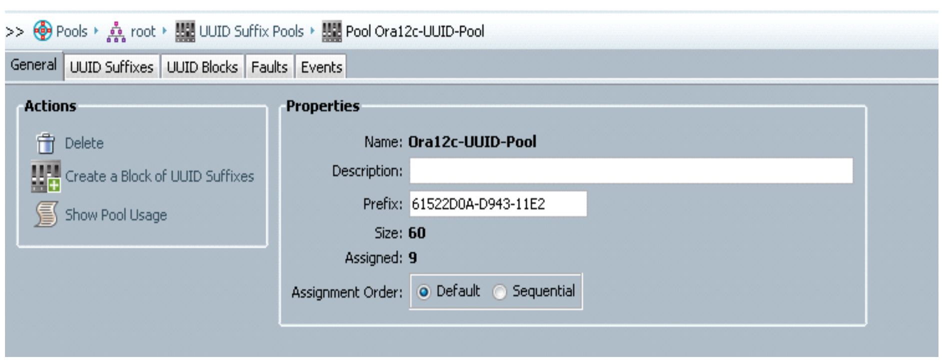

IP and MAC Pools

Click the LAN tab, filter on pools and create IP and MAC pools as shown below:

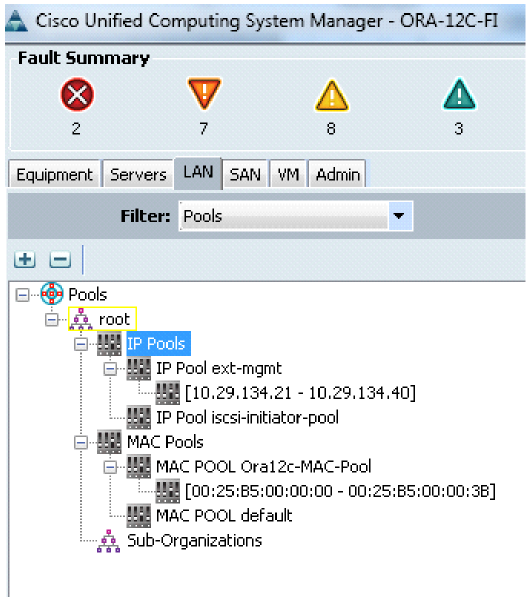

Note The IP pools will be used for console management, while MAC addresses for the vNICs will be addressed later in this document.

WWNN and WWPN pools

Click the SAN tab filter on pools and create the pools as shown below:

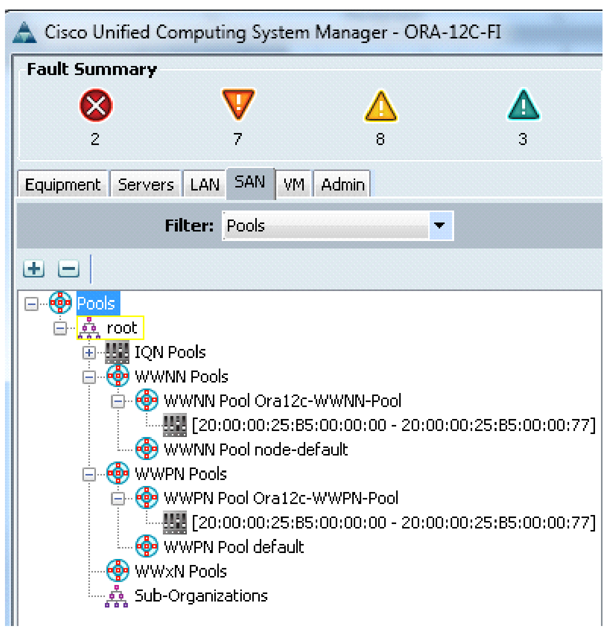

Configure VNIC Templates

Click the LAN tab, filter on policies and select vNIC templates as shown below. Two templates are created one for Public network and one for Private network. The Private network is for the internal Heart Beat and message transfers between Oracle Nodes while Public network for external clients like middle tiers and ssh sessions to the Oracle database hosts.

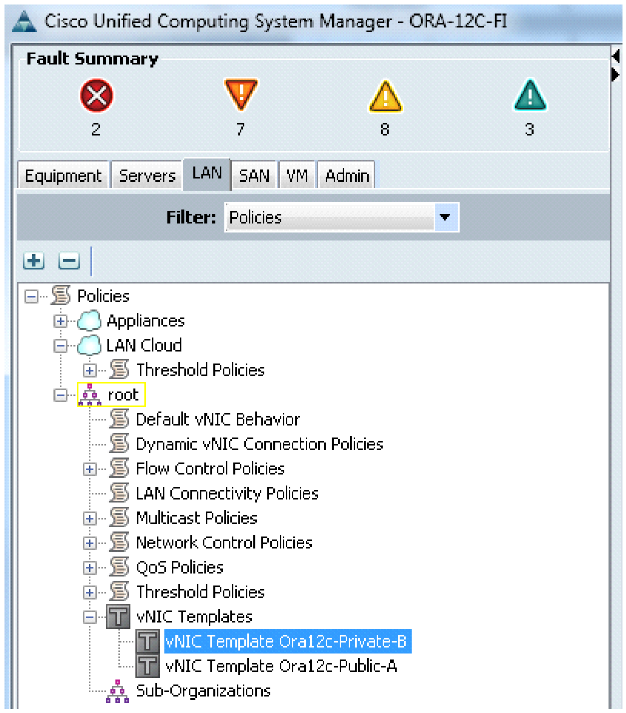

The vNIC template for the Oracle Private link is set at 9000 MTU and pinned to Fabric B and the failover is enabled. This allows the vNIC to failover to Fabric A, in case of failures of Fabric B.

Create a Private vNIC Template

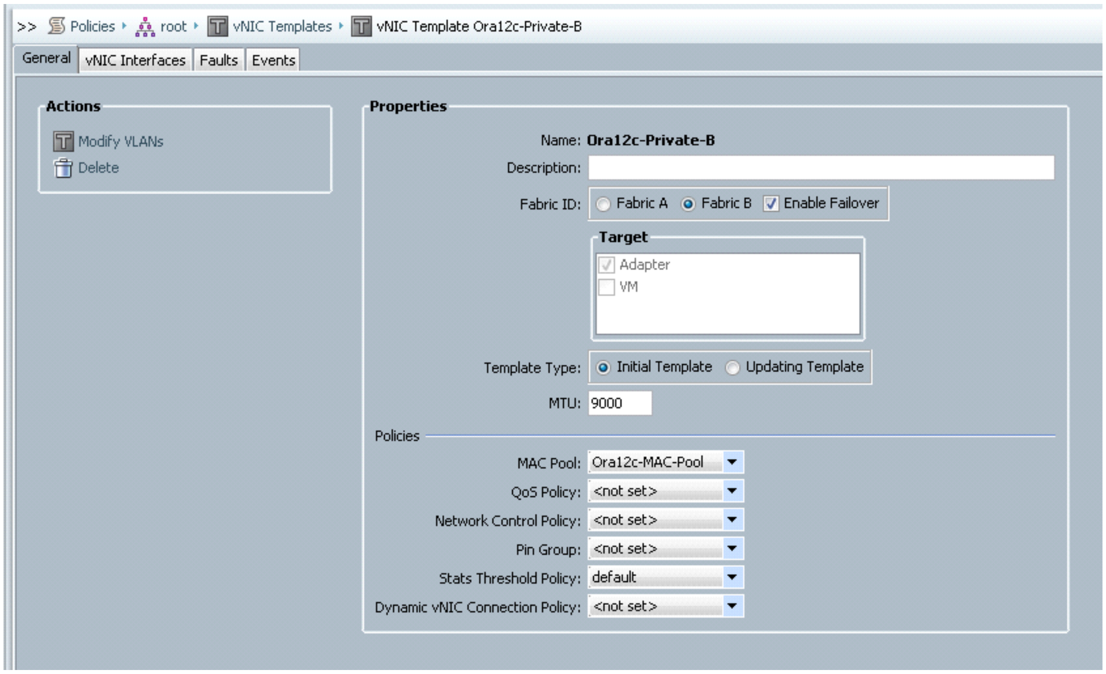

Create a Public vNIC Template

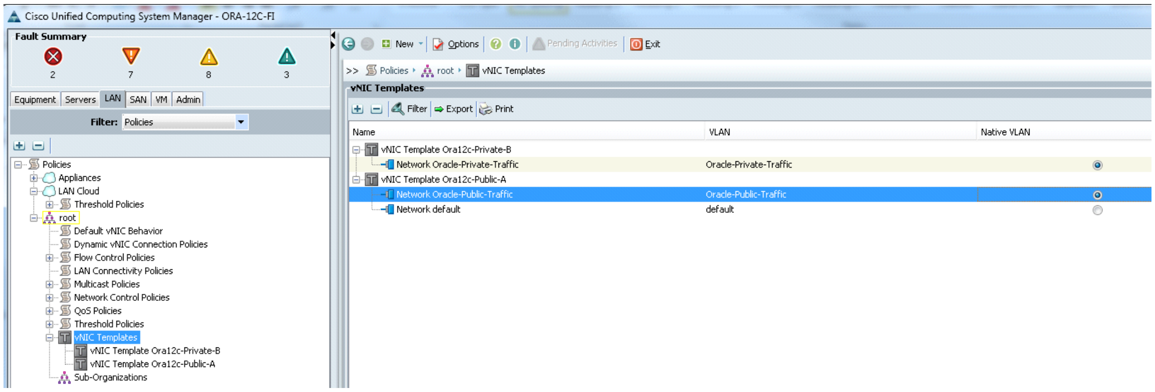

Create HBA templates

Click the SAN tab, filter out policies, right-click the vHBA templates and create a template as shown below:

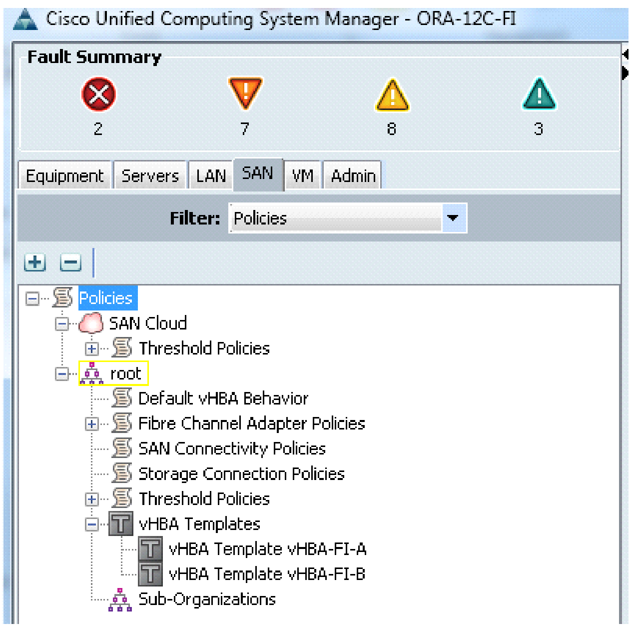

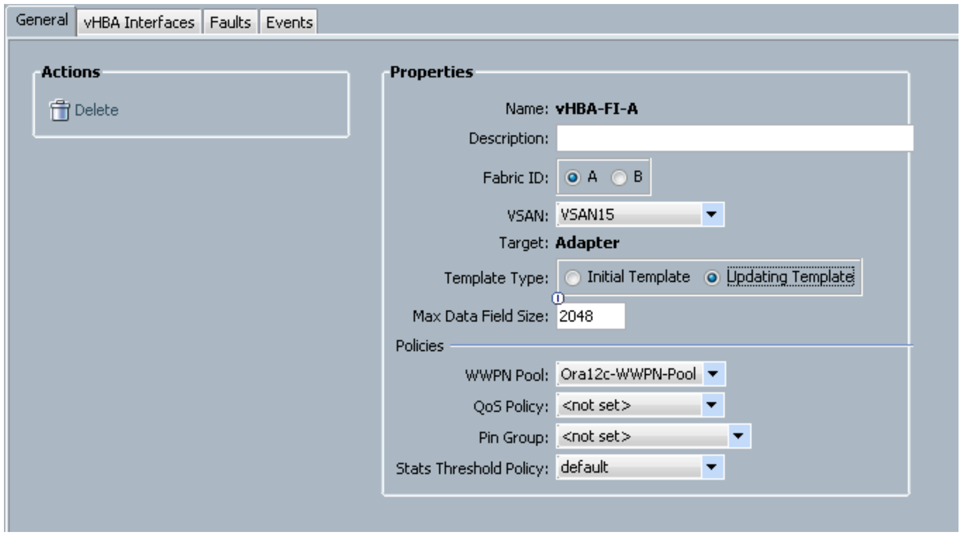

Create another vHBA_FI_B template as shown below:

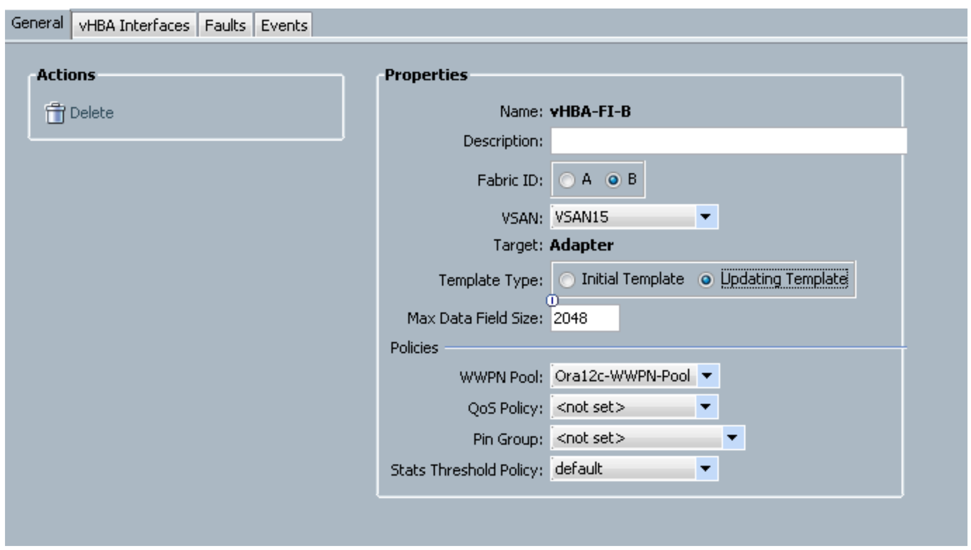

When the above preparatory steps are complete you can create a service template from which the service profiles can be easily created.

Create Service Profile Template

Create a service profile template as shown below before forking service profiles to be allocated to the servers later. Click the Servers tab in the Cisco UCS manager, filter out on Service Profile Templates as below and select Create Service Profile Template.

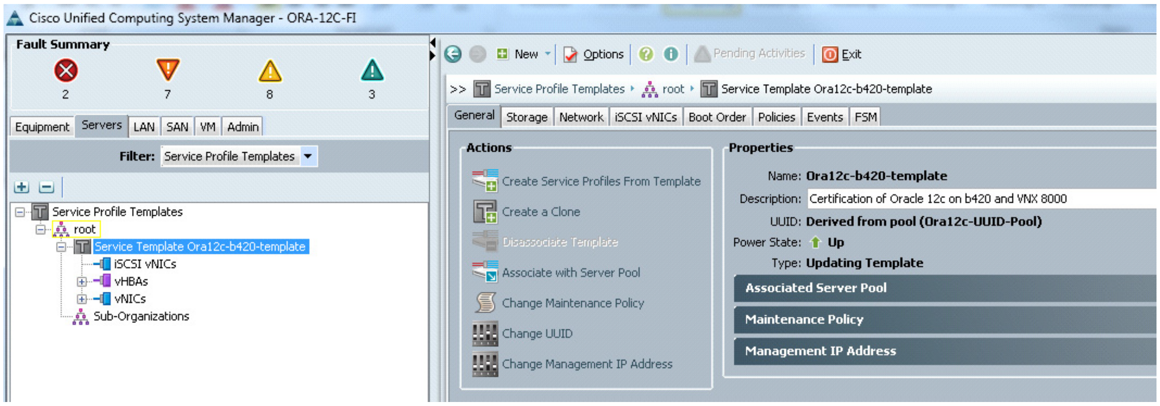

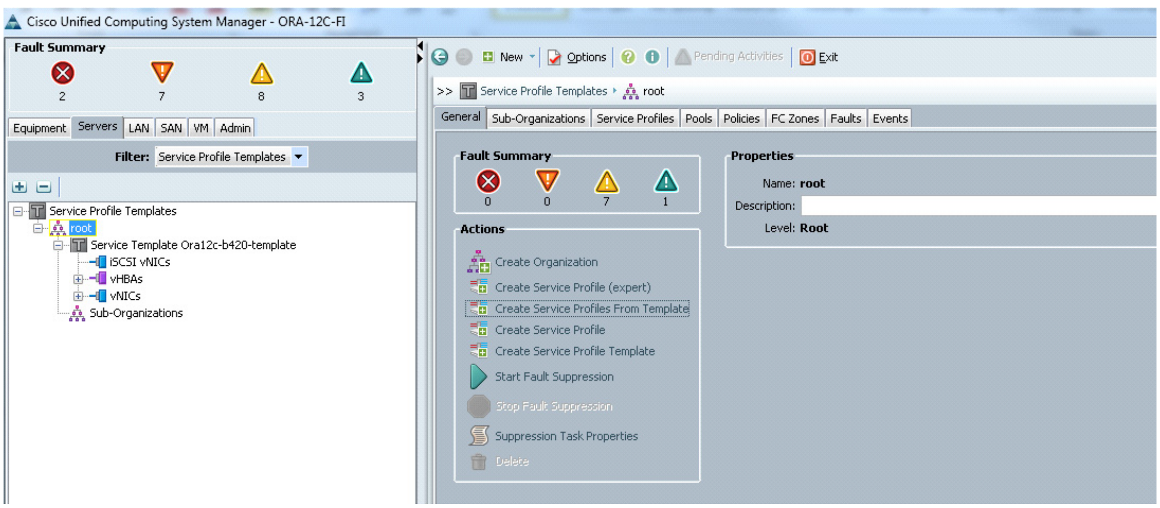

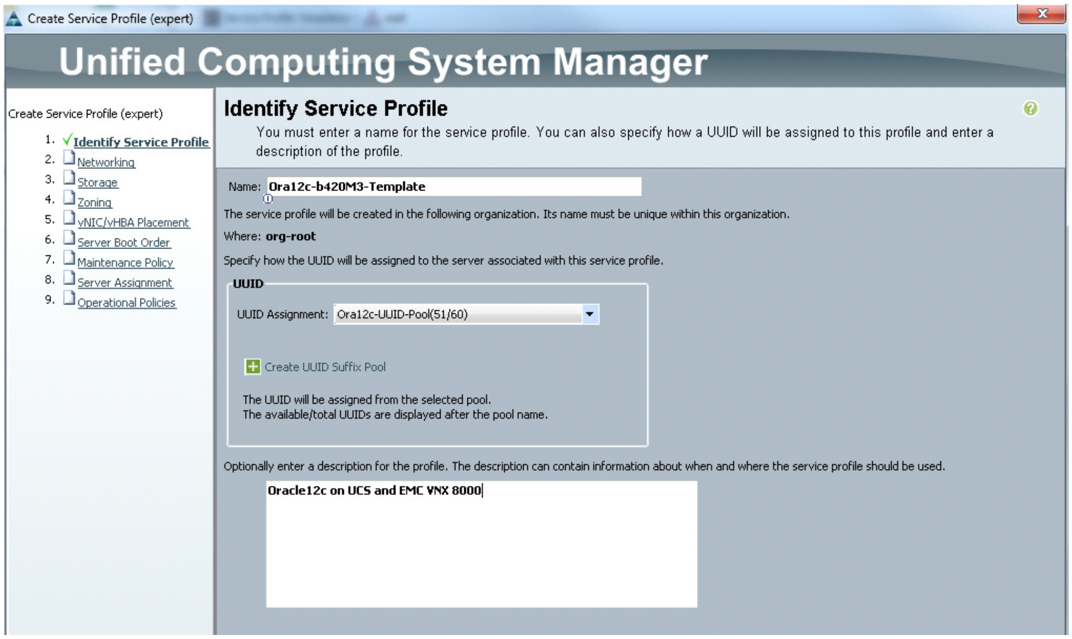

Enter a name, select the default UUID created earlier and move to the next screen.

In the Networking page, create vNICs; one on each fabric and associate them with the VLAN policies created earlier.

Select expert mode, click add in the section that specifies add one or more vNICs that the server should use to connect to the LAN.

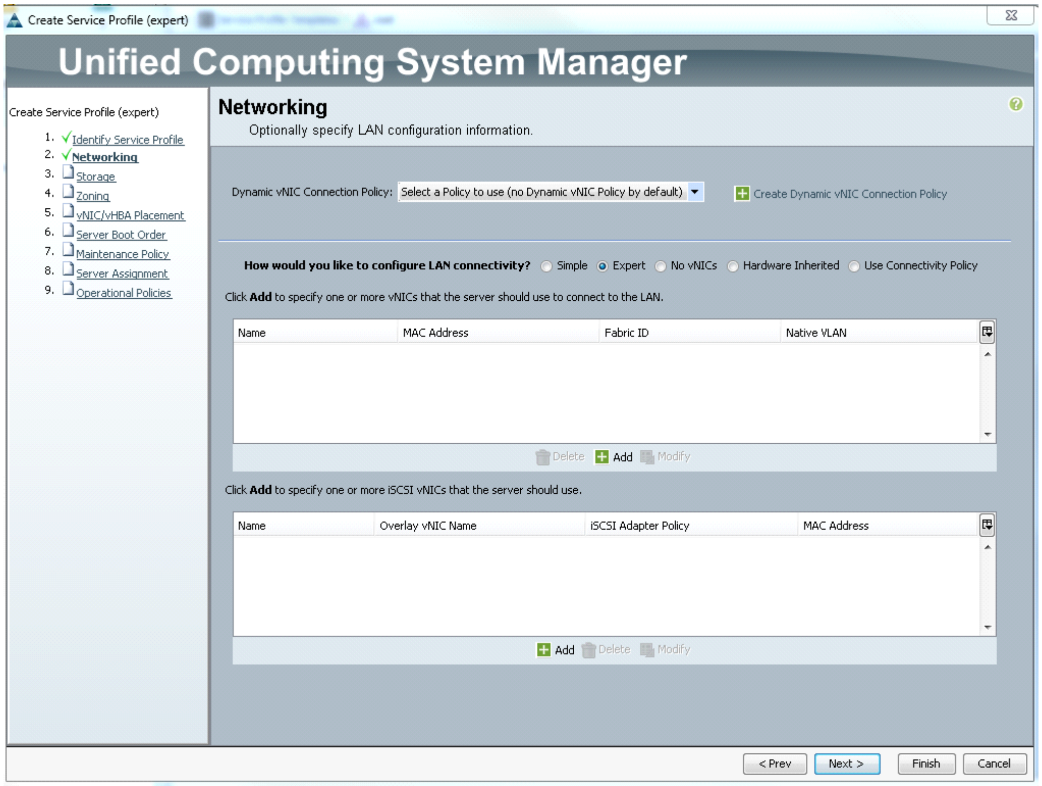

From the Create vNIC page, select "Use vNIC template" and adapter policy as Linux. vNIC1 was selected for Oracle private network as shown below:

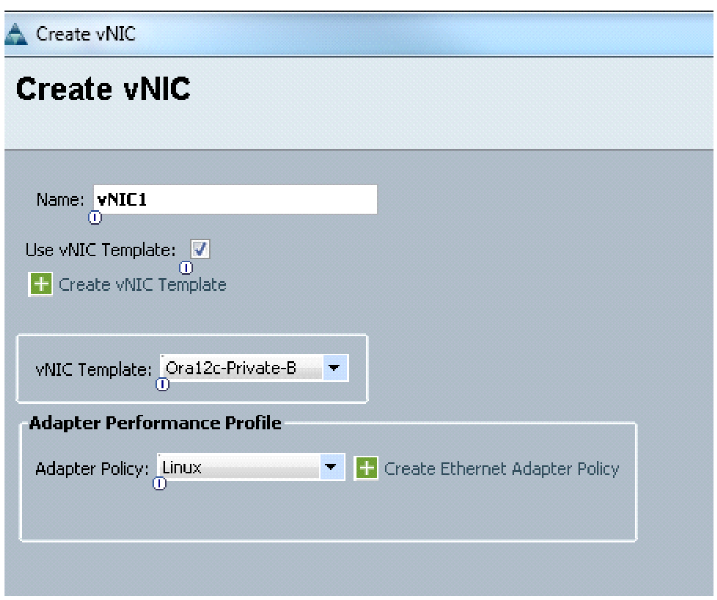

Create a vNIC2 for public.

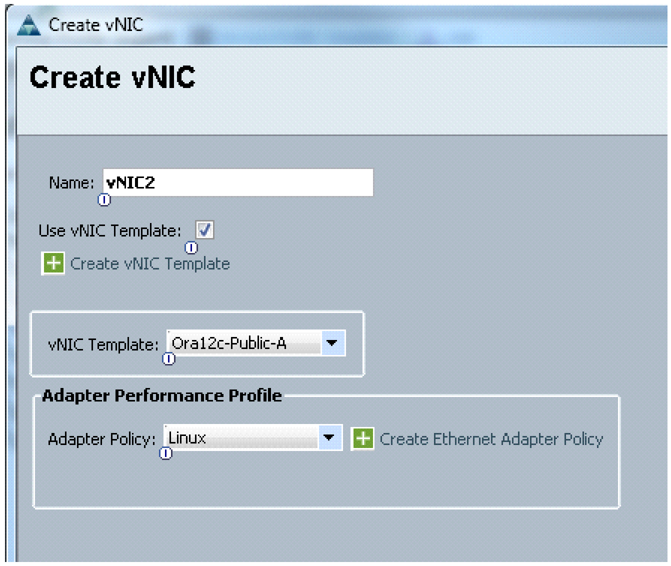

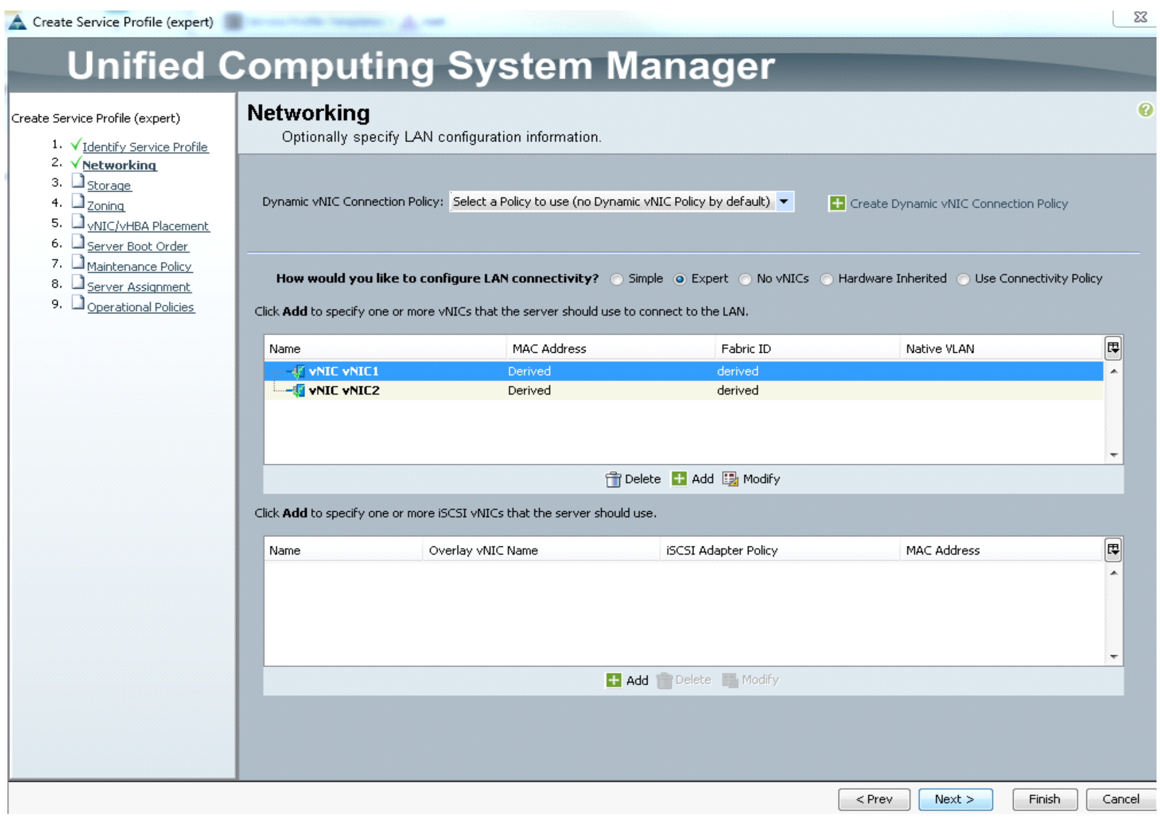

From the Storage page, similar to vNICs select expert mode in adapter, choose the WWNN pool created earlier and click Add to create vHBAs. We selected 4xvHBA's as shown below:

Create vHBA1 using template vHBA_FI_A.

Create vHBA2 using template vHBA_FI_B.

Create vHBA3 using template vHBA_FI_A.

Create vHBA4 using template vHBA_FI_B.

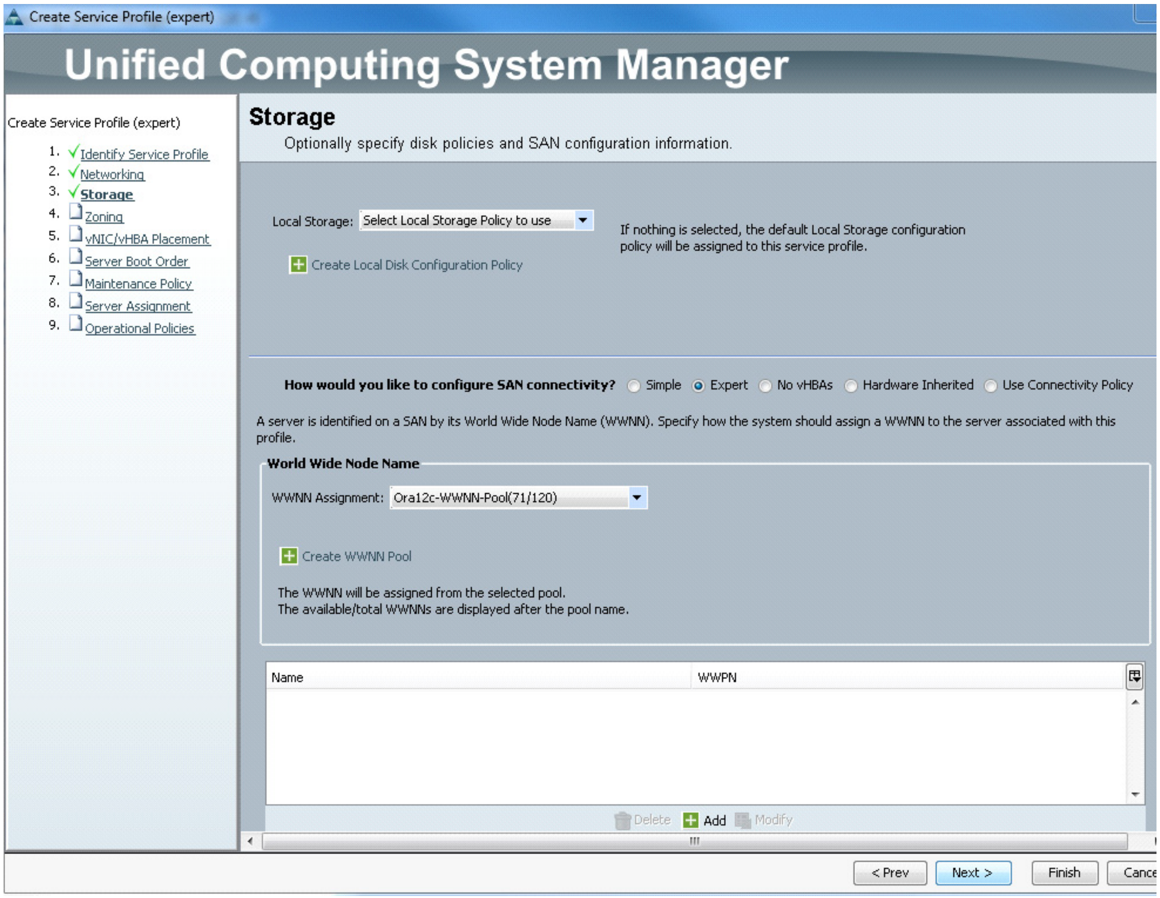

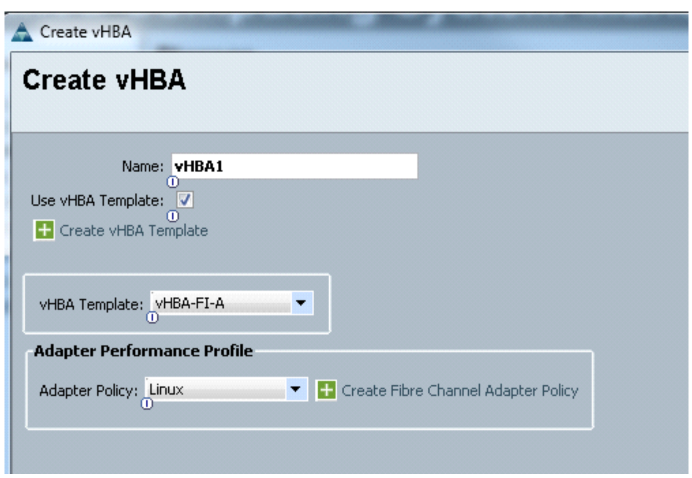

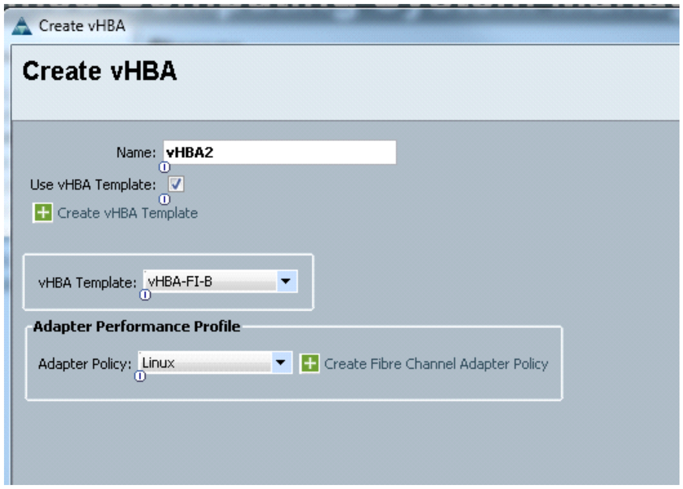

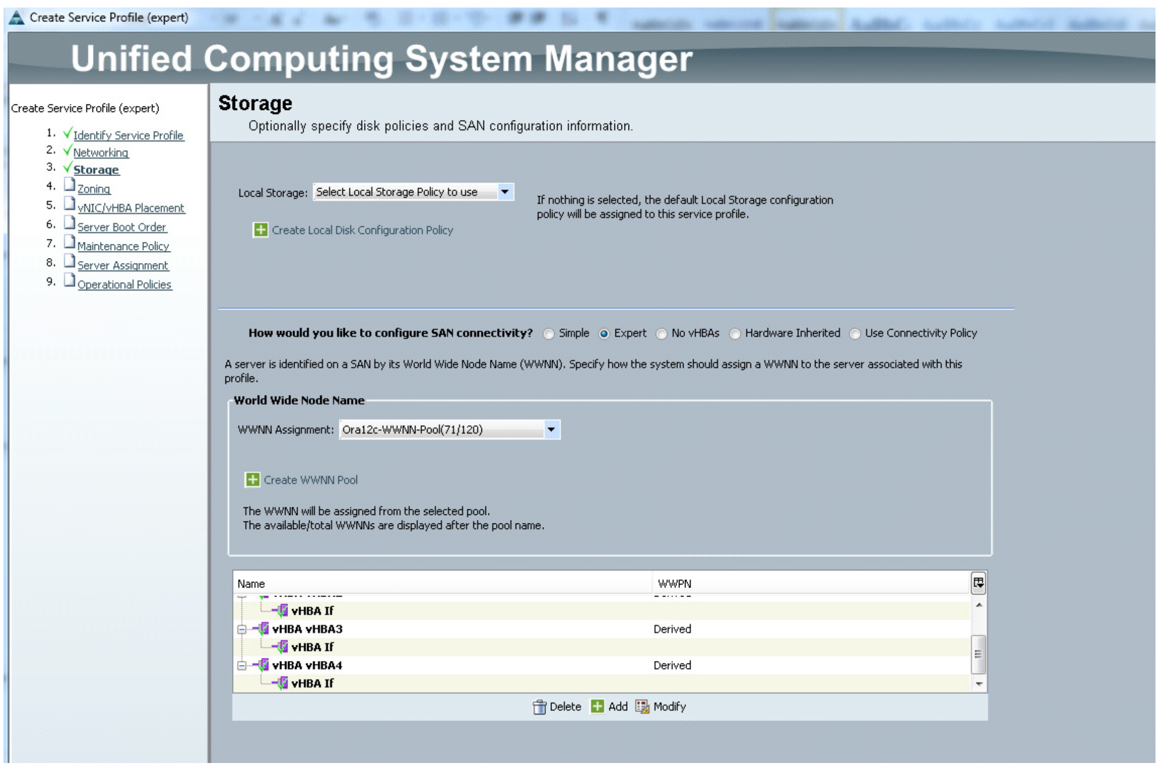

Skip the zoning section and go to vNIC/vHBA placement.

You can retain the system defaults or you can specify them manually as shown below:

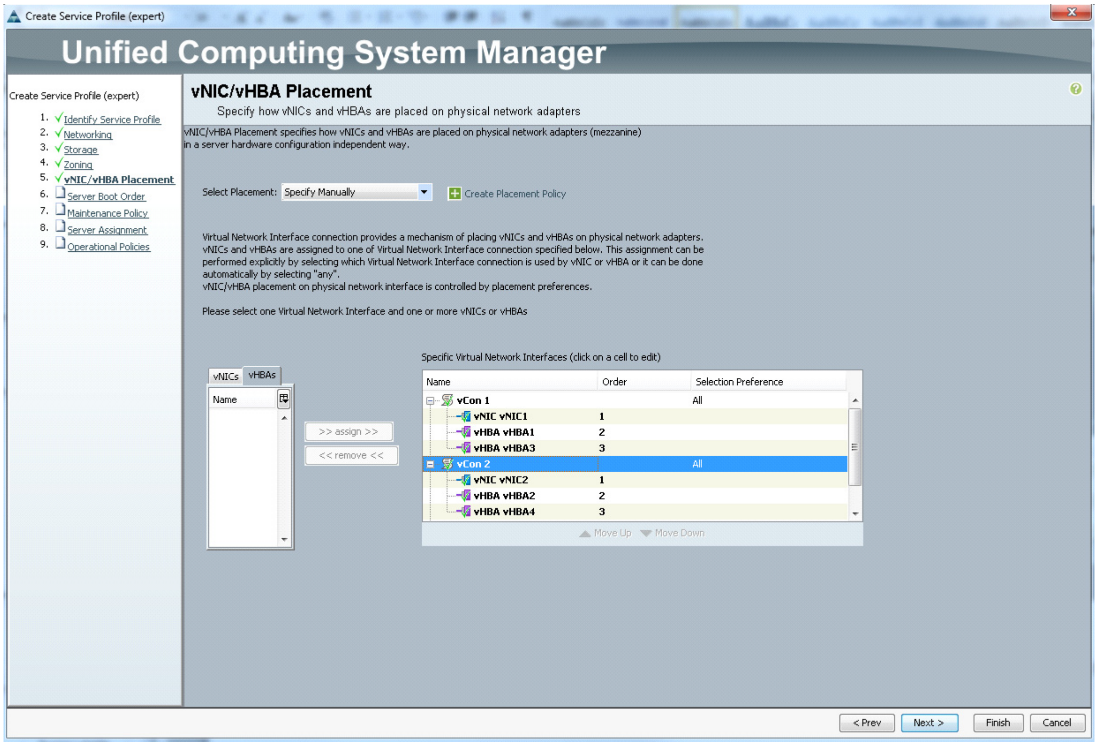

Allocate vNIC1, vHBA1 and vHBA3 to the first vic1240, while the rest of vNIC2, vHBA2 and vHBA4 to the second vNIC 1280

Server Boot Policy:

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

The rest of the maintenance and assignment policies were left to default in the test bed. Other policies may be selected and will vary from site to site, depending on your workloads, best practices and policies.

Create Service Profiles from Service Profile Templates

Click the Servers tab, right-click on the root and select Create Service Profile from Template as shown below:

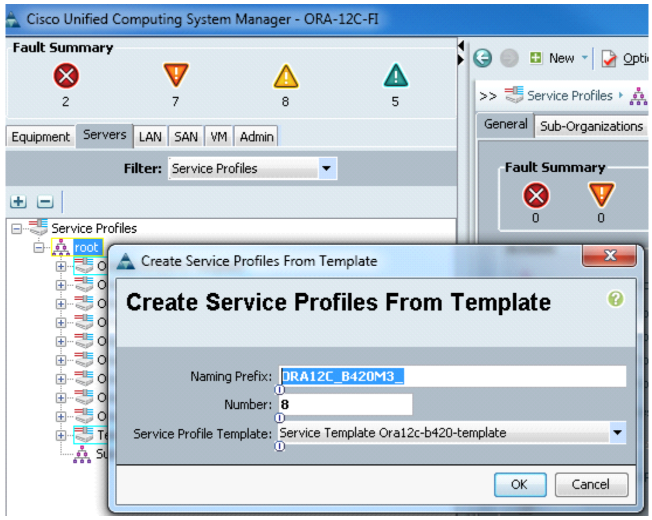

This creates 8 service profiles with the following names:

ORA12C_B420M3_1, ORA12C_B420M3_2, ORA12C_B420M3_3, ORA12C_B420M3_4, ORA12C_B420M3_5, ORA12C_B420M3_6, ORA12C_B420M3_7, ORA12C_B420M3_8,

Associating Service Profile to the Servers

Make sure that a few of the entries in the service profile appear as expected and as shown below before associating them to a server.

In order to associate this service profile to a server, perform the following steps.

From the Servers tab, select the desired service profile and select change service profile association.

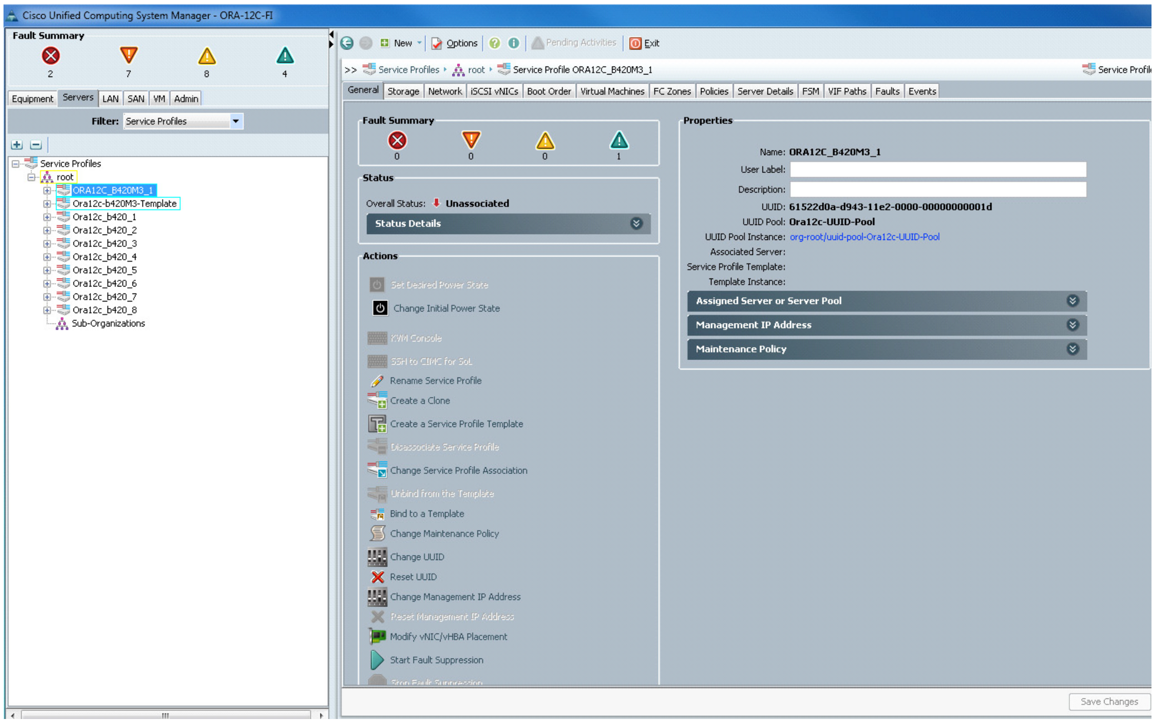

The service profile is unassociated as of now and can be assigned to a server in the pool.

Click Change Service Profile Association under the General tab, from the drop-down of Server Assignment, select an existing server that you would like to assign, and click OK.

Setting Up EMC VNX Storage

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

The following are some generic recommendations for EMC VNX storage configuration with mixed drives.

Turn off the read and write caches for flash drive-based LUNs. In most situations, it is better to turn off both the read and write caches on all the LUNs that reside on flash drives, for the following reasons:

- The flash drives are extremely fast: When the read cache is enabled for the LUNs residing on them, the read cache lookup for each read request adds more overhead compared to SAS drives. This scenario occurs in an application profile that is not expected to get many read cache hits at any rate. It is generally much faster to directly read the block from the flash drives.
- Typically, the storage array is also shared by several other applications along with the database. In some situations, the write cache may become fully saturated, placing the flash drives in a force-flush situation. This adds unnecessary latency. This typically occurs particularly when storage deploys mixed drives and consists of slower Near Line SAS drives. Therefore, it is better in these situations to write the block directly to the flash drives than to the write cache of the storage system.

Distribute database files for flash drives. Refer to Table 2 below for recommendations about distributing database files based on the type of workload.

Four databases were created.

- OLTP - Online Transaction system completely on flash drives.
- DSS - Decision support system on mixed pools.
- CPU - Database to stress the CPU on SAS drives
- CRS - CRS database on SAS drives.

Table 2 illustrates the distribution of LUNs from a VNX8000 for the setup.

Storage Configuration

Table 2 Storage Configuration

Purpose	OLTP Database data and temp files	DSS Database data and temp files	CPU and CRS Database data and temp files	Redo Log Files for all the databases
Disk Type	Flash	Mixed (SAS and Flash)	SAS	SAS
RAID Type	RAID 5 Storage Pool	RAID 5 Storage Pool	RAID 5 Storage Pool	RAID 1/0 Storage Pool
SAS Disks	0	150	20	32
Flash Disks	50	30	0	0
Total LUNs	16	16	16	16
LUN Size	400GB	1000GB	250GB	300GB

Purpose	Boot LUNs and Oracle RAC
Disk Type	SAS
RAID Type	RAID 5
SAS Disks	5
Flash Disks	0
Total LUNs	8 Boot LUNs
LUN Size	Boot LUNs - 50GB

Purpose	Oracle RAC code Tree and OCR/Voting LUNs
Disk Type	SAS
RAID Type	RAID 5
SAS Disks	5
Flash Disks	0
Total LUNs	8 RAC LUNs and 5 OCR/Voting LUNs
LUN Size	RAC LUNs - 60GB OCR/Voting LUNs – 20GB

Purpose	ACFS File system for RDBMS Code Tree

Disk Type	SAS
RAID Type	RAID 5
SAS Disks	3
Flash Disks	0
Total LUNs	3 ACFS LUNs
LUN Size	ACFS LUNs - 175GB

Hardware Storage Processors Configuration

A total of 32 ports were used from storage processors and were equally distributed between SPA and SPB as shown in Table 3 and were connected to the respective N5K's.

Table 3 Service Processor Distribution

Processor	Slot/Port	WWPN
SPA	A0P0	50:06:01:60:36:60:05:e7
	A0P1	50:06:01:61:36:60:05:e7
	A0P2	50:06:01:62:36:60:05:e7
	A0P3	50:06:01:63:36:60:05:e7
	A1P0	50:06:01:64:36:60:05:e7
	A1P1	50:06:01:65:36:60:05:e7
	A1P2	50:06:01:66:36:60:05:e7
	A1P3	50:06:01:67:36:60:05:e7
	A7P0	50:06:01:60:36:64:05:e7
	A7P1	50:06:01:61:36:64:05:e7
	A7P2	50:06:01:62:36:64:05:e7
	A7P3	50:06:01:63:36:64:05:e7
	A8P0	50:06:01:64:36:64:05:e7
	A8P1	50:06:01:65:36:64:05:e7
	A8P2	50:06:01:66:36:64:05:e7
	A8P3	50:06:01:67:36:64:05:e7
SPB	B0P0	50:06:01:68:36:60:05:e7
	B0P1	50:06:01:69:36:60:05:e7
	B0P2	50:06:01:6a:36:60:05:e7
	B0P3	50:06:01:6b:36:60:05:e7
	B1P0	50:06:01:6c:36:60:05:e7
	B1P1	50:06:01:6d:36:60:05:e7
	B1P2	50:06:01:6e:36:60:05:e7
	B1P3	50:06:01:6f:36:60:05:e7
	B7P0	50:06:01:68:36:64:05:e7
	B7P1	50:06:01:69:36:64:05:e7
	B7P2	50:06:01:6a:36:64:05:e7
	B7P3	50:06:01:6b:36:64:05:e7
	B8P0	50:06:01:6c:36:64:05:e7
	B8P1	50:06:01:6d:36:64:05:e7
	B8P2	50:06:01:6e:36:64:05:e7
	B8P3	50:06:01:6f:36:64:05:e7

In the sections of N5K zoning, we will cover how these WWPNs will be used in zoning, boot policies and in achieving high availability in case of failures.

Configuring SAN Zoning on N5K 5548 UP Switches

Note Two numbers of N5K 5548 UP switches were configured.

Figure 12 *N5K Configurations with vPC*

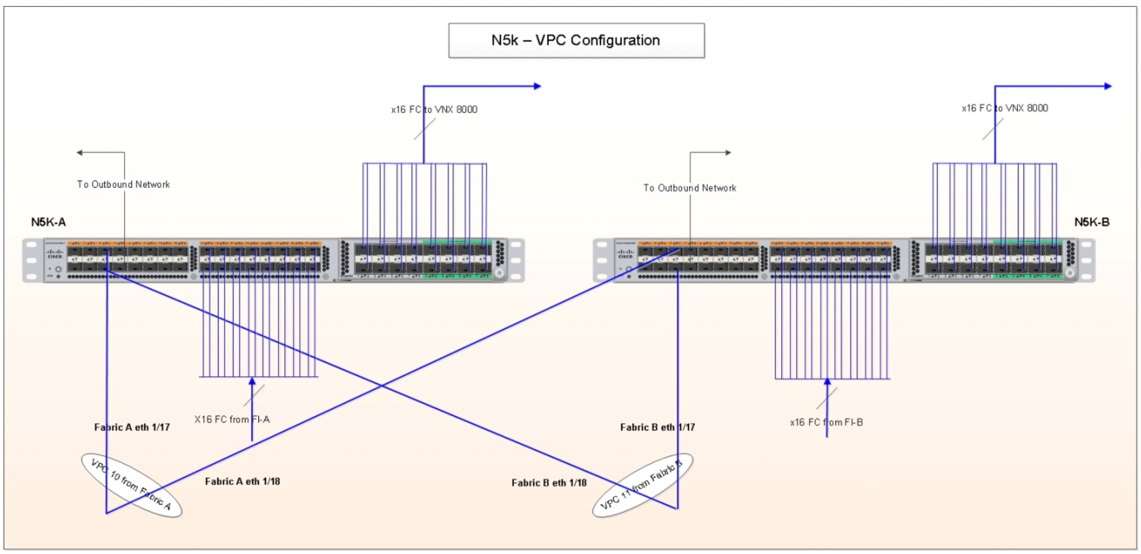

Figure 12 illustrates how the N5K UP switches are connected to North bound switches and storage, while connected to the underlying Cisco UCS Fabrics. The N5K switches form a core group in controlling SAN zoning.

Fibre Channel Zoning

Prior to configuring zoning details, decide how many paths are needed for each LUN and extract the WWPN numbers for each of the HBAs.

To access the WWPNs for each of the HBAs, login to the Cisco UCS manager.

Click Equipment, chassis, servers and select the desired server. On the right-hand menu, click the Inventory tab and HBAs, as shown below:

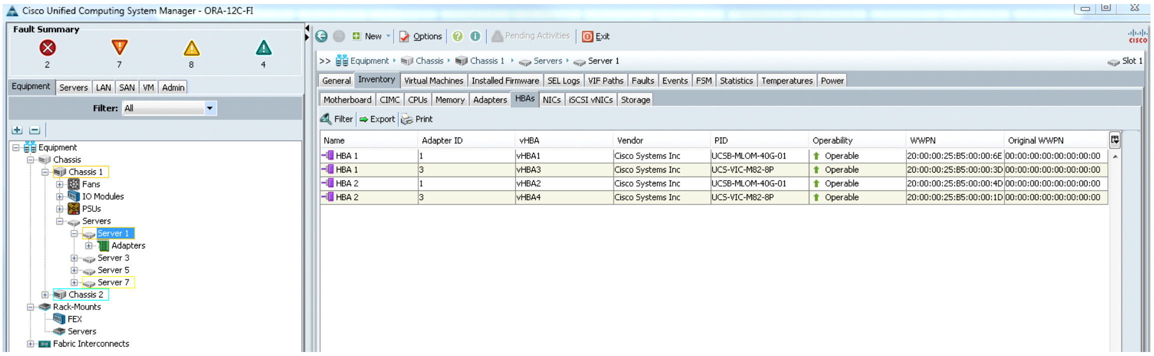

The WWPN numbers for all the 4 HBAs for server 1 as an example is illustrated above. In the current setup, it was decided to have a total of 4 paths, 2 paths from each Fabrics and N5K's to the storage.

The zoning for Server1, HBA1 can be setup as shown below:

zone name ora12c_b420_1_hba1 vsan 15

* fcid 0x7413ef [device-alias A0P0] <-SPA

* fcid 0x7405ef [device-alias B0P0] <- SPB

* fcid 0x740001 [pwwn 20:00:00:25:b5:00:00:6e] <- Extracted from the above figure for HBA1.

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

The following is an example for server 1:

N5K-A		
	zone ora12c_b420_1_hba1	[pwwn 20:00:00:25:b5:00:00:6e]
		[pwwn 50:06:01:60:36:60:05:e7] [A0P0]
		[pwwn 50:06:01:68:36:60:05:e7] [B0P0]
	zone ora12c_b420_1_hba3	[pwwn 20:00:00:25:b5:00:00:3d]
		[pwwn 50:06:01:61:36:60:05:e7] [A0P1]
		[pwwn 50:06:01:69:36:60:05:e7] [B0P1]
N5K-B		
	zone ora12c_b420_1_hba2	[pwwn 20:00:00:25:b5:00:00:4d]
		[pwwn 50:06:01:60:36:64:05:e7] [A7P0]
		[pwwn 50:06:01:68:36:64:05:e7] [B7P0]
	zone ora12c_b420_1_hba4	[pwwn 20:00:00:25:b5:00:00:1d]
		[pwwn 50:06:01:61:36:64:05:e7] [A7P1]
		[pwwn 50:06:01:69:36:64:05:e7] [B7P1]

Log in to each N5K through ssh and issue the following:

Setup VLAN and VSANs on both N5Ks

```
conf term
vlan 134
  name Oracle_RAC_Public_Traffic
exit
vlan 10
  name Oracle_RAC_Private_Traffic
  no ip igmp snooping
exit
vsan database
vsan 15
exit
```

Setting Up Device Aliases for Storage Initiators

```
N5K-A
device-alias database
device-alias name A0P0 pwwn 50:06:01:60:36:60:05:e7
device-alias name A0P1 pwwn 50:06:01:61:36:60:05:e7
device-alias name A0P2 pwwn 50:06:01:62:36:60:05:e7
device-alias name A0P3 pwwn 50:06:01:63:36:60:05:e7
device-alias name A1P0 pwwn 50:06:01:64:36:60:05:e7
device-alias name A1P1 pwwn 50:06:01:65:36:60:05:e7
device-alias name A1P2 pwwn 50:06:01:66:36:60:05:e7
device-alias name A1P3 pwwn 50:06:01:67:36:60:05:e7
device-alias name B0P0 pwwn 50:06:01:68:36:60:05:e7
device-alias name B0P1 pwwn 50:06:01:69:36:60:05:e7
device-alias name B0P2 pwwn 50:06:01:6a:36:60:05:e7
device-alias name B0P3 pwwn 50:06:01:6b:36:60:05:e7
device-alias name B1P0 pwwn 50:06:01:6c:36:60:05:e7
device-alias name B1P1 pwwn 50:06:01:6d:36:60:05:e7
device-alias name B1P2 pwwn 50:06:01:6e:36:60:05:e7
device-alias name B1P3 pwwn 50:06:01:6f:36:60:05:e7
exit
device-alias commit
exit


N5K-B
device-alias database
device-alias name A7P0 pwwn 50:06:01:60:36:64:05:e7
device-alias name A7P1 pwwn 50:06:01:61:36:64:05:e7
device-alias name A7P2 pwwn 50:06:01:62:36:64:05:e7
device-alias name A7P3 pwwn 50:06:01:63:36:64:05:e7
device-alias name A8P0 pwwn 50:06:01:64:36:64:05:e7
device-alias name A8P1 pwwn 50:06:01:65:36:64:05:e7
device-alias name A8P2 pwwn 50:06:01:66:36:64:05:e7
device-alias name A8P3 pwwn 50:06:01:67:36:64:05:e7
device-alias name B7P0 pwwn 50:06:01:68:36:64:05:e7
device-alias name B7P1 pwwn 50:06:01:69:36:64:05:e7
device-alias name B7P2 pwwn 50:06:01:6a:36:64:05:e7
device-alias name B7P3 pwwn 50:06:01:6b:36:64:05:e7
device-alias name B8P0 pwwn 50:06:01:6c:36:64:05:e7
device-alias name B8P1 pwwn 50:06:01:6d:36:64:05:e7
device-alias name B8P2 pwwn 50:06:01:6e:36:64:05:e7
device-alias name B8P3 pwwn 50:06:01:6f:36:64:05:e7
exit
```

```
device-alias commit
exit
```

Setting Up Zones

The following is an example for one zone on one N5K:

```
conf term
zoneset name ora12c_FI_A vsan 15
zone name orarac1_hba1
member device-alias A2P0
member device-alias A2P2
member device-alias B2P0
member device-alias B2P0
member pwwn 20:00:00:25:b5:00:00:1f ( ora12crac1 hba1 wwpn )
exit
exit
zoneset activate name ora12c_FI_A vsan 15
copy running-config startup-config
```

Optionally, configure the device aliases for simpler maintenance as above, or use WWPN for the storage initiators.

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

Setting Up VPC on N5Ks

As shown in Figure 12, both N5K's port 17 receives traffic from UCS Fabric A, that has port-channel 10 defined. Similarly both N5K's port 18 receives traffic from UCS Fabric B, that has port-channel 11 configured.

Login into N5K-A as admin.

```
conf term
feature vpc

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

interface port-channel1
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type network
  speed 10000
  vpc peer-link

interface port-channel10
  description oracle 12c port-channel
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  vpc 10

interface port-channel11
```

```
  description oracle 12c port-channel
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  vpc 11

interface Ethernet1/5
  description connected to UCS-Fab-A-Port17
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  channel-group 10 mode active

interface Ethernet1/6
  description connected to UCS-Fab-B-Port18
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  channel-group 11 mode active

interface Ethernet1/8
  description "Public Traffic to 3750"
  switchport mode trunk
  switchport access vlan 134
  switchport trunk native vlan 134
  switchport trunk allowed vlan 1,134
  speed 1000

copy running-config startup-config
```

Repeat the above on both N5K's.

Show vpc status should show the following for a successful configuration:

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

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active
vlans
------ ----------- ------ ----------- -------------------------- -----------
10     Po10        up     success     success                    1,10,134
11     Po11        up     success     success                    1,10,134

show interface port-channel 10-11 brief

----------------------------------------------------------------------------
----
Port-channel VLAN    Type Mode   Status  Reason                    Speed
Protocol
Interface
----------------------------------------------------------------------------
----
Po10         1       eth  trunk  up      none                      a-10G(D)
lacp
```

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

Setting Up Jumbo Frames on N5K

Jumbo frames with an mtu=9000 have to be setup on n5k. Oracle Interconnect traffic under normal conditions does not go to the northbound switch like N5K's as all the private vNICs are configured in Fabric B. However if there is a partial link or IOM failure, the private interconnect traffic has to go to the immediate northbound switch (N5K in our case) to reach Fabric B.

The command shown below details how to configure Jumbo frames Nexus 5K Fabric A Switch:

```
ora12c-n5k-a# conf terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ora12c-n5k-a (config)# class-map type network-qos class-platinum
ora12c-n5k-a (config-cmap-nq)# exit
ora12c-n5k-a (config)# policy-map type network-qos jumbo
ora12c-n5k-a (config-pmap-nq)# class type network-qos class-default
ora12c-n5k-a (config-pmap-nq-c)# mtu 9216
ora12c-n5k-a (config-pmap-nq-c)# multicast-optimize
ora12c-n5k-a (config-pmap-nq-c)# exit
ora12c-n5k-a (config-pmap-nq)# system qos
ora12c-n5k-a (config-sys-qos)# service-policy type network-qos jumbo
ora12c-n5k-a (config-sys-qos)# exit
ora12c-n5k-a (config)# copy running-config startup-config
[########################################] 100%
ora12c-n5k-a (config)#
```

Enable this on both N5K setups.

Installing the Operating System, Additional RPMs and Preparing the System for Oracle 12c RAC, Flex ASM, ACFS and RDBMS Installation

Note Oracle Linux 6.3 was installed.

Preparatory Steps

A few changes may have to be done on the storage and on N5K in order to install Oracle Linux 6.3 with boot LUNs, configured on EMC PowerPath. Detailed steps are provided in EMC PowerPath for Linux version 5.7 Installation and Administration guide.

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

Storage Boot LUNs Configuration

Make sure that the boot LUN for the server is presented to the host first from the storage side. Eight storage groups were defined, one for each Cisco UCS B420. For server 1, the boot LUN was added to the first storage group. Also make a note of the host ID (preferably 0 as this is the first LUN presented to the host)

Figure 13 Storage Group Properties

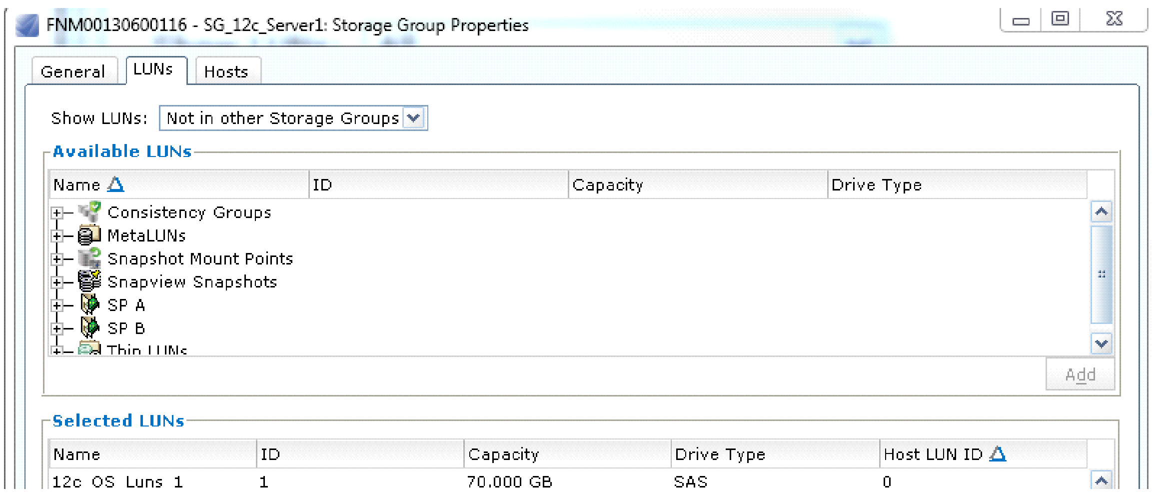

Verify the connectivity status and the host initiators.

SAN Zoning Changes on N5K for Boot

Change the zoning policy on N5K's so that only one path is available during the boot time. Disable the zones say on N5K-B and enable only on N5K-A. Also make sure that only one path is available before install. The Linux installer should show the paths as sda/sdb etc and not as mpatha/mpathb in order to make boot LUNs with PowerPath. When the installation is complete and PowerPath is completely setup, this may be reverted back to it's full paths. As an example for server 1 (ora12crac1) only one zone is made available before install as below.

```
  zone name ora12c_b420_1_hba1 vsan 15
* fcid 0x7413ef [device-alias A0P0]
* fcid 0x740001 [pwwn 20:00:00:25:b5:00:00:6e]
```

Configure Boot Policies on Cisco UCS Servers

Define boot policy for the server 1

Log in to Cisco UCS Manager, go to the Servers tab, filter on Policies, right-click Boot Policy to create a policy as shown in Figure 14.

Figure 14 Creating Boot Policy

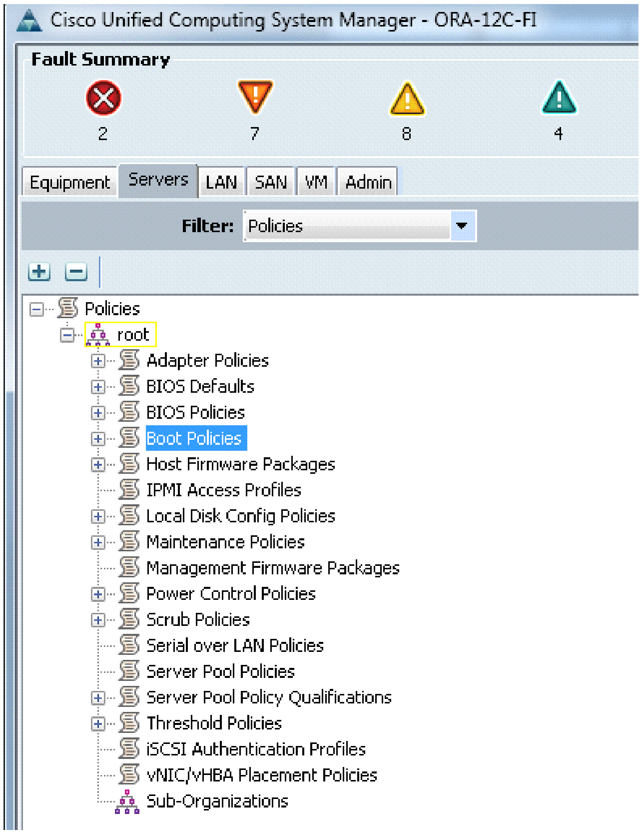

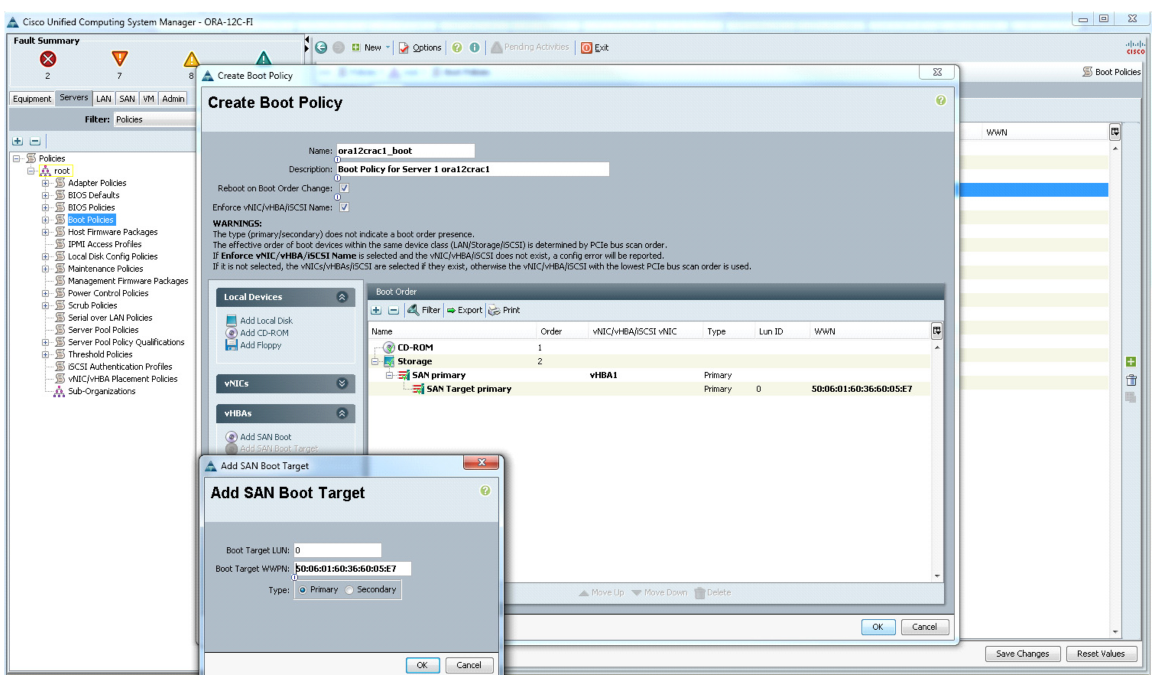

Click OK to add the SAN Boot Target. Next, add the secondary target for SAN primary

For both SAN Primary and SAN secondary, add the SAN Boot targets as shown below. The Boot Target LUN ID should match the Host ID from VNX as mentioned earlier.

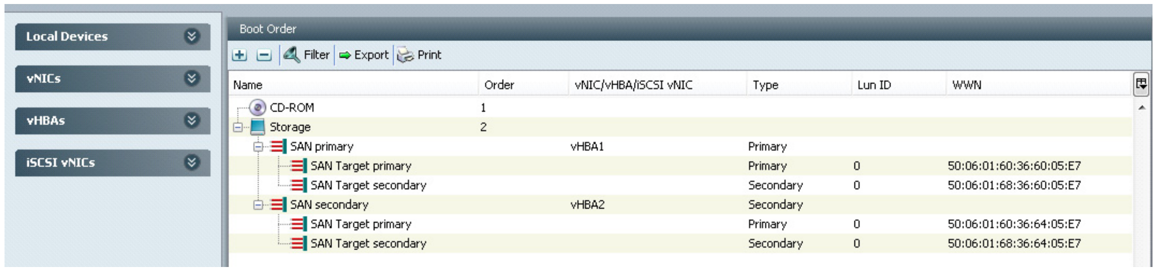

The policy will look like the screenshot above after adding the targets to both Primary and Secondary. These steps have to be repeated for all the Oracle 12c servers.

While the screenshot illustrates how to set up the paths during the runtime operation, changes have to be made to a single path before the OS install.

To make sure that you do not have multiple paths during boot time, temporarily disable all the paths and enable only one as shown below:

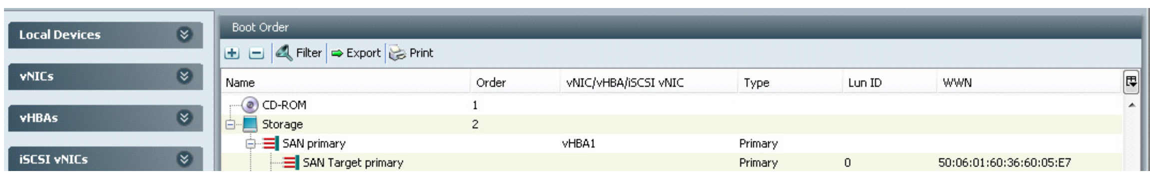

This completes the preparatory step for the OS installation.

Installing Oracle Linux 6.3 Image

Download Oracle Linux 6.3 images from https://edelivery.oracle.com/linux or as appropriate. Mount the image and launch the installer.

Launch the KVM console for the desired server, click Virtual Media, add image and reset the server. When the server comes up, it launches the Oracle Linux Installer.

Only a few of the screen shots for the install are provided below.

Select your language and installation.

Figure 15 Installing Linux 1

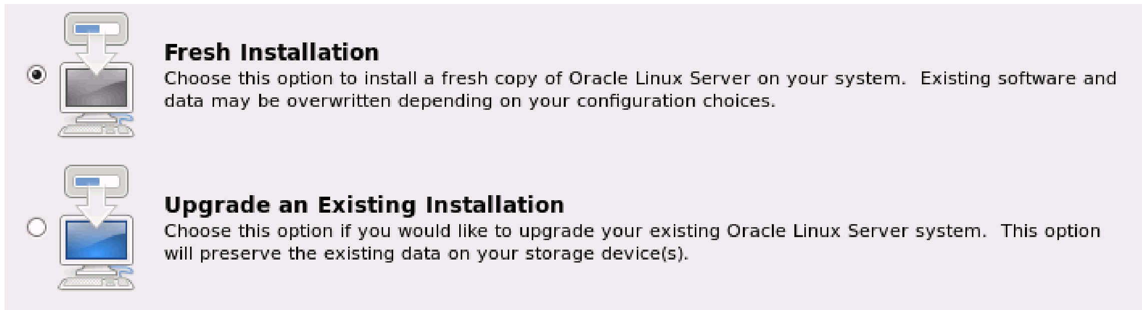

Select the hostname and click Configure Network to configure both your private and public networks. Edit each network interface and populate with the appropriate entries.

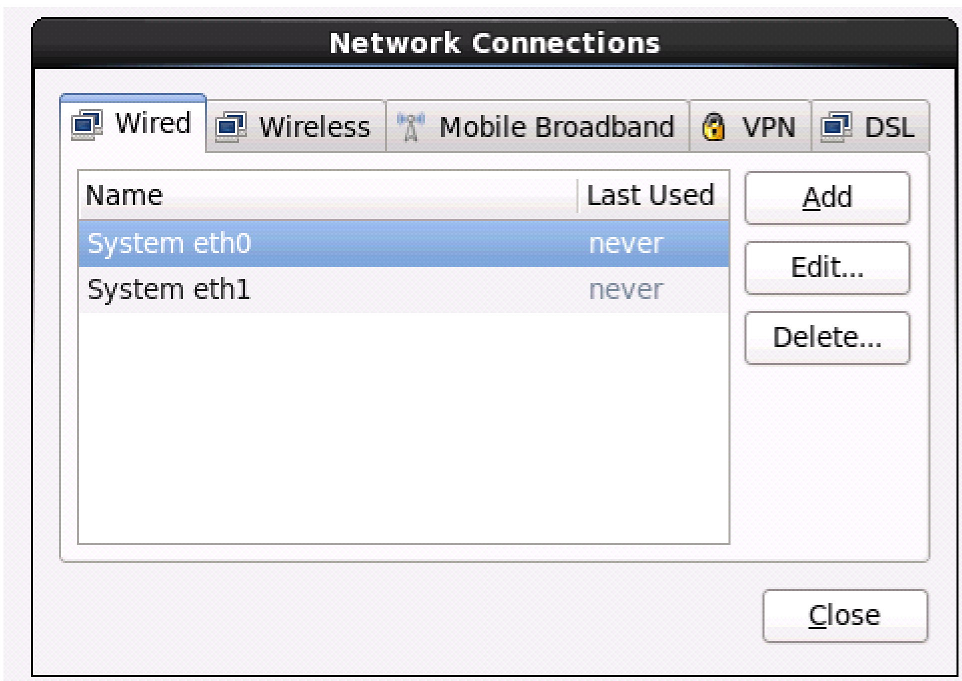

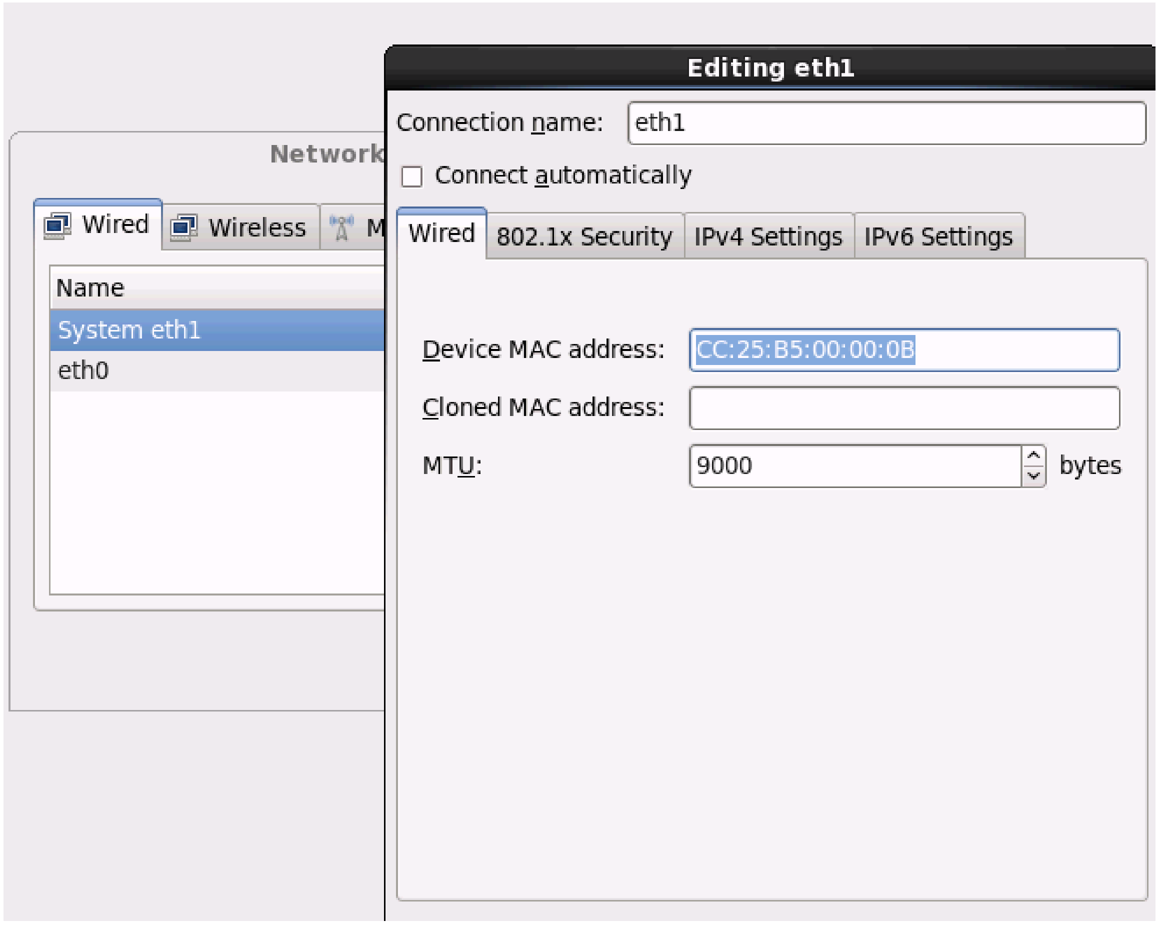

Figure 16 Installing Linux 2

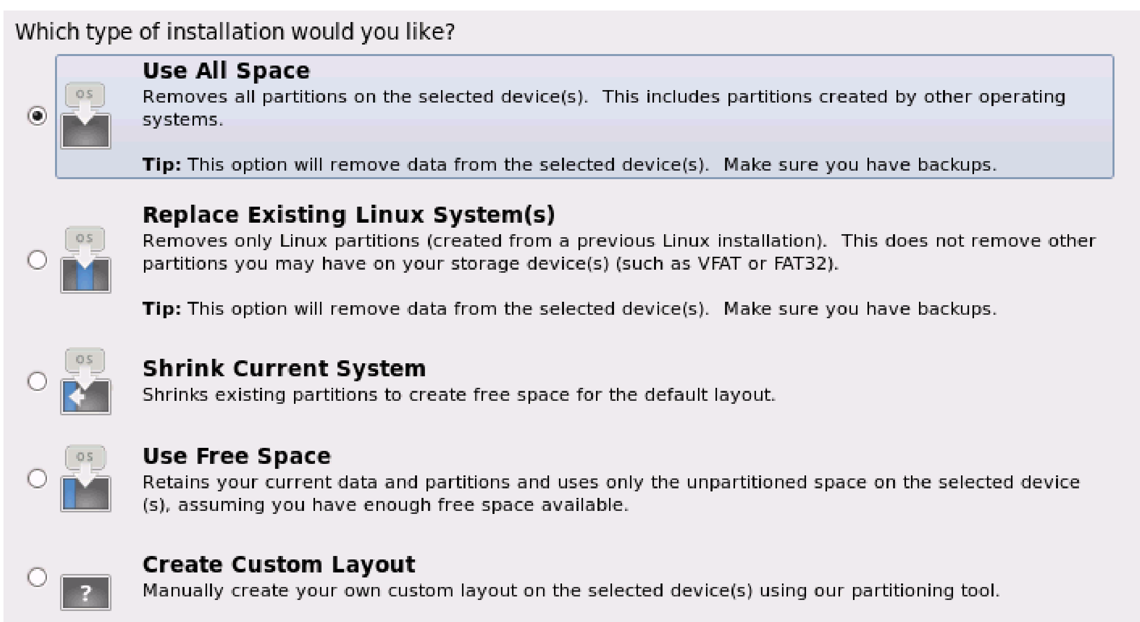

Before clicking Next, select Review and modify partitioning layout.

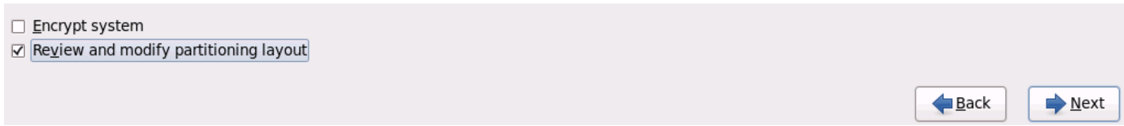

Select the appropriate devices and size.

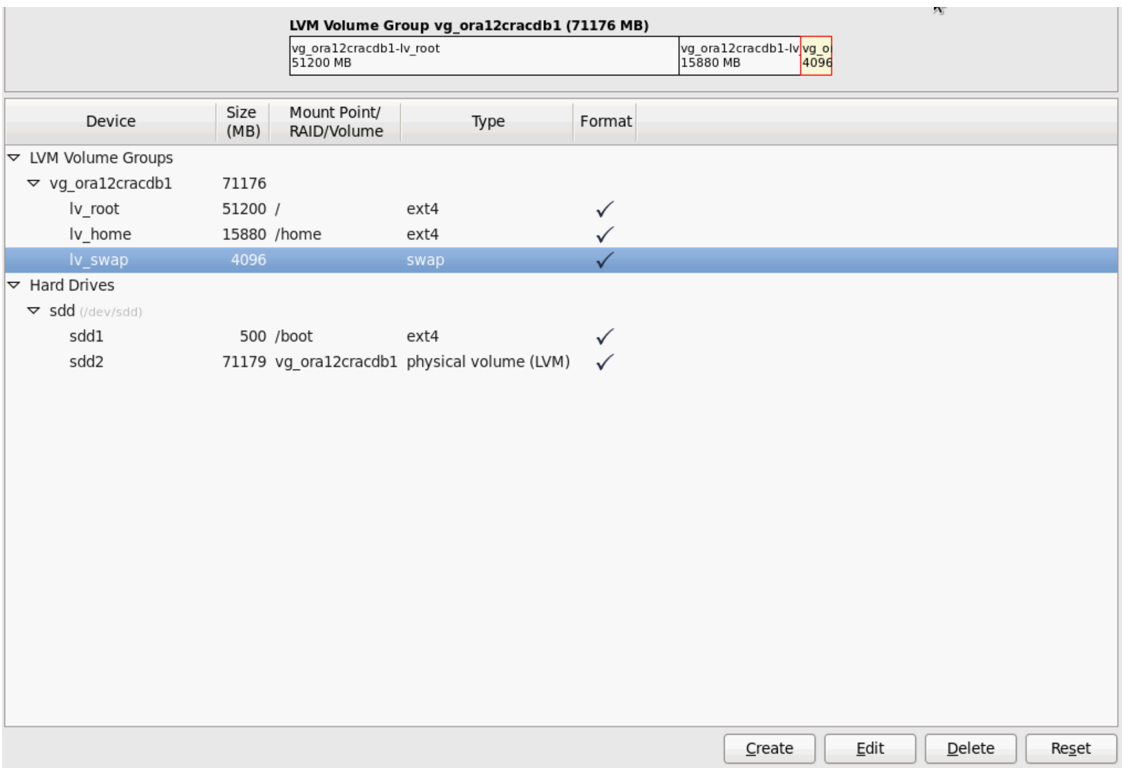

Click Change device and select First sector of boot partition as shown below:

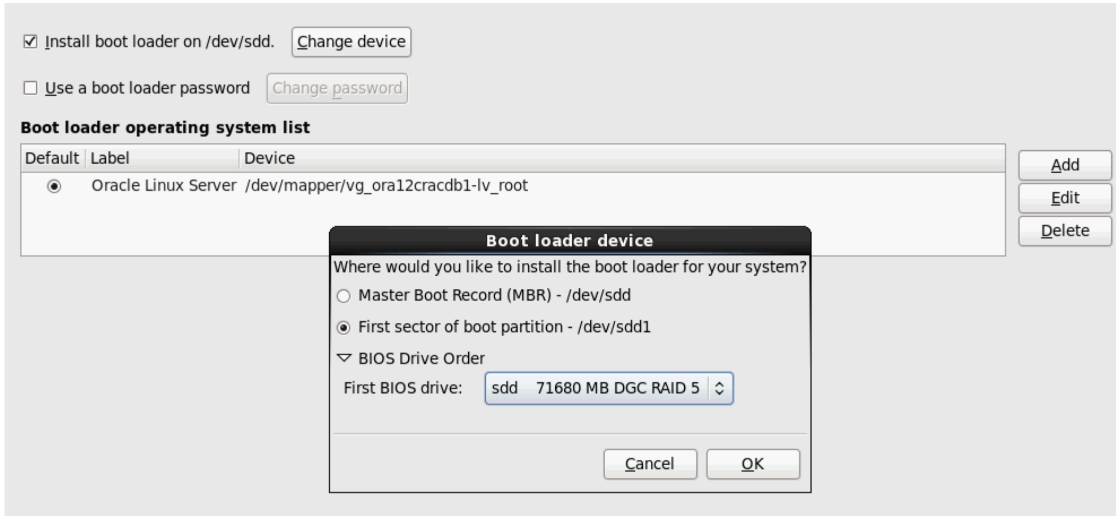

Select Customize now to add additional packages to the existing install.

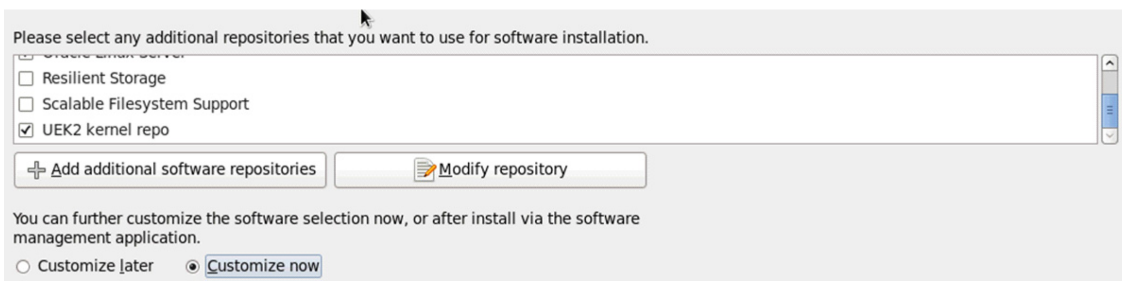

From the Servers menu, select System administration tools and then select oracleasm-support tools.

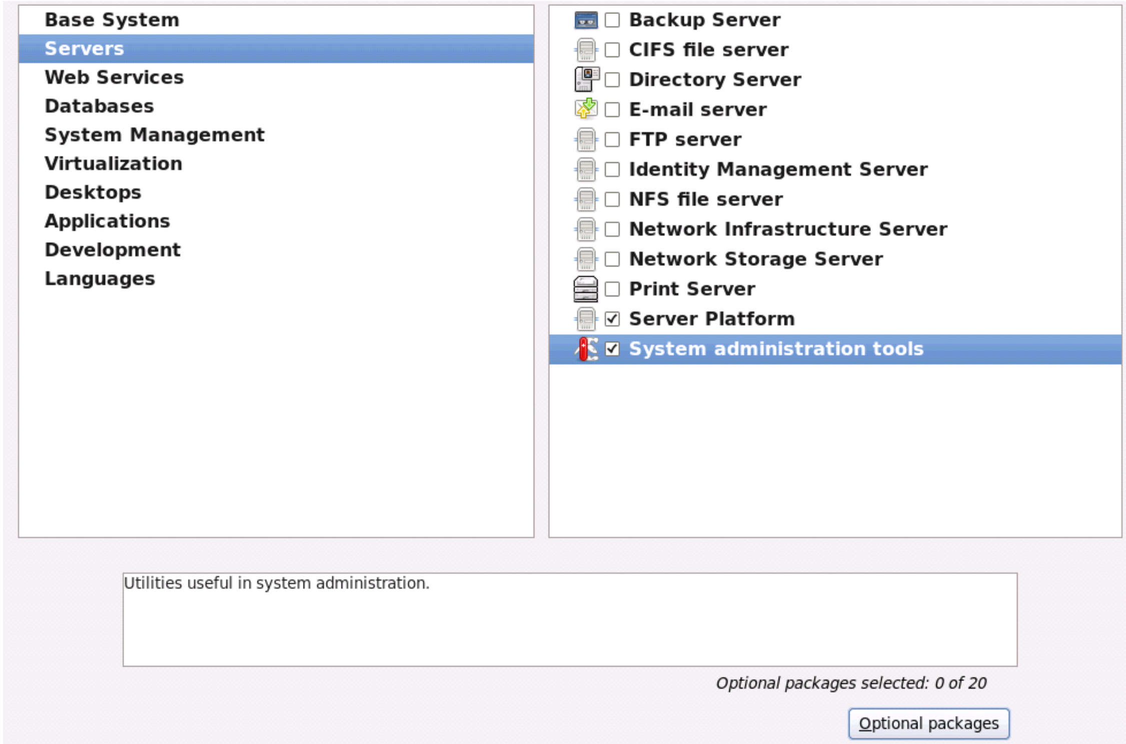

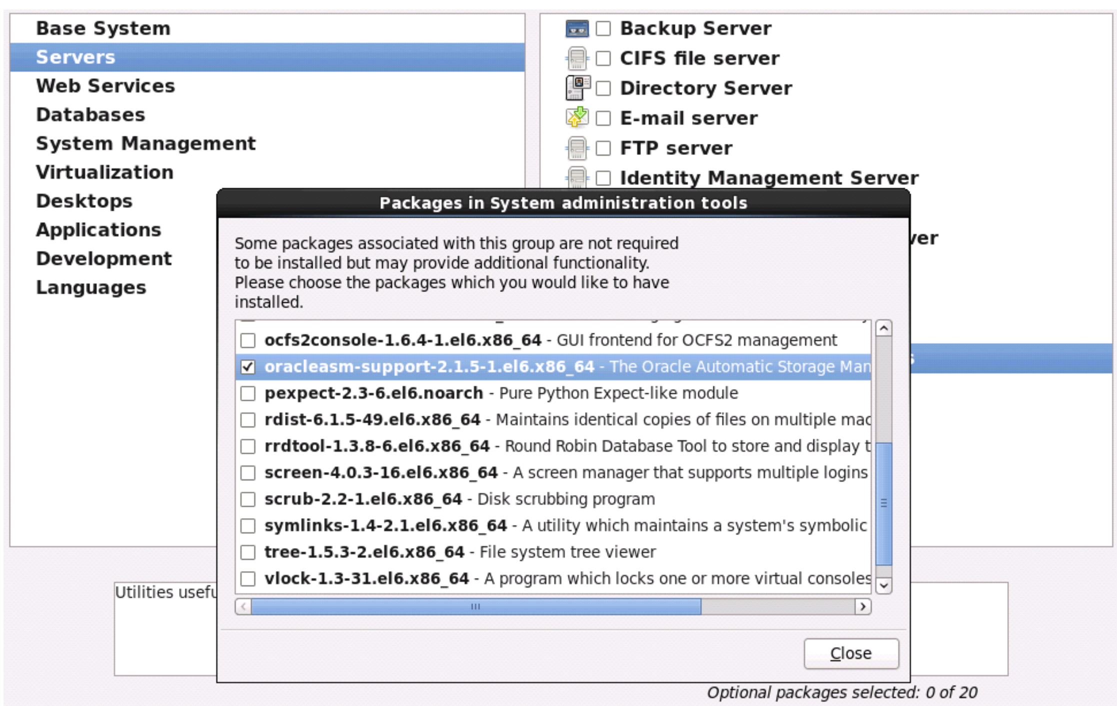

Select Desktops and then select the X Windows System.

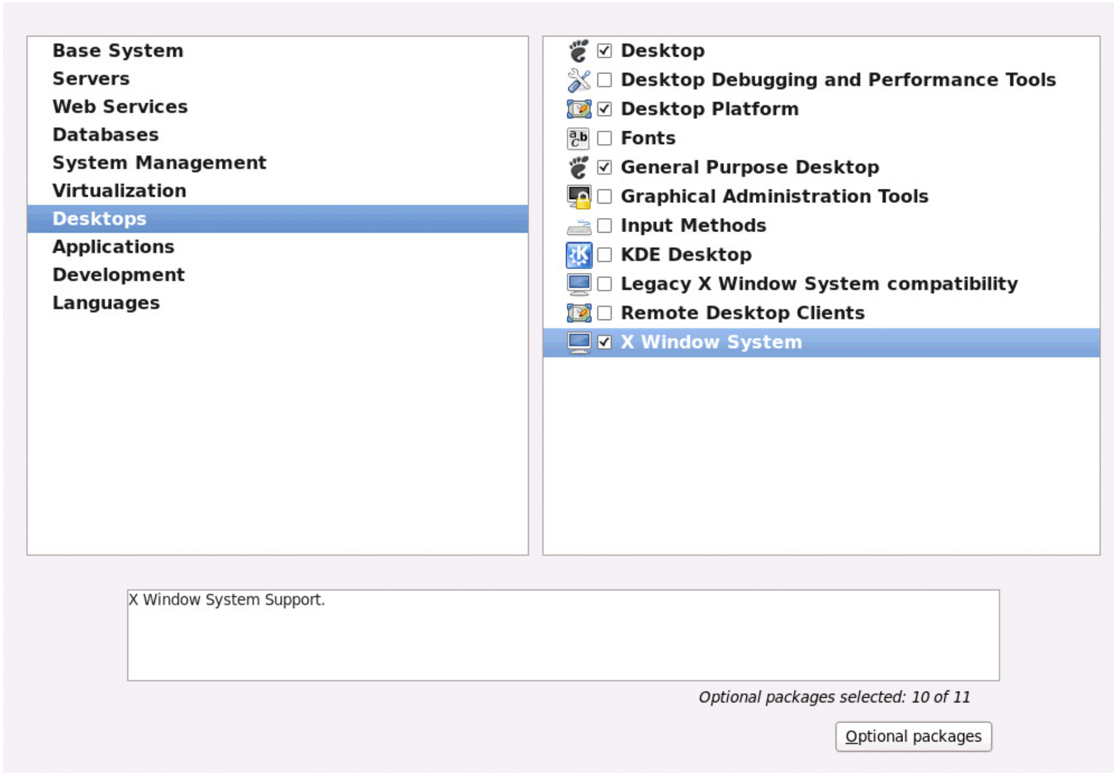

After few minutes, the system installation completes.

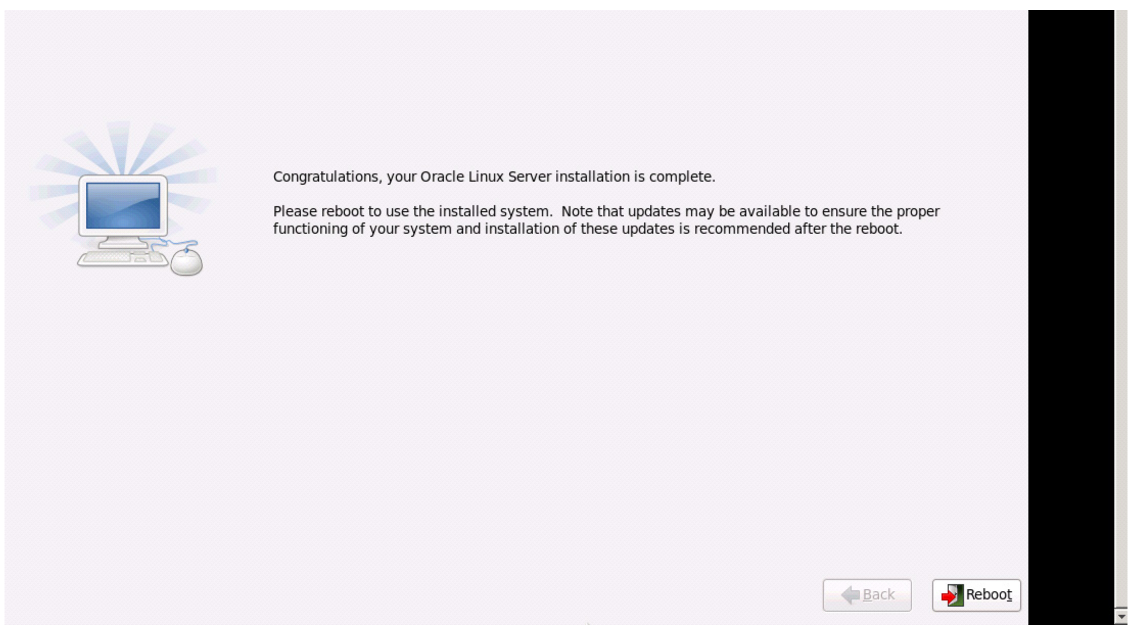

Reboot the server and accept license information, register the system as needed and synchronize the time with NTP. If NTP is not configured, Oracle RAC cluster synchronization daemon kicks in on a Oracle RAC node to sync up the time between the cluster nodes and maintaining the mean cluster time. Both NTP and OCSSD are mutually exclusive.

This completes the OS Installation.

Miscellaneous Post-Installation Steps

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

Disable selinux

It is recommended to disable selinux.

Edit /etc/selinux/config and change to

SELINUX=disabled

#SELINUXTYPE=targeted

Modify/Create the DBA Group if Needed

groupmod -g 500 oinstall

Change sshd_config File

RSAAuthentication yes

PubkeyAuthentication yes

AuthorizedKeysFile   .ssh/authorized_keys

AuthorizedKeysCommand none

UsePAM yes

X11Forwarding yes

Subsystem   sftp   /usr/libexec/openssh/sftp-server

Disable Firewalls

service iptables stop

service ip6tables stop

chkconfig iptables off

chkconfig ip6tables off

Make sure /etc/sysconfig/network has an entry for hostname. Preferably add NETWORKING_IPV6=no

Configure SSH Trust for Oracle User

Configure trust between nodes for Oracle user. This can also be done by the Oracle Installer during run time.

ssh-keygen -t rsa

cd $HOME/.ssh.

cat id_rsa.pub >> authorized_keys

ssh <server name > should login back to the host.

Setup yum.repository

cd /etc/yum.repos.d

wget http://public-yum.oracle.com/public-yum-ol6.repo

edit the downloaded file public-yum-ol6.repo and change status as enabled=1

Run yum update.

You may have to set up the http_proxy environment variable in case the server accesses the Internet through a proxy.

Please make sure that the following RPM packages are available after the yum update. Alternatively install with yum install.

oracleasmlib-2.0.4-1.el6.x86_64

oracleasm-support-2.1.5-1.el6.x86_64

oracle-rdbms-server-12cR1-preinstall-1.0-1.el6.x86_64.rpm

The exact version of the packages could be different on the uek kernel being used.

Install the Linux Driver for Cisco 10G FCOE HBA

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

In the download page, select servers-Unified computing. On the right menu select your class of servers say Cisco UCS B-series Blade server software and then select Unified Computing System (UCS) Drivers in the following page.

Select your firmware version under All Releases, say 2.1 and download the ISO image of

UCS-related drivers for your matching firmware, for example ucs-bxxx-drivers.2.1.1a.iso.

Extract the fnic rpm from the iso.

Alternatively you can also mount the iso file. You can use KVM console too and map the iso.

After mapping virtual media - Login to host to copy the rpm

```
[root@ora12crac1 ~]# mount -o loop /dev/cdrom  /mnt
[root@ora12crac1 ~]# cd /mnt
[root@ora12crac1 ~]# cd /mnt/Linux/Storage/Cisco/1280/Oracle/OL6.3
[root@ora12crac1 ~]# ls
dd-fnic-1.5.0.18-oracle-uek-6.3.iso
README-Oracle Linux Driver for Cisco 10G FCoE HBA.docx
```

Extract the rpm from iso.

Follow the instructions in README-Oracle Linux Driver for Cisco 10G FCoE HBA. In case you are running this on Oracle Linux Redhat compatible kernel, the appropriate driver for your Linux version should be installed.

Below are the steps followed for uek2 kernel:

```
[root@ora12crac1 ~]# rpm -ivh kmod-fnic-1.5.0.18-1.el6uek.x86_64.rpm
Preparing...                ###########################################
[100%]
```

```
   1:kmod-fnic              ###########################################
[100%]

[root@rac2 fnic]# modprobe fnic
[root@rac2 fnic]# modinfo fnic
filename:
/lib/modules/2.6.39-200.24.1.el6uek.x86_64/extra/fnic/fnic.ko
version:        1.5.0.18
license:        GPL v2
author:         Abhijeet Joglekar <abjoglek@cisco.com>, Joseph R. Eykholt
<jeykholt@cisco.com>
description:    Cisco FCoE HBA Driver
srcversion:     24F8E443F0EEDBDF4802F20
alias:          pci:v00001137d00000045sv*sd*bc*sc*i*
depends:        libfc,libfcoe,scsi_transport_fc
vermagic:       2.6.39-200.24.1.el6uek.x86_64 SMP mod_unload modversions
parm:           fnic_log_level:bit mask of fnic logging levels (int)
parm:           fnic_trace_max_pages:Total allocated memory pages for fnic
trace buffer (uint)
```

For more details on the install, follow the README document found in the iso above.

In general it is good practice to install the latest drivers. In case you are planning to run RHEL compatible kernel, you may have to check for any additional drivers in enic/fnic category to be installed.

Reboot the host after making the changes and verify.

Configure PowerPath

After reboot, configure PowerPath as it is only with single path now. Please contact EMC for the appropriate version of PowerPath for the operating system.

The Oracle Linux 6.3 installs with 2 kernels

```
Uek2 kernel - 2.6.39-200.24.1.el6uek.x86_64 which is the default.
Red Hat binary compatible kernel - 2.6.32-279.el6.x86_64.
```

Obtain the following rpm's from EMC directly.

```
HostAgent-Linux-64-x86-en_US-1.0.0.1.0474-1.x86_64
EMCpower.LINUX-5.7.1.02.00-004.ol6_uek2_r2.x86_64 ( power path rpm for uek2
kernel )
```

For the actual list of PowerPath and Linux Kernel versions, please check at http://powerlink.emc.com

Make sure that multipath is not running.

```
[root@ora12crac1 ~]# service --status-all  | grep multipath
[root@ora12crac1 ~]# multipath -ll
-bash: multipath: command not found

[root@ora12crac1 ~]# rpm -ivh
HostAgent-Linux-64-x86-en_US-1.0.0.1.0474-1.x86_64.rpm
Preparing...                ###########################################
[100%]
   1:HostAgent-Linux-64-x86-###########################################
[100%]

[root@ora12crac1 ~]# rpm -ivh
EMCpower.LINUX-5.7.1.02.00-004.ol6_uek2_r2.x86_64.rpm
```

```
Preparing...                ###########################################
[100%]
   1:EMCpower.LINUX         ###########################################
[100%]
All trademarks used herein are the property of their respective owners.

[root@ora12crac1 ~]# service hostagent start
Starting Navisphere agent:                                 [  OK  ]

[root@ora12crac1 ~]# service PowerPath start
Starting PowerPath:  done

[root@ora12crac1 ~]# powermt check_registration
There are no license keys now registered.

[root@ora12crac1 ~]#  emcpreg  -add  < power path key here >
1 key(s) successfully added.

[root@ora12crac1 ~]# powermt set policy=co
[root@ora12crac1 ~]# powermt config
[root@ora12crac1 ~]# powermt save

[root@ora12crac1 ~]# powermt display dev=all
Pseudo name=emcpowera
VNX ID=FNM00130600116 [SG_12c_Server1]
Logical device ID=6006016013203300E3AEB4EE32E0E211 [12c_OS_Luns_1]
state=alive; policy=CLAROpt; queued-IOs=0
Owner: default=SP A, current=SP A       Array failover mode: 4
============================================================================
==
--------------- Host ---------------   - Stor -  -- I/O Path --   -- Stats
---
###  HW Path               I/O Paths    Interf.  Mode     State   Q-IOs
Errors
============================================================================
==
   3 fnic                   sda         SP A0    active   alive      0
0
```

Note Only one path is active right now.

Reconfigure Zoning and Boot Policies

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

The zoning attributes for each HBA (hba1 as an example below) needs to be reverted back to what was planned earlier

```
zone name ora12c_b420_1_hba1 vsan 15
    device-alias A0P0
    device-alias B0P0
    pwwn 20:00:00:25:b5:00:00:6e
```

Similarly, change the boot policy of the server to multiple paths as below

Figure 17 Configure San Boot Order in Cisco UCS

Reboot the server.

After reboot, all the paths should be active as shown below.

After activating, powermt should display text as shown below:

```
[root@ora12crac1 ~]# powermt display dev=all
Pseudo name=emcpowera
VNX ID=FNM00130600116 [SG_12c_Server1]
Logical device ID=600601601320330 0E3AEB4EE32E0E211 [12c_OS_Luns_1]
state=alive; policy=CLAROpt; queued-IOs=0
Owner: default=SP A, current=SP A       Array failover mode: 4
============================================================================
==
--------------- Host ---------------   - Stor -  -- I/O Path --   -- Stats
---
###  HW Path               I/O Paths    Interf.  Mode     State   Q-IOs
Errors
============================================================================
==
   2 fnic                   sduo        SP B8    active   alive      0
0
   2 fnic                   sdrm        SP A8    active   alive      0
0
   4 fnic                   sdok        SP B9    active   alive      0
0
   4 fnic                   sdkk        SP A9    active   alive      0
0
   1 fnic                   sdig        SP B0    active   alive      0
0
   1 fnic                   sdco        SP A0    active   alive      0
0
   3 fnic                   sdcc        SP A1    active   alive      0
0
   3 fnic                   sda         SP B1    active   alive      0
0
```

Configuring Boot LUN

Please follow the instructions from the EMC PowerPath Install and Administration guide. A few of the steps are mentioned below:

The Powermt command shown above details that emcpowera is the pseudo device for 12c_OS_Luns_1 lun.

- Capture the partitions from /proc/partitions

```
[root@ora12crac1 ~]# cat /proc/partitions | grep emcpowera
 120        0   73400320 emcpowera
 120        1     512000 emcpowera1 -> Boot partition
 120        2   72887296 emcpowera2
```

- Backup /etc/fstab file and change the entries

```
/dev/mapper/vg_ora12crac1-lv_root /                    ext4    defaults
1 1
#UUID=e7b411c6-815e-4196-a755-3187529c3554 /boot                  ext4
defaults        1 2
/dev/emcpowera1 /boot                  ext4    defaults,_netdev        1 0
# fsck disabled for /boot partition
/dev/mapper/vg_ora12crac1-lv_home /home                ext4    defaults
1 0
/dev/mapper/vg_ora12crac1-lv_swap swap                 swap    defaults
0 0
tmpfs                   /dev/shm                tmpfs   size=132326088
0 0
devpts                  /dev/pts                devpts  gid=5,mode=620  0 0
sysfs                   /sys                    sysfs   defaults        0 0
proc                    /proc                   proc    defaults        0 0
```

Change to pseudo devices entries in fstab

- Unmount and mount boot partition

```
[root@ora12crac1 ~]# umount /boot
[root@ora12crac1 ~]# mount /boot
```

- Check emcpower devices for system partitions now

```
[root@ora12crac1 ~]#  df -k
Filesystem         1024-blocks      Used Available Capacity Mounted on
/dev/mapper/vg_ora12crac1-lv_root  51606140  18937112  30047588      39% /
tmpfs                115778344    596064 115182280       1% /dev/shm
/dev/mapper/vg_ora12crac1-lv_home  16005540   2317200  12875284      16%
/home
/dev/emcpowera1          495844     80353    389891      18% /boot
```

- Make lvm changes

```
Take backup of /etc/lvm.conf  and make changes to filter as below.
#  filter = [ "a/.*/" ] -> Comment out the existing entry
    filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ] # New values
```

- Run vgscan and lvmdiskscan to flush out cache

```
[root@ora12crac1 ~]#  vgscan -v
Wiping cache of LVM-capable devices
Wiping internal VG cache
Reading all physical volumes.  This may take a while...

[root@ora12crac1 ~]#  lvmdiskscan
  /dev/ram0       [      16.00 MiB]
  /dev/ram1       [      16.00 MiB]
  /dev/emcpowera1 [     500.00 MiB]
  /dev/ram2       [      16.00 MiB]
  /dev/emcpowera2  [      69.51 GiB]
    ……………………………..
```

- Create new image file

```
cd /boot
[root@ora12crac1 boot]# dracut /boot/initramfs-PP-$(uname -r).img $(uname
-r)

[root@ora12crac1 boot]# ls -l initramfs*
-rw-r--r--. 1 root root 16155005 Mar 17 09:25
initramfs-2.6.32-279.el6.x86_64.img
-rw-r--r--  1 root root 20732599 Mar 19 12:30
initramfs-2.6.39-200.24.1.el6uek.x86_64.img
-rw-r--r--  1 root root 20666728 Mar 17 10:32
initramfs-PP-2.6.39-200.24.1.el6uek.x86_64.img
```

- Backup grub.conf and replace the entries pointing to new PowerPath initramfs.
- Reboot the server

This completes the SAN boot installation items.

Note Repeat these steps (above) on all the hosts.

Configure Oracle ASM

Oracle ASM is installed as part of the installation in OEL 6 but needs to be configured:

```
[root@ora12crac1 ~]# /etc/init.d/oracleasm configure
Configuring the Oracle ASM library driver.

This will configure the on-boot properties of the Oracle ASM library
driver.  The following questions will determine whether the driver is
loaded on boot and what permissions it will have.  The current values
will be shown in brackets ('[]').  Hitting <ENTER> without typing an
answer will keep that current value.  Ctrl-C will abort.

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

[root@ora12crac1 ~]#cat /etc/sysconfig/oracleasm | grep -v '^#'
ORACLEASM_ENABLED=true
ORACLEASM_UID=oracle
ORACLEASM_GID=oinstall
ORACLEASM_SCANBOOT=true
ORACLEASM_SCANORDER="emcpower" • Add this entry
ORACLEASM_SCANEXCLUDE="sd"  • Add this entry
```

This will create a mount point /dev/oracleasm/disks

Configure ASM LUNs and Create Disks

Mask the LUNs and create partitions

- Configure Storage LUNs

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

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

- Partition LUNs

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

  Use host utilities like fdisk to create a partition on the disk.

  Create a input file, fdisk.input as shown below:

```
d
n
p
1


  <- Leave a double space here
x
b
1
2048 <- 2048 for EMC VNX.
p
w
```

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

Now all the pseudo partitions should be available in /dev as emcpowera1, emcpowerb1, emcpowerab1 etc.

Create ASM Disks

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

oracleasm createdisk -v DSS_1 /dev/emc[partition name]

This will create a disk label as DSS_1 on the partition. This can be queried with oracle supplied kfed/kfod utilities as well.

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

Scan the disks with oracleasm and these will be visible under /dev/oracleasm/disks mount point created by oracleasm earlier as shown below:

```
[root@ora12crac1 ~]# oracleasm scandisks
```

```
Reloading disk partitions: done
Cleaning any stale ASM disks...
Scanning system for ASM disks...
[root@ora12crac1 ~]# cd /dev/oracleasm/disks/
[root@ora12crac1 disks]# ls
ACFSORAHOME1  CPUCRS_12  CPUCRS_3  CPUCRS_9  DSS_14  DSS_5     OCRVOTE2
OLTP_11  OLTP_2  OLTP_8   REDO_13  REDO_4
ACFSORAHOME2  CPUCRS_13  CPUCRS_4  DSS_1     DSS_15  DSS_6     OCRVOTE3
OLTP_12  OLTP_3  OLTP_9   REDO_14  REDO_5
ACFSORAHOME3  CPUCRS_14  CPUCRS_5  DSS_10    DSS_16  DSS_7     OCRVOTE4
OLTP_13  OLTP_4  REDO_1   REDO_15  REDO_6
CPUCRS_1      CPUCRS_15  CPUCRS_6  DSS_11    DSS_2   DSS_8     OCRVOTE5
OLTP_14  OLTP_5  REDO_10  REDO_16  REDO_7
CPUCRS_10     CPUCRS_16  CPUCRS_7  DSS_12    DSS_3   DSS_9     OLTP_1
OLTP_15  OLTP_6  REDO_11  REDO_2   REDO_8
CPUCRS_11     CPUCRS_2   CPUCRS_8  DSS_13    DSS_4   OCRVOTE1  OLTP_10
OLTP_16  OLTP_7  REDO_12  REDO_3   REDO_9
```

The system is ready for the Oracle installation.

Oracle 12c Install

12c Grid Infrastructure Install

The steps to create the 8 Node Oracle RAC 12c database are not detailed this section. A few of the screenshots are provided for reference. While Oracle 12c has several features, only a few of them have been validated as part of the Oracle Certification effort. The most notable of them are the Oracle flex ASM along with pluggable databases.

Launch the Installer:

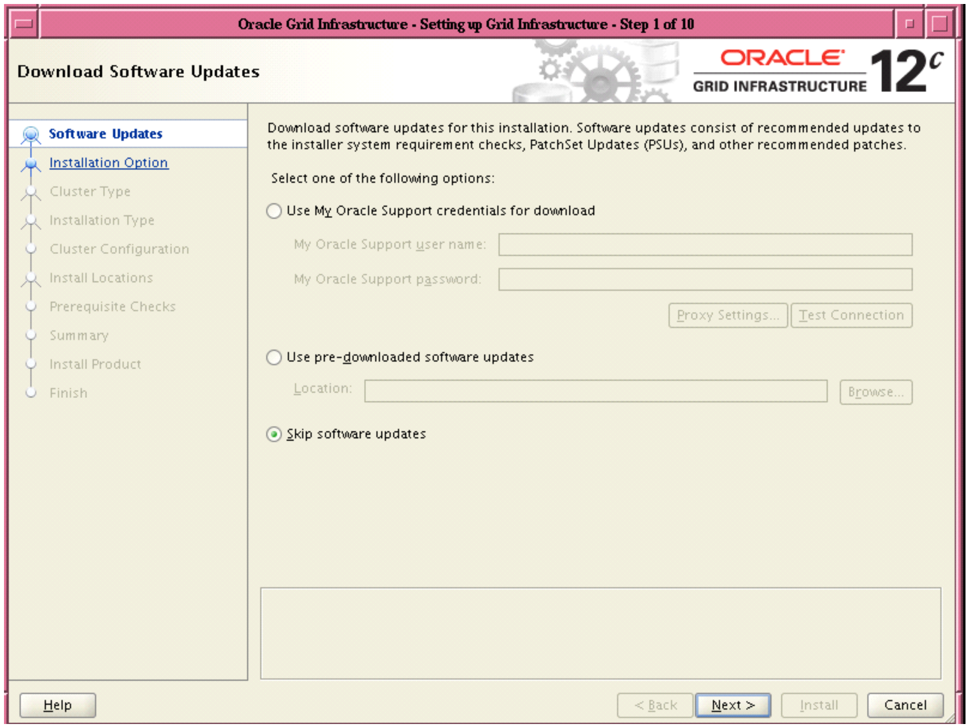

Select Install and configure Oracle Grid Infrastructure.

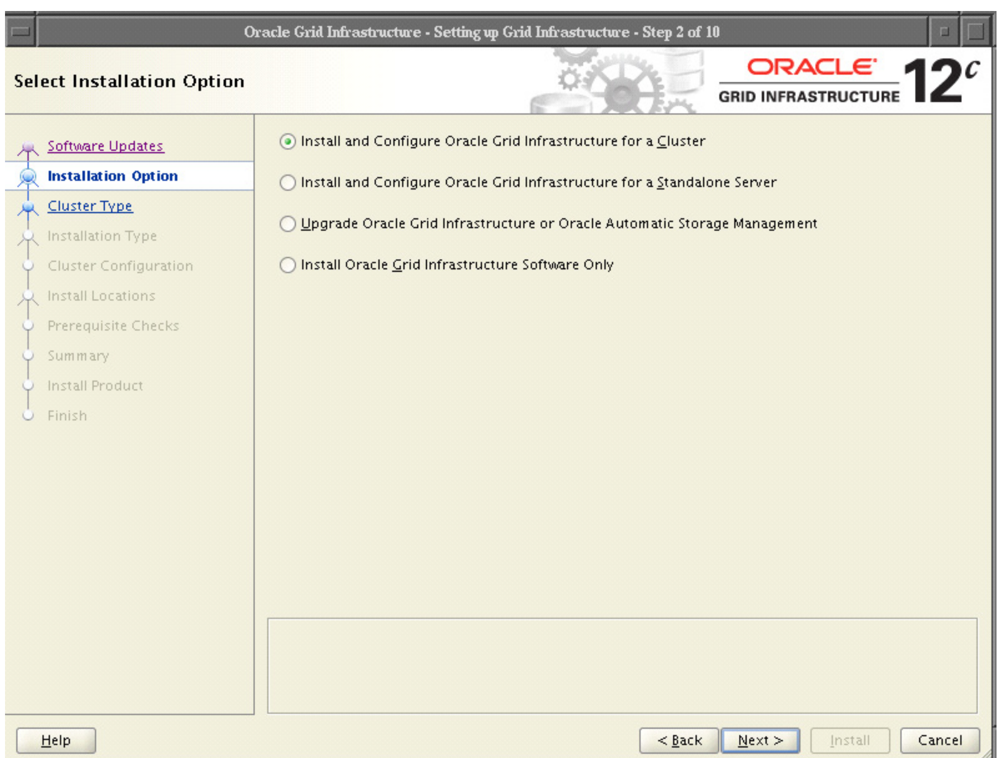

Select Configure Flex Cluster.

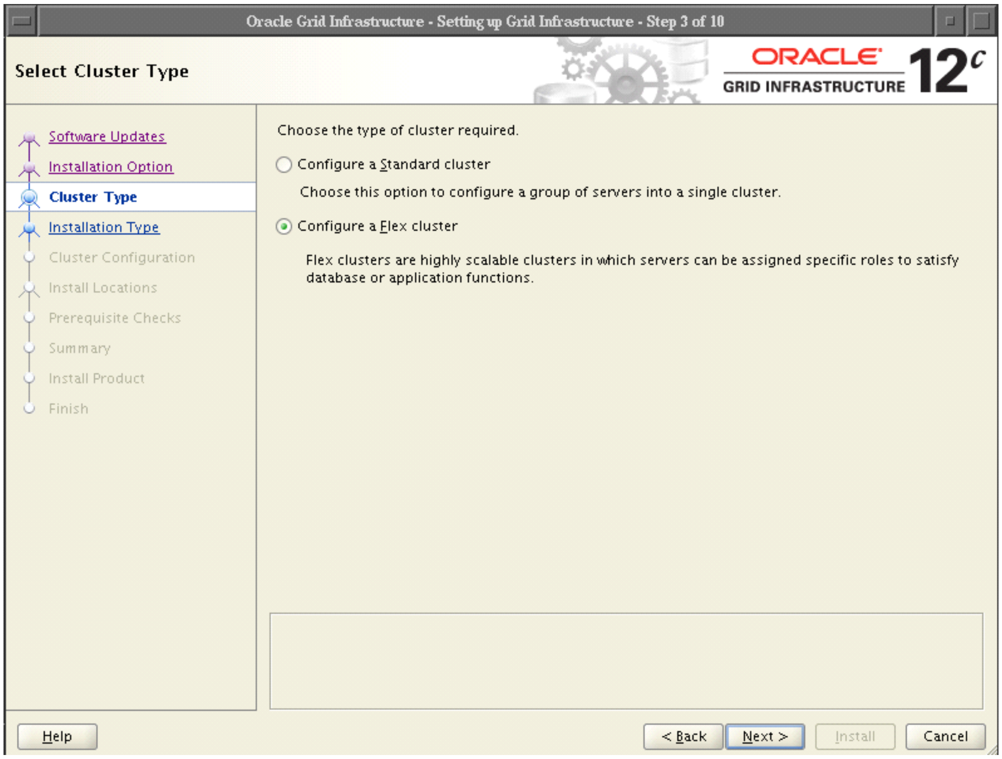

After selecting the language, select Grid Plug and play information. To use Flex ASM, configuring GNS is mandatory. The Appendix includes a few of the configuration files for GNS.

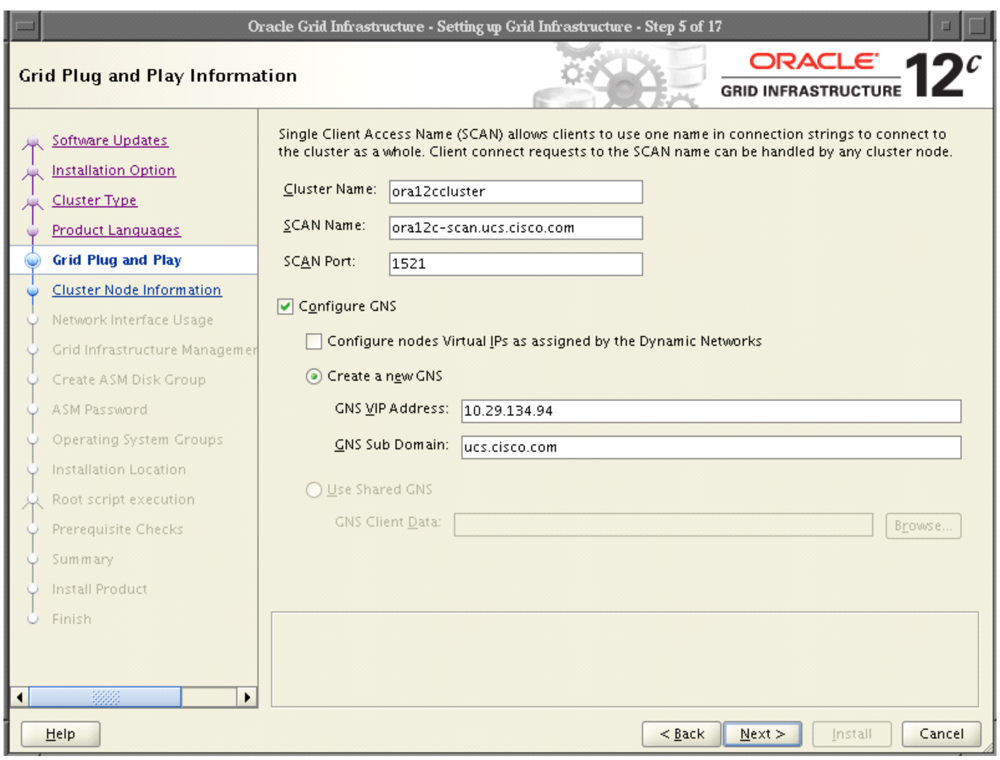

Select only HUB nodes in the setup.

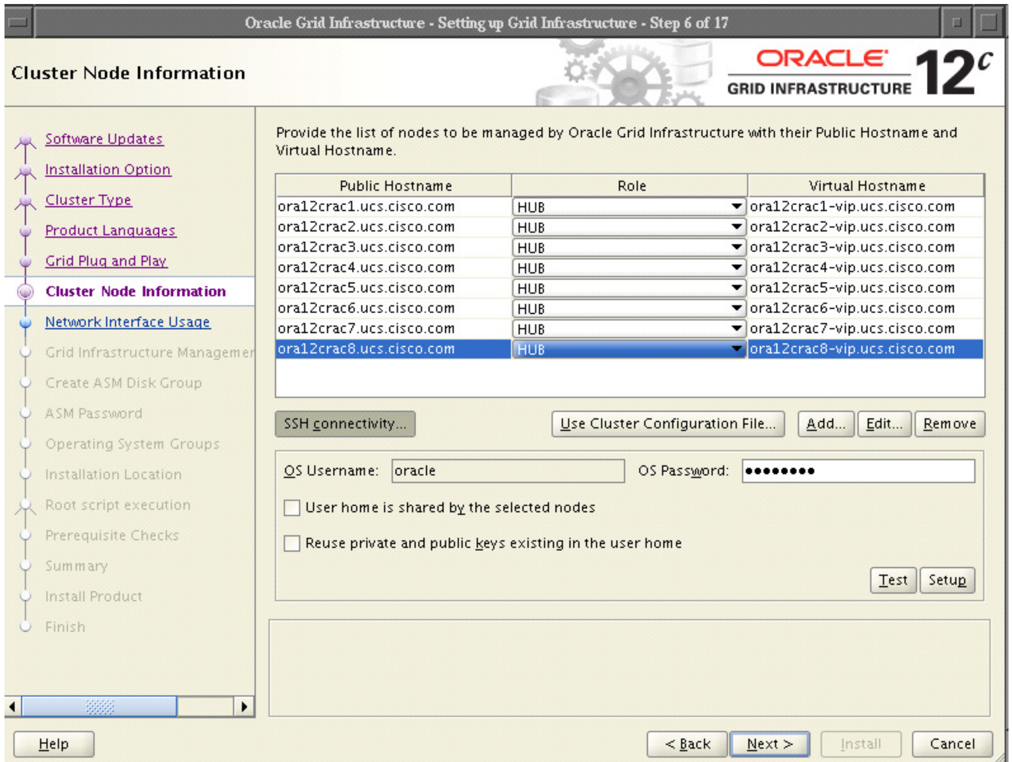

Configure the private and public networks as needed for your setup.

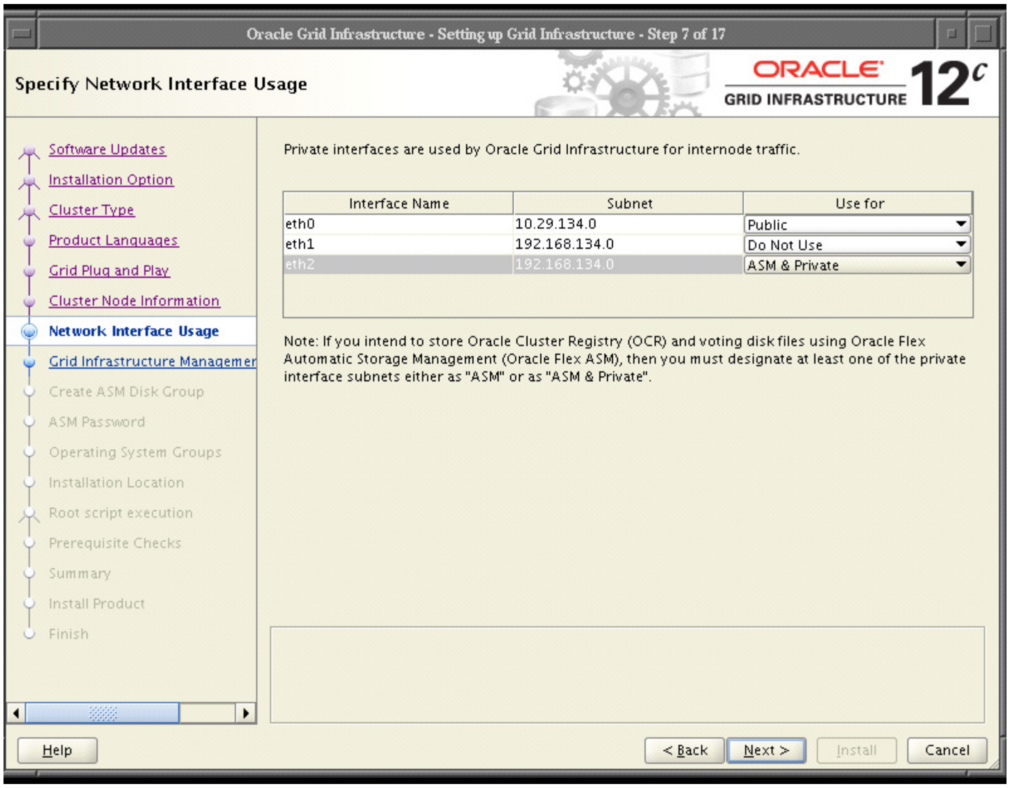

Configure OCR/Voting disks with Normal redundancy.

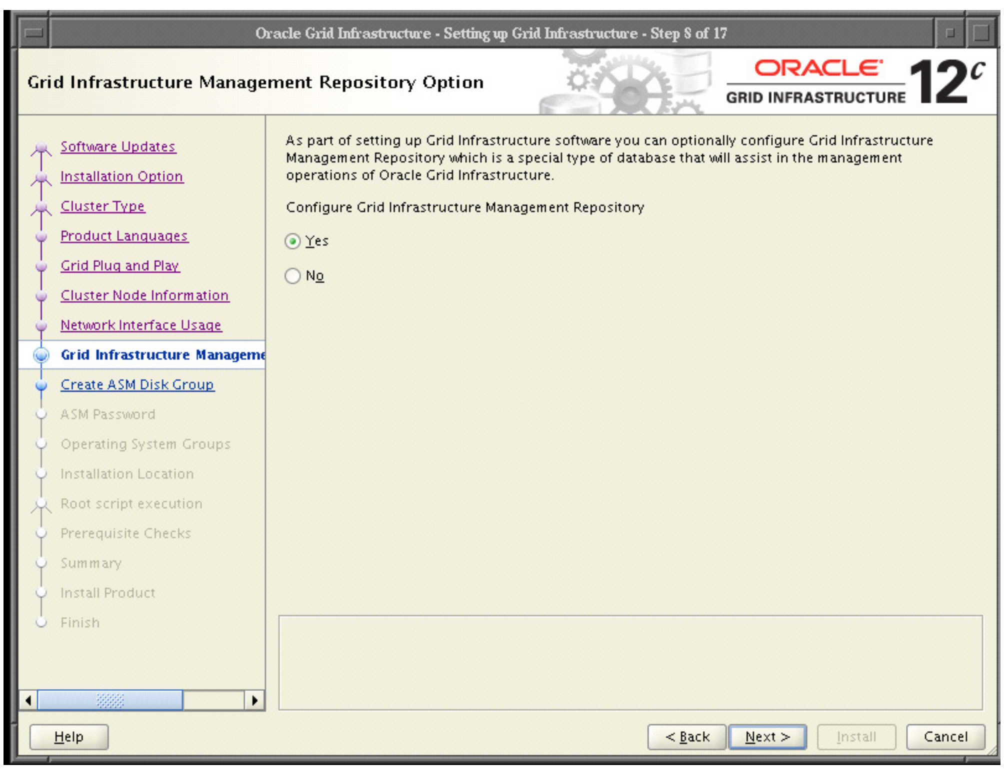

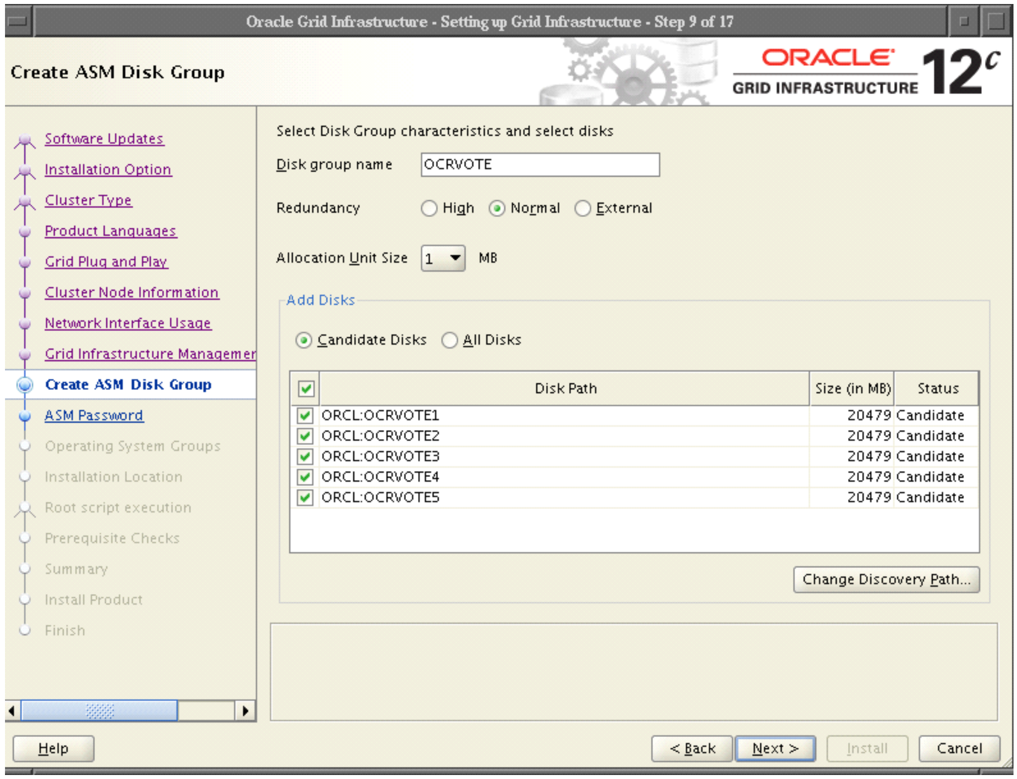

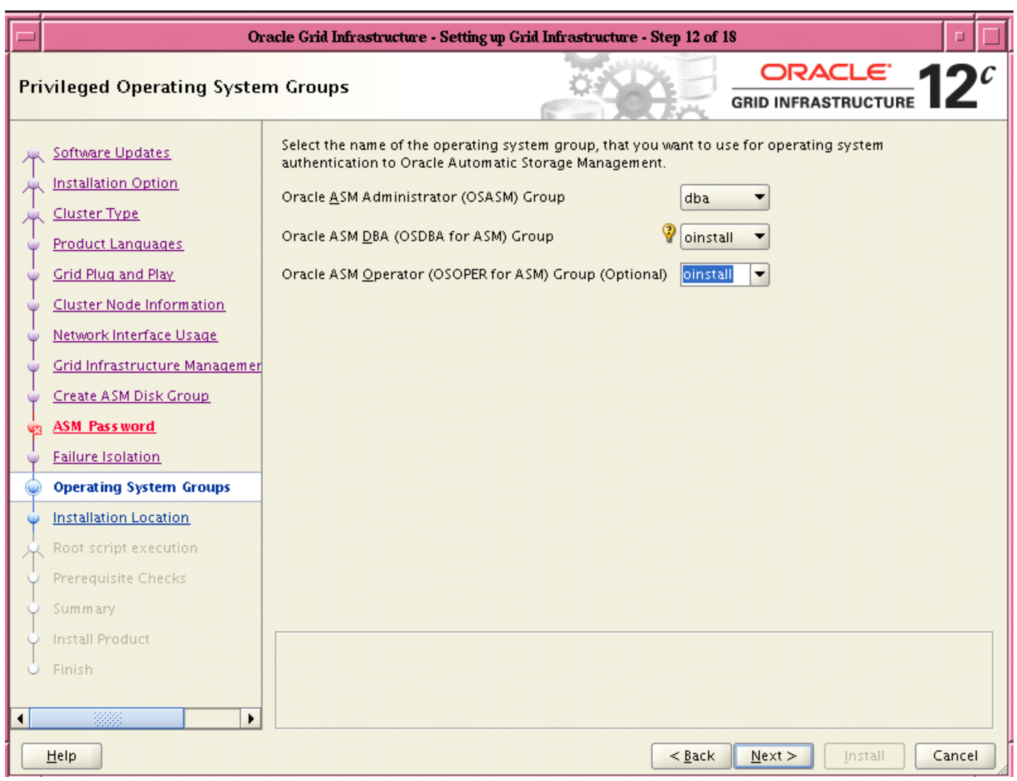

After selecting the software location, confirm the root script execution by providing the root password.

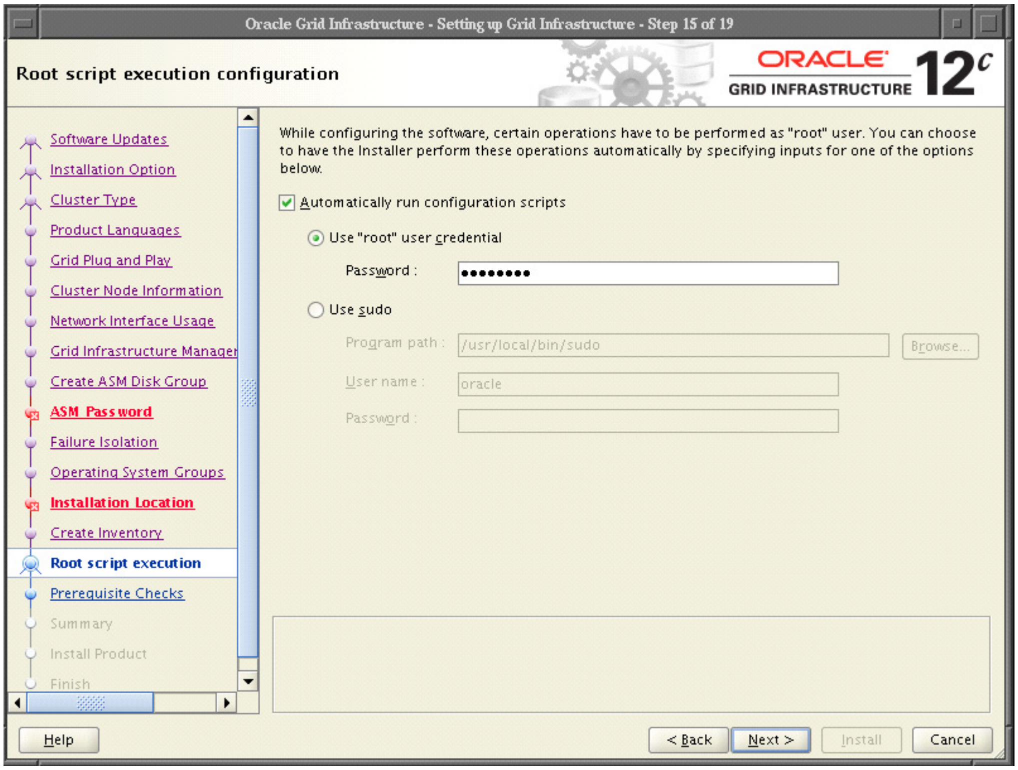
Oracle Grid Infrastructure - Setting up Grid Infrastructure - Step 15 of 19
Root script execution configuration
ORACLE
GRID INFRASTRUCTURE
12c
Software Updates
Installation Option
Cluster Type
Product Languages
Grid Plug and Play
Cluster Node Information
Network Interface Usage
Grid Infrastructure Manager
Create ASM Disk Group
ASM Password
Failure Isolation
Operating System Groups
Installation Location
Create Inventory
Root script execution
Prerequisite Checks
Summary
Install Product
Finish
While configuring the software, certain operations have to be performed as "root" user. You can choose to have the Installer perform these operations automatically by specifying inputs for one of the options below.
Automatically run configuration scripts
Use "root" user credential
Password :
Use sudo
Program path :
/usr/local/bin/sudo
Browse...
User name :
oracle
Password :
Help
< Back
Next >
Install
Cancel

Click Yes.

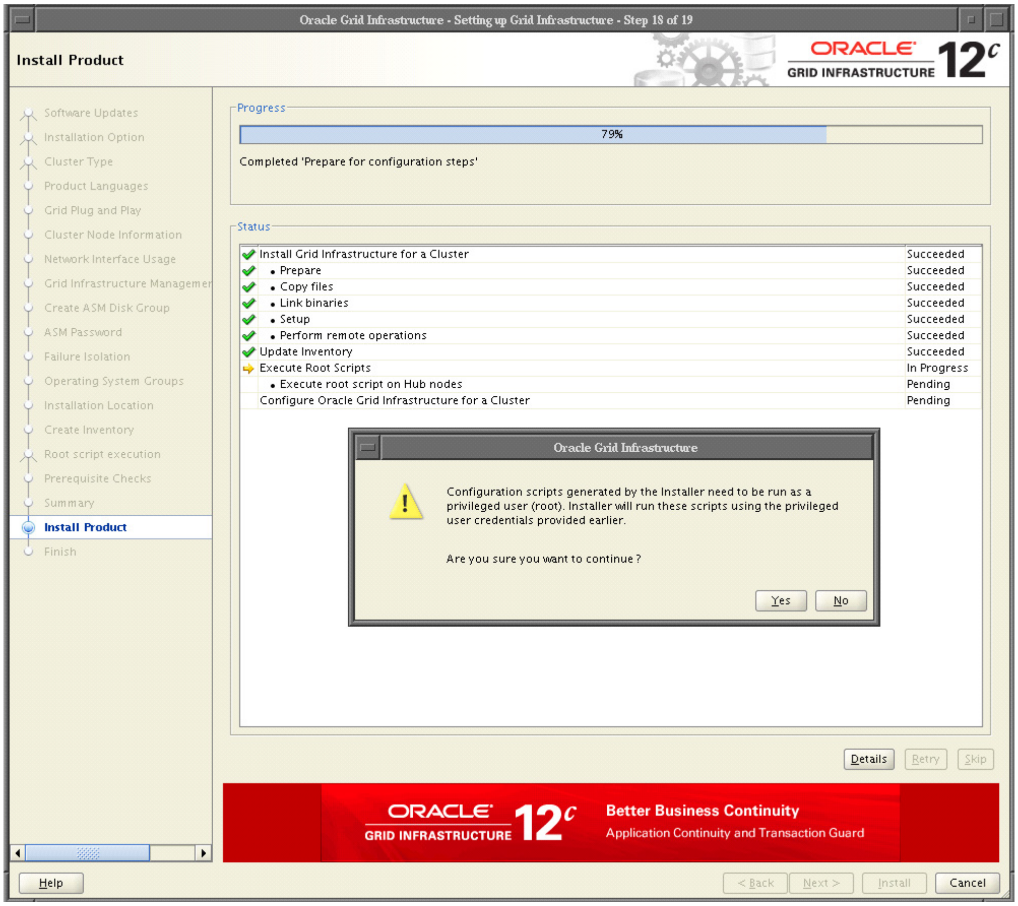
Oracle Grid Infrastructure - Setting up Grid Infrastructure - Step 18 of 19
Install Product
ORACLE
GRID INFRASTRUCTURE 12c
Software Updates
Installation Option
Cluster Type
Product Languages
Grid Plug and Play
Cluster Node Information
Network Interface Usage
Grid Infrastructure Managemer
Create ASM Disk Group
ASM Password
Failure Isolation
Operating System Groups
Installation Location
Create Inventory
Root script execution
Prerequisite Checks
Summary
Install Product
Finish
Progress
79%
Completed 'Prepare for configuration steps'
Status
Install Grid Infrastructure for a Cluster Succeeded
Prepare Succeeded
Copy files Succeeded
Link binaries Succeeded
Setup Succeeded
Perform remote operations Succeeded
Update Inventory Succeeded
Execute Root Scripts In Progress
Execute root script on Hub nodes Pending
Configure Oracle Grid Infrastructure for a Cluster Pending
Oracle Grid Infrastructure
Configuration scripts generated by the Installer need to be run as a privileged user (root). Installer will run these scripts using the privileged user credentials provided earlier.
Are you sure you want to continue ?
Yes
No
Details
Retry
Skip
ORACLE
GRID INFRASTRUCTURE 12c
Better Business Continuity
Application Continuity and Transaction Guard
Help
< Back
Next >
Install
Cancel

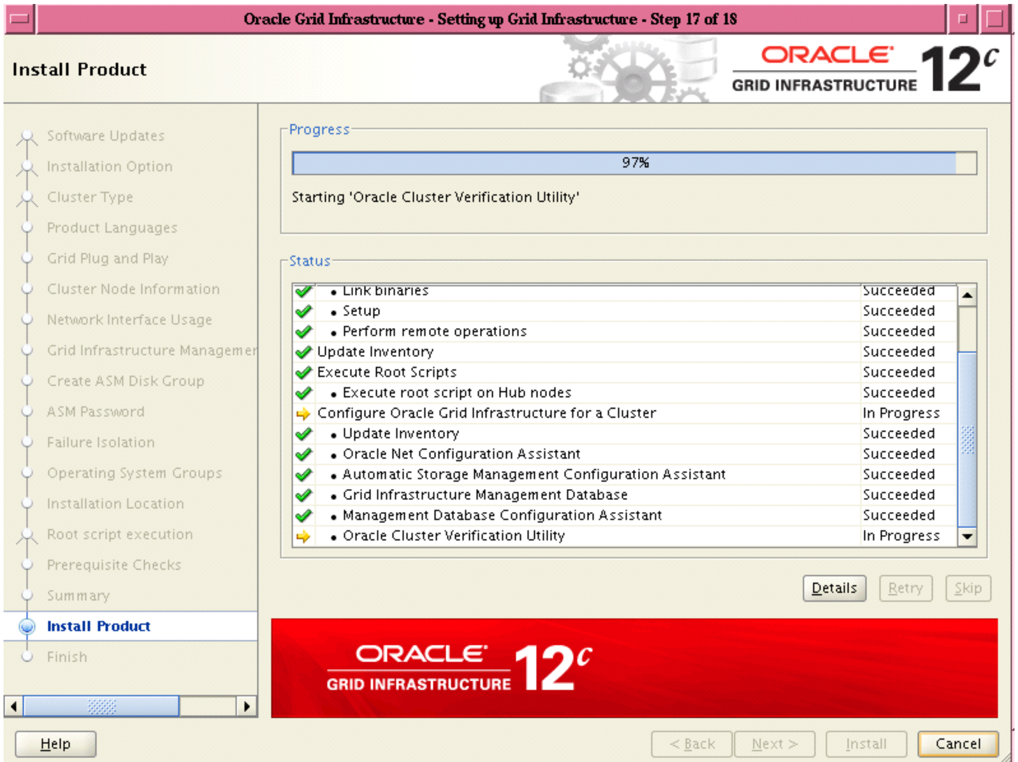
Oracle Grid Infrastructure - Setting up Grid Infrastructure - Step 17 of 18
Install Product
ORACLE
GRID INFRASTRUCTURE 12c
Software Updates
Installation Option
Cluster Type
Product Languages
Grid Plug and Play
Cluster Node Information
Network Interface Usage
Grid Infrastructure Managemen
Create ASM Disk Group
ASM Password
Failure Isolation
Operating System Groups
Installation Location
Root script execution
Prerequisite Checks
Summary
Install Product
Finish
Progress
97%
Starting 'Oracle Cluster Verification Utility'
Status
Link binaries Succeeded
Setup Succeeded
Perform remote operations Succeeded
Update Inventory Succeeded
Execute Root Scripts Succeeded
Execute root script on Hub nodes Succeeded
Configure Oracle Grid Infrastructure for a Cluster In Progress
Update Inventory Succeeded
Oracle Net Configuration Assistant Succeeded
Automatic Storage Management Configuration Assistant Succeeded
Grid Infrastructure Management Database Succeeded
Management Database Configuration Assistant Succeeded
Oracle Cluster Verification Utility In Progress
Details
Retry
Skip
ORACLE
GRID INFRASTRUCTURE 12c
Help
< Back
Next >
Install
Cancel

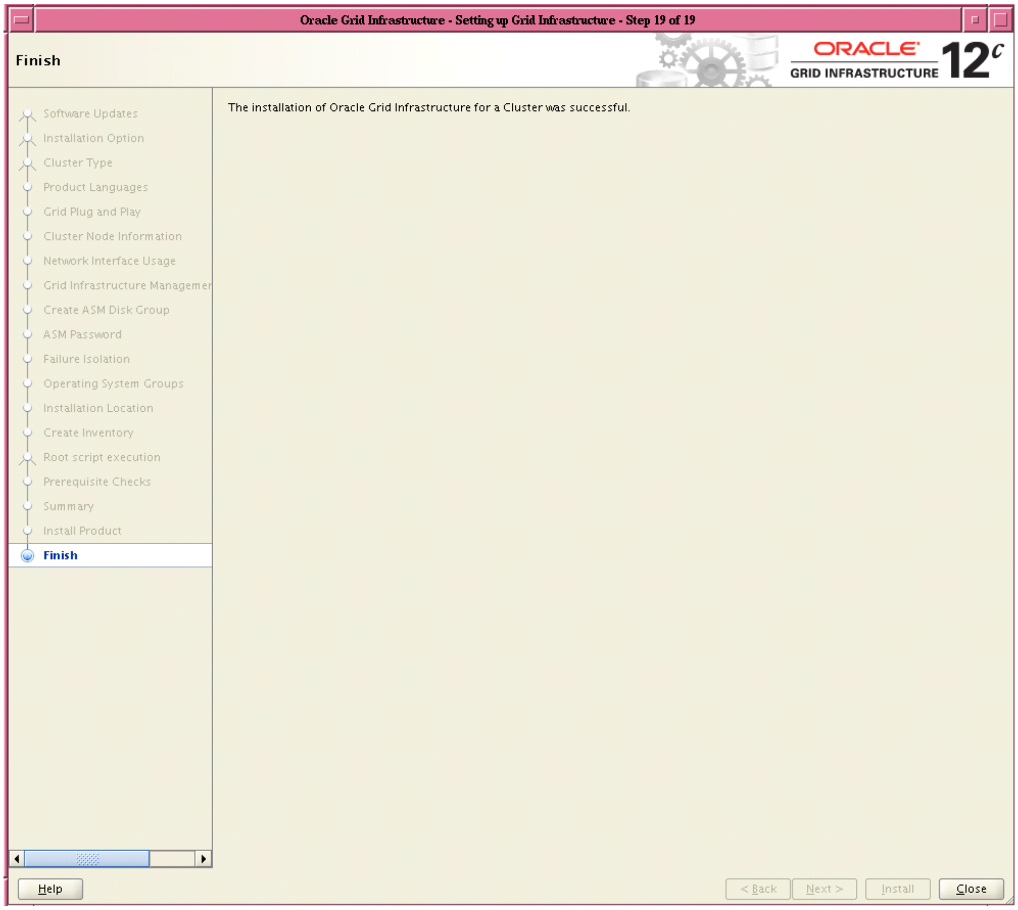

Login to each of the nodes and perform the minimum health checks to make sure that the Cluster is healthy on all the eight nodes.

12c ACFS Installation

Create a diskgroup named ACFSOH from the command line followed by launching ASMCA from grid nodes.

ACFS file system will be used as shared oracle home for 12c RDBMS binaries.

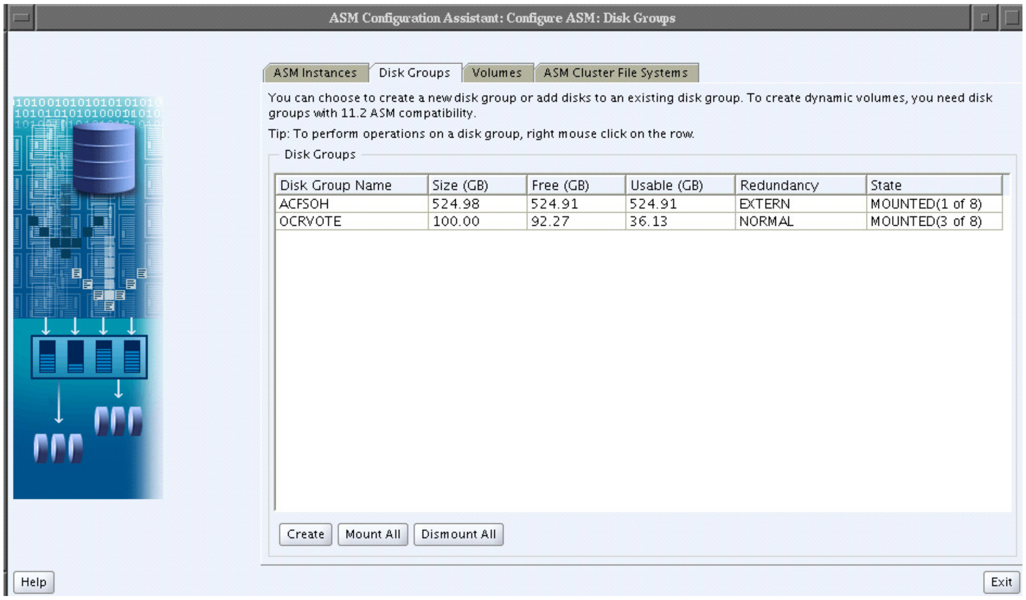
ASM Configuration Assistant: Configure ASM: Disk Groups
ASM Instances
Disk Groups
Volumes
ASM Cluster File Systems
You can choose to create a new disk group or add disks to an existing disk group. To create dynamic volumes, you need disk groups with 11.2 ASM compatibility.
Tip: To perform operations on a disk group, right mouse click on the row.
Disk Groups
Disk Group Name | Size (GB) | Free (GB) | Usable (GB) | Redundancy | State
ACFSOH | 524.98 | 524.91 | 524.91 | EXTERN | MOUNTED(1 of 8)
OCRVOTE | 100.00 | 92.27 | 36.13 | NORMAL | MOUNTED(3 of 8)
Create
Mount All
Dismount All
Help
Exit

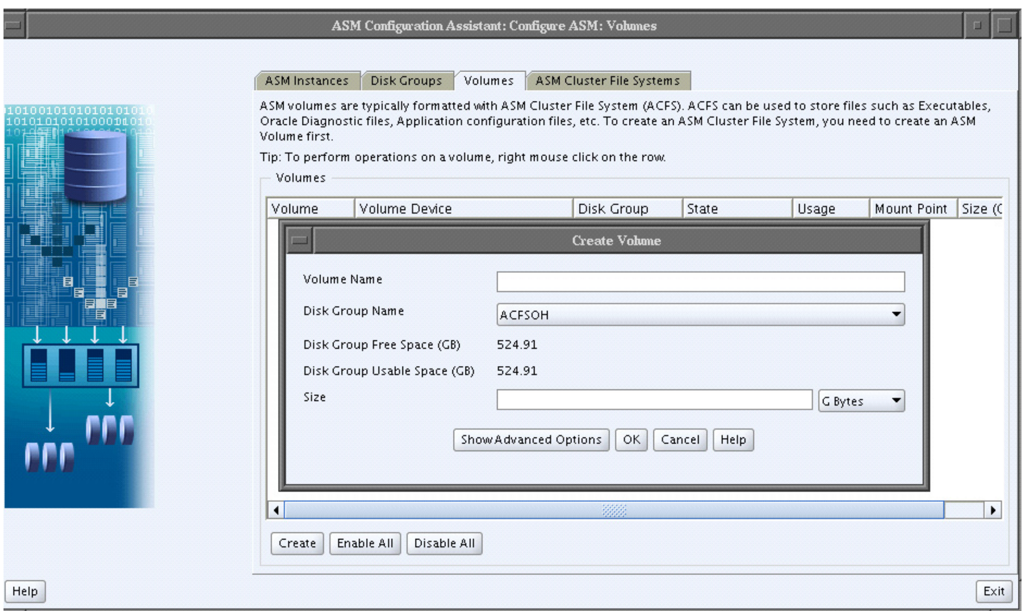
ASM Configuration Assistant: Configure ASM: Volumes
ASM Instances
Disk Groups
Volumes
ASM Cluster File Systems
ASM volumes are typically formatted with ASM Cluster File System (ACFS). ACFS can be used to store files such as Executables, Oracle Diagnostic files, Application configuration files, etc. To create an ASM Cluster File System, you need to create an ASM Volume first.
Tip: To perform operations on a volume, right mouse click on the row.
Volumes
Volume | Volume Device | Disk Group | State | Usage | Mount Point | Size (C
Create Volume
Volume Name
Disk Group Name
ACFSOH
Disk Group Free Space (GB)
524.91
Disk Group Usable Space (GB)
524.91
Size
G Bytes
Show Advanced Options
OK
Cancel
Help
Create
Enable All
Disable All
Help
Exit

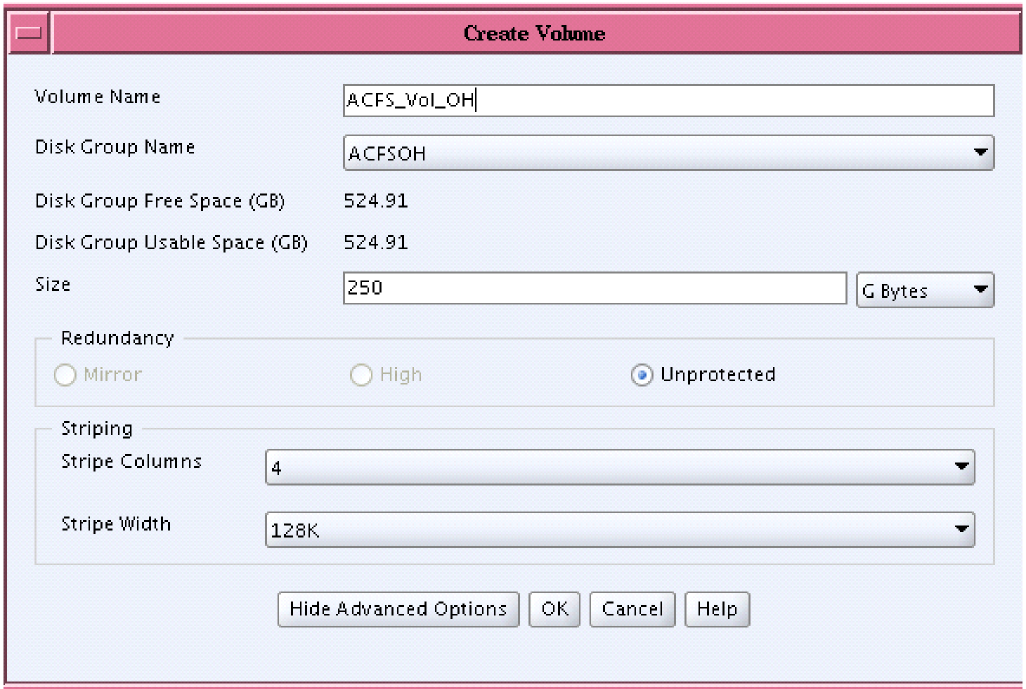
Create Volume
Volume Name
ACFS_Vol_OH
Disk Group Name
ACFSOH
Disk Group Free Space (GB)
524.91
Disk Group Usable Space (GB)
524.91
Size
250
G Bytes
Redundancy
Mirror
High
Unprotected
Striping
Stripe Columns
4
Stripe Width
128K
Hide Advanced Options
OK
Cancel
Help

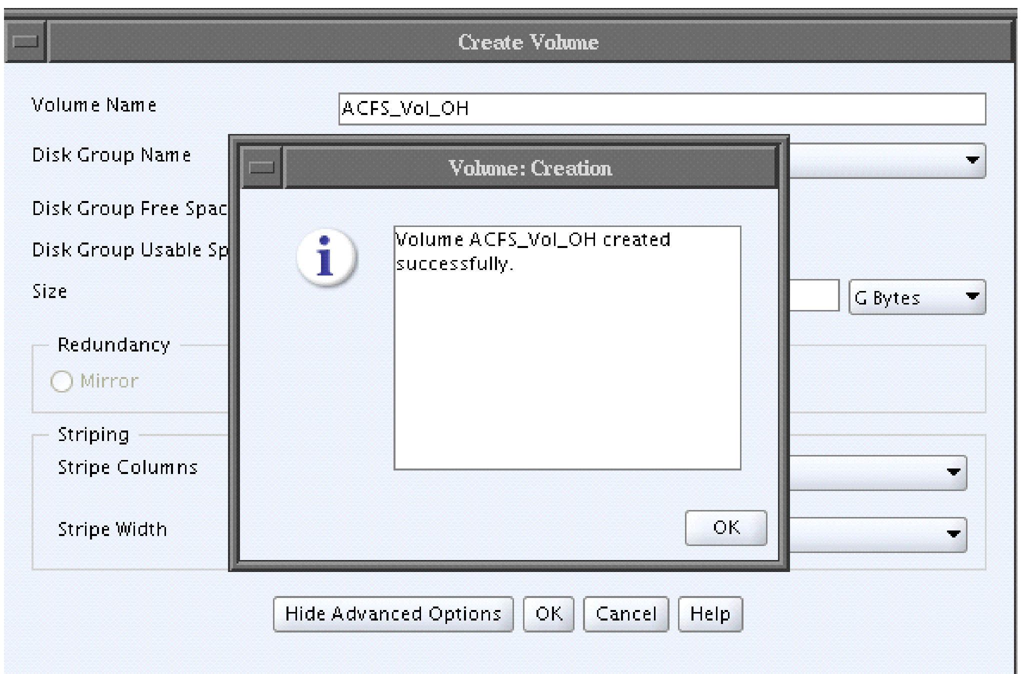
Create Volume
Volume Name
ACFS_Vol_OH
Disk Group Name
Volume: Creation
Volume ACFS_Vol_OH created successfully.
OK
G Bytes
Redundancy
Mirror
Striping
Stripe Columns
Stripe Width
Hide Advanced Options
OK
Cancel
Help

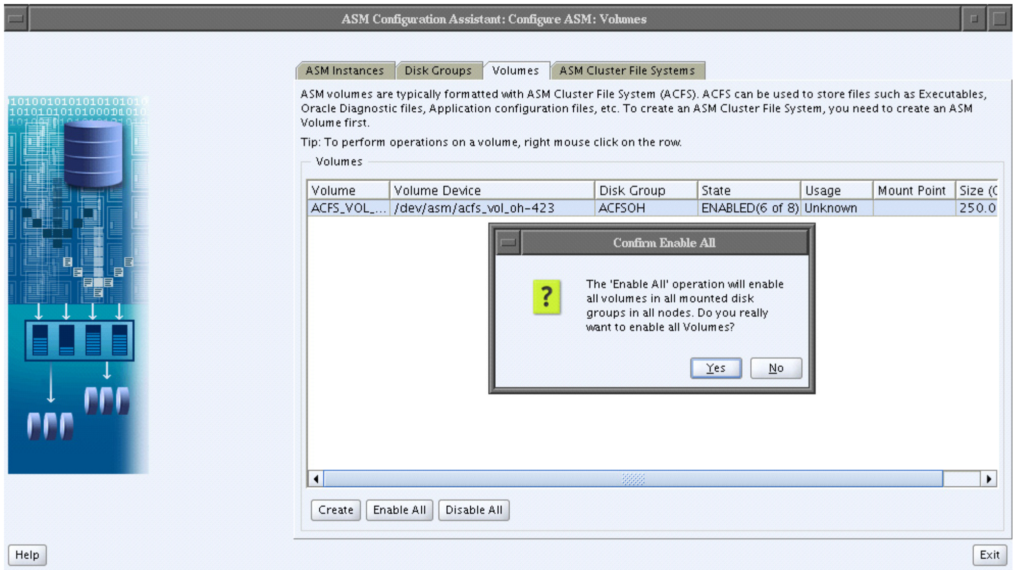

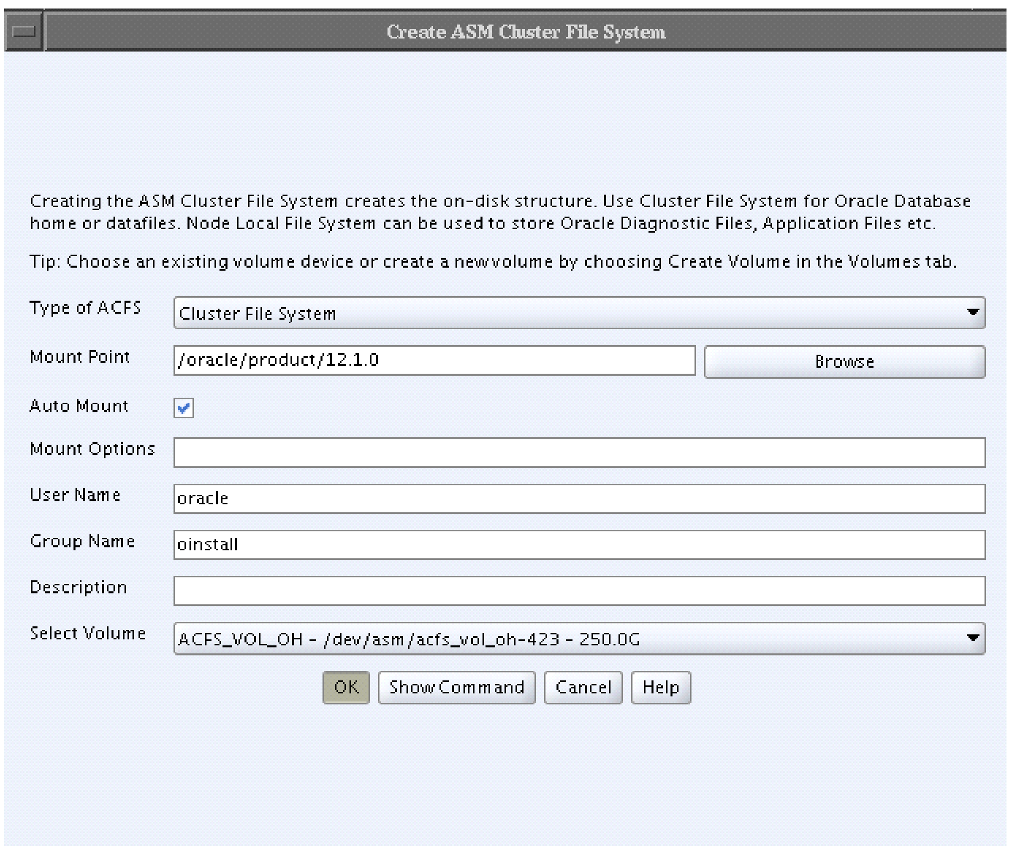

Run the registration scripts for auto-start if prompted.

When configured, reboot the nodes to validate that the ACFS Cluster file system is running automatically along with the other cluster services. Querying the cluster resources will reveal information as shown below:

```
ora.acfsoh.acfs_vol_oh.acfs
      1         ONLINE  ONLINE       ora12crac4               mounted on
/oracle/product/12.1.0,STABLE
      2         ONLINE  ONLINE       ora12crac2               mounted on
/oracle/product/12.1.0,STABLE
      3         ONLINE  ONLINE       ora12crac3               mounted on
/oracle/product/12.1.0,STABLE
      4         ONLINE  ONLINE       ora12crac1               mounted on
/oracle/product/12.1.0,STABLE
      5         ONLINE  ONLINE       ora12crac5               mounted on
/oracle/product/12.1.0,STABLE
      6         ONLINE  ONLINE       ora12crac6               mounted on
/oracle/product/12.1.0,STABLE
      7         ONLINE  ONLINE       ora12crac7               mounted on
/oracle/product/12.1.0,STABLE
      8         ONLINE  ONLINE       ora12crac8               mounted on
/oracle/product/12.1.0,STABLE
```

12c RDBMS Installation

The details of RDBMS install are not covered in this document. Using the shared Oracle Home mount point created above from ACFS, launch Oracle Installer and install the software. The database in the certification test bed was created with OAST tool kit and not through DBCA. DBCA could as well be used for creating the databases.

12c Client Installation

Oracle client for generating the load on the databases was done by installing 12c OAST toolkit. OAST is a load generating tool kit from oracle that is used for certifications. For details on this tool please contact oracle for details. However, any other testing tool openly available in the market and compatible with Oracle 12c could be used to stress the tool.

OAST tool was used to create the database and see data in it. Four databases were created.

Oastcpu -> Stress the CPU of all the nodes.

Oastcrs -> Stress the interconnect.

Oastiop -> OLTP type of database to stress the storage mostly on IOPS

Oastdss -> DSS type of database to stress the bandwidth of the system.

Oracle 12c New Features and Enhancements

Table 4lists some of the Oracle 12c new features and enhancements.

Table 4 Oracle 12c New Features and Enhancements

Multitenant Architecture (CDB's and PDB's)	Introduced in 12c
RAC (Real Application Clusters)	Application Continuity
Advanced Compression	Smart Compression
Data Guard	Far sync, Global Data Services
Real Application Testing	Consolidated Replays
Advanced Security	Data Redaction

Flex Clusters	Several Hub and Spoke nodes
Flex ASM	No dependency on ASM processes any more
RMAN enhancements	CDB/PDB backup and Recovery, duplicating Pluggable databases
Heat Map and Automatic Data Optimization	ILM Implementation Strategy
Redaction Policies	Enhanced Data Masking
Resource Manager	DBRM for Pluggable databases

Some of the above features either introduced or enhanced with 12c enable rapid provisioning, plug/unplug databases with minimal or no application level changes. The multitenacy with resource isolation reduces CAPEX and OPEX expenses along with faster provisioning. It's beyond the scope to explain all the features of 12c in this document. However few of them that were used in the certification exercise are mentioned below. Please refer Oracle online documentation for further details on all of these. A glimpse of these features provided here to give an overview before jumping into the stress and performance tests done with these features on the test bed.

Oracle Flex Clusters

Figure 18 Overview of Flex Clusters

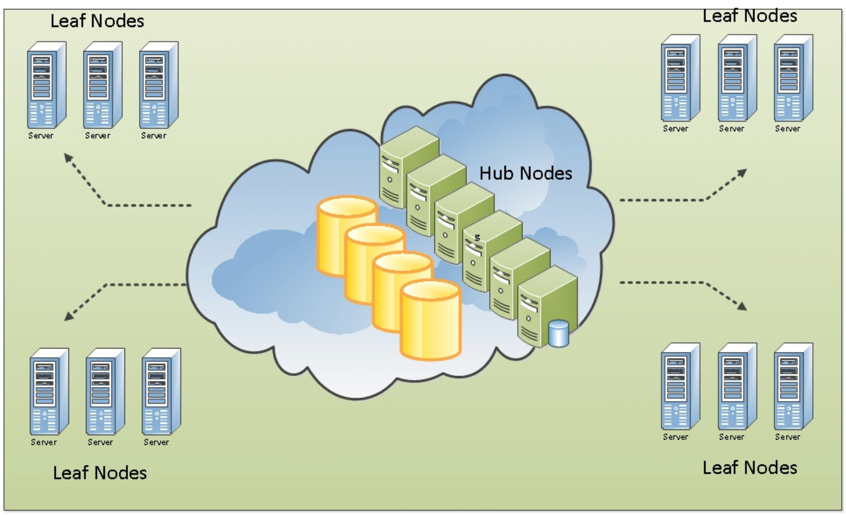

The traditional Oracle RAC clusters have been extended to flex clusters that could serve several applications including Oracle RAC for databases. By embracing flex clusters policy decisions are centralized to Central Grid Infrastructure. The hub and leaf nodes are connected and a leaf is attached to at least to one Hub. Hub nodes are same as earlier generation of RAC nodes and have access to shared storage while leaf nodes do not require access to shared storage and can run applications that can be controlled with Grid infrastructure.

Oracle Flex ASM

Figure 19 Overview of Flex ASM

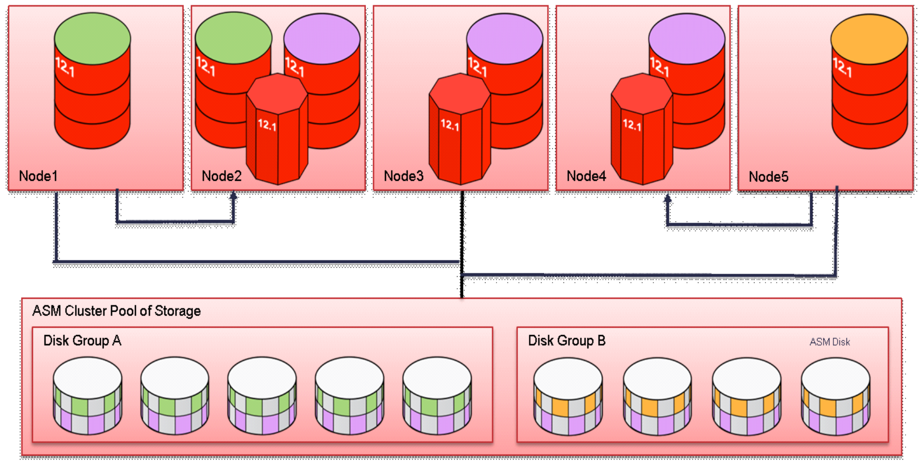

Prior to Oracle 12c, if ASM instance on one of the RAC nodes crashes, all the instances running on that node will crash too. This issue has been addressed in Flex ASM; Flex ASM can be used even if all the nodes are hub nodes. However, GNS configuration is mandatory for enabling Flex ASM.

Figure 19 illustrates ASM running only on nodes 1, 3 and 5 while nodes 2 and 4 are dependent on other nodes.

You can check what instances are connected with a simple query as shown below:

```
[oracle@ora12crac3 ~]$ srvctl status asm -detail
ASM is running on ora12crac1,ora12crac2,ora12crac5
ASM is enabled.
On Node1 - ora12crac1
SQL> @apxinfo

GROUP_NUMBER INSTANCE_NAME    DB_NAME    STATUS               CON_ID
------------ ---------------- ---------- --------------- ----------
           5 +ASM1            +ASM       CONNECTED                0
           2 +ASM1            +ASM       CONNECTED                0
           1 +APX1            +APX       CONNECTED                0
           1 +APX6            +APX       CONNECTED                0
           1 +APX8            +APX       CONNECTED                0
```

APX1(ora12crac1), APX6(ora12crac3) and APX8(ora12crac4) are connected to ASM1

The above result is from a simple query as below. Change the instance_name per your setup.

Select group_number,instance_name,db_name,status,con_id from v$asm_client where instance_name not like '%oast%';

The output indicates that the query that was run on node1 having +ASM1 running is connected to node3 and node4. Hence if +ASM1 instance on node dies, the APX connections will migrate to another node in the cluster running ASM processes. This was tested as part of ASM failover tests.

Oracle 12c Multitenant Architecture

Figure 20 Oracle 12c's Multitenant Architecture

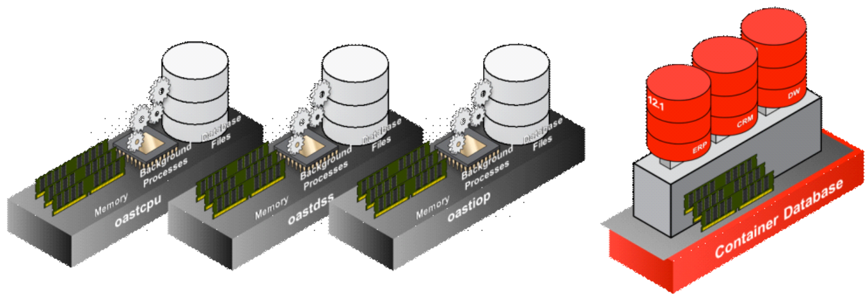

Oracle 12c Multitenant architecture encapsulates one or more databases that are contained in a single container. Each database in a container is called as Pluggable database that can be plugged and unplugged into/from a container whenever needed.

The diagram below, from Oracle's web site, provides an overview of Multitenant architecture in Oracle 12c.

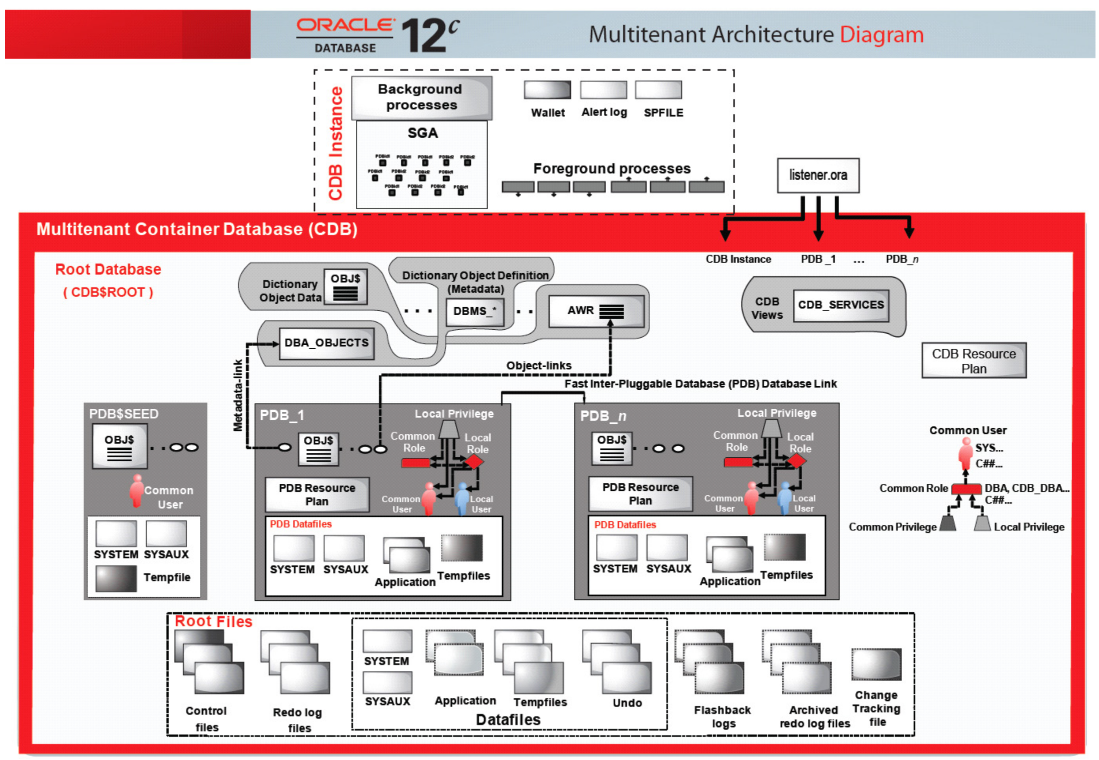

Each container has the following:

- Root - that stores Oracle supplied metadata. The root container is named cdb$root.
- Seed PDB - that is a seed PDB, a system supplied template and can be used to create new PDB's
- PDB's - One or more Pluggable databases which are user created databases

Container and Pluggable databases help consolidation of smaller databases to a single container; ease manageability and consolidation of the databases. While pluggable databases are not exactly like virtualized database, they provide similar functionality. We can have a single container database, with many databases plugged into them. There is a single operating system unlike virtualized mode where we have multiple OS for each VM. If a PDB database grows or shrinks or the server runs out of the resources it can be moved to another container database on another server. The backup and restore policies can be centralized and privilege controlled too and the Dev/QA/Test systems can be consolidated as well. May be using some of the rolling upgrade features, patches can be centralized and applied too. While cpu_count has been there for a while in Oracle RDBMS, Oracle 12c Resource Manager provides fine grain control of resources for PDB's. However, like any other Oracle feature, use cases could be different from one Organization to other and needs to be validated, checking both the functionality and compliance and best practices if any, with internal IT procedures, while implementing pluggable databases.

Creation of CDB

Container database is created first and is very similar to a conventional database of earlier versions. Create database …enable pluggable database instructs Oracle to create a CDB database.

Creation of PDB

Create pluggable database statement creates a PDB. This includes all the metadata and links necessary with cdb$root.

Figure 21 illustrates the creation of PDB.

Figure 21 PDB Database Creation Methods

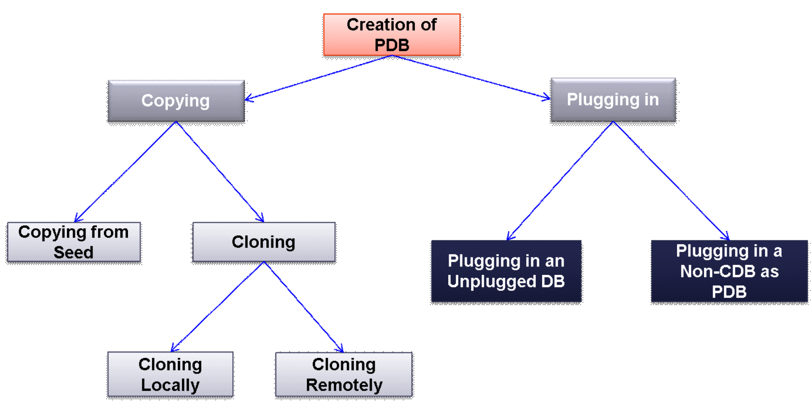

There are several methods you can use to migrate to a PDB database as shown in Figure 21. In the test bed for Oracle 12c certification, the cpu and dss databases were plugged into iops. Originally the database oastcpu, oastdss and oastiops were created as a single container databases with one PDB each. The CRS database was created as a Non-CDB. CPU and DSS were pluged to IOPS container. The following is a snippet of code used to create pluggable databases on the testbed:

```
Shutdown pluggable databases  - oastdss_pdb, oastiop_pdb, oastcpu_pdb
oastcpu - 40G sga and 27,264 MB buffer cache at 4096 block size
oastdss - 16G sga and 9,888 MB buffer cache at 16384 block size
oastiop - 24G sga and 14,656 MB buffer cache at 8192 block size

Connect to dss database
```

```
=======================
alter pluggable database oastdss_pdb unplug into
'/DataLoad/PDB/oastdss_pdb.xml';
drop pluggable database oastdss_pdb keep datafiles;

Connect to  cpu database
=======================
alter pluggable database oastcpu_pdb unplug into
'/DataLoad/PDB/oastcpu_pdb.xml';
drop pluggable database oastcpu_pdb keep datafiles;

Shutdown both dss and cpu databases
===================================
srvctl stop database -d oastdss
srvctl stop database -d oastcpu

Modify the init parameters accordingly
===================================
alter system set db_files=4000 scope=spfile sid='*';
alter system set db_4k_cache_size=27264M scope=spfile sid='*';
alter system set db_16k_cache_size=9888M scope=spfile sid='*';
alter system set db_cache_size=21504M scope=spfile sid='*';
alter system set sga_target=128G scope=spfile sid='*';
alter system set sga_max_size=128G scope=spfile sid='*';
alter system set memory_target=0 scope=spfile sid='*';
alter system set memory_max_target=0 scope=spfile sid='*';
alter system set shared_pool_size=24g scope=spfile sid='*';
alter system set pga_aggregate_target=30g scope=spfile sid='*';

Restart the oastiop database to reflect the spfile
==================================================
srvctl stop database -d oastiop
srvctl start database -d oastiop

Plug the dss and cpu databases to iop database
==============================================
create pluggable database oastdss_pdb using '/DataLoad/PDB/oastdss_pdb.xml'
nocopy tempfile reuse;
create pluggable database oastcpu_pdb using '/DataLoad/PDB/oastcpu_pdb.xml'
nocopy tempfile reuse;
```

Using Resource Manager with PDB

When PDB's are plugged into a single container, there is a need to effectively control the resources used by each PDB. Oracle 12c Resource Manager was used to control the resources of each PDB in the iops container.

```
alter system set resource_manager_plan='' scope=both sid='*';
exec DBMS_RESOURCE_MANAGER.CREATE_PENDING_AREA();
exec DBMS_RESOURCE_MANAGER.DELETE_CDB_PLAN(plan => 'cpu_oltp_dss_plan');
EXEC DBMS_RESOURCE_MANAGER.VALIDATE_PENDING_AREA();
EXEC DBMS_RESOURCE_MANAGER.SUBMIT_PENDING_AREA();

exec DBMS_RESOURCE_MANAGER.CREATE_PENDING_AREA();
 BEGIN
```

```
  DBMS_RESOURCE_MANAGER.CREATE_CDB_PLAN(
    plan => 'cpu_oltp_dss_plan',
    comment => 'CDB Resource plan for OAST stress minus crs');
END;
/
BEGIN
  DBMS_RESOURCE_MANAGER.CREATE_CDB_PLAN_DIRECTIVE(
    plan => 'cpu_oltp_dss_plan',
    pluggable_database => 'oastcpu_pdb',
    shares => 2,
    utilization_limit => 60,
    parallel_server_limit => 0);
END;
/
BEGIN
  DBMS_RESOURCE_MANAGER.CREATE_CDB_PLAN_DIRECTIVE(
    plan => 'cpu_oltp_dss_plan',
    pluggable_database => 'oastiop_pdb',
    shares => 3,
    utilization_limit => 35,
    parallel_server_limit => 0);
END;
/
BEGIN
  DBMS_RESOURCE_MANAGER.CREATE_CDB_PLAN_DIRECTIVE(
    plan => 'cpu_oltp_dss_plan',
    pluggable_database => 'oastdss_pdb',
    shares => 1,
    utilization_limit =>5,
    parallel_server_limit => 8);
END;
/
EXEC DBMS_RESOURCE_MANAGER.VALIDATE_PENDING_AREA();
EXEC DBMS_RESOURCE_MANAGER.SUBMIT_PENDING_AREA();
alter system set resource_manager_plan='cpu_oltp_dss_plan' scope=both
sid='*';

The following query can be leverage to check the resource utilization limits
in the database.
COLUMN PLAN HEADING 'Plan' FORMAT A26
COLUMN PLUGGABLE_DATABASE HEADING 'Pluggable|Database' FORMAT A25
COLUMN SHARES HEADING 'Shares' FORMAT 999
COLUMN UTILIZATION_LIMIT HEADING 'Utilization|Limit' FORMAT 999
COLUMN PARALLEL_SERVER_LIMIT HEADING 'Parallel|Server|Limit' FORMAT 999

SELECT PLAN,
         PLUGGABLE_DATABASE,
         SHARES,
         UTILIZATION_LIMIT,
         PARALLEL_SERVER_LIMIT
    FROM DBA_CDB_RSRC_PLAN_DIRECTIVES
   ORDER BY PLAN;
```

Oracle 12c Performance, Stress and Destructive Scenarios

Performance Tests

Separate Databases	Find out Peak Values through iteration				
	CPU	Max CPU			
	CRS		Max Interconnect		
	IOPS			Max IOPS	
	DSS				Max MBPS
	Simultaneously load all of the above	CPU Usage%	Interconnect in MBPS	IOPS	MBPS
Container Databases	Combine IOPS, DSS and CPU in a single Container as pluggable databases				
	Control resources through Resource Manager	CPU Usage%	Interconnect in MBPS	IOPS	MBPS

Stress Tests

48 hours stress test	Single container with 3 pluggable databases and one CRS database	Measure any drop in performance output when stressed at the peak values obtained from Performance tests.

Destructive Tests

A detailed list of destructive test cases attempted with the system running in full load, at its peak, is documented in the following sections.

Performance Data from the Test Bed

Orion Tests

Before conducting the Oracle Read Write tests using the OAST workload, orion tests were run, to capture the performance data from the test bed. A separate suite of tests were run on OLTP and DSS LUNs to understand the performance characteristics and guesstimate the IO throughput we can get from the test bed.

Figure 22 Orion Tests 100% Reads on OLTP LUNs

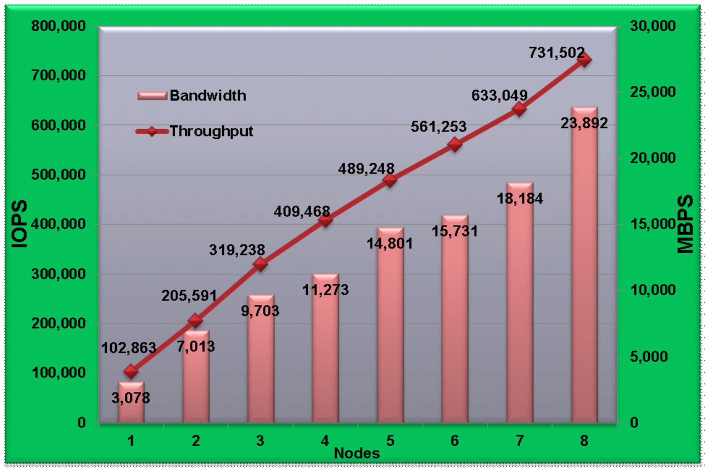

Figure 23 Orion Tests IOPS and Latency vs Read% on OLTP LUNs

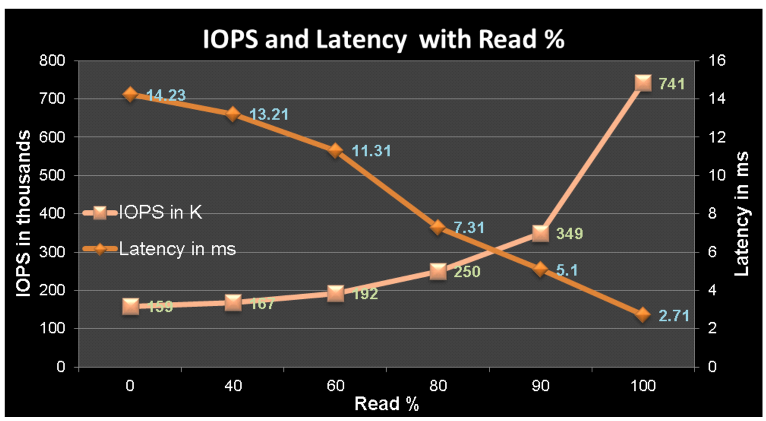

Figure 24 Orion Tests 30% Writes on OLTP LUNs

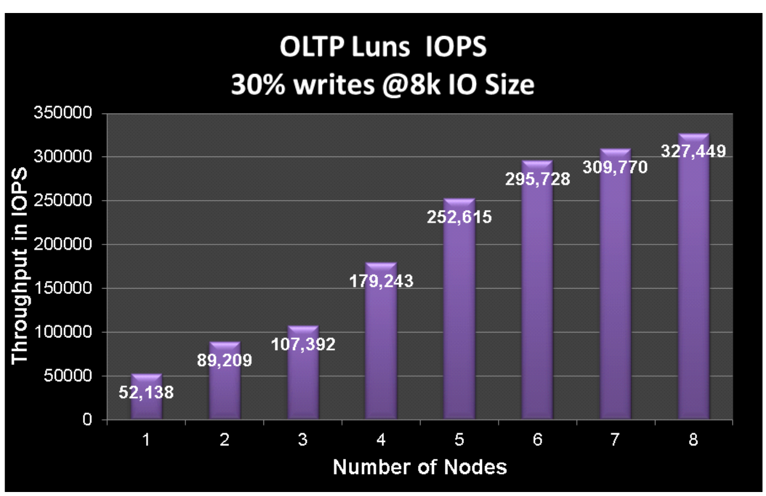

Figure 25 Orion Tests 10% Writes on DSS LUNs

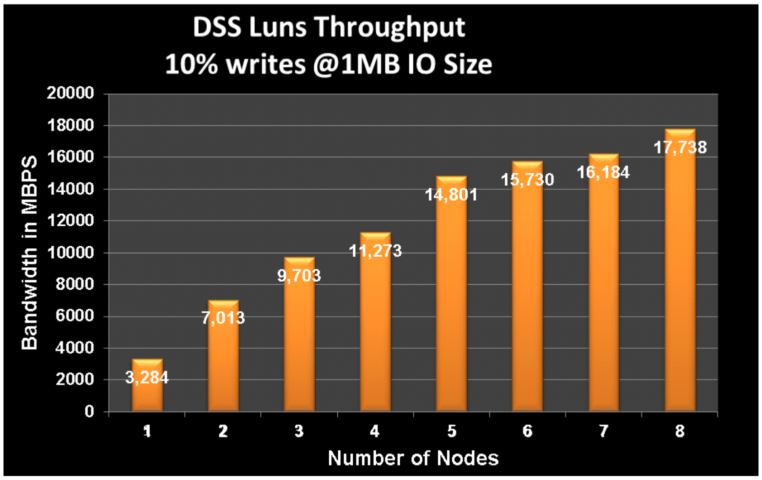

Summary

Orion tests revealed that system can go up to 750k IOPS with OLTP LUNs at 100% Reads, while it drops to around 300k with 30% writes. The latency increases considerably with write %. On the other hand DSS LUNs gave a bandwidth of 17GB/s of output. It was expected that there will be 30% writes on OLTP and 10% writes on DSS before starting the tests and hence the percentages chosen. While the above tests were run separately, a combined work load on both OLTP and DSS is expected to bring down the IOPS and MBPS figures further.

OAST Workloads

Each database was stressed to obtain the peak point before reducing them in a combined work load so that system sustains the load. The purpose was to max out the system on IO/CPU/Interconnect resources to the maximum balancing all of them. While this involved quite a number of iterations to achieve, only the final results are presented below. Out of the individual database tests only OLTP databases is presented here.

OLTP Workload

Data from Enterprise Manager 12c Grid Control

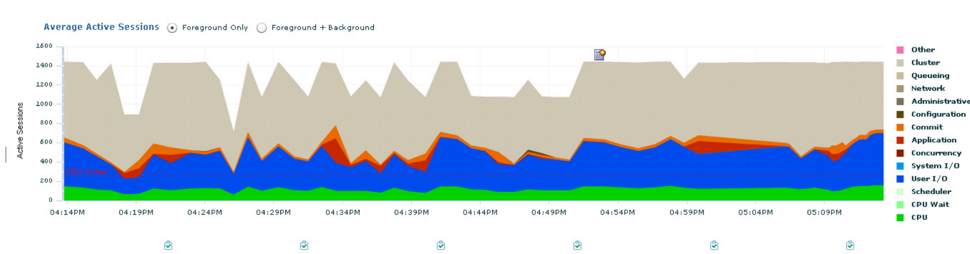

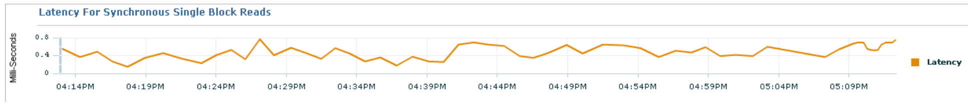

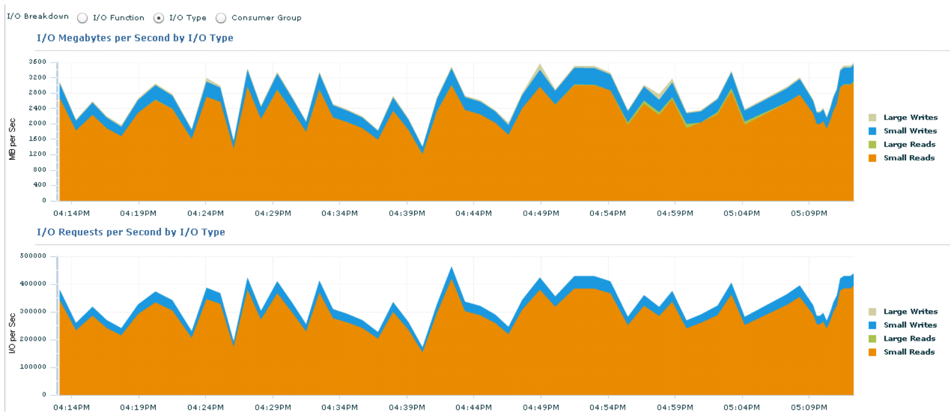

Iostat Data from One Out-of-Eight Nodes

Device:	r/s	w/s	rMB/s	wMB/s	avgrq-sz	avgqu-sz	await	svctm	%util
emcpoweraa1	3,412	313	26.66	3.00	16.3	4.95	1.33	0.26	96.4
emcpowerab1	3,379	301	26.40	3.05	16.39	4.56	1.24	0.26	96.07
emcpowerac1	3,406	327	26.61	3.21	16.36	4.99	1.34	0.26	96.8
emcpowerad1	3,480	324	27.19	3.29	16.41	4.78	1.26	0.25	96.7
emcpowerae1	3,503	317	27.36	3.00	16.28	4.99	1.31	0.25	96.7
emcpoweraf1	3,463	311	27.05	2.90	16.25	4.72	1.25	0.26	97.13
emcpowerag1	3,416	324	26.69	3.23	16.38	5.04	1.35	0.26	96.27
emcpowerah1	3,460	287	27.03	2.61	16.2	4.65	1.24	0.26	96.57
emcpowerai1	3,476	318	27.16	3.16	16.37	5.03	1.33	0.26	96.97
emcpoweraj1	3,448	311	26.94	3.16	16.4	4.72	1.26	0.26	96.37
emcpowerak1	3,514	327	27.45	3.48	16.49	5.15	1.34	0.25	96.43
emcpoweral1	3,489	317	27.26	3.12	16.35	4.79	1.26	0.25	96.17
emcpoweram1	3,424	320	26.75	3.04	16.29	5	1.34	0.26	96.73
emcpoweran1	3,508	309	27.41	3.36	16.5	4.75	1.25	0.25	97.17
emcpowerao1	3,528	329	27.56	3.11	16.29	5.18	1.34	0.25	96.77
emcpowerap1	3,440	325	26.88	3.09	16.3	4.77	1.27	0.26	96.23
Total	**55,347**	**5,059**	**432.40**	**49.81**					

AWR Data

System Statistics - Per Second

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	388,598.48	48,826.11	5,942.44	10,590.08	41,182.57	1,650.25	5,626.67	10.07	0.74	754.18
2	343,303.22	38,135.89	5,131.02	9,379.12	36,464.53	1,458.12	4,962.98	6.93	0.76	665.82
3	429,543.52	48,894.73	6,383.38	11,770.42	45,817.56	1,825.47	6,228.78	8.51	0.77	835.02
4	489,979.66	49,552.22	7,104.00	13,403.16	52,110.05	2,085.29	7,100.56	7.10	0.75	953.39
5	390,682.37	47,975.94	5,959.25	10,725.86	41,794.09	1,664.92	5,670.82	7.26	0.78	761.56
6	422,139.41	48,267.50	6,341.00	11,585.16	45,026.99	1,794.91	6,115.79	8.14	0.75	820.71
7	349,212.60	42,395.61	5,288.12	9,540.31	37,039.26	1,480.26	5,055.48	11.65	0.80	676.48
8	420,084.79	51,688.54	6,412.30	11,539.70	44,713.68	1,785.07	6,081.97	7.08	0.77	816.07
Sum	3,233,544.06	**375,736.54**	**48,561.53**	**88,533.83**	344,148.74	13,744.29	46,843.04	66.75	6.11	**6,283.23**
Avg	404,193.01	46,967.07	6,070.19	11,066.73	43,018.59	1,718.04	5,855.38	8.34	0.76	785.40
Std	47,408.19	4,435.96	641.38	1,307.56	5,084.63	202.95	689.89	1.70	0.02	93.03

	Reads MB/sec			Writes MB/sec				Reads requests/sec			Writes requests/sec			
I#	Total	Buffer Cache	Direct Reads	Total	DBWR	Direct Writes	LGWR	Total	Buffer Cache	Direct Reads	Total	DBWR	Direct Writes	LGWR
1	404.09	372.64	0.01	56.35	45.49	0.00	10.67	47,820.12	47,651.24	0.94	5,381.75	4,535.46	0.11	845.28
2	293.29	293.03	0.01	48.99	39.46	0.00	9.52	37,339.12	37,322.03	1.37	4,631.14	3,870.18	0.22	760.03
3	377.23	376.10	0.01	61.22	49.27	0.00	11.93	48,122.36	48,048.73	1.07	5,685.95	4,866.28	0.15	818.84
4	384.22	383.97	0.00	68.94	55.24	0.00	13.69	49,136.10	49,120.55	0.00	6,343.94	5,380.98	0.00	962.48
5	367.57	367.31	0.01	56.96	45.75	0.00	10.86	46,996.51	46,979.08	1.73	5,461.78	4,563.98	0.22	896.44
6	373.67	373.34	0.00	61.01	49.18	0.00	11.82	47,708.30	47,686.13	0.00	5,680.27	4,831.02	0.00	848.74
7	326.92	323.80	0.02	67.59	40.50	0.00	9.63	41,474.51	41,411.60	2.44	4,852.43	4,042.61	0.26	790.20
8	395.35	395.10	0.01	60.97	49.16	0.00	11.63	50,551.03	50,534.30	1.15	5,636.70	4,848.58	0.10	787.08
Sum	2,922.34	2,885.29	0.07	482.02	374.06	0.01	89.74	369,148.05	368,753.66	8.69	43,673.97	36,939.09	1.05	6,709.08
Avg	365.29	360.66	0.01	60.25	46.76	0.00	11.22	46,143.51	46,094.21	1.09	5,459.25	4,617.39	0.13	838.63

CPU Utilization Across Nodes

ora12crac1	Cpu(s):	**11.7%us,**	4.4%sy,	0.0%ni,	75.0%id,	**7.5%wa,**	0.0%hi,	1.4%si,	0.0%st
ora12crac2	Cpu(s):	**21.2%us,**	8.0%sy,	0.0%ni,	52.5%id,	**15.3%wa,**	0.0%hi,	2.9%si,	0.0%st
ora12crac3	Cpu(s):	**33.7%us,**	12.2%sy,	0.0%ni,	27.5%id,	**22.3%wa,**	0.0%hi,	4.3%si,	0.0%st
ora12crac4	Cpu(s):	**26.2%us,**	9.8%sy,	0.0%ni,	38.3%id,	**22.1%wa,**	0.0%hi,	3.7%si,	0.0%st
ora12crac5	Cpu(s):	**31.5%us,**	12.6%sy,	0.0%ni,	24.1%id,	**27.2%wa,**	0.0%hi,	4.6%si,	0.0%st
ora12crac6	Cpu(s):	**19.7%us,**	7.1%sy,	0.0%ni,	55.3%id,	**15.4%wa,**	0.0%hi,	2.5%si,	0.0%st
ora12crac7	Cpu(s):	**21.1%us,**	7.1%sy,	0.0%ni,	55.8%id,	**13.9%wa,**	0.0%hi,	2.1%si,	0.0%st
ora12crac8	Cpu(s):	**31.3%us,**	11.9%sy,	0.0%ni,	28.9%id,	**23.7%wa,**	0.0%hi,	4.3%si,	0.0%st

Interconnect Across Nodes

dstat	Time	recv	send	usr	sys	idl	wai	hiq	siq
ora12crac1	17:26:25	**56M**	**55M**	21	8	51	17	0	3
ora12crac2	17:26:25	**55M**	**55M**	25	9	43	19	0	3
ora12crac3	17:26:25	**55M**	**54M**	23	8	48	17	0	3
ora12crac4	17:26:25	**56M**	**57M**	23	9	47	18	0	3
ora12crac5	17:26:25	**56M**	**57M**	27	10	36	22	0	4
ora12crac6	17:26:25	**52M**	**50M**	20	7	55	16	0	3
ora12crac7	17:26:25	**53M**	**50M**	19	7	60	12	0	2
ora12crac8	17:26:25	**61M**	**61M**	31	12	32	21	0	5
	Total	**444M**	**439M**						

Summary

	Enterprise Manager	Iostat	AWR
Txn/s			6283 txn/s
Throughput	400,000 iops	483,248 iops	412,822 iops
Bandwidth	3200 MB/s	3,858 MB/s	3404 MB/s

Note The iostat is a point in time, while AWR is averaged over 10mts interval.

Discrete Databases

There are four databases one for CPU (oastcpu), one for IOPS (oastiop), one for DSS (oastdss) and one for CRS (oastcrs). All the four databases were loaded simultaneously to stress the system in all the fronts and in order to capture the peak performance from the system.

Oastcpu

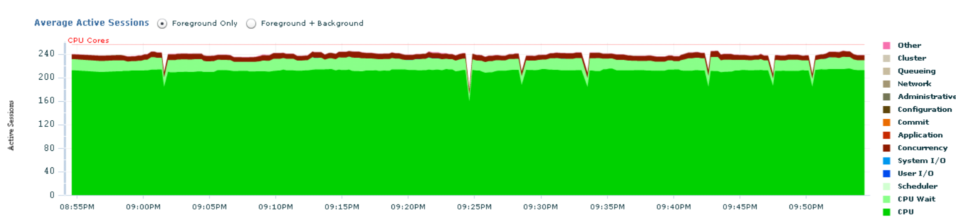

Oastiop

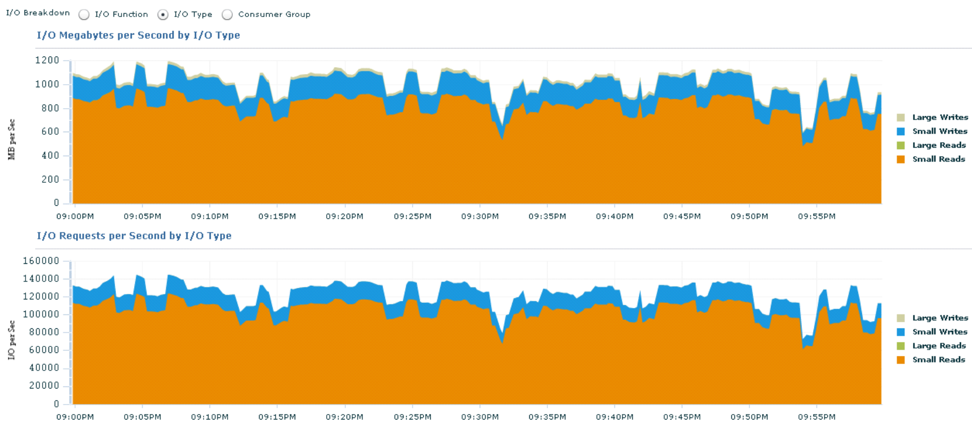

Oastdss

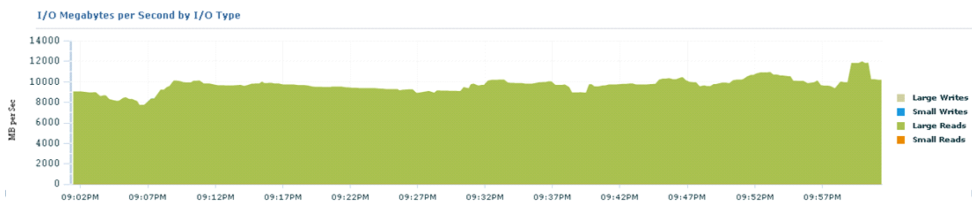

Oastcrs

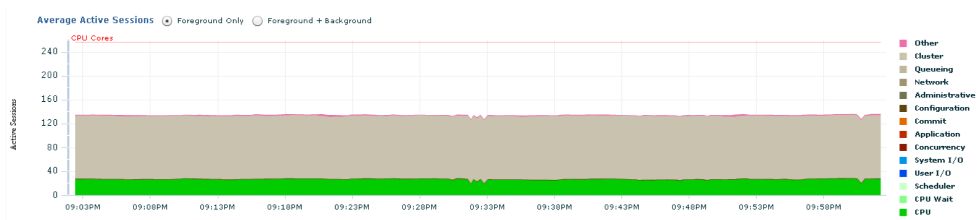

AWR Data - oastiop

	Reads MB/sec			Writes MB/sec				Reads requests/sec			Writes requests/sec			
I#	Total	Buffer Cache	Direct Reads	Total	DBWR	Direct Writes	LGWR	Total	Buffer Cache	Direct Reads	Total	DBWR	Direct Writes	LGWR
1	152.46	151.25	0.00	35.78	28.19	0.00	7.15	19,393.06	19,351.66	0.31	3,365.06	2,629.36	0.02	734.60
2	92.40	91.45	0.01	22.72	18.11	0.00	4.35	11,702.94	11,682.58	1.30	2,054.20	1,679.86	0.15	373.17
3	92.12	91.87	0.01	22.95	18.27	0.00	4.41	11,770.98	11,753.25	1.37	2,106.65	1,693.62	0.14	411.79
4	139.68	134.53	0.01	32.10	25.37	0.00	6.33	17,237.28	17,196.03	0.73	2,997.99	2,365.36	0.05	631.41
5	95.59	94.31	0.01	23.48	18.70	0.00	4.53	12,087.40	12,065.17	1.29	2,184.53	1,755.38	0.12	427.96
6	112.71	107.76	0.01	26.78	21.09	0.00	5.17	13,836.10	13,796.18	1.14	2,431.40	1,953.18	0.11	476.82
7	127.49	127.24	0.01	30.50	24.14	0.00	6.01	16,294.77	16,277.69	0.87	2,813.35	2,242.48	0.09	569.68
8	116.83	111.88	0.01	26.67	21.21	0.00	5.20	14,353.95	14,313.93	1.06	2,475.03	1,972.47	0.09	501.47
Sum	**929.29**	**910.28**	**0.06**	**220.97**	**175.08**	**0.01**	**43.16**	**116,676.49**	**116,436.49**	**8.08**	**20,428.21**	**16,291.72**	**0.76**	**4,126.92**
Avg	116.16	113.79	0.01	27.62	21.88	0.00	5.39	14,584.56	14,554.56	1.01	2,553.53	2,036.46	0.10	515.86

AWR Data - oastdss

	Reads MB/sec			Writes MB/sec				Reads requests/sec			Writes requests/sec			
I#	Total	Buffer Cache	Direct Reads	Total	DBWR	Direct Writes	LGWR	Total	Buffer Cache	Direct Reads	Total	DBWR	Direct Writes	LGWR
1	0.34	0.01	0.00	0.02	0.01	0.00	0.00	22.29	0.35	0.00	1.19	0.52	0.01	0.33
2	1,563.12	0.01	1,558.72	0.02	0.01	0.00	0.00	12,761.63	0.46	12,484.09	1.22	0.55	0.01	0.32
3	1,573.67	0.01	1,569.48	0.02	0.01	0.00	0.00	12,830.42	0.39	12,570.42	1.27	0.57	0.01	0.34
4	570.03	0.01	565.84	0.02	0.01	0.00	0.00	4,795.81	0.45	4,531.68	1.20	0.53	0.01	0.32
5	1,534.77	0.01	1,530.81	0.02	0.01	0.00	0.00	12,497.76	0.44	12,247.76	1.31	0.62	0.01	0.34
6	1,122.61	0.01	1,114.25	0.03	0.01	0.00	0.00	9,447.06	0.49	8,924.14	1.21	0.53	0.01	0.32
7	1,200.57	0.02	1,196.18	0.02	0.01	0.00	0.00	9,845.54	0.78	9,574.15	1.39	0.56	0.01	0.43
8	1,171.53	0.02	1,170.89	0.03	0.01	0.00	0.01	9,411.82	0.90	9,377.75	1.30	0.59	0.01	0.32
Sum	**8,736.62**	**0.10**	**8,706.17**	**0.19**	**0.09**	**0.00**	**0.04**	**71,612.33**	**4.26**	**69,710.00**	**10.11**	**4.47**	**0.05**	**2.70**
Avg	1,092.08	0.01	1,088.27	0.02	0.01	0.00	0.00	8,951.54	0.53	8,713.75	1.26	0.56	0.01	0.34

AWR Data – Interconnect

	Sent (MB/s)						Received (MB/s)					
I#	Total	Cache	IPQ	DLM	PNG	Misc	Total	Cache	IPQ	DLM	PNG	Misc
1	19.44	0.09	0.10	19.24	0.00	0.01	194.47	183.55	0.12	10.78	0.00	0.02
2	125.10	110.99	0.06	14.03	0.00	0.01	132.01	118.19	0.03	13.78	0.00	0.01
3	328.40	310.89	0.06	17.44	0.00	0.01	217.55	194.75	0.03	22.75	0.00	0.01
4	205.76	187.71	0.16	17.87	0.00	0.01	208.63	190.71	0.26	17.64	0.00	0.01
6	217.87	203.74	0.06	14.05	0.00	0.01	160.60	143.80	0.03	16.74	0.00	0.02
7	19.04	0.77	0.06	18.20	0.00	0.01	232.29	224.38	0.03	7.86	0.00	0.01
8	348.06	330.88	0.10	17.07	0.00	0.01	252.73	231.06	0.11	21.53	0.00	0.01
Sum	**1,263.66**	**1,145.08**	**0.61**	**117.91**	**0.02**	**0.05**	**1,398.28**	**1,286.44**	**0.63**	**111.09**	**0.02**	**0.10**
Avg	180.52	163.58	0.09	16.84	0.00	0.01	199.75	183.78	0.09	15.87	0.00	0.01
Std	133.58	134.09	0.04	2.03	0.00	0.00	41.65	40.73	0.09	5.44	0.00	0.00

AWR Data – Interconnect 2

The interconnect configuration and internode communication will influence the performance of cluster databases. The tables below show network interfaces on all hosts and network interfaces currently in use by cluster databases. It is important that cluster databases are configured to use a private interconnect for message and block transfers.

Private Interconnect Transfer Rate (MB/Sec) 2,577.466
Transfer rate on the private network in the last 5 minutes.

Interfaces by Hosts

View Private

Expand All | Collapse All

Name	Type	Subnet	Interface Type	Total I/O Rate (MB/Sec) (Last 5 Minutes)	Total Error Rate (%) (Last 5 Minutes)
ora12cCluster	Cluster				
ora12crac1.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	362.916	0
ora12crac2.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	599.503	0
ora12crac3.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	661.482	0
ora12crac4.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	438.928	0
ora12crac5.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	876.672	0
ora12crac6.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	936.324	0
ora12crac7.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	438.759	0
ora12crac8.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	819.366	0

OS Watcher

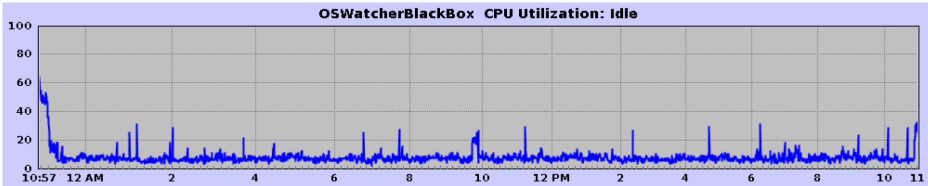

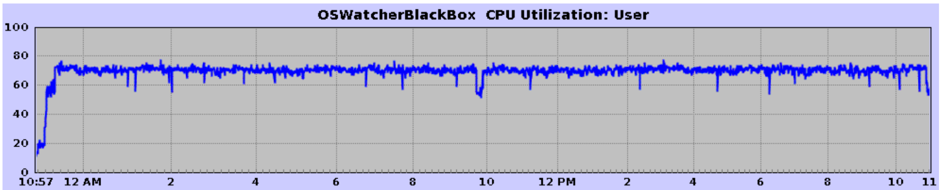

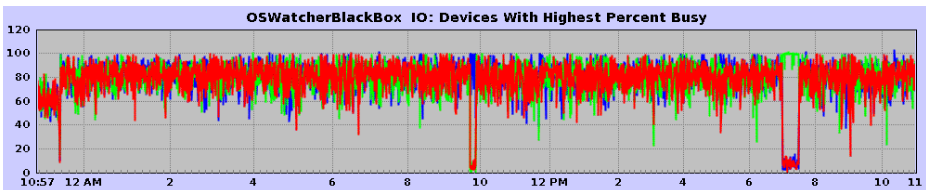

Summary

Mining the oast files for txn/s:

	Date	Time	Users	TPM
oastiops	1/29/2014	09:40:03:959354	150	164,551
	1/29/2014	09:40:18:780963	150	165,057
	1/29/2014	09:40:10:259751	150	163,541
	1/29/2014	09:40:00:255790	150	165,420
oastcpu	1/29/2014	09:40:04:143137	200	183,463
	1/29/2014	09:40:06:426822	200	183,099
	1/29/2014	09:40:18:839190	100	88,145
Total				**1,113,276**

System performance:

TPM from OAST Tool	**1,113,276**
IOPS	208,726
MBPS	9,887 MB/s
Interconnect Traffic	2,662 MB/s

Performance of the Pluggable Database

CPU, IOPS and DSS databases were plugged to single container database oastiop. The CRS database was left as separate non-CDB database.

For controlling the resources used across the pluggable databases, resource manager was used. It was bit iterative to arrive at the optimal stress of CPU and IO resources.

The performance data was collected with these two databases, as shown below:

EM Graphs

The EM graph clearly depicts the fall in IOPS and increase in MBPS when DSS workload kicks in. Data collected from ASM is shown below.

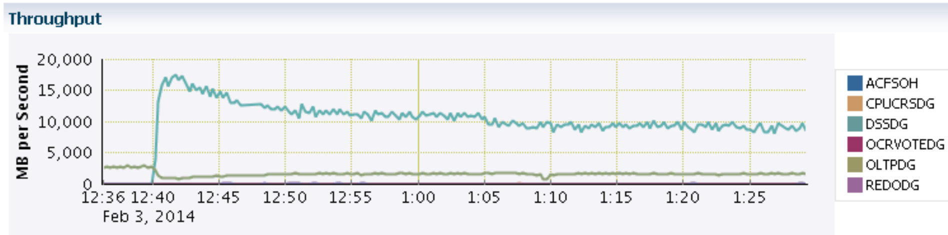
Throughput
MB per Second
20,000
15,000
10,000
5,000
0
12:36 12:40 12:45 12:50 12:55 1:00 1:05 1:10 1:15 1:20 1:25
Feb 3, 2014
ACFSOH
CPUCRSDG
DSSDG
OCRVOTEDG
OLTPDG
REDODG

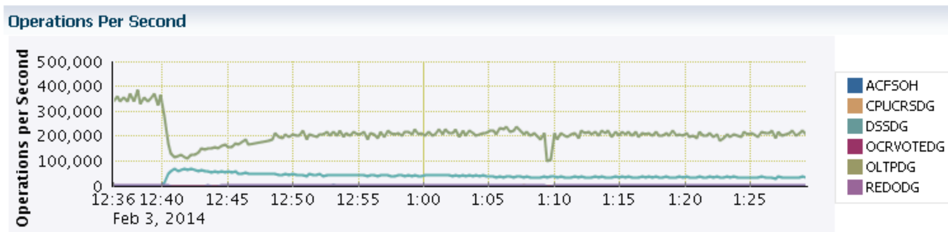
Operations Per Second
Operations per Second
500,000
400,000
300,000
200,000
100,000
0
12:36 12:40 12:45 12:50 12:55 1:00 1:05 1:10 1:15 1:20 1:25
Feb 3, 2014
ACFSOH
CPUCRSDG
DSSDG
OCRVOTEDG
OLTPDG
REDODG

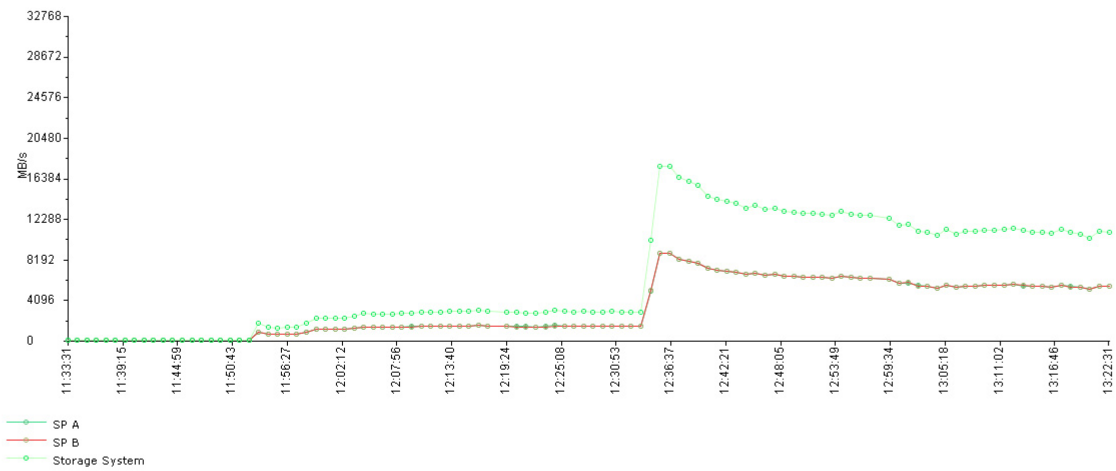
MB/s
32768
28672
24576
20480
16384
12288
8192
4096
0
11:33:31
11:39:15
11:44:59
11:50:43
11:56:27
12:02:12
12:07:56
12:13:40
12:19:24
12:25:08
12:30:53
12:36:37
12:42:21
12:48:05
12:53:49
12:59:34
13:05:18
13:11:02
13:16:46
13:22:31
SP A
SP B
Storage System

Bandwidth from VNX ~ 10,908 MB/s

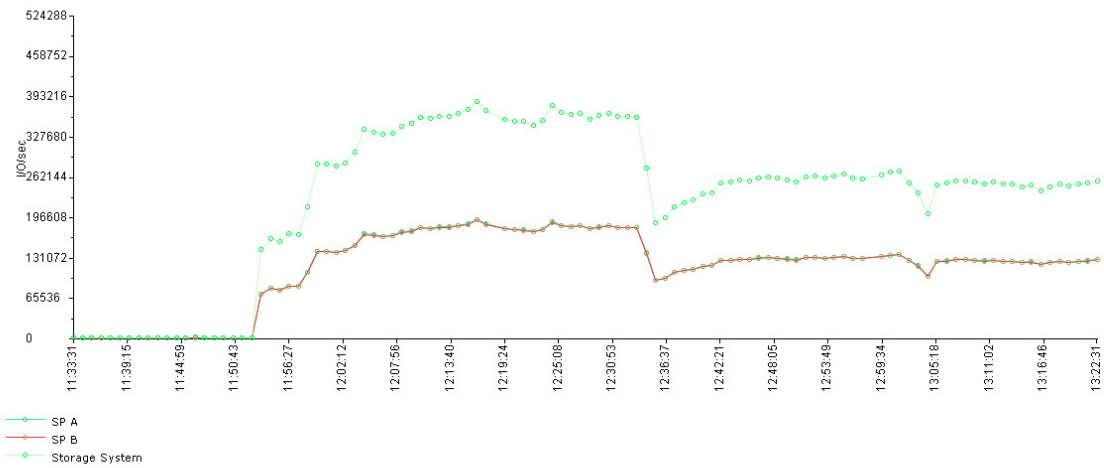

Throughput from VNX ~ 256,067

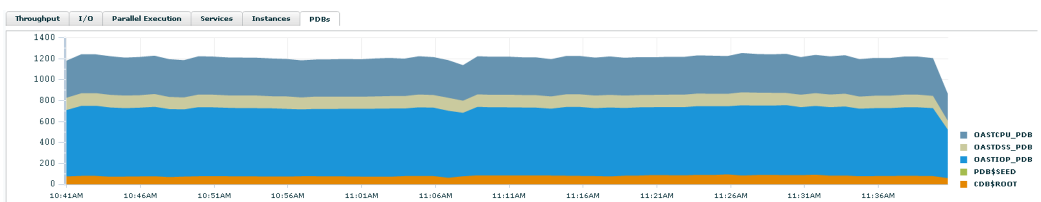

oratop 1: 1414 oast 12:57:57 up 2.0h, 8 ins, 813G mt, 779 sn, 4 us, 133% db

ID	%CU	HLD	MBPS	IOPS	%FR	PGAU	ASC	ASI	ASW	ASP	AAS	USN	TPS	UCPS	SSRT	DBC	DBW
1	40	12	1633	19k	35	3G	16	82	70	26	158.0	175	543	2675	19m	16	84
5	42	7	1508	22k	37	3G	16	74	68	16	154.4	165	574	2710	18m	18	82
4	43	8	1421	23k	37	3G	24	46	94	16	155.9	162	587	2685	19m	18	82
6	47	6	1062	25k	35	3G	31	53	66	10	142.1	153	651	3112	14m	21	79
2	46	7	186	23k	37	3G	20	34	95	1	137.6	144	681	3388	13m	21	79

EVENT (Real-Time)	AVG: TOT WAITS	TIME(s)	AVG MS	PCT	WAIT CLASS
DB CPU		798510		91	
wait for EMON to process ntf	7698	36977	4803.4	4	Configuration
log file parallel write	6779688	21003	3.1	2	System I/O
db file sequential read	7185876	12856	1.8	1	User I/O
db file parallel write	2014361	11244	5.6	1	System I/O

Interconnect

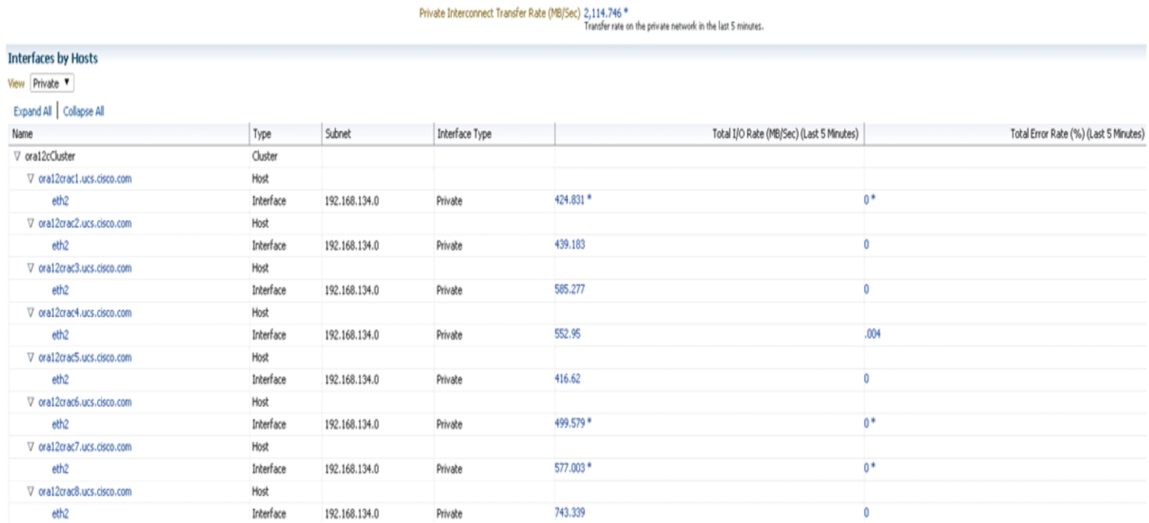
Private Interconnect Transfer Rate (MB/Sec) 2,114.746 *
Transfer rate on the private network in the last 5 minutes.

Interfaces by Hosts

View Private

Expand All | Collapse All

Name	Type	Subnet	Interface Type	Total I/O Rate (MB/Sec) (Last 5 Minutes)	Total Error Rate (%) (Last 5 Minutes)
ora12cCluster	Cluster				
ora12crac1.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	424.831 *	0 *
ora12crac2.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	439.183	0
ora12crac3.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	585.277	0
ora12crac4.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	552.95	.004
ora12crac5.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	416.62	0
ora12crac6.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	499.579 *	0 *
ora12crac7.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	577.003 *	0 *
ora12crac8.ucs.cisco.com	Host				
eth2	Interface	192.168.134.0	Private	743.339	0

System Performance

TPM	975,402
IOPS	**256,067**
MBPS	**10,908 MB/s**
Interconnect Traffic	**2,114 MB/s**

Compare Performance of Discrete and Pluggable Databases

Below is a snippet of the performance characteristics:

	Discrete	Pluggable
TPM	**1,113,276**	**975,402**
IOPS	**208,726**	**256,067**
MBPS	**9,887 MB/s**	**10,908 MB/s**
Interconnect Traffic	**2,662 MB/s**	**2,114 MB/s**

It appears that both Discrete and Pluggable databases performed very close in terms of performance. All the four variables listed above depend on each other. A drop in IOPS may be accompanied by an increase in TPM and vice-versa. It was very difficult to obtain and tune for consistency in the tests. Also, in order to control the resources, resource manager was used for pluggable databases which in turn had its own tuning effort.

Figure 26 CDB vs PDB Performance Comparison

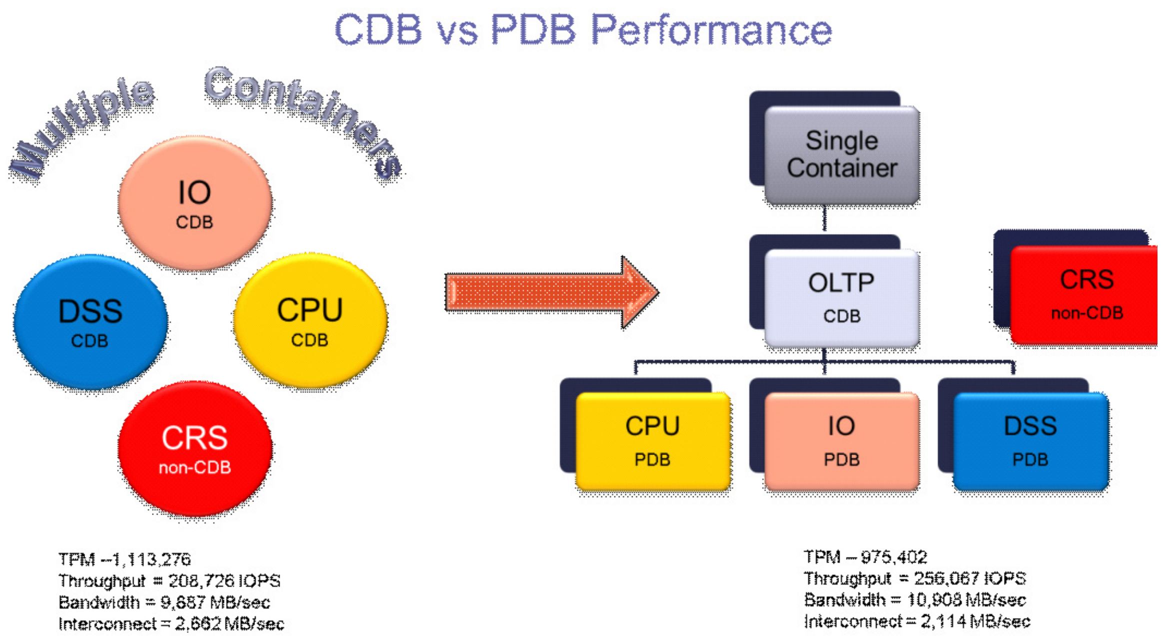

Stress Tests

A 36-48 hrs stress tests were run on the system in order to ascertain, that there is no substantial drop in performance and also to check for any errors in the alert log and other Oracle trace files, system message files etc.

Data was collected from Oracle, EMC frame and also from system files and no errors were observed.

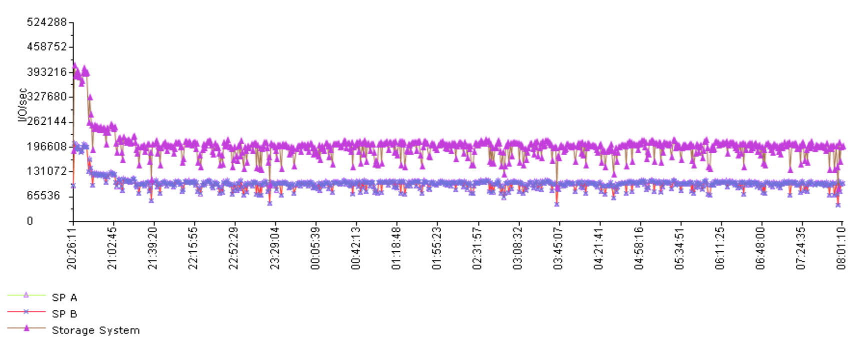

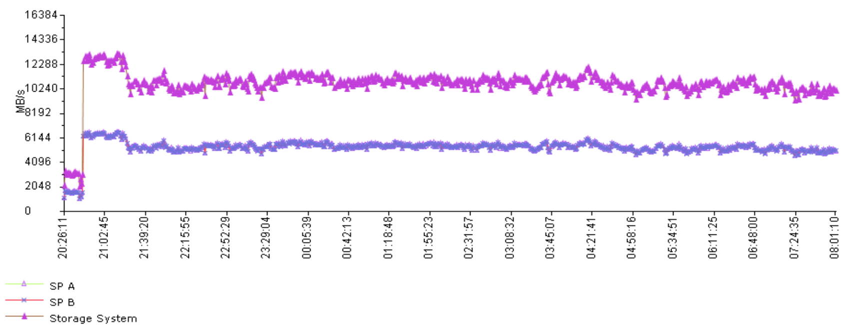

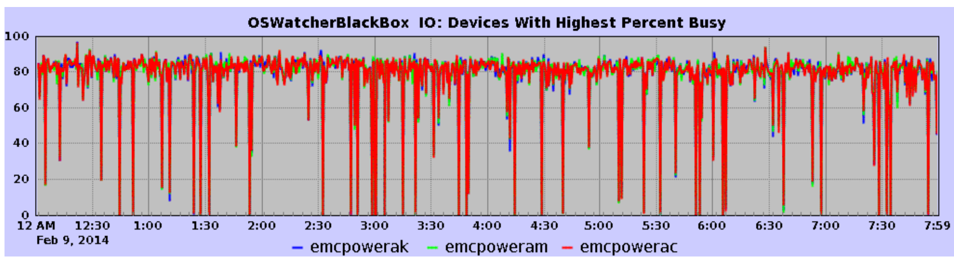

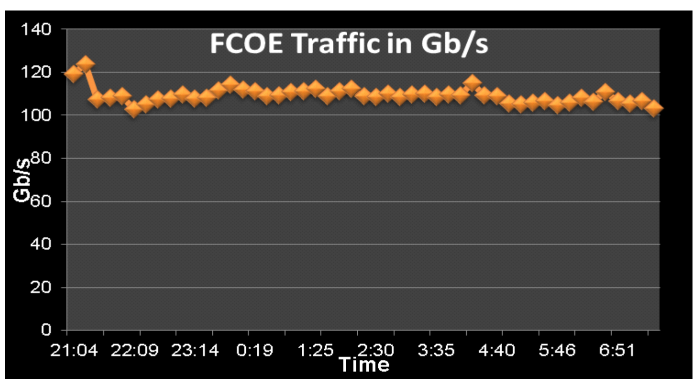

Database LUNs throughput during the 48 hrs run:

DB	r/s	w/s	rMB/s	wMB/s	Lun util%
OLTP	170,533	15,057	1,332	150	83
DSS	35,379	0	8,835	0	76
REDO	56	2,930	3	30	4
Total	**205,968**	**17,987**	**10,170**	**180**	

The performance of the system was consistent throughout without any errors in the log files. Care was of course taken to make sure that system is not running out of space during the 48 hrs stress run. The IOPS remained around 200k while the throughput around 10GB/s. The FCoE traffic captured from Cisco UCS Fabrics above, the iostat data from OS, the AWR data from Oracle and SP[A,B] data from EMC frame showed consistent results throughout the test.

Destructive and Hardware Failover Tests

Some of the destructive and hardware failover tests in this section were conducted on a fully loaded system (with both OLTP and DSS workload running) to check on the high availability and recoverability of the system when faults were injected. The test cases are listed in Figure 27.

Figure 27 Destructive Test Cases

	Test	Status
Test 1 – Multiple Network Connection Failures	Run the system on full work load, stressing all the components to their max. Disconnect the public links from first chassis and private links from second chassis one after other and reconnect each of them after 5 minutes.	Only second chassis servers rebooted. They joined the cluster back with a successful reconfigurations.
Test 2 – Network failure between FI and Corporate network	Run the system on full work load as above. Disconnect connection to N5K-A from Fabric A, wait 5 minutes, connect it back and repeat for N5K-B.	No disruption to the client or server traffic.
Test 3 – Fabric Failover tests	Run the system on full load as above. Reboot Fabric A, let it join the cluster back and then Reboot Fabric B.	Fabric failovers did not cause any disruption to ethernet and/or FC traffic.
Test 4 – Disconnect all FC Storage paths	Run the system on full load and disconnect the storage paths on the fabrics.	All nodes went for a reboot because of inaccessibility of voting files. All instances joined the cluster back later.
Test 5 – Blades Swap Test	Run the system on full load and swap out the blades one from each chassis.	The blades were configured with boot over SAN. The blades were rediscovered by the chassis and the nodes joined the cluster back.
Test 6 – ACFS Test 1	Kill the volume driver process vdbg under full load, on couple of nodes, that handles extents locks and other volume management functions.	DB instances crash as the underlying acfs files system is not available.
Test 7 – ACFS Test 2	Run the system on full load and enlarge the ACFS file system.	No disruption to Oracle client or BG processes.
Test 8 – ACFS Test 3	Run the system on full load and kill the ASM CKPT process.	Process re-spawned and workload continued.
Test 9 – Flex ASM failures	Run the system on full load and send sigkill to the ASM pmon process.	ASM process migrates to another node and there is no interruption to Oracle Work load.

Figure 28 illustrates some of the destructive tests.

Figure 28 ASM Destructive Test

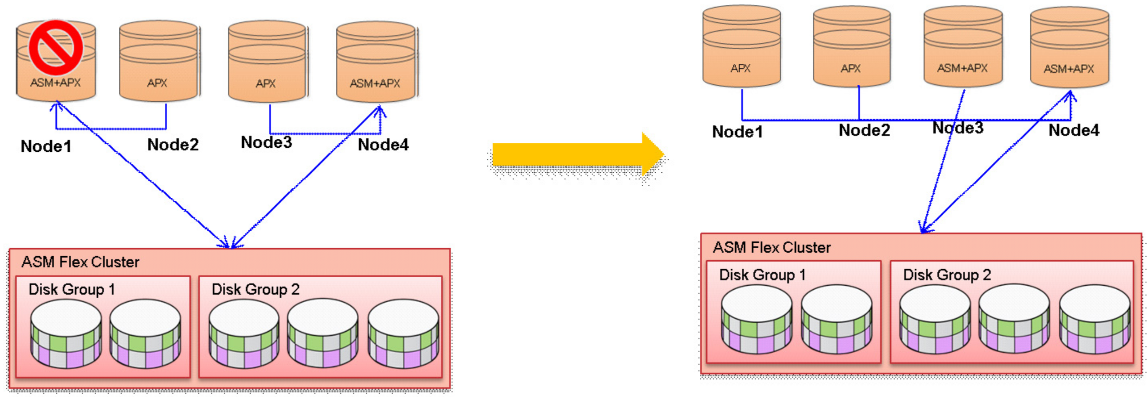

Test 2 - Network Connectivity Failures

Figure 29 N5K Failure Tests

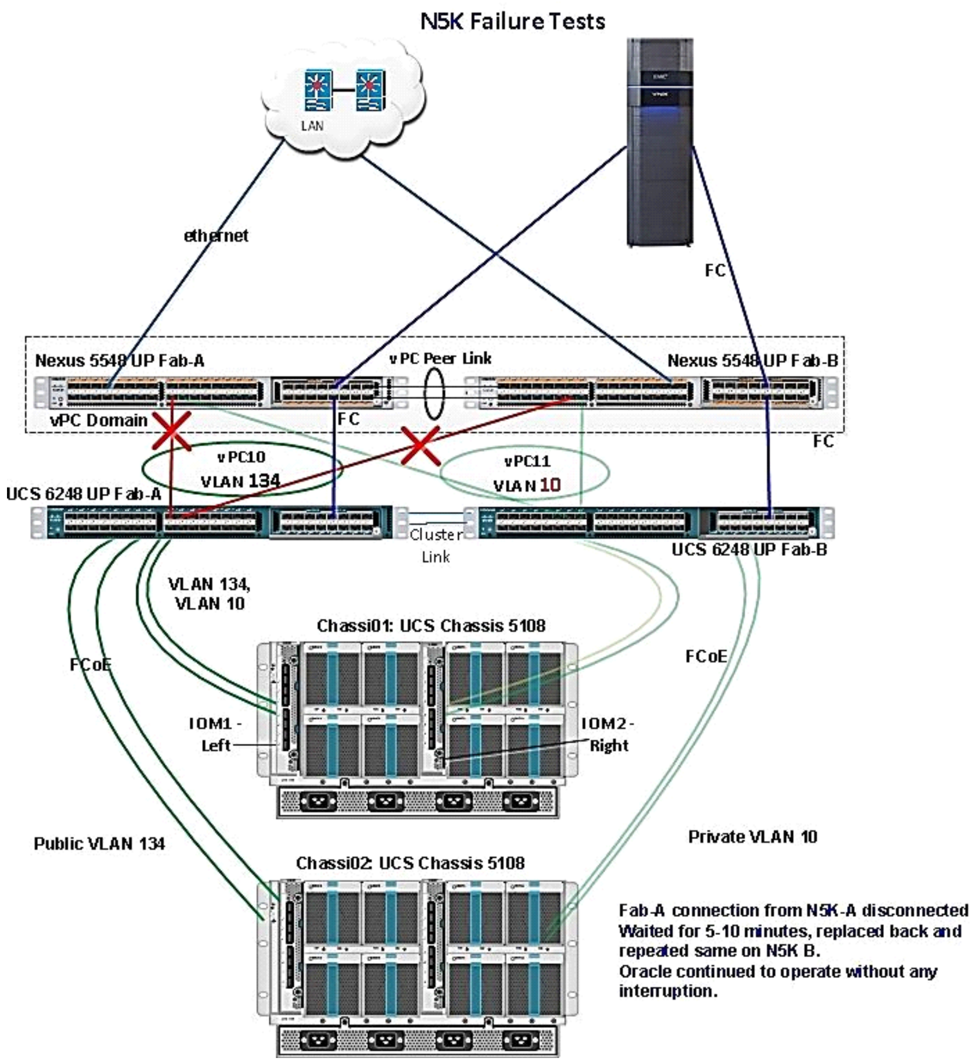

Test 3 - Fabric Failover Tests

Figure 30 Fabric Failover Tests

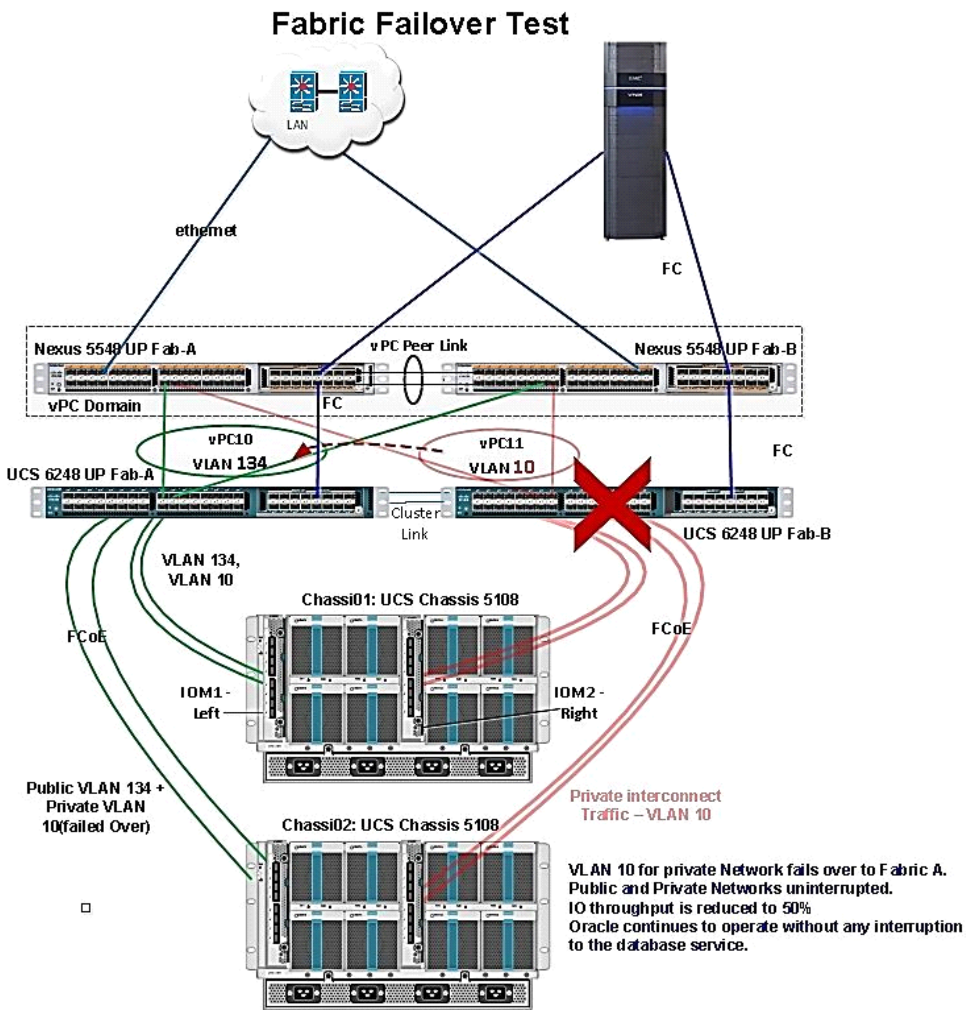

Test 4: Storage Path Failure Tests

Figure 31 Voting Disk Failures

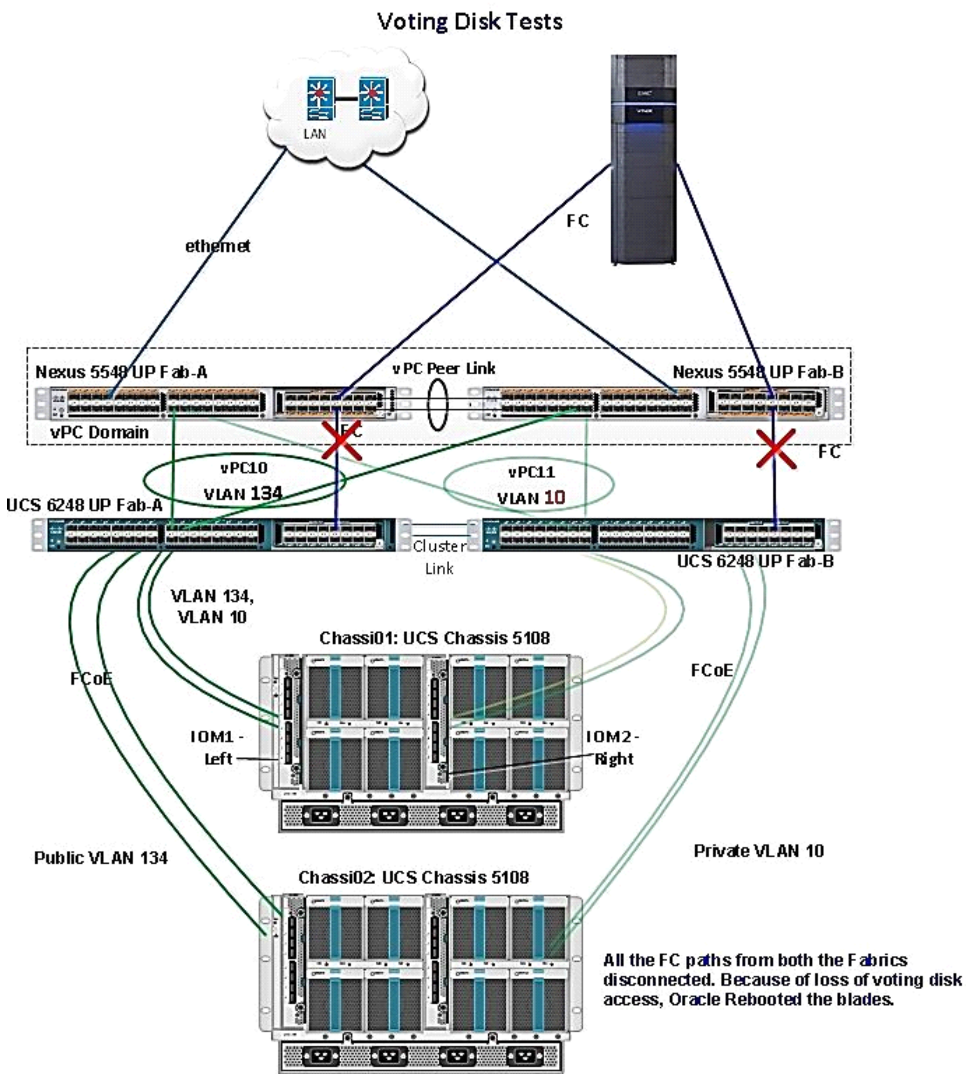

Lessons Learned and Best Practices for Running Oracle 12c on VNX 8000

Linux	· **If using yum update (latest) on uek kernel please make sure to check cisco Inter-operability matrix. Cisco Certified matrix should list the uek versions.** · **Use huge pages instead of AMM (Automatic Memory Management). Disable transparent huge pages preferably. Refer metalink note** · **1606759.1 for details.** · **Make sure you are using the right/latest Power Path for the kernel.**
Oracle	· May have to increase few parameters in flex ASM depending on the load (sudden bursts). · Controlling the resources across PDB's through resource manager is iterative. · Create PDB's does not do a space check. Being worked on by Oracle. Take due care while plugging and unplugging databases.
VNX	Upgrade to the latest firmware.
Cisco UCS	Discovery issues observed few times on Cisco UCS B420 M3. The blade gets hung up during the discovery phase; this issue is fixed in Cisco UCS Manager 2.2.

Appendix A: Cisco UCS Service Profiles

```
ORA-12C-FI-B# show fabric-interconnect

Fabric Interconnect:
    ID   OOB IP Addr     OOB Gateway     OOB Netmask     Operability
    ---- --------------- --------------- --------------- -----------
    A    10.29.134.8     10.29.134.1     255.255.255.0   Operable
    B    10.29.134.9     10.29.134.1     255.255.255.0   Operable
ORA-12C-FI-B# show fabric version
Fabric Interconnect A:
    Running-Kern-Vers: 5.0(3)N2(2.11f)
    Running-Sys-Vers: 5.0(3)N2(2.11f)
    Package-Vers: 2.1(1f)A
    Startup-Kern-Vers: 5.0(3)N2(2.11f)
    Startup-Sys-Vers: 5.0(3)N2(2.11f)
    Act-Kern-Status: Ready
    Act-Sys-Status: Ready
    Bootloader-Vers:   v3.5.0(02/03/2011)

Fabric Interconnect B:
    Running-Kern-Vers: 5.0(3)N2(2.11f)
    Running-Sys-Vers: 5.0(3)N2(2.11f)
    Package-Vers: 2.1(1f)A
    Startup-Kern-Vers: 5.0(3)N2(2.11f)
    Startup-Sys-Vers: 5.0(3)N2(2.11f)
    Act-Kern-Status: Ready
    Act-Sys-Status: Ready
    Bootloader-Vers:   v3.5.0(02/03/2011)

ORA-12C-FI-B# show server inventory
```

```
Server  Equipped PID Equipped VID Equipped Serial (SN) Slot Status      Ackd
Memory (MB) Ackd Cores
------- ------------ ------------ -------------------- ----------------
---------------- ----------
1/1     UCSB-B420-M3 V01         FCH16447HH2         Equipped
245760         32
1/2                                                     Equipped Not Pri
1/3     UCSB-B420-M3 V01         FCH16437BSM         Equipped
262144         32
1/4                                                     Equipped Not Pri
1/5     UCSB-B420-M3 V01         FCH161374EC         Equipped
262144         32
1/6                                                     Equipped Not Pri
1/7     UCSB-B420-M3 V01         FCH16437BTK         Equipped
262144         32
1/8                                                     Equipped Not Pri
2/1     UCSB-B420-M3 V01         FCH16447HEV         Equipped
262144         32
2/2                                                     Equipped Not Pri
2/3     UCSB-B420-M3 V01         FCH16447HF4         Equipped
262144         32
2/4                                                     Equipped Not Pri
2/5     UCSB-B420-M3 V01         FCH1623700J         Equipped
262144         32
2/6                                                     Equipped Not Pri
2/7     UCSB-B420-M3 V01         FCH16437BRS         Equipped
262144         32
2/8                                                     Equipped Not Pri


ORA-12C-FI-B# show service-profile inventory
Service Profile Name Type              Server  Assignment Association
-------------------- ----------------- ------- ---------- -----------

Ora12c_b420_1        Instance          1/1     Assigned   Associated
Ora12c_b420_2        Instance          1/3     Assigned   Associated
Ora12c_b420_3        Instance          1/5     Assigned   Associated
Ora12c_b420_4        Instance          1/7     Assigned   Associated
Ora12c_b420_5        Instance          2/1     Assigned   Associated
Ora12c_b420_6        Instance          2/3     Assigned   Associated
Ora12c_b420_7        Instance          2/5     Assigned   Associated
Ora12c_b420_8        Instance          2/7     Assigned   Associated
ORA12C_B420M3_1      Instance                  Unassigned Unassociated

ORA-12C-FI-B# show service-profile inventory expand
Service Profile Name: Ora12c-b420-template
Type: Updating Template
Server:
Description: Certification of Oracle 12c on b420 and VNX 8000
Assignment: Unassigned
Association: Unassociated

Service Profile Name: Ora12c-b420M3-Template
Type: Instance
Server:
Description: Oracle12c on UCS and EMC VNX 8000
Assignment: Unassigned
Association: Unassociated
```

```
Service Profile Name: Ora12c_b420_1
Type: Instance
Server: 1/1
Description:
Assignment: Assigned
Association: Associated
    Server 1/1:
        Name:
        Acknowledged Serial (SN): FCH16447HH2
        Acknowledged Product Name: Cisco UCS B420 M3
        Acknowledged PID: UCSB-B420-M3
        Acknowledged VID: V01
        Acknowledged Memory (MB): 245760
        Acknowledged Effective Memory (MB): 229376
        Acknowledged Cores: 32
        Acknowledged Adapters: 2

        Bios:
            Model: UCSB-B420-M3
            Revision: 0
            Serial:
            Vendor: Cisco Systems, Inc.

        Motherboard:
            Product Name: Cisco UCS B420 M3
            PID: UCSB-B420-M3
            VID: V01
            Vendor: Cisco Systems Inc
            Serial (SN): FCH16447HH2
            HW Revision: 0

            Array 1:
                DIMM Location   Presence         Overall Status
Type         Capacity (MB) Clock
                ---- ---------- ---------------- -----------------------
------------ ------------ -----
                   1 A0         Equipped         Operable
DDR3         16384         1600
                   2 A1         Missing          Removed
Undisc       Unknown       Unknown
                   3 A2         Missing          Removed
Undisc       Unknown       Unknown
                   4 B0         Equipped         Operable
DDR3         16384         1600
                   5 B1         Missing          Removed
Undisc       Unknown       Unknown
                   6 B2         Missing          Removed
Undisc       Unknown       Unknown
                   7 C0         Equipped         Operable
DDR3         16384         1600
                   8 C1         Missing          Removed
Undisc       Unknown       Unknown
                   9 C2         Missing          Removed
Undisc       Unknown       Unknown
                  10 D0         Equipped         Operable
DDR3         16384         1600
```

```
                 11 D1          Missing          Removed
Undisc       Unknown        Unknown
                 12 D2          Missing          Removed
Undisc       Unknown        Unknown
                 13 E0          Equipped         Operable
DDR3         16384          1600
                 14 E1          Missing          Removed
Undisc       Unknown        Unknown
                 15 E2          Missing          Removed
Undisc       Unknown        Unknown
                 16 F0          Equipped         Operable
DDR3         16384          1600
                 17 F1          Missing          Removed
Undisc       Unknown        Unknown
                 18 F2          Missing          Removed
Undisc       Unknown        Unknown
                 19 G0          Equipped         Operable
DDR3         16384          1600
                 20 G1          Missing          Removed
Undisc       Unknown        Unknown
                 21 G2          Missing          Removed
Undisc       Unknown        Unknown
                 22 H0          Equipped         Operable
DDR3         16384          1600
                 23 H1          Missing          Removed
Undisc       Unknown        Unknown
                 24 H2          Missing          Removed
Undisc       Unknown        Unknown
                 25 I0          Missing          Removed
Undisc       Unknown        Unknown
                 26 I1          Equipped         Disabled
Unknown      Unknown        Unknown
                 27 I2          Missing          Removed
Undisc       Unknown        Unknown
                 28 J0          Equipped         Operable
DDR3         16384          1600
                 29 J1          Missing          Removed
Undisc       Unknown        Unknown
                 30 J2          Missing          Removed
Undisc       Unknown        Unknown
                 31 K0          Equipped         Operable
DDR3         16384          1600
                 32 K1          Missing          Removed
Undisc       Unknown        Unknown
                 33 K2          Missing          Removed
Undisc       Unknown        Unknown
                 34 L0          Equipped         Operable
DDR3         16384          1600
                 35 L1          Missing          Removed
Undisc       Unknown        Unknown
                 36 L2          Missing          Removed
Undisc       Unknown        Unknown
                 37 M0          Equipped         Operable
DDR3         16384          1600
                 38 M1          Missing          Removed
Undisc       Unknown        Unknown
                 39 M2          Missing          Removed
Undisc       Unknown        Unknown
```

```
                    40 N0          Equipped         Operable
DDR3         16384         1600
                    41 N1          Missing          Removed
Undisc       Unknown       Unknown
                    42 N2          Missing          Removed
Undisc       Unknown       Unknown
                    43 O0          Missing          Removed
Undisc       Unknown       Unknown
                    44 O1          Missing          Removed
Undisc       Unknown       Unknown
                    45 O2          Missing          Removed
Undisc       Unknown       Unknown
                    46 P0          Equipped         Operable
DDR3         16384         1600
                    47 P1          Missing          Removed
Undisc       Unknown       Unknown
                    48 P2          Missing          Removed
Undisc       Unknown       Unknown

            CPUs:
                ID: 1
                Presence: Equipped
                Architecture: Xeon
                Socket: CPU1
                Cores: 8
                Speed (GHz): 2.700000
                Stepping: 7
                Product Name: Intel(R) Xeon(R) E5-4650
                PID: UCS-CPU-E5-4650
                VID: 01
                Vendor: Intel(R) Corporation
                HW Revision: 0

                ID: 2
                Presence: Equipped
                Architecture: Xeon
                Socket: CPU2
                Cores: 8
                Speed (GHz): 2.700000
                Stepping: 7
                Product Name: Intel(R) Xeon(R) E5-4650
                PID: UCS-CPU-E5-4650
                VID: 01
                Vendor: Intel(R) Corporation
                HW Revision: 0

                ID: 3
                Presence: Equipped
                Architecture: Xeon
                Socket: CPU3
                Cores: 8
                Speed (GHz): 2.700000
                Stepping: 7
                Product Name: Intel(R) Xeon(R) E5-4650
                PID: UCS-CPU-E5-4650
                VID: 01
                Vendor: Intel(R) Corporation
                HW Revision: 0
```

```
            ID: 4
            Presence: Equipped
            Architecture: Xeon
            Socket: CPU4
            Cores: 8
            Speed (GHz): 2.700000
            Stepping: 7
            Product Name: Intel(R) Xeon(R) E5-4650
            PID: UCS-CPU-E5-4650
            VID: 01
            Vendor: Intel(R) Corporation
            HW Revision: 0

        RAID Controller 1:
            Type: SAS
            Vendor: LSI Logic   Symbios Logic
            Model: LSI MegaRAID SAS 2208 ROMB
            Serial: LSIROMB-0
            HW Revision: 0
            PCI Addr: 01:00.0
            Raid Support: RAID0, RAID1, RAID5, RAID10

           .........
            Local Disk Config Definition:
                Mode: No Local Storage
                Description:
                Protect Configuration: No

    Adapter:

    Adapter PID          Vendor            Serial       Overall Status
    ------- ------------ ----------------- ------------ --------------
          1 UCSB-MLOM-40G-01
                         Cisco Systems Inc FCH1645731H  Operable
          3 UCS-VIC-M82-8P
                         Cisco Systems Inc FCH164471JT  Operable
```

Appendix B: N5K Zone Definitions

N5K-A

```
ora12c-n5k-a(config)# show zones active
zoneset name ora12c_FI_A vsan 15
  zone name ora12c_b420_1_hba1 vsan 15
  * fcid 0x7413ef [device-alias A0P0]
  * fcid 0x7405ef [device-alias B0P0]
  * fcid 0x740001 [pwwn 20:00:00:25:b5:00:00:6e]

  zone name ora12c_b420_1_hba3 vsan 15
  * fcid 0x740fef [device-alias A0P1]
  * fcid 0x7407ef [device-alias B0P1]
  * fcid 0x740063 [pwwn 20:00:00:25:b5:00:00:3d]
```

```
zone name ora12c_b420_2_hba1 vsan 15
* fcid 0x7413ef [device-alias A0P0]
* fcid 0x7405ef [device-alias B0P0]
* fcid 0x740021 [pwwn 20:00:00:25:b5:00:00:5d]

zone name ora12c_b420_2_hba3 vsan 15
* fcid 0x740041 [pwwn 20:00:00:25:b5:00:00:0d]
* fcid 0x740fef [device-alias A0P1]
* fcid 0x7407ef [device-alias B0P1]

zone name ora12c_b420_3_hba1 vsan 15
* fcid 0x7411ef [device-alias A0P2]
* fcid 0x7406ef [device-alias B0P2]
* fcid 0x740005 [pwwn 20:00:00:25:b5:00:00:3c]

zone name ora12c_b420_4_hba1 vsan 15
* fcid 0x7406ef [device-alias B0P2]
* fcid 0x7411ef [device-alias A0P2]
* fcid 0x740082 [pwwn 20:00:00:25:b5:00:00:1b]

zone name ora12c_b420_4_hba3 vsan 15
* fcid 0x7404ef [device-alias B0P3]
* fcid 0x7414ef [device-alias A0P3]
* fcid 0x740081 [pwwn 20:00:00:25:b5:00:00:0e]

zone name ora12c_b420_6_hba3 vsan 15
* fcid 0x7410ef [device-alias B1P1]
* fcid 0x7400ef [device-alias A1P1]
* fcid 0x740161 [pwwn 20:00:00:25:b5:00:00:0f]

zone name ora12c_b420_5_hba1 vsan 15
* fcid 0x7401ef [device-alias A1P0]
* fcid 0x7415ef [device-alias B1P0]
* fcid 0x740003 [pwwn 20:00:00:25:b5:00:00:59]

zone name ora12c_b420_3_hba3 vsan 15
* fcid 0x7404ef [device-alias B0P3]
* fcid 0x7414ef [device-alias A0P3]
* fcid 0x7400c1 [pwwn 20:00:00:25:b5:00:00:6d]

zone name ora12c_b420_5_hba3 vsan 15
* fcid 0x7400ef [device-alias A1P1]
* fcid 0x7410ef [device-alias B1P1]
* fcid 0x7400a1 [pwwn 20:00:00:25:b5:00:00:2e]

zone name ora12c_b420_6_hba1 vsan 15
* fcid 0x7415ef [device-alias B1P0]
* fcid 0x7401ef [device-alias A1P0]
* fcid 0x740181 [pwwn 20:00:00:25:b5:00:00:38]

zone name ora12c_b420_7_hba1 vsan 15
* fcid 0x7403ef [device-alias A1P2]
* fcid 0x7412ef [device-alias B1P2]
* fcid 0x740183 [pwwn 20:00:00:25:b5:00:00:27]

zone name ora12c_b420_7_hba3 vsan 15
* fcid 0x7402ef [device-alias A1P3]
```

```
    * fcid 0x7416ef [device-alias B1P3]
    * fcid 0x740141 [pwwn 20:00:00:25:b5:00:00:2f]

    zone name ora12c_b420_8_hba1 vsan 15
    * fcid 0x7412ef [device-alias B1P2]
    * fcid 0x7403ef [device-alias A1P2]
    * fcid 0x740007 [pwwn 20:00:00:25:b5:00:00:16]

    zone name ora12c_b420_8_hba3 vsan 15
    * fcid 0x7416ef [device-alias B1P3]
    * fcid 0x7402ef [device-alias A1P3]
    * fcid 0x740002 [pwwn 20:00:00:25:b5:00:00:4f]
```

N5K-B

```
ora12c-n5k-b(config)# show zones active
zoneset name ora12c_FI_B vsan 15
  zone name ora12c_b420_1_hba4 vsan 15
  * fcid 0x7f0023 [pwwn 20:00:00:25:b5:00:00:1d]
  * fcid 0x7f0eef [device-alias A7P1]
  * fcid 0x7f0aef [device-alias B7P1]

  zone name ora12c_b420_1_hba2 vsan 15
  * fcid 0x7f0021 [pwwn 20:00:00:25:b5:00:00:4d]
  * fcid 0x7f18ef [device-alias A7P0]
  * fcid 0x7f0bef [device-alias B7P0]

  zone name ora12c_b420_2_hba2 vsan 15
  * fcid 0x7f0001 [pwwn 20:00:00:25:b5:00:00:1f]
  * fcid 0x7f18ef [device-alias A7P0]
  * fcid 0x7f0bef [device-alias B7P0]

  zone name ora12c_b420_4_hba4 vsan 15
  * fcid 0x7f19ef [device-alias A7P3]
  * fcid 0x7f08ef [device-alias B7P3]
  * fcid 0x7f0101 [pwwn 20:00:00:25:b5:00:00:3e]

  zone name ora12c_b420_3_hba2 vsan 15
  * fcid 0x7f09ef [device-alias B7P2]
  * fcid 0x7f11ef [device-alias A7P2]
  * fcid 0x7f0004 [pwwn 20:00:00:25:b5:00:00:2c]

  zone name ora12c_b420_6_hba4 vsan 15
  * fcid 0x7f0fef [device-alias A8P1]
  * fcid 0x7f10ef [device-alias B8P1]
  * fcid 0x7f0181 [pwwn 20:00:00:25:b5:00:00:3f]

  zone name ora12c_b420_4_hba2 vsan 15
  * fcid 0x7f11ef [device-alias A7P2]
  * fcid 0x7f09ef [device-alias B7P2]
  * fcid 0x7f00c2 [pwwn 20:00:00:25:b5:00:00:0b]

  zone name ora12c_b420_5_hba2 vsan 15
  * fcid 0x7f0002 [pwwn 20:00:00:25:b5:00:00:69]
  * fcid 0x7f15ef [device-alias A8P0]
  * fcid 0x7f16ef [device-alias B8P0]

  zone name ora12c_b420_3_hba4 vsan 15
```

```
* fcid 0x7f08ef [device-alias B7P3]
* fcid 0x7f19ef [device-alias A7P3]
* fcid 0x7f00c1 [pwwn 20:00:00:25:b5:00:00:1e]

zone name ora12c_b420_2_hba4 vsan 15
* fcid 0x7f0008 [pwwn 20:00:00:25:b5:00:00:6c]
* fcid 0x7f0eef [device-alias A7P1]
* fcid 0x7f0aef [device-alias B7P1]

zone name ora12c_b420_6_hba2 vsan 15
* fcid 0x7f01a1 [pwwn 20:00:00:25:b5:00:00:48]
* fcid 0x7f15ef [device-alias A8P0]
* fcid 0x7f16ef [device-alias B8P0]

zone name ora12c_b420_5_hba4 vsan 15
* fcid 0x7f0fef [device-alias A8P1]
* fcid 0x7f10ef [device-alias B8P1]
* fcid 0x7f0081 [pwwn 20:00:00:25:b5:00:00:5e]

zone name ora12c_b420_7_hba2 vsan 15
* fcid 0x7f12ef [device-alias A8P2]
* fcid 0x7f13ef [device-alias B8P2]
* fcid 0x7f01a3 [pwwn 20:00:00:25:b5:00:00:37]

zone name ora12c_b420_7_hba4 vsan 15
* fcid 0x7f0141 [pwwn 20:00:00:25:b5:00:00:5f]
* fcid 0x7f1aef [device-alias A8P3]
* fcid 0x7f17ef [device-alias B8P3]

zone name ora12c_b420_8_hba2 vsan 15
* fcid 0x7f12ef [device-alias A8P2]
* fcid 0x7f13ef [device-alias B8P2]
* fcid 0x7f0183 [pwwn 20:00:00:25:b5:00:00:26]

zone name ora12c_b420_8_hba4 vsan 15
* fcid 0x7f1aef [device-alias A8P3]
* fcid 0x7f17ef [device-alias B8P3]
* fcid 0x7f0161 [pwwn 20:00:00:25:b5:00:00:6f]
```

Appendix C: Oracle Patches

Grid Infrastructure Patches	Description	Comments
Patch 17552800	"Database Patch Set Update : 12.1.0.1.2 (17552800)"	Patchset 12.1.0.1.2
Patch 17303297	"ACFS Patch Set Update 12.1.0.1.1"	Patchset 12.1.0.1.1
Patch 17077442	"Oracle Clusterware Patch Set Update 12.1.0.1.1"	Patchset 12.1.0.1.1

DB Patches	Description	Comments
Patch 18148383	AWR Snapshots not working – Bug 18225695	No Snapshots
Patch 17039360	ORA-600 [OCIKDBLinkConn-8] quering CDB views	Perf issues
Patch 17552800	"Database Patch Set Update : 12.1.0.1.2 (17552800)"	Patchset 12.1.0.1.2
Patch 17077442	"Oracle Clusterware Patch Set Update 12.1.0.1.1"	Patchset 12.1.0.1.1

Appendix D: Oracle spfile Parameters

Note Linux Huge Pages were setup on each host

ASM

```
asm_diskgroups='OCRVOTEDG','OLTPDG','DSSDG','CPUCRSDG','ACFSOH','REDODG'
asm_power_limit=1
memory_target=1076M
sessions=600

Container Database
sga_max_size=128G
sga_target=128G
db_name='oltp'
cluster_database=TRUE
log_buffer= 183623680
processes=3000
db_files=4000
db_4k_cache_size=26G //CPU PDB
db_16_cache_size=10G //DSS PDB
memory_max_target=0
memory_target=0 //use huge pages
pga_aggregate_target=30G
resource_manager_plan=cpu_oltp_dss_plan //Resource Manager
```

Appendix E: Key Linux Parameters

```
sysctl.conf
kernel.sem = 8192 48000 8192 8192
net.core.rmem_default = 4194304
net.core.rmem_max = 16777216
net.core.wmem_default = 4194304
net.core.wmem_max = 16777216
vm.nr_hugepages = 72100
limits.conf
/etc/security/limits.d/oracle-rdbms-server-12cR1-preinstall.conf
oracle soft nofile 4096
oracle hard nofile 65536
oracle soft nproc 32767
oracle hard nproc 32767
oracle soft stack 10240
oracle hard stack 32768
```

Appendix F: Cluster Information

```
Local Time Now :2014-05-17 09:28:24

The Cluster Nodes are :                ora12crac1, ora12crac2, ora12crac3,
ora12crac4, ora12crac5, ora12crac6, ora12crac7, ora12crac8
The Local Node is :                    ora12crac1
The Remote Nodes are :                 ora12crac2, ora12crac3, ora12crac4,
ora12crac5, ora12crac6, ora12crac7, ora12crac8

Major Clusterware Software Version is :12.1.0.1.0
Major Clusterware Active Version is :  12.1.0.1.0
Major Clusterware Release Version is : 12.1.0.1.0

CRS_HOME is installed at :             /oracle/product/12cgrid
CRS_BASE is installed at :             /oracle/product/grid
CRS_OWNER is :                         oracle
CRS_GROUP is :                         oinstall



All databases created :oastcpu, oastcrs, oastdss, oastiop

DB_NAME     MANAGEMENT  DB_TYPE     DB_VERSION  DB_HOME
DG_USED
=======     ==========  =======     ==========  =======
=======
oastcpu     administratorRAC         12.1.0.1.0
/oracle/product/12.1.0/ucs12c             ''
oastcrs     administratorRAC         12.1.0.1.0
/oracle/product/12.1.0/ucs12c             '+CPUCRSDG'
oastdss     administratorRAC         12.1.0.1.0
/oracle/product/12.1.0/ucs12c             +DSSDG'
oastiop     administratorRAC         12.1.0.1.0
/oracle/product/12.1.0/ucs12c             '+OLTPDG'


NODE_NAME           NODE_IDNODE_STATENODE_ROLE
=========           ==========================
ora12crac1          1      Active    Hub
ora12crac4          2      Active    Hub
ora12crac2          3      Active    Hub
ora12crac8          4      Active    Hub
ora12crac5          5      Active    Hub
ora12crac7          6      Active    Hub
ora12crac3          7      Active    Hub
ora12crac6          8      Active    Hub



Cluster Name :          ora12ccluster
SCAN Name :             ora12c-scan.ucs.cisco.com
SCAN Listeners :        LISTENER_SCAN1 (Port: TCP:1521)
                        LISTENER_SCAN2 (Port: TCP:1521)
                        LISTENER_SCAN3 (Port: TCP:1521)
GNS Status :            configured and enabled
GNS Version :           12.1.0.1.0
```

```
GNS Subdomain :         ucs.cisco.com
GNS-to-DNS  Port :      53
GNS-to-mDNS Port :      5,353



Node VIP Version :      12.1.0.1.0
Local Node VIPs :       ora.ora12crac1.vipora12crac1-vip.ucs.cisco.com
(static)
                        ora.ora12crac2.vipora12crac2-vip.ucs.cisco.com
(static)
                        ora.ora12crac3.vipora12crac3-vip.ucs.cisco.com
(static)
                        ora.ora12crac4.vipora12crac4-vip.ucs.cisco.com
(static)
                        ora.ora12crac5.vipora12crac5-vip.ucs.cisco.com
(static)
                        ora.ora12crac6.vipora12crac6-vip.ucs.cisco.com
(static)
                        ora.ora12crac7.vipora12crac7-vip.ucs.cisco.com
(static)
                        ora.ora12crac8.vipora12crac8-vip.ucs.cisco.com
(static)



Oracle Interfaces :     eth010.29.134.0globalpublic
                        eth1192.168.134.0globalcluster_interconnect,asm



OCR Location :          '+OCRVOTEDG'
Voting Disk Location :  '+OCRVOTEDG'

Cluster Mode :          Flex Cluster

Hub Node                            attaches                   Leaf Node
========                            ========                   =========
ora12crac1(1,Active)                <---                       None
ora12crac4(2,Active)                <---                       None
ora12crac2(3,Active)                <---                       None
ora12crac8(4,Active)                <---                       None
ora12crac5(5,Active)                <---                       None
ora12crac7(6,Active)                <---                       None
ora12crac3(7,Active)                <---                       None
ora12crac6(8,Active)                <---                       None



MGMTDB Status :         enabled and is running on ora12crac1
MGMTDB HOME :           <CRS home>
MGMTDB Spfile :         '+OCRVOTEDG/_mgmtdb/spfile-MGMTDB.ora'
MGMTDB Instance :       '-MGMTDB'


DISKGROUPREDUNDANCYAU   COMPATIBILITYDB_COMPATIBILITYSIZE_MBFREE_MBUSABLE_MB
PATH
```

```
==========================================================================
====
ACFSOH   EXTERN    1MB 12.1.0.0.0   10.1.0.0.0      537576215712215712
ORCL:ACFSORAHOME3

ORCL:ACFSORAHOME1

ORCL:ACFSORAHOME2
CPUCRSDG EXTERN    1MB 12.1.0.0.0   10.1.0.0.0      409592038400653840065
ORCL:CPUCRS_10

ORCL:CPUCRS_1

ORCL:CPUCRS_14

ORCL:CPUCRS_15

ORCL:CPUCRS_16

ORCL:CPUCRS_2

ORCL:CPUCRS_3

ORCL:CPUCRS_4

ORCL:CPUCRS_5

ORCL:CPUCRS_6

ORCL:CPUCRS_7

ORCL:CPUCRS_8

ORCL:CPUCRS_11

ORCL:CPUCRS_9

ORCL:CPUCRS_13

ORCL:CPUCRS_12
DSSDG    EXTERN    4MB 12.1.0.0.0   10.1.0.0.0      1638387250326325032632
ORCL:DSS_16

ORCL:DSS_15

ORCL:DSS_14

ORCL:DSS_13

ORCL:DSS_12

ORCL:DSS_11

ORCL:DSS_10

ORCL:DSS_1

ORCL:DSS_2
```

```
ORCL:DSS_9

ORCL:DSS_8

ORCL:DSS_7

ORCL:DSS_6

ORCL:DSS_5

ORCL:DSS_4

ORCL:DSS_3
OCRVOTEDGNORMAL    1MB 12.1.0.0.0   10.1.0.0.0      1023959346336492
ORCL:OCRVOTE5

ORCL:OCRVOTE4

ORCL:OCRVOTE3

ORCL:OCRVOTE2

ORCL:OCRVOTE1
OLTPDG   EXTERN    1MB 12.1.0.0.0   10.1.0.0.0      655348813582171358217
ORCL:OLTP_9

ORCL:OLTP_15

ORCL:OLTP_16

ORCL:OLTP_8

ORCL:OLTP_7

ORCL:OLTP_6

ORCL:OLTP_5

ORCL:OLTP_4

ORCL:OLTP_3

ORCL:OLTP_2

ORCL:OLTP_13

ORCL:OLTP_12

ORCL:OLTP_11

ORCL:OLTP_10

ORCL:OLTP_1

ORCL:OLTP_14
REDODG   EXTERN    1MB 12.1.0.0.0   10.1.0.0.0      491512040795854079585
ORCL:REDO_8
```

```
ORCL:REDO_9

ORCL:REDO_6

ORCL:REDO_7

ORCL:REDO_4

ORCL:REDO_5

ORCL:REDO_2

ORCL:REDO_3

ORCL:REDO_1

ORCL:REDO_10

ORCL:REDO_11

ORCL:REDO_12

ORCL:REDO_13

ORCL:REDO_14

ORCL:REDO_15

ORCL:REDO_16



VOLUME_NAME    VOLUME_DEVICE DG_NAME      MOUNTPOINT USAGE VOL_SIZE_MB
=========== =============            =======     ==========          =====
===========
ACFS_VOL_OH /dev/asm/acfs_vol_oh-423 ACFSOH         /oracle/product/12.1ACFS
321536


ASM Host                           connects                   Client
========                           ========                   ======
ora12crac1.ucs.cisco.com(+ASM1)    <---
'+APX1(ora12crac1)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'+APX7(ora12crac3)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'+APX8(ora12crac6)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'+ASM1(ora12crac1)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'-MGMTDB(ora12crac1)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastcrs1(ora12crac1)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastcrs2(ora12crac4)'
```

```
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastcrs3(ora12crac8)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastcrs4(ora12crac3)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastcrs7(ora12crac7)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastcrs8(ora12crac6)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastiop1(ora12crac1)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastiop4(ora12crac7)'
ora12crac1.ucs.cisco.com(+ASM1)    <---
'oastiop6(ora12crac3)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'+APX3(ora12crac2)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'+APX4(ora12crac8)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'+ASM3(ora12crac2)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'oastcrs6(ora12crac2)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'oastiop2(ora12crac8)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'oastiop5(ora12crac2)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'oastiop7(ora12crac6)'
ora12crac2.ucs.cisco.com(+ASM3)    <---
'oastiop8(ora12crac4)'
ora12crac5.ucs.cisco.com(+ASM5)    <---
'+APX2(ora12crac4)'
ora12crac5.ucs.cisco.com(+ASM5)    <---
'+APX5(ora12crac5)'
ora12crac5.ucs.cisco.com(+ASM5)    <---
'+APX6(ora12crac7)'
ora12crac5.ucs.cisco.com(+ASM5)    <---
'+ASM5(ora12crac5)'
ora12crac5.ucs.cisco.com(+ASM5)    <---
'oastcrs5(ora12crac5)'
ora12crac5.ucs.cisco.com(+ASM5)    <---
'oastiop3(ora12crac5)'

OCR/CRSD Master :ora12crac1
CRSD PE Master : ora12crac1
CRSD PE Standby :ora12crac5

<ora12crac1> 05-14 20:01:12.587: My R:3 = Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:2   OCR on ASM:1   ASM mode:2
<ora12crac4> 05-14 20:04:03.879: My R:1 < Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:0   OCR on ASM:1   ASM mode:2
<ora12crac2> 05-14 20:03:06.693: My R:3 = Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:2   OCR on ASM:1   ASM mode:2
<ora12crac8> 05-14 20:02:25.735: My R:1 < Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:0   OCR on ASM:1   ASM mode:2
<ora12crac5> 05-14 20:03:14.102: My R:4 > Min R:3   R of OCR:1   R of CRS
Standby:1   R of ASM Inst:2   OCR on ASM:1   ASM mode:2
```

```
<ora12crac7> 05-14 20:03:42.636: My R:1 < Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:0   OCR on ASM:1   ASM mode:2
<ora12crac3> 05-14 20:02:24.307: My R:1 < Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:0   OCR on ASM:1   ASM mode:2
<ora12crac6> 05-14 20:02:26.508: My R:1 < Min R:3   R of OCR:1   R of CRS
Standby:0   R of ASM Inst:0   OCR on ASM:1   ASM mode:2
<ora12crac1> 05-14 20:01:14.086: PE Role|State Update: old role [MASTER] new
[MASTER]; old state [Starting] new [Running]
<ora12crac4> 05-14 20:04:18.178: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]
<ora12crac2> 05-14 20:02:44.621: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]
<ora12crac8> 05-14 20:02:44.008: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]
<ora12crac5> 05-14 20:02:43.420: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]
<ora12crac7> 05-14 20:03:56.493: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]
<ora12crac3> 05-14 20:02:45.298: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]
<ora12crac6> 05-14 20:02:42.100: PE Role|State Update: old role [SLAVE] new
[SLAVE]; old state [Starting] new [Running]



CTSS Master :    ora12crac1
UI Master :      ora12crac1
ONS Master :     ora12crac6
CHM Master :     ora12crac1
CHM Replica :    REPLICA has been deprecated from 12c

OCR Local/Writer                  connects                 ASM Instance
================                  ========               ============
ora12crac1(Hub,OCR Writer)        --->
ora12crac1(+ASM1)
ora12crac4(Hub,OCR Local)         --->
ora12crac1(+ASM1)
ora12crac2(Hub,OCR Local)         --->
ora12crac1(+ASM1)
ora12crac8(Hub,OCR Local)         --->
ora12crac1(+ASM1)
ora12crac5(Hub,OCR Local)         --->
ora12crac1(+ASM1)
ora12crac7(Hub,OCR Local)         --->
ora12crac2(+ASM3)
ora12crac3(Hub,OCR Local)         --->
ora12crac1(+ASM1)
ora12crac6(Hub,OCR Local)         --->
ora12crac1(+ASM1)
```

Appendix G: FCoE Throughput from Fabric Interconnects

The following is a sample script that was used to mine FCoE throughput from fabrics. This is Cisco UCS Manager version dependent and may need some changes.

Write an expect script to access the details from each IOM as shown below:

```
#!/usr/bin/expect
set timeout -1
spawn -noecho ssh -o StrictHostKeyChecking=no admin@<ip address of UCSM>
expect {
   "Password:" {
      send "UCSMPassword\r"
   }
}
expect "# "
send "connect local-mgmt a\r"
expect "# "
send "connect iom 1\r"
expect "# "
log_file "a-iom-1.lst"
send "show platform software woodside rate | grep 0-NI\r"
expect "# "
log_file
send "exit 0\r"
```

Repeat the same for each of the other IOM's like Fab A iom 2, Fab B iom 1, Fab B iom 2.

Mine the log files created to sum up the IO throughput at a given time. The output captured could be Kbps or Mbps or Gbps. A little Math may be needed in scripting, to covert all of them to one units and proceed. The script below is provided as a reference only:

```
if [ $units == "Kbps" ]; then
test1=`echo $value/1024/1024 | bc -l`….
```