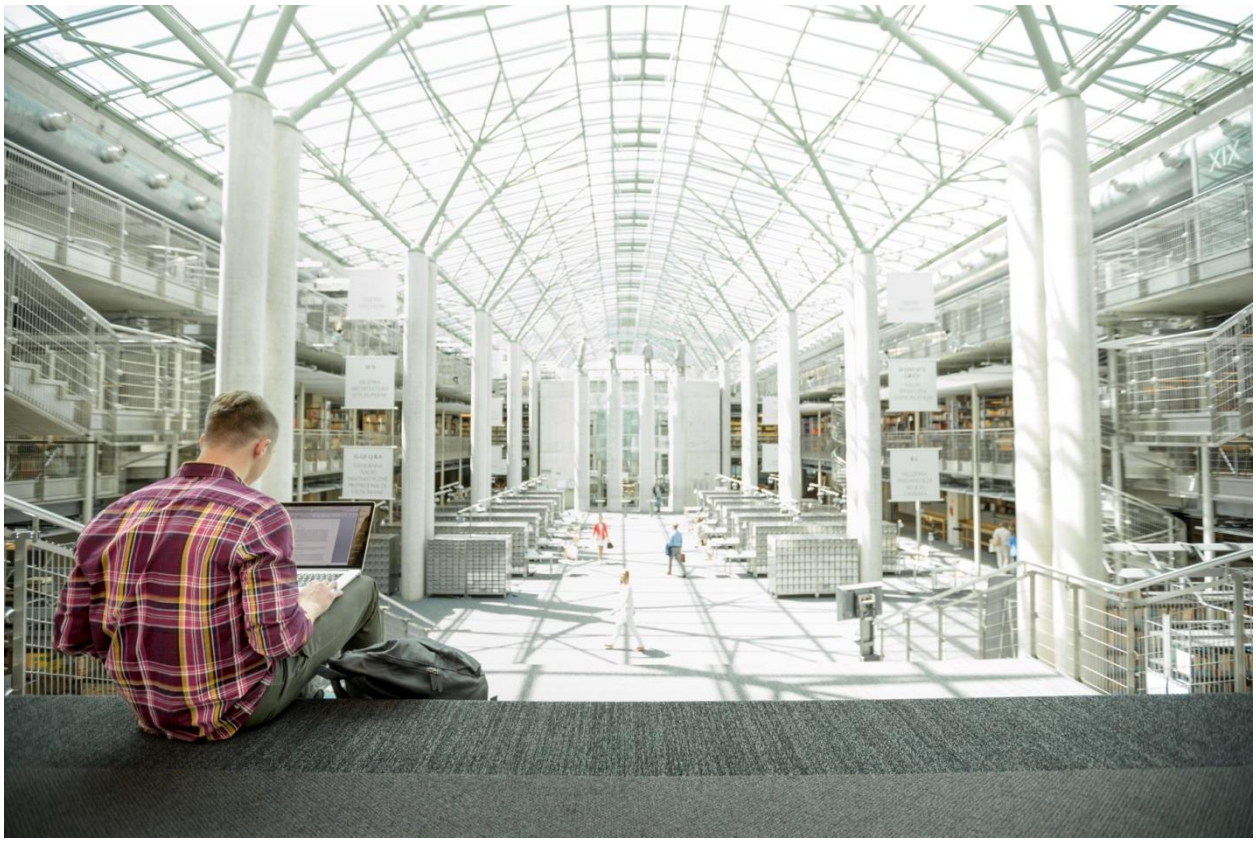

FlashStack Data Center with Oracle RAC 12cR2 Database

Deployment Guide for Oracle RAC Database 12cR2 on Cisco Unified Computing System and Pure Storage FlashArray//X Series

Last Updated: March 1, 2018

About the Cisco Validated Design Program

The Cisco Validated Design (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

Table of Contents

Executive Summary

The Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, network, storage access, and virtualization into a single cohesive system. Cisco UCS is an ideal platform for the architecture of mission critical database workloads. The combination of Cisco UCS platform, Pure Storage® and Oracle Real Application Cluster (RAC) architecture can accelerate your IT transformation by enabling faster deployments, greater flexibility of choice, efficiency, high availability and lower risk.

Cisco® Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers.

This Cisco Validated Design (CVD) describes a FlashStack reference architecture for deploying a highly available Oracle RAC Databases environment on Pure Storage FlashArray//X using Cisco UCS Compute Servers, Cisco Fabric Interconnect Switches, Cisco MDS Switches, Cisco Nexus Switches and Oracle Linux. Cisco and Pure Storage have validated the reference architecture with OLTP (On-line Transaction Processing) and Data Warehouse workload in Cisco's lab. This document presents the hardware and software configuration of the components involved, results of various tests and offers implementation and best practices guidance.

FlashStack is a converged infrastructure solution that brings the benefits of an all-flash storage platform to your converged infrastructure deployments. Built on best of breed components from Cisco UCS Systems and Pure Storage, FlashStack provides a converged infrastructure solution that is simple, flexible, efficient, and costs less than legacy converged infrastructure solutions based on traditional disk.

FlashStack embraces the latest technology and efficiently simplifies data center workloads that redefine the way IT delivers value:

- Guarantee customer success with prebuilt, pre-tested drivers and Oracle database software
- A cohesive, integrated system that is managed, serviced and tested as a whole
- Faster Time to Deployment – Leverage a pre-validated platform to minimize business disruption, improve IT agility, and reduce deployment time from months to weeks.
- Reduces Operational Risk – Highly available architecture with no single point of failure, non-disruptive operations, and no downtime.

Solution Overview

Introduction

Database administrators and their IT departments face many challenges that demand a simplified Oracle deployment and operation model providing high performance, availability and lower TCO. The current industry trend in data center design is towards shared infrastructures featuring multitenant workload deployments. Cisco® and Pure Storage have partnered to deliver FlashStack, which uses best-in-class storage, server, and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

FlashStack solution provides the advantage of having the compute, storage, and network stack integrated with the programmability of the Cisco Unified Computing System (Cisco UCS). This Cisco Validated Design (CVD) describes how Cisco UCS System can be used in conjunction with Pure Storage FlashArray//X System to implement an Oracle Real Application Clusters (RAC) 12c R2 Database solution.

Audience

The target audience for this document includes but is not limited to storage administrators, data center architects, database administrators, field consultants, IT managers, Oracle solution architects and customers who want to implement Oracle RAC database solutions with Linux on a FlashStack Converged Infrastructure solution. A working knowledge of Oracle RAC Database, Linux, Storage technology, and Network is assumed but is not a prerequisite to read this document.

Purpose of this Document

Oracle RAC databases deployments are extremely complicated in nature and customers face enormous challenges in maintaining these landscapes in terms of time, efforts and cost. Oracle RAC databases often manage the mission critical components of a customer's IT department, ensuring availability while also lowering the IT TCO is always their top priority.

The goal of this CVD is to highlight the performance, scalability, manageability, and simplicity of the FlashStack Converged Infrastructure solution for deploying mission critical applications such as Oracle RAC databases.

The following are the objectives of this reference architecture document:

1. Provide reference architecture design guidelines for the FlashStack based Oracle RAC Databases.
2. Build, validate, and predict performance of Server, Network, and Storage platform on a per workload basis.
3. Seamless scalability of performance and capacity to meet growth needs of Oracle Database.
4. High availability of DB instances without performance compromise through software and hardware upgrades.

We will demonstrate the scalability and performance of this solution by running SwingBench and SLOB (Silly Little Oracle Benchmark) on OLTP (On-line Transaction Processing) and DSS (Decision Support System), such as benchmarking with varying users, nodes and read/write workload characteristics.

FlashStack System Overview

The FlashStack platform, developed by Cisco and Pure Storage, is a flexible, integrated infrastructure solution that delivers pre-validated storage, networking, and server technologies. Cisco and Pure Storage have carefully validated and verified the FlashStack solution architecture and its many use cases while creating a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model.

This portfolio includes, but is not limited to, the following items:

- Best practice architectural design
- Implementation and deployment instructions and provides application sizing based on results

Figure 1 FlashStack System Components

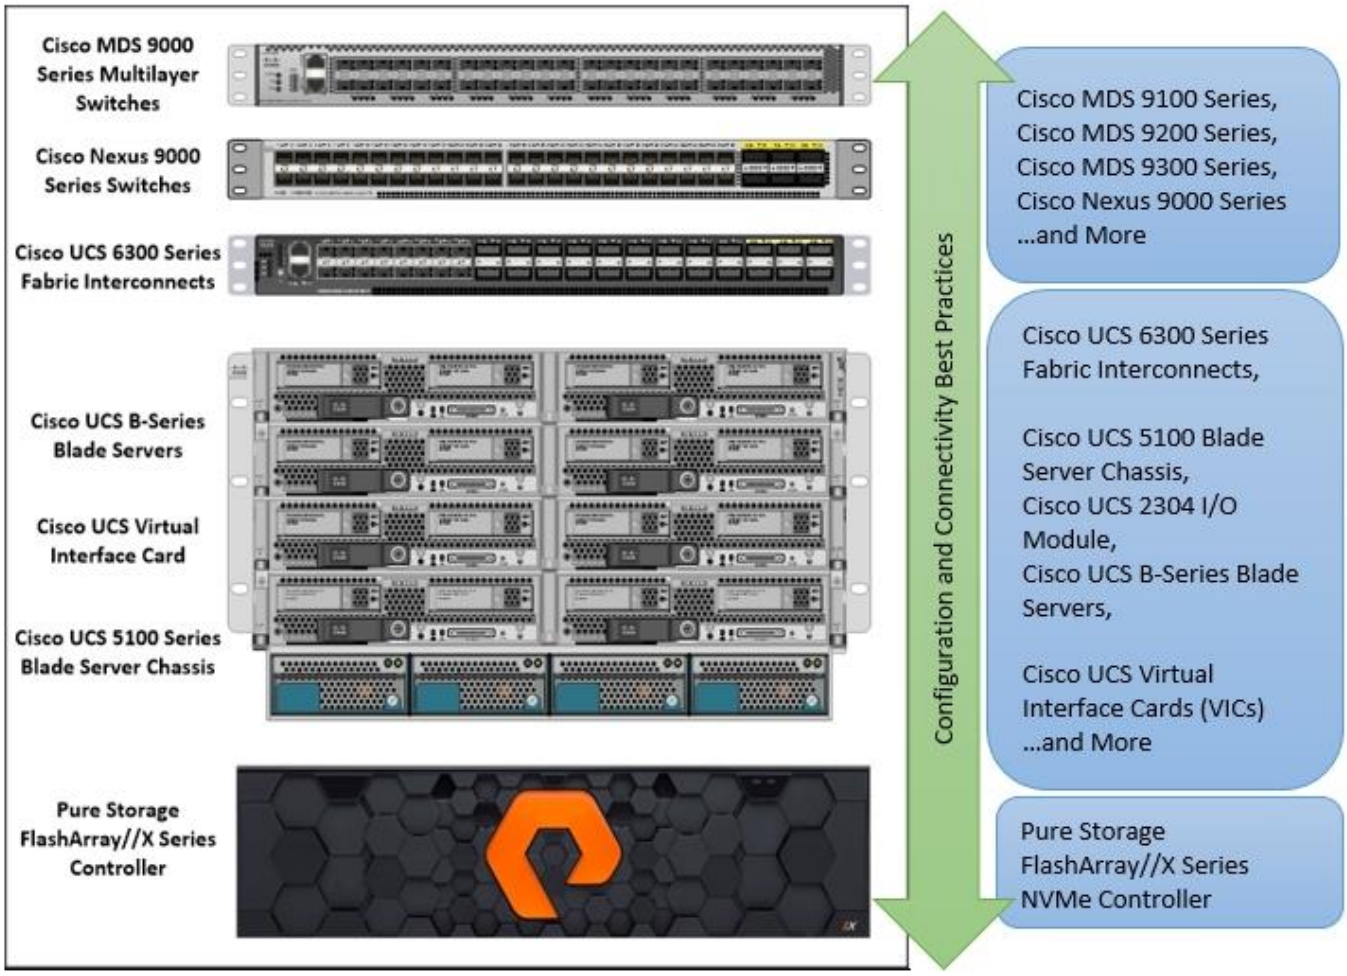

As shown in Figure 1, these components are connected and configured according to best practices of both Cisco and Pure Storage and provide the ideal platform for running a variety of enterprise database workloads with confidence. FlashStack can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed), or it can scale out for environments that require multiple consistent deployments.

The reference architecture covered in this document leverages the Pure Storage FlashArray//X70 Controller with NVMe based DirectFlash modules for Storage, Cisco UCS B200 M5 Blade Server for Compute, Cisco Nexus 9000 and Cisco MDS 9100 series for the switching element and Cisco Fabric Interconnects 6300 series for System Management. As shown in Figure 1, FlashStack Architecture can maintain consistency at scale. Each of the component families shown in (Cisco UCS, Cisco Nexus, Cisco MDS, Cisco FI and Pure Storage) offers platform and resource options to scale the infrastructure up or down, while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlashStack.

FlashStack Solution Benefits

FlashStack provides a jointly supported solution by Cisco and Pure Storage. Bringing a carefully validated architecture built on superior compute, world-class networking, and the leading innovations in all flash storage. The portfolio of validated offerings from FlashStack includes but is not limited to the following:

- Consistent Performance and Scalability
  - Consistent sub-millisecond latency with 100 percent NVMe enterprise flash storage
  - Consolidate hundreds of enterprise-class applications in a single rack
  - Scalability through a design for hundreds of discrete servers and thousands of virtual machines, and the capability to scale I/O bandwidth to match demand without disruption
  - Repeatable growth through multiple FlashStack CI deployments.
- Operational Simplicity
  - Fully tested, validated, and documented for rapid deployment
  - Reduced management complexity
  - No storage tuning or tiers necessary
  - 3x better data reduction without any performance impact
- Lowest TCO
  - Dramatic savings in power, cooling and space with Cisco UCS and 100 percent Flash
  - Industry leading data reduction
  - Free FlashArray controller upgrades every three years with Forever Flash™
- Mission Critical and Enterprise Grade Resiliency
  - Highly available architecture with no single point of failure
  - Non-disruptive operations with no downtime
  - Upgrade and expand without downtime or performance loss
  - Native data protection: snapshots and replication

Cisco and Pure Storage have also built a robust and experienced support team focused on FlashStack solutions, from customer account and technical sales representatives to professional services and technical support engineers. The support alliance between Pure Storage and Cisco gives customers and channel services partners direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

What's New in this FlashStack Release

This version of the FlashStack CVD introduces new hardware with the Pure Storage FlashArray//X, that is 100 percent NVMe enterprise class all-flash array along with Cisco UCS B200 M5 Blade Servers featuring the Intel Xeon Scalable Family of CPUs. This is the second Oracle RAC Database deployment Cisco Validated Design with Pure Storage. It incorporates the following features:

- Pure Storage FlashArray//X

- Cisco UCS B200 M5 Blade Servers
- Oracle RAC Database 12c Release 2

Solution Components

This section provides a list of all the components used in this solution.

Cisco UCS 6332-16UP Fabric Interconnect

The 6332-16UP Fabric Interconnect is the management and communication backbone for Cisco UCS B-Series Blade Servers, C-Series Rack Servers, and 5100 Series Blade Server Chassis. It implements 20x40 Gigabit Ethernet and Fibre Channel over Ethernet ports, with additional support for 16 unified ports that can be configured to 1 or 10 Gbps Ethernet, or 4/8/16 Gbps Fibre Channel.

The Fabric Interconnect provides high-speed upstream connectivity to the network, or converged traffic to servers through its 40 Gbps ports, but also allows for Fibre Channel connectivity to SAN switches like the MDS, or alternately directly attached Fibre Channel to storage arrays like the Pure Storage FlashArray through its unified ports.

Cisco UCS MDS 9148S Fabric Switch

The Cisco® MDS 9148S 16G Multilayer Fabric Switch is the next generation of highly reliable, flexible and low-cost Cisco MDS 9100 Series Switches. It provides up to 48 auto-sensing Fibre Channel ports, which are capable of speeds of 2, 4, 8, and 16 Gbps, with 16 Gbps of dedicated bandwidth for each port.

In all, the Cisco MDS 9148S is a powerful and flexible switch that delivers high performance and comprehensive Enterprise-class features at an affordable price.

Cisco Nexus 9372PX-E Switch

The Cisco Nexus 9372PX-E Switches are 1RU switches that support 1.44 Tbps of bandwidth and over 1150 mpps across 48 fixed 10-Gbps SFP+ ports and 6 fixed 40-Gbps QSFP+ ports.

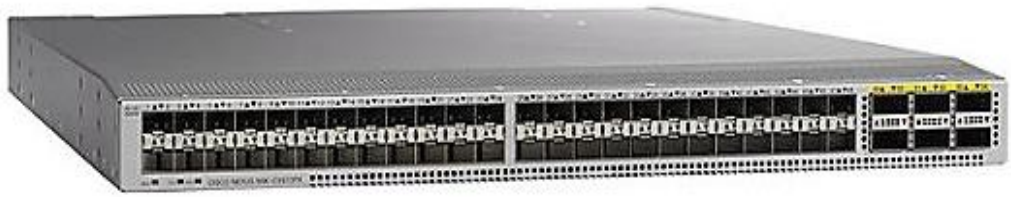

Cisco UCS B200 M5 Blade Servers

The Cisco UCS B200 M5 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads including Virtual Desktop Infrastructure (VDI), web infrastructure, distributed databases, converged infrastructure, and enterprise applications such as Oracle and SAP HANA.

The Cisco UCS B200 M5 server can quickly deploy stateless physical and virtual workloads through programmable, easy-to-use Cisco UCS Manager Software and simplified server access through Cisco Single-Connect technology.

Cisco UCS 5108 Blade Server Chassis

Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high, can mount in an industry-standard 19-inch rack, and uses standard front-to-back cooling. A chassis can accommodate up to eight half-width or four full-width Cisco UCS B-Series Blade Servers form factors within the same chassis.

By incorporating unified fabric and fabric-extender technology, the Cisco Unified Computing System eliminates the need for dedicated chassis management and blade switches, reduces cabling, and allowing scalability to 20 chassis without adding complexity. The Cisco UCS 5108 Blade Server Chassis is a critical component in delivering the simplicity and IT responsiveness for the data center as part of the Cisco Unified Computing System.

Cisco UCS 2304 Fabric Extender

Cisco UCS 2304 Fabric Extender brings the unified fabric into the blade server enclosure, providing multiple 40 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

The Cisco UCS 2304 connects the I/O fabric between the Cisco UCS 6300 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic Fibre Channel over Ethernet (FCoE) fabric to connect all blades and chassis together.

Cisco UCS Virtual Interface Card (VIC) 1340

The Cisco UCS Virtual Interface Card (VIC) 1340 is a 2-port, 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) mezzanine adapter.

Cisco UCS 1340 VIC delivers 80 Gbps throughput to the Server and helps reduce TCO by consolidating the overall number of NICs, HBAs, cables, and switches; LAN and SAN traffic runs over the same mezzanine card and fabric.

Pure Storage FlashArray //X70

The Pure Storage FlashArray family delivers purpose-built, software-defined all-flash power and reliability for businesses of every size. FlashArray is all-flash enterprise storage that is up to 10X faster, more space and power efficient, more reliable, and far simpler than other available solutions. Critically, FlashArray also costs less, with a TCO that's typically 50% lower than traditional performance disk arrays.

At the top of the FlashArray line is FlashArray//X – the first mainstream, 100 percent NVMe, enterprise-class all-flash array. //X represents a higher performance tier for mission-critical databases, top-of-rack flash deployments, and Tier 1 application consolidation. It is optimized for the lowest-latency workloads and delivers an unprecedented level of performance density that makes possible previously unattainable levels of consolidation.

FlashArray//X provides microsecond latency, 1PB in 3U, and GBs of bandwidth, with rich data services, proven 99.9999 percent availability (inclusive of maintenance and generational upgrades), 2X better data reduction versus alternative all-flash solutions, and DirectFlash™ global flash management. Further, //X is self-managing and plug-

n-play, thanks to unrivalled Pure1® Support and the cloud-based, machine-learning predictive analytics of Pure1 Meta. Finally, FlashArray//X, like the rest of the FlashArray line, has revolutionized the 3-5 year storage refresh cycle by eliminating it: Pure's Evergreen™ Storage model provides a subscription to hardware and software innovation that enables organizations to expand and enhance their storage for 10 years or more.

Figure 2 Pure Storage FlashArray //X70

At the heart of FlashArray//X is the Purity Operating Environment software. Purity enables organizations to enjoy Tier 1 data services for all workloads, completely non-disruptive operations, and the power and efficiency of DirectFlash. Moreover, Purity includes enterprise-grade data security, comprehensive data protection options, and complete business continuity via ActiveCluster multi-site stretch cluster – all included with every array.

Figure 3 Pure Storage FlashArrays

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers the following:

- Consistent Performance FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.
- Less Cost than Disk Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure. With VDI workloads data reduction is typically > 10:1
- Disaster Recovery Built-In FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. In addition, recover instantly from local, space-efficient snapshots or remote replicas.
- Mission-Critical Resiliency FlashArray delivers >99.999 percent proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

Solution Architecture

The FlashStack architecture brings together the proven data center strengths of the Cisco UCS and Cisco Nexus network switches with the Fibre Channel delivered storage of the leading visionary in all flash arrays. This collaboration creates a simple, yet powerful and resilient data center footprint for the modern enterprise. The FlashStack Data Center with Oracle RAC database on Oracle Linux solution provides an end-to-end architecture with Cisco, Oracle, and Pure Storage technologies and demonstrates the FlashStack configuration benefits for running highly available Oracle RAC Database 12c R2 with Cisco VICs (Virtual Interface Cards).

Design Topology

This section describes the design considerations for the Oracle RAC Database 12c Release 2 on FlashStack deployment. In this solution design, we used two Cisco UCS Blade Server Chassis with 8 identical Intel Xeon CPU based Cisco UCS B-Series B200 M5 Blade Servers for hosting the 8-Node Oracle RAC Databases. The Cisco UCS B200 M5 Server has Virtual Interface Card (VIC) 1340 with port expander and they were connected four ports from each Cisco Fabric extender of the Cisco UCS Chassis to the Cisco Fabric Interconnects, which were in turn connected to the Cisco MDS Switches for upstream connectivity to access the Pure Storage FlashArray//X70.

The following table list the inventory of the components used in the FlashStack solution.

Table 1 Inventory and Bill of Material

Vendor	Name	Model	Description	Qty
Cisco	Cisco Nexus 9372PX-E Switch	N9K-C9372PX-E	Cisco Nexus 9300 Series Switches	2
Cisco	Cisco MDS 9148S 16G Fabric Switch	DS-C9148S-12PK9	Cisco MDS 9100 Series Multilayer Fabric Switches	2
Cisco	Cisco UCS 6332-16UP Fabric Interconnect	UCS-FI-6332-16UP	Cisco 6300 Series Fabric Interconnects	2
Cisco	Cisco UCS Fabric Extender	UCS-IOM-2304	Cisco UCS 2304XP I/O Module (4 External, 8 Internal 40Gb Ports)	4
Cisco	Cisco UCS 5108 Blade Server Chassis	UCSB-5108-AC2	Cisco UCS 5100 Series Blade Server AC2 Chassis	2
Cisco	Cisco UCS B200 M5 Blade Servers	UCSB-B200-M5	Cisco UCS B-Series Blade Servers	8
Cisco	Cisco UCS VIC 1340	UCSB-MLOM-40G-03	Cisco UCS Virtual Interface Card 1340	8
Cisco	Cisco UCS Port Expander Card	UCSB-MLOM-PT-01	Port Expander Card for Cisco UCS MLOM	8
Pure Storage	Pure FlashArray //X70 Controller	Purity 4.10.6	Pure Storage FlashArray (FA //X70)	1

The following table list the server configuration used in the FlashStack solution.

Table 2 Cisco UCS B200 M5 Blade Server Configuration

Server Configuration	
Processor	2 x Intel® Xeon® Gold 6152 Processor (2.10 GHz, 140W, 22C, 30.25MB Cache, DDR4 2666MHz 768GB)
Memory	16 x 32GB DDR4-2666-MHz RDIMM/dual rank/x4/1.2v
Cisco UCS VIC 1340	Cisco UCS VIC 1340 Blade MLOM
Cisco UCS Port Expander Card	Port Expander Card for Cisco UCS MLOM

For this FlashStack solution design, we configured two VLANs and two VSANs as described in the table below.

Table 3 VLAN and VSAN Configuration

Name	ID	Description
VLANs		
• **Default VLAN**	1	Native VLAN
• **Public VLAN**	134	VLAN for Public Network Traffic
• **Private VLAN**	10	VLAN for Private Network Traffic
VSANs		
• **VSAN – A**	201	SAN Communication through for Fabric Interconnect A
• **VSAN – B**	202	SAN Communication through for Fabric Interconnect B

The FlashStack design comprises of Pure Storage FlashArray //X70 with NVMe enterprise class all-flash for increased scalability and throughput. The table below lists the components of the array.

Table 4 Pure Storage FlashArray Configuration

Storage Components	Description
FlashArray	// X70
Capacity	23 TB
Connectivity	8 x 16 Gb/s redundant Fibre Channel 1 Gb/s redundant Ethernet (Management port)
Physical	3U

For this FlashStack solution design, we used the following Software and Firmware.

Table 5 Software and Firmware Configuration

Software and Firmware	Version
Oracle Linux Server 7.4 (64 bit) Operating System	Linux 4.1.12-94.3.9.el7uek.x86_64
Oracle 12c Release 2 GRID	12.2.0.1.0
Oracle 12c Release 2 Database Enterprise Edition	12.2.0.1.0
Cisco Nexus 9372PX-E NXOS Version	6.1(2) I2 (2a)
Cisco MDS 9148S System Version	6.2 (9)
Cisco UCS Manager System	3.2 (2c)
Cisco UCS Adapter VIC 1340	4.2 (2b)
Cisco eNIC (modinfo enic)	2.3.0.31
Cisco fNIC (modinfo fnic)	1.6.0.24
Pure Storage Purity Version	4.10.6
Oracle Swingbench	2.5.971
Oracle DBMS_RESOURCE_MANAGER_CALIBRATE_IO	
SLOB	2.4.2

Physical Topology

FlashStack consists of a combined stack of hardware (storage, network and compute) and software (Cisco UCS Manager, Oracle Database, Pure Storage GUI, Purity, and Oracle Linux).

- **Network**: Cisco Nexus 9372PX-E, Cisco MDS 9148S and Cisco UCS Fabric Interconnect 6332-16UP for external and internal connectivity of IP and FC network.
- **Storage**: Pure Storage FlashArray//X70 with 16Gb Fibre Channel connectivity
- **Compute**: Cisco UCS B200 M5 Blade Server

Figure 4 illustrates the FlashStack solution physical infrastructure.

Figure 4 Physical Topology

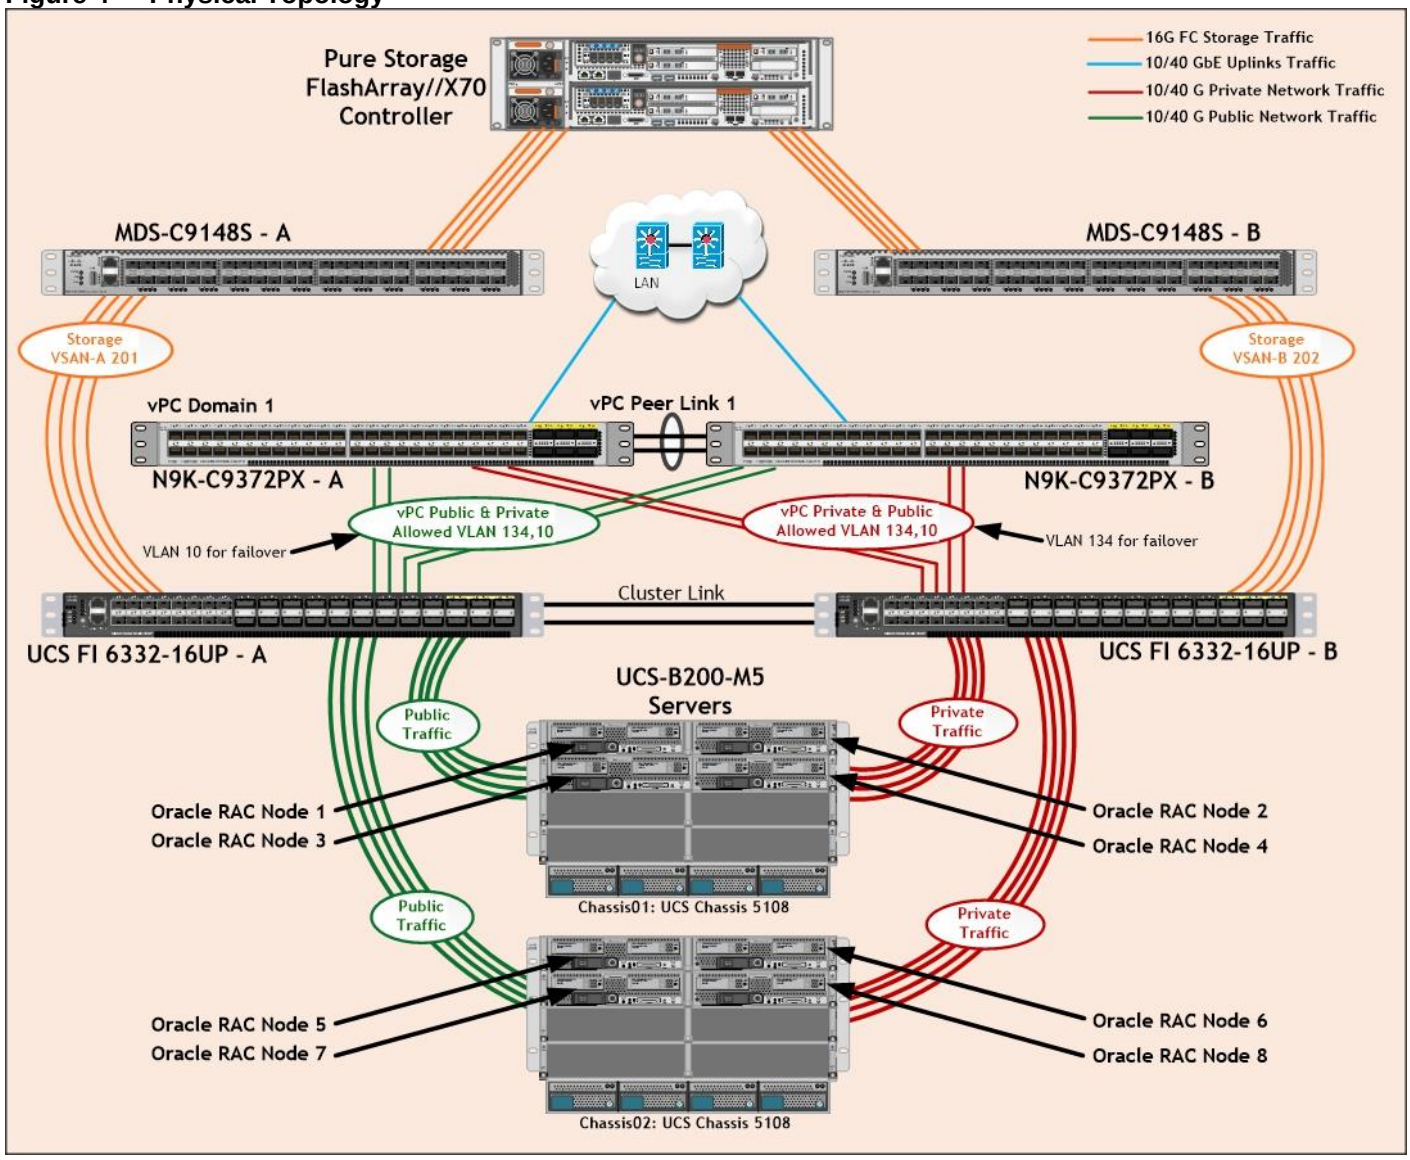

Figure 4 is a typical network configuration that can be deployed in a customer's environment. The best practices and setup recommendations are described later in this document.

As shown in Figure 4, a pair of Cisco UCS 6332-16UP fabric interconnects carries both storage and network traffic from the server blades with the help of Cisco Nexus 9372PX-E and Cisco MDS 9148S switches. Both the fabric interconnect and the Cisco Nexus switch are clustered with the peer link between them to provide high availability. Two virtual Port-Channels (vPCs) are configured to provide public network and private network paths for the server blades to northbound switches. Each vPC has VLANs created for application network data and management data paths.

As illustrated in Figure 4, eight (4 x 40G link per chassis) links go to Fabric Interconnect "A". Similarly, eight links go to Fabric Interconnect B. Fabric Interconnect-A links are used for Oracle Public network traffic shown as green lines. Fabric Interconnect-B links are used for Oracle private interconnect traffic shown as red lines. FC Storage access from Fabric Interconnect-A and Fabric Interconnect-B show as an orange line.

For Oracle RAC configuration on Cisco Unified Computing System, we recommend to keep all private interconnects local on a single Fabric interconnect. In this case, the private traffic stays local to that fabric intercon-

nect and will not be routed via northbound network switch. All inter-server blade (or RAC node private) communication will be resolved locally at the fabric interconnect and this significantly reduces latency for Oracle Cache Fusion traffic.

Solution Configuration

Cisco UCS Configuration Overview

This section details the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power, and installation of the chassis are described in the install guide, see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html). It is beyond the scope of this document to cover detailed information about UCS infrastructure setup and connectivity. The documentation guides and examples are available at http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html

All of the tasks to configure Cisco UCS are detailed in this document, but only some of the screenshots are included.

Cisco UCS Manager Software Version 3.2 (2c)

This document assumes the use of Cisco UCS Manager Software version 3.2(2c). To upgrade the Cisco UCS Manager software and the Cisco UCS 6332-16UP Fabric Interconnect software to a higher version of the firmware, refer to Cisco UCS Manager Install and Upgrade Guides.

Configure Base Cisco Unified Computing System

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

1. Configure Fabric Interconnects for a Cluster Setup.
2. Set Fabric Interconnects to Fibre Channel End Host Mode.
3. Synchronize Cisco UCS to NTP.
4. Configure Fabric Interconnects for Chassis and Blade Discovery:
    a. Configure Global Policies
    b. Configure Server Ports
5. Configure LAN and SAN on Cisco UCS Manager:
    a. Configure Ethernet LAN Uplink Ports
    b. Create Uplink Port Channels to Cisco Nexus Switches
    c. Configure FC SAN Uplink Ports
    d. Configure VLAN
    e. Configure VSAN
6. Configure IP, UUID, Server, MAC, WWNN and WWPN Pools:
    a. IP Pool Creation
    b. UUID Suffix Pool Creation
    c. Server Pool Creation
    d. MAC Pool Creation
    e. WWNN and WWPN Pool Creation

7. Set Jumbo Frames in both the Cisco Fabric Interconnect.
8. Configure Server BIOS Policy.
9. Create Adapter Policy.
10. Configure Update Default Maintenance Policy.
11. Configure vNIC and vHBA Template:
    a. Create Public vNIC Template
    b. Create Private vNIC Template
    c. Create Storage vHBA Template
12. Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

Configure Fabric Interconnects for a Cluster Setup

To configure the Cisco UCS Fabric Interconnects, complete the following steps:

1. Verify the following physical connections on the fabric interconnect:
    a. The management Ethernet port (mgmt0) is connected to an external hub, switch, or router
    b. The L1 ports on both fabric interconnects are directly connected to each other
    c. The L2 ports on both fabric interconnects are directly connected to each other

For more information, refer to the *Cisco UCS Hardware Installation Guide* for your fabric interconnect.

2. Connect to the console port on the first Fabric Interconnect.

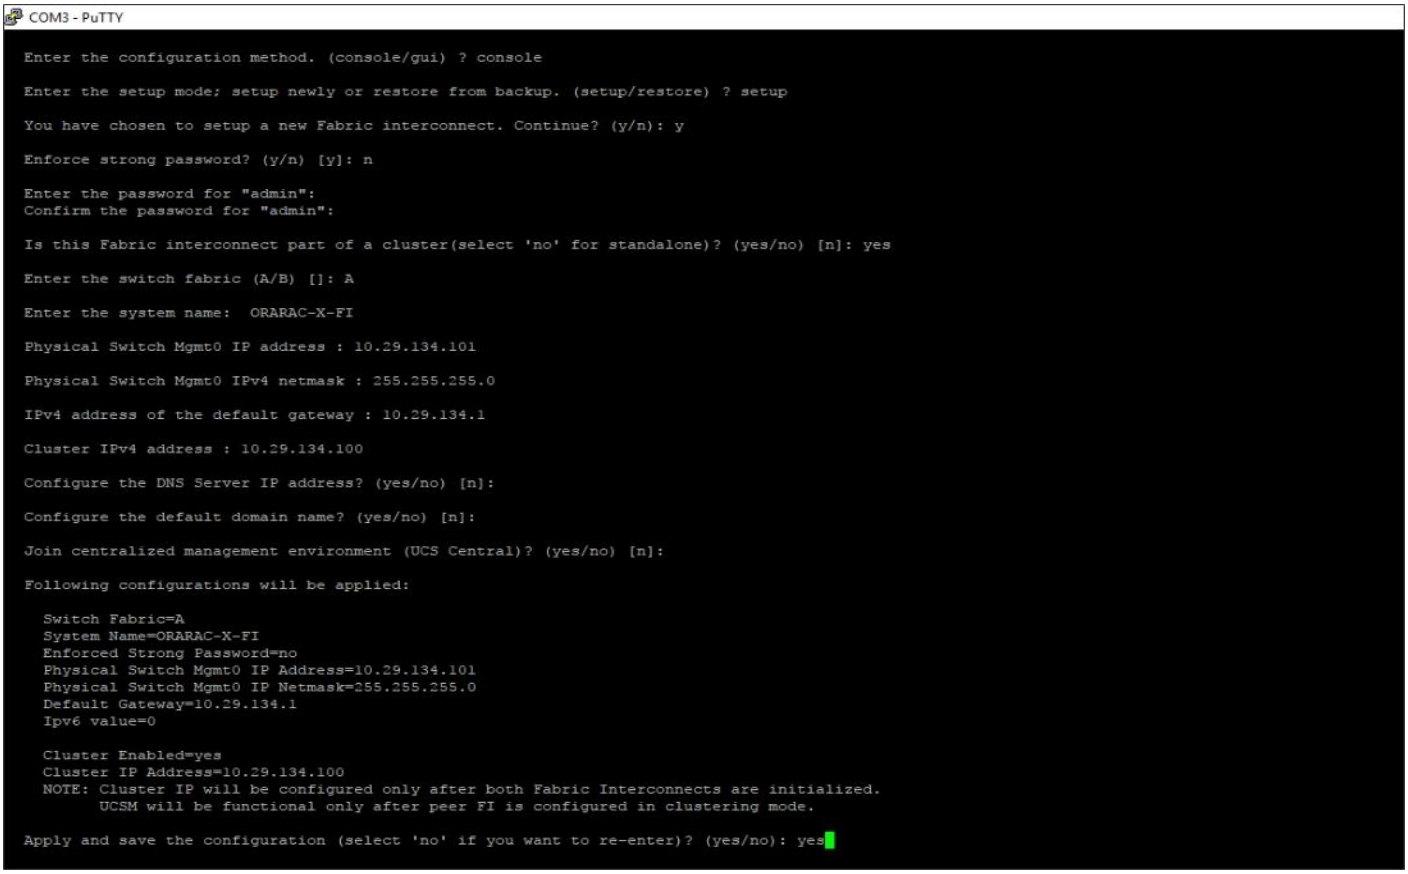

3. Review the settings on the console. Answer yes to Apply and Save the configuration.
4. Wait for the login prompt to make sure the configuration has been saved to Fabric Interconnect A.
5. Connect the console port on the second Fabric Interconnect and do as follows:

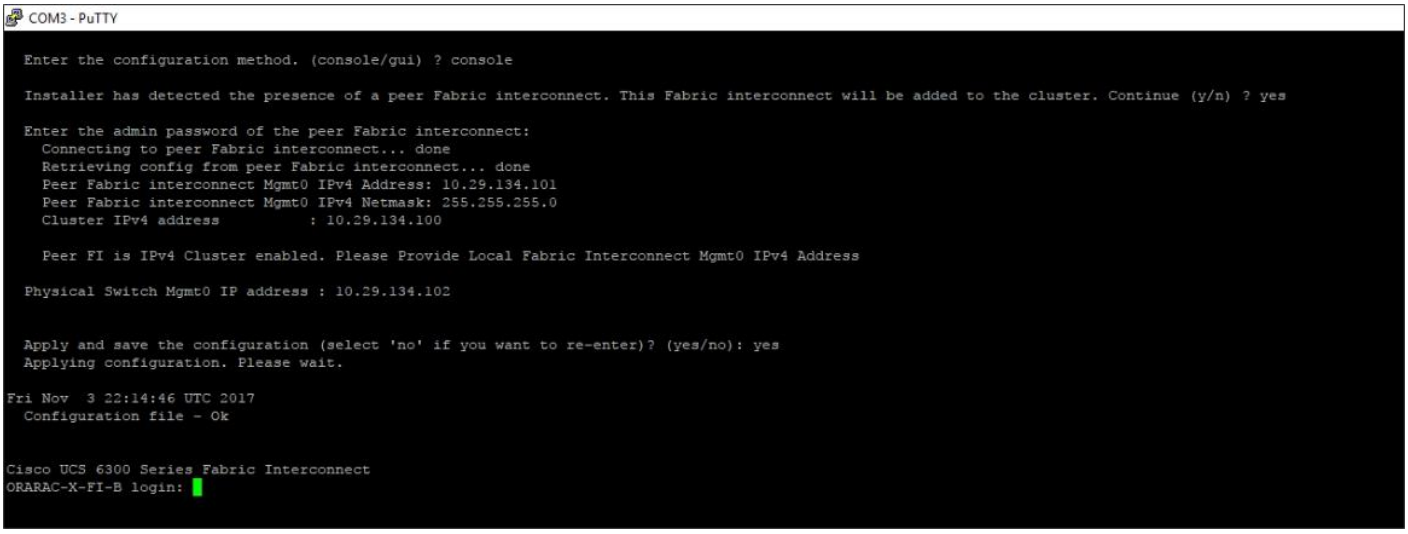

6. Review the settings on the console. Answer yes to Apply and Save the configuration.
7. Wait for the login prompt to make sure the configuration has been saved to Fabric Interconnect B.

To log into the Cisco Unified Computing System (Cisco UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address configured above.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted, accept the security certificates.

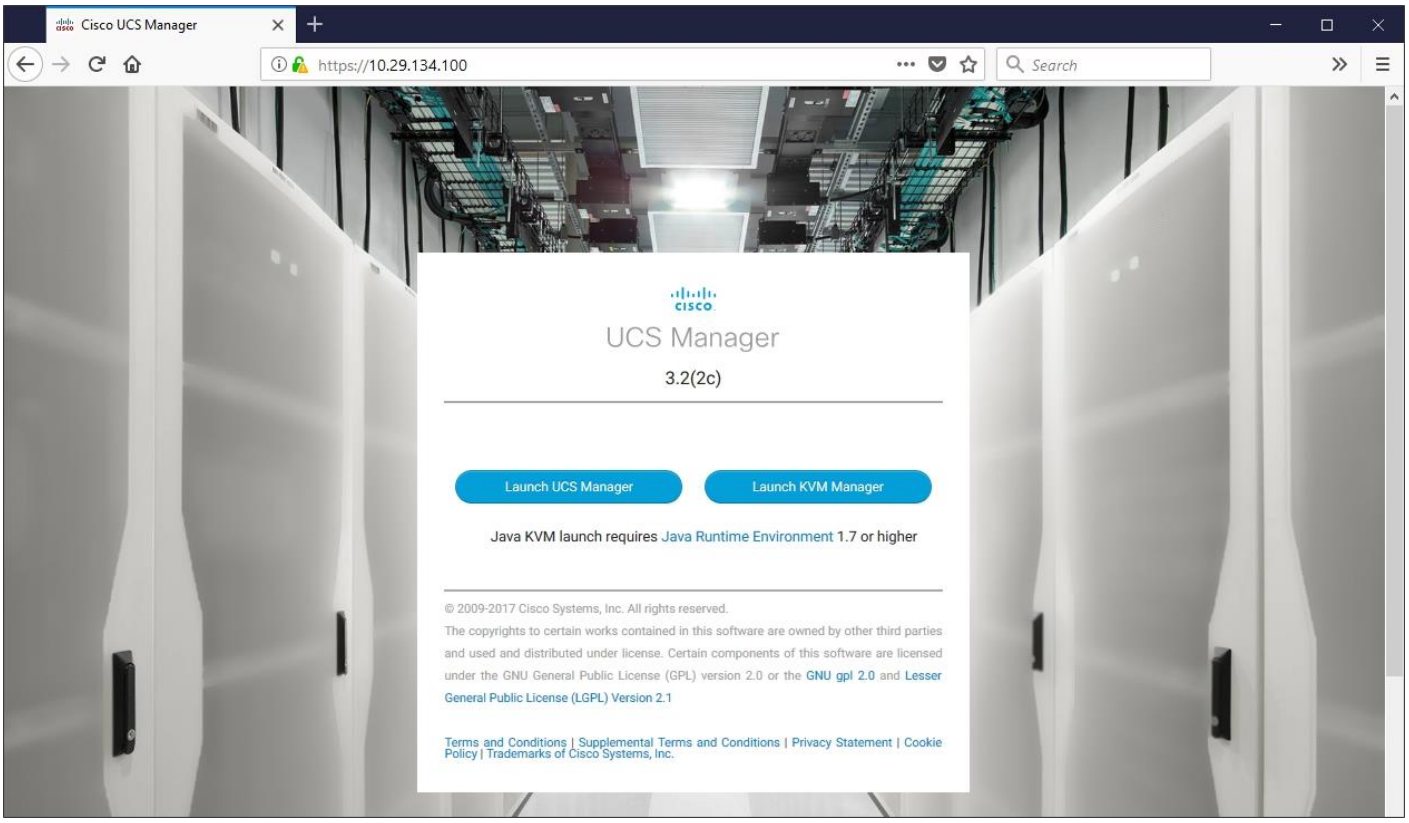

4. When prompted, enter the user name and password enter the password.
5. Click “Log In” to login to Cisco UCS Manager.

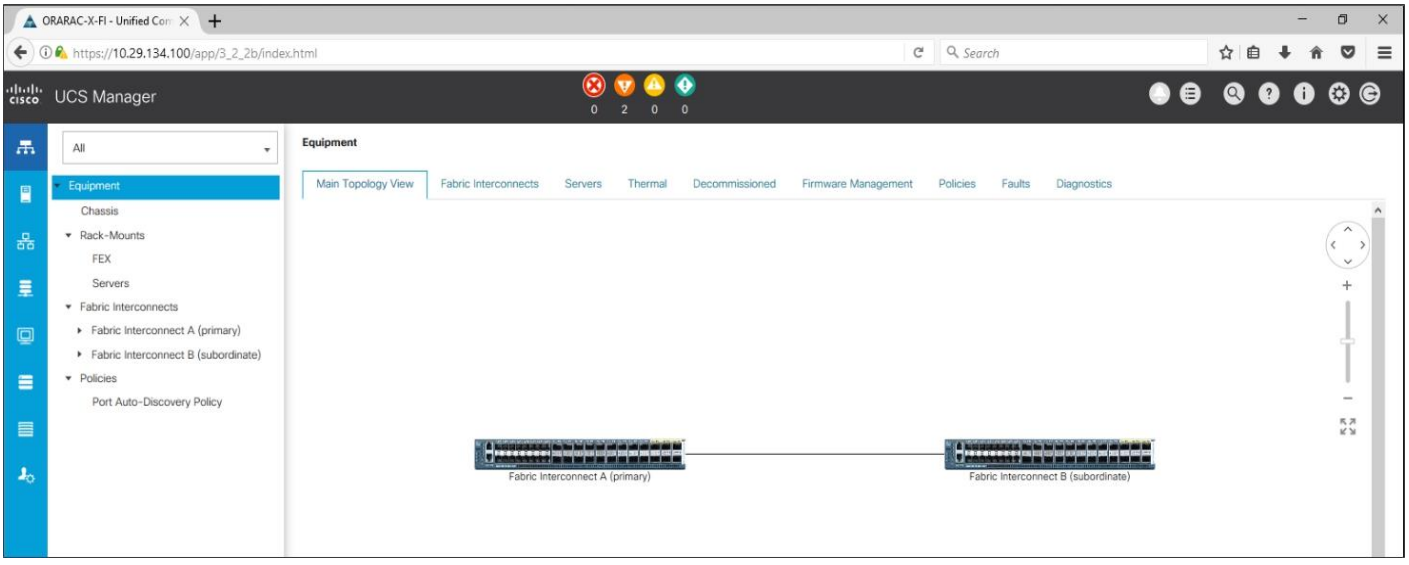

Set Fabric Interconnects to Fibre Channel End Host Mode

To set the Fabric Interconnects to the Fibre Channel End Host Mode, complete the following steps:

1. On the Equipment tab, expand the Fabric Interconnects node and click Fabric Interconnect A.
2. On the General tab in the Actions pane, click Set FC End Host mode.
3. Follow the dialogs to complete the change.

Both Fabric Interconnects automatically reboot sequentially when you confirm you want to operate in this mode.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Admin tab.
2. Select All > Time zone Management.
3. In the Properties pane, select the appropriate time zone in the Time zone menu.
4. Click Save Changes and then click OK.
5. Click Add NTP Server.
6. Enter the NTP server IP address and click OK.
7. Click OK to finish.

Configure Fabric Interconnects for Chassis and Blade Discovery

Cisco UCS 6332-16UP Fabric Interconnects are configured for redundancy. It provides resiliency in case of failures. The first step is to establish connectivity between blades and Fabric Interconnects.

Configure Global Policies

The chassis discovery policy determines how the system reacts when you add a new chassis. We recommend using the platform max value as shown. Using platform max helps ensure that Cisco UCS Manager uses the maximum number of IOM uplinks available.

To configure global policies, complete the following steps:

1. Go to Equipment > Policies (right pane) > Global Policies > Chassis/FEX Discovery Policies. As shown in the screenshot below, select Action as “Platform Max” from the drop-down list and set Link Grouping to Port Channel.
2. Click Save Changes.
3. Click OK.

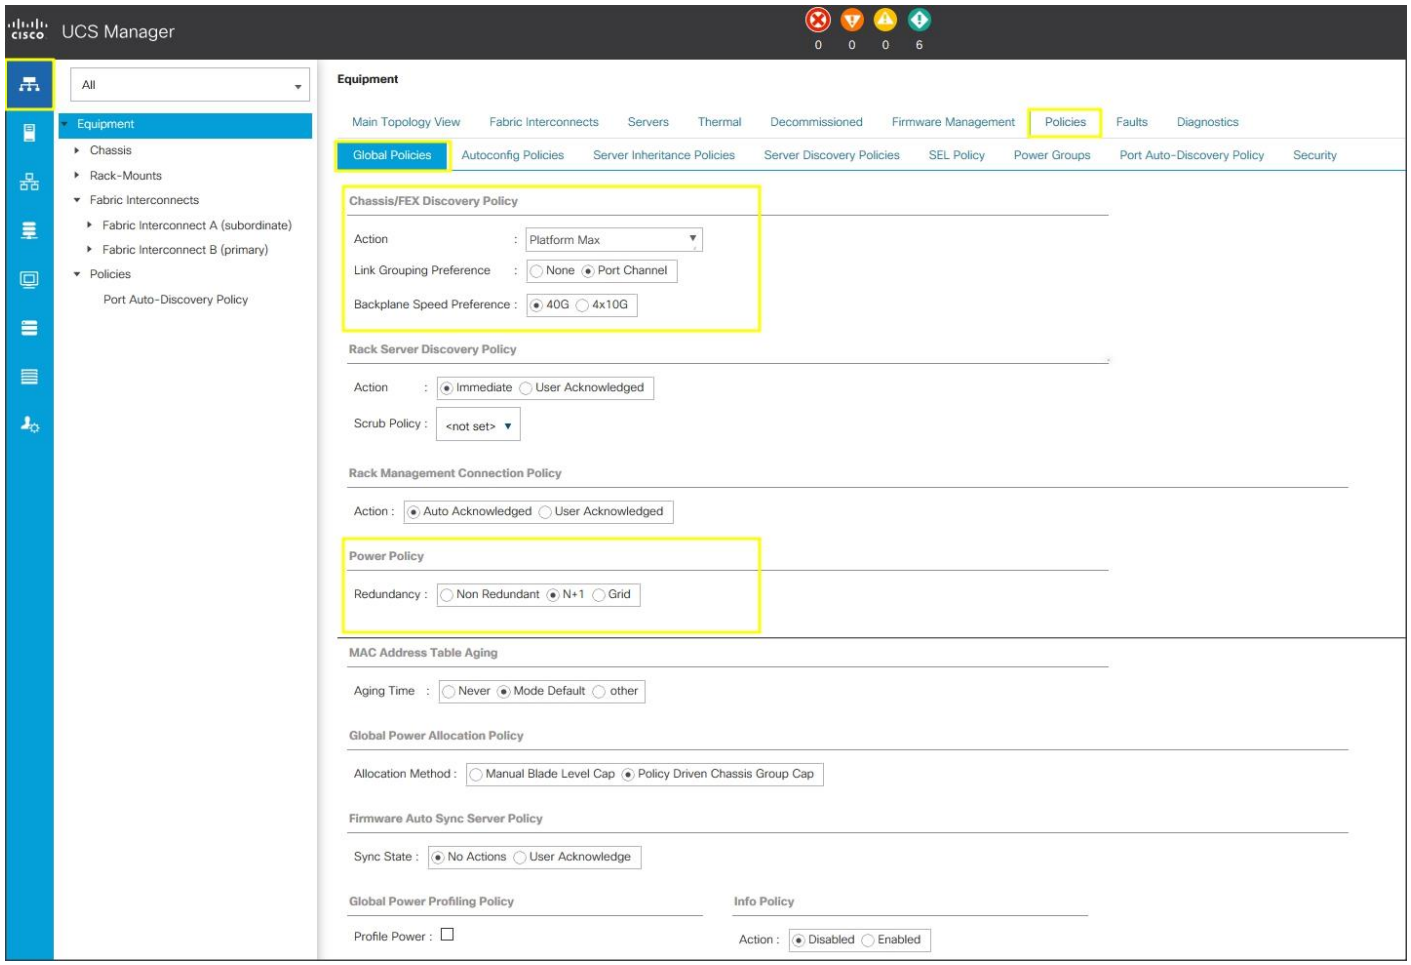

Figure 5 illustrates the advantage of having Discrete mode versus Port Channel mode.

Figure 5 Fabric Ports: Discrete vs. Port Channel Mode

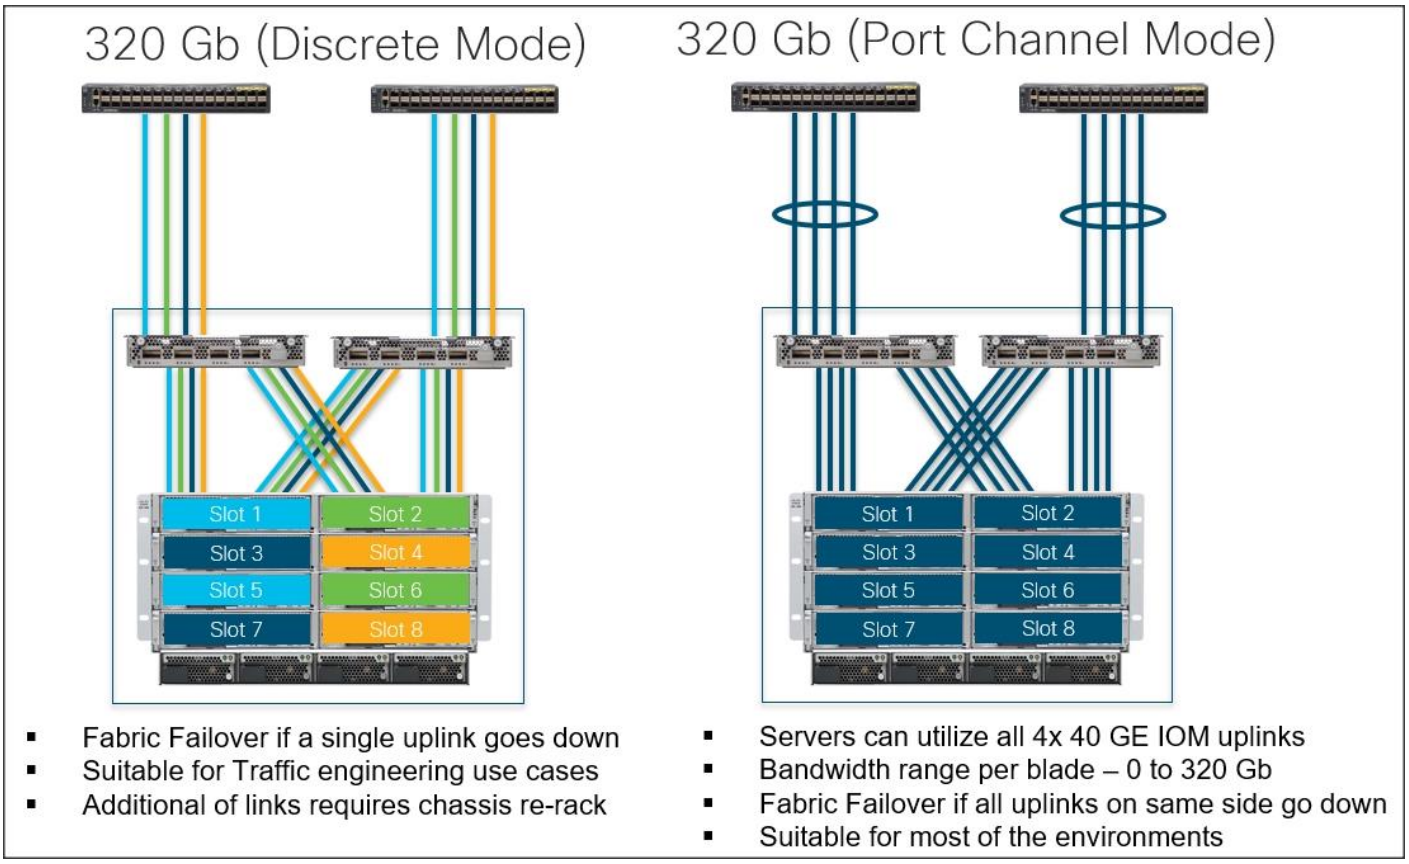

Configure Server Ports

Configure Server Ports to initiate Chassis and Blade discovery. To configure server ports, complete the following steps:

1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > Ethernet Ports.
2. Select the ports (for this solution ports are 17-24) which are connected to the Cisco IO Modules of the two B-Series 5108 Chassis.
3. Right-click and select “Configure as Server Port”.
4. Click Yes to confirm and click OK.

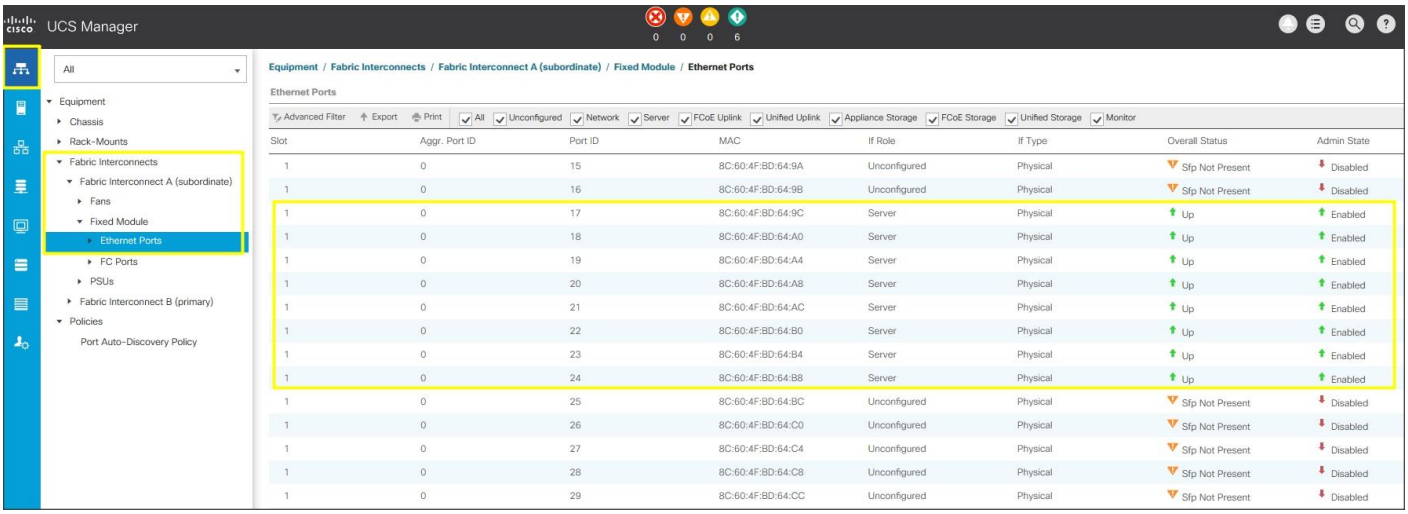

5. Repeat the steps above for Fabric Interconnect B.

6. After configuring Server Ports, acknowledge both the Chassis. Go to Equipment >Chassis > Chassis 1 > General > Actions > select “Acknowledge Chassis”. Similarly, acknowledge the chassis 2.

7. After acknowledging both the chassis, re-acknowledge all the servers placed in the chassis. Go to Equipment > Chassis 1 > Servers > Server 1 > General > Actions > select Server Maintenance > select option “Re-acknowledge” and click OK. Repeat this process to re-acknowledge all eight Servers.

8. When the acknowledgement of the Servers is completed, verify the Port-channel of Internal LAN. Go to the LAN tab > Internal LAN > Internal Fabric A > Port Channels as shown in the screenshot below.

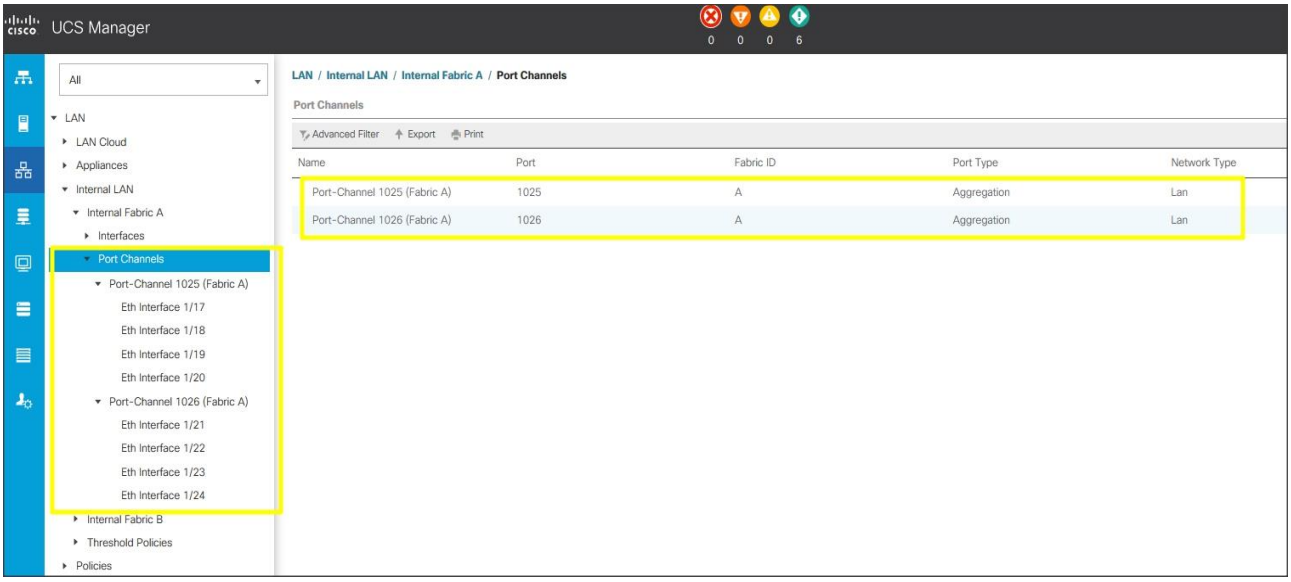

9. Repeat these steps for Internal Fabric B.

Configure LAN and SAN on Cisco UCS Manager

Configure Ethernet Uplink Ports and Fibre Channel (FC) Storage ports as explained in the following section.

Configure Ethernet LAN Uplink Ports

To configure network ports used to uplink the Fabric Interconnects to the Cisco Nexus switches, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports (for this solution ports are 11-14) that are connected to the Nexus switches, right-click them, and select Configure as Network Port.
5. Click Yes to confirm ports and click OK.
6. Verify the Ports connected to Cisco Nexus upstream switches are now configured as network ports.
7. Repeat the above steps for Fabric Interconnect B. The screenshot below shows the network uplink ports for Fabric A.

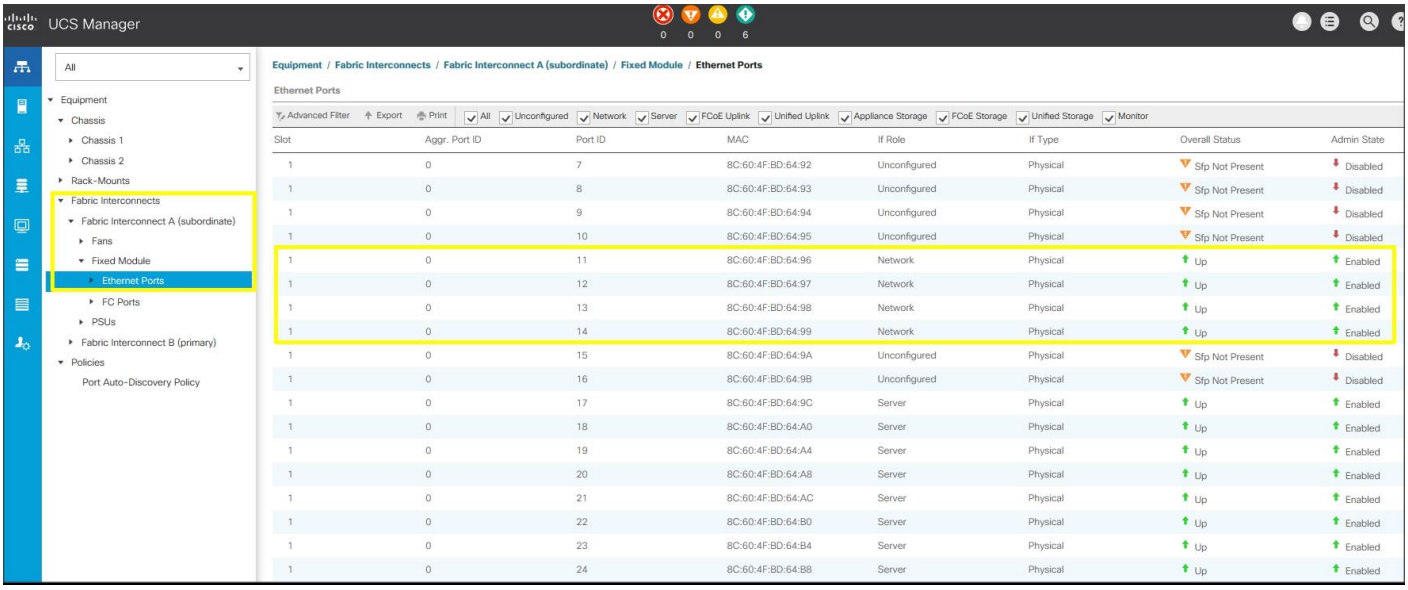

You have now created four uplink ports on each Fabric Interconnect as shown above. These ports will be used to create Virtual Port Channel in the next section.

Create Uplink Port Channels to Cisco Nexus Switches

In this procedure, two port channels were created; one from Fabric A to both Cisco Nexus 9372PX-E switches and one from Fabric B to both Cisco Nexus 9372PX-E switches. To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, expand node Fabric A tree:
   a. Right-click Port Channels.
   b. Select Create Port Channel.
   c. Enter 21 as the unique ID of the port channel.

d. Enter Oracle-Public as the name of the port channel.

e. Click Next.

f. Select Ethernet ports 11-14 for the port channel.

3. Click Finish.

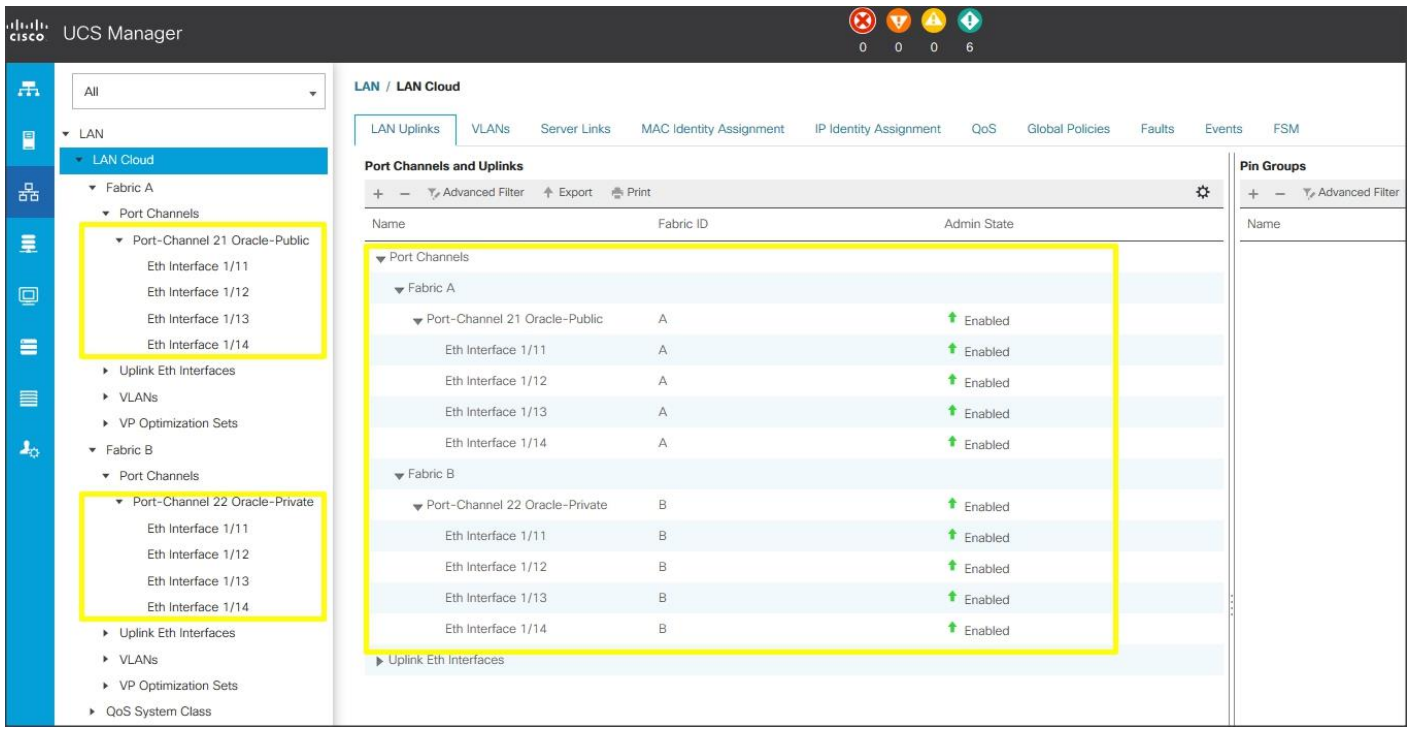

4. Repeat steps 1-3 for Fabric Interconnect B, substituting 22 for the port channel number and Oracle-Private for the name. Your resulting configuration should look like the screenshot above.

Configure FC SAN Uplink Ports

To configure Fibre Channel Uplink ports, complete the following steps:

1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > General tab > Actions pane, click Configure Unified ports.

2. Click Yes to confirm in the pop-up window.

3. Move the slider to the right.

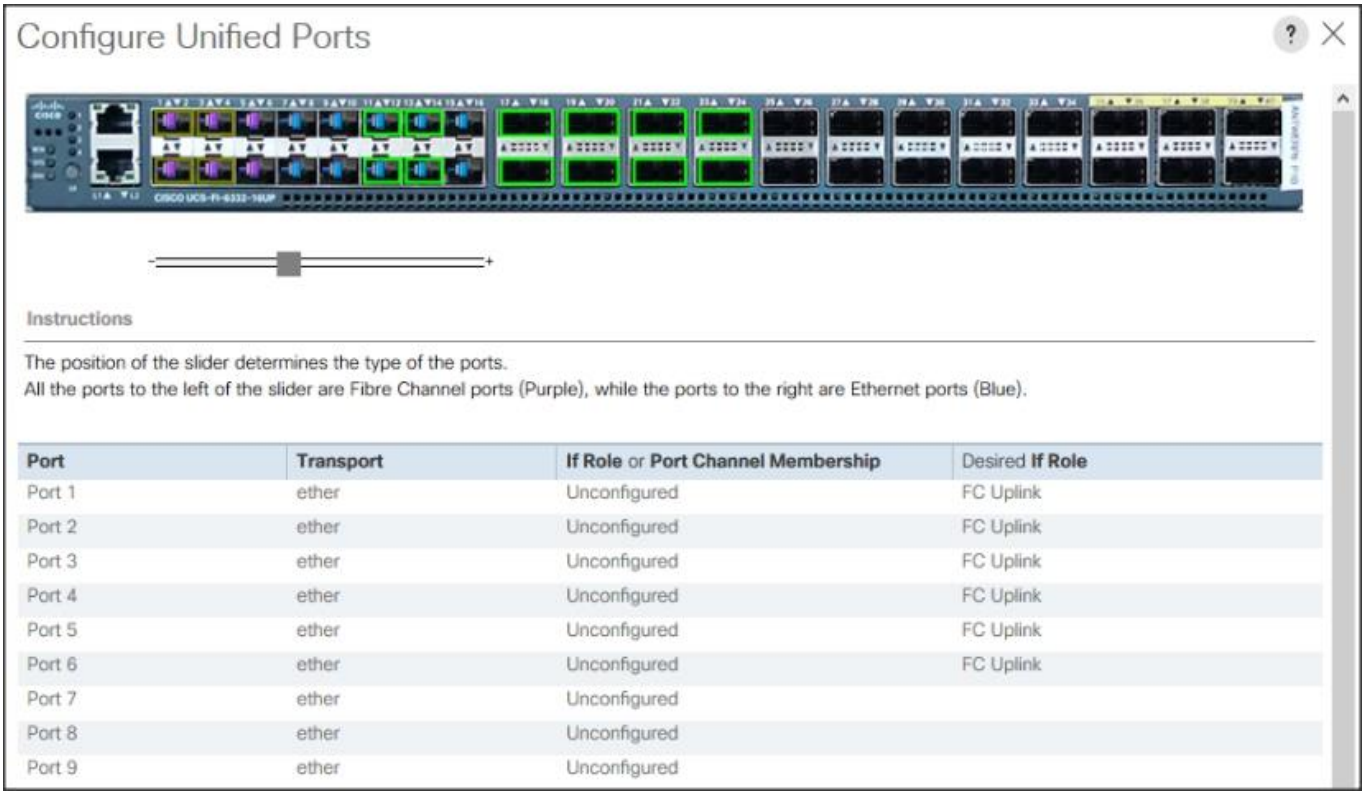

Ports to the right of the slider will become FC ports. For our study, we configured the first six ports on the FI as FC Uplink ports.

4. Click OK.

5. Click Yes to apply the changes.

Applying this configuration will cause the immediate reboot of Fabric Interconnect and/or Expansion Module(s)

6. After the FI reboot, your FC Ports configuration should look like the screenshot below:

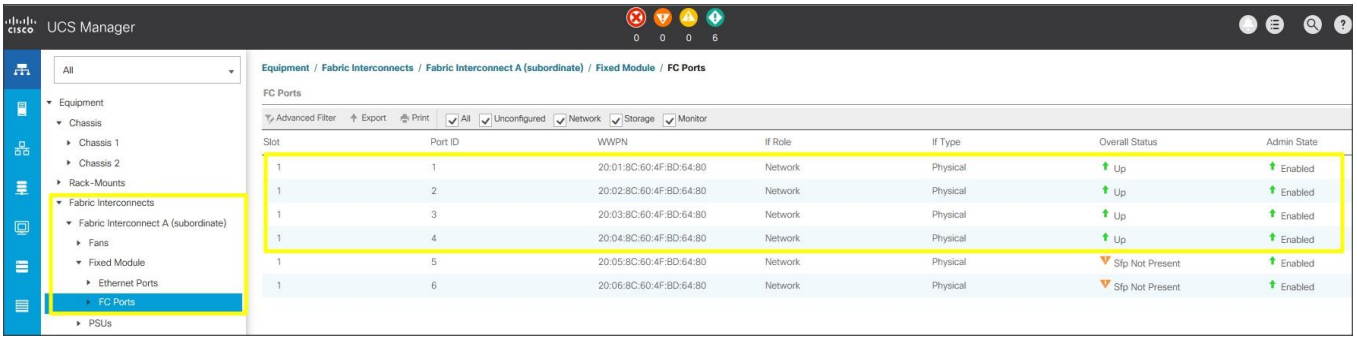

Configure VLAN

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

In this solution, we created two VLANs: one for private network (VLAN 10) traffic and one for public network (VLAN 134) traffic. These two VLANs will be used in the vNIC templates that are discussed later.

It is very important to create both VLANs as global across both fabric interconnects. This way, VLAN identity is maintained across the fabric interconnects in case of NIC failover.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.

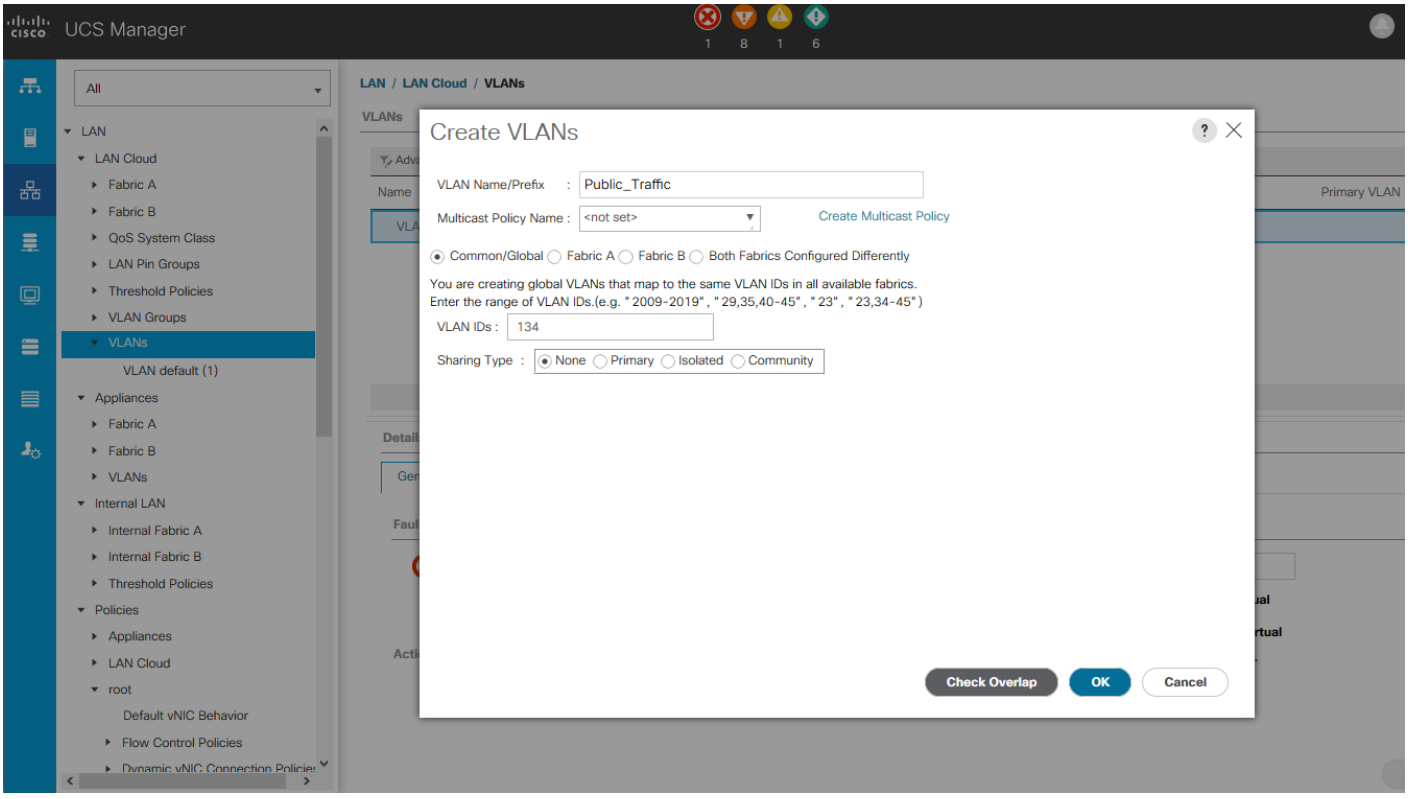

5. Enter Public_Traffic as the name of the VLAN to be used for Public Network Traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 134 as the ID of the VLAN ID.
8. Keep the Sharing Type as None.

9. Click OK and then click OK again.

Similarly, we also created the second VLAN: for private network (VLAN 10) traffic.

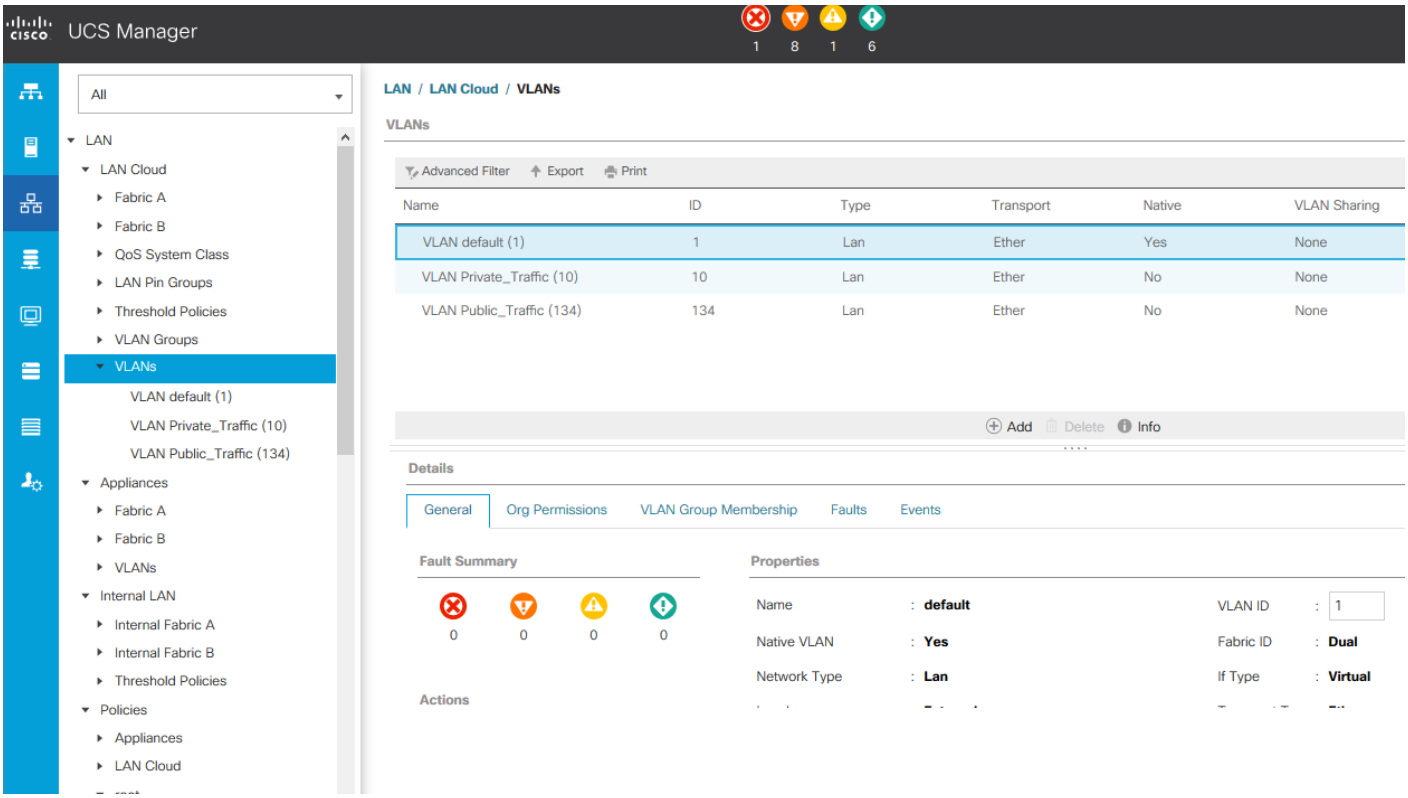

These two VLANs will be used in the vNIC templates that are discussed later.

Configure VSAN

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.

In this solution, we have created two VSANs. VSAN-A 101 and VSAN-B 102 for SAN Boot and Storage Access.

2. Select SAN > SAN Cloud.
3. Under VSANs, right-click VSANs.
4. Select Create VSANs.
5. Enter VSAN 201 as the name of the VSAN.

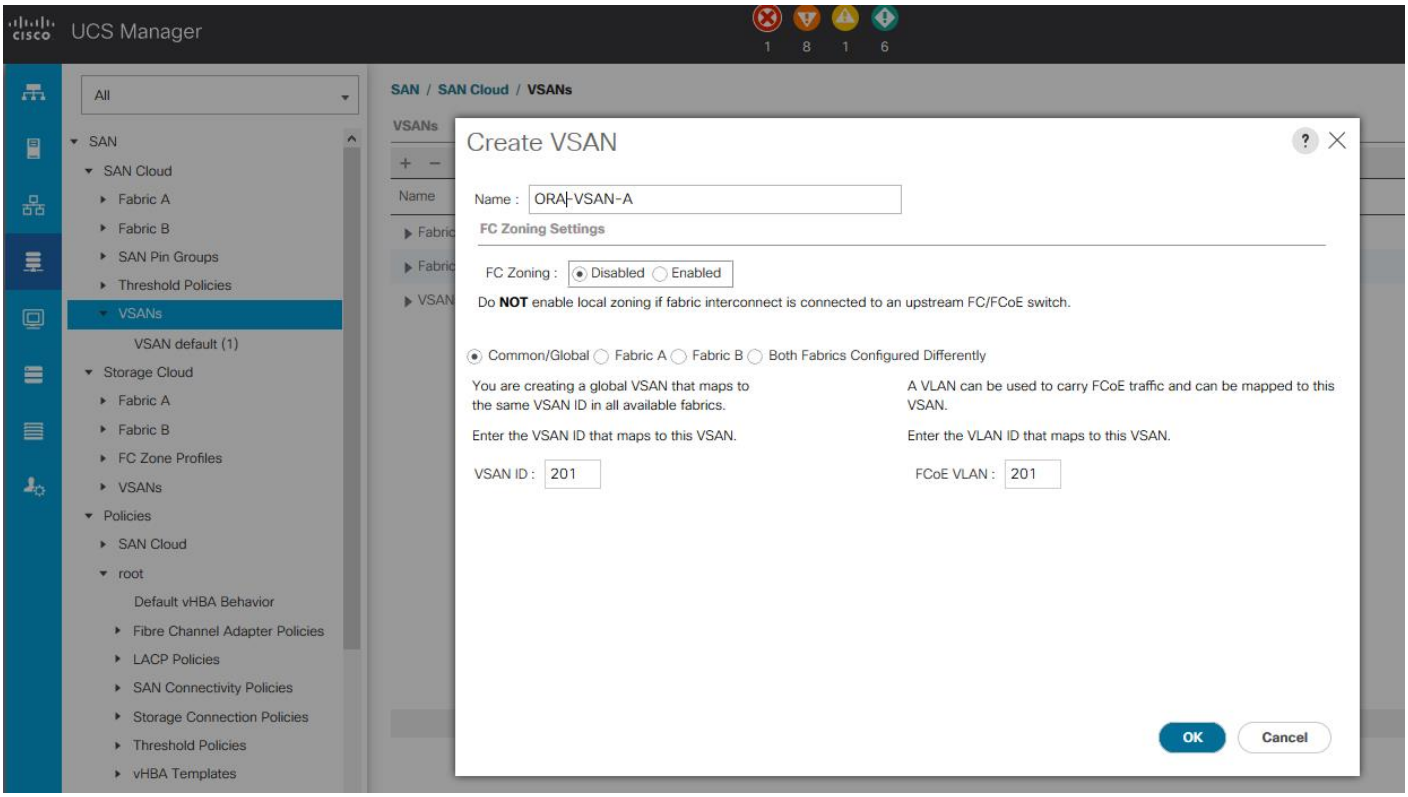

6. Select Fabric A for the scope of the VSAN.
7. Enter 20 as the ID of the VSAN.
8. Click OK and then click OK again.
9. Repeat the above steps to create the VSANs necessary for this solution. VSAN 201 and 202 are configured as shown below:

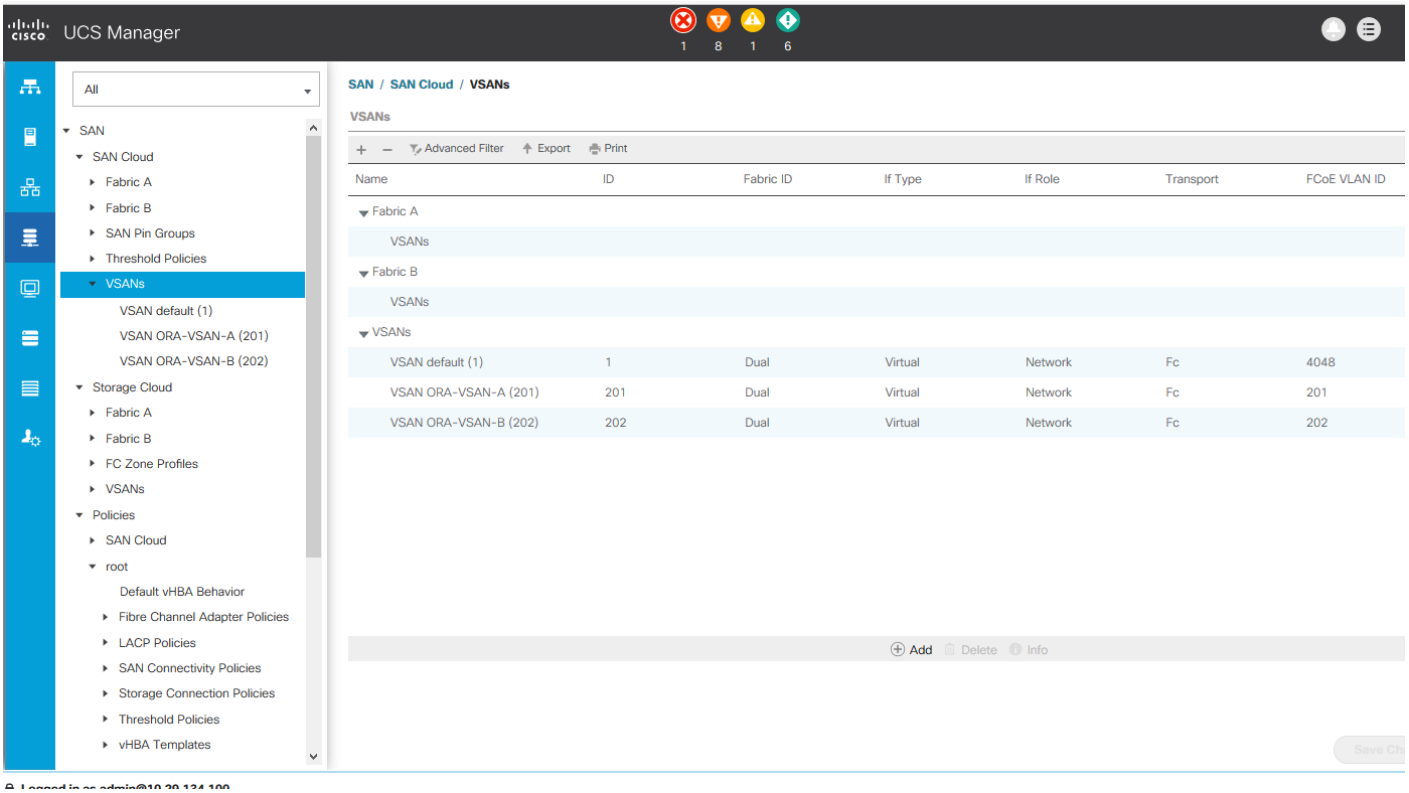

Configure IP, UUID, Server, MAC, WWNN, and WWPN Pools

IP Pool Creation

An IP address pool on the out of band management network must be created to facilitate KVM access to each compute node in the Cisco UCS domain. To create a block of IP addresses for server KVM access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.

2. Select Pools > root > IP Pools > click Create IP Pool.

We named the IP Pool as ORA-IP-Pool for this solution.

3. Select option Sequential to assign IP in sequential order then click Next.

4. Click Add IPv4 Block.

5. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gate-way information as shown in the screenshot.

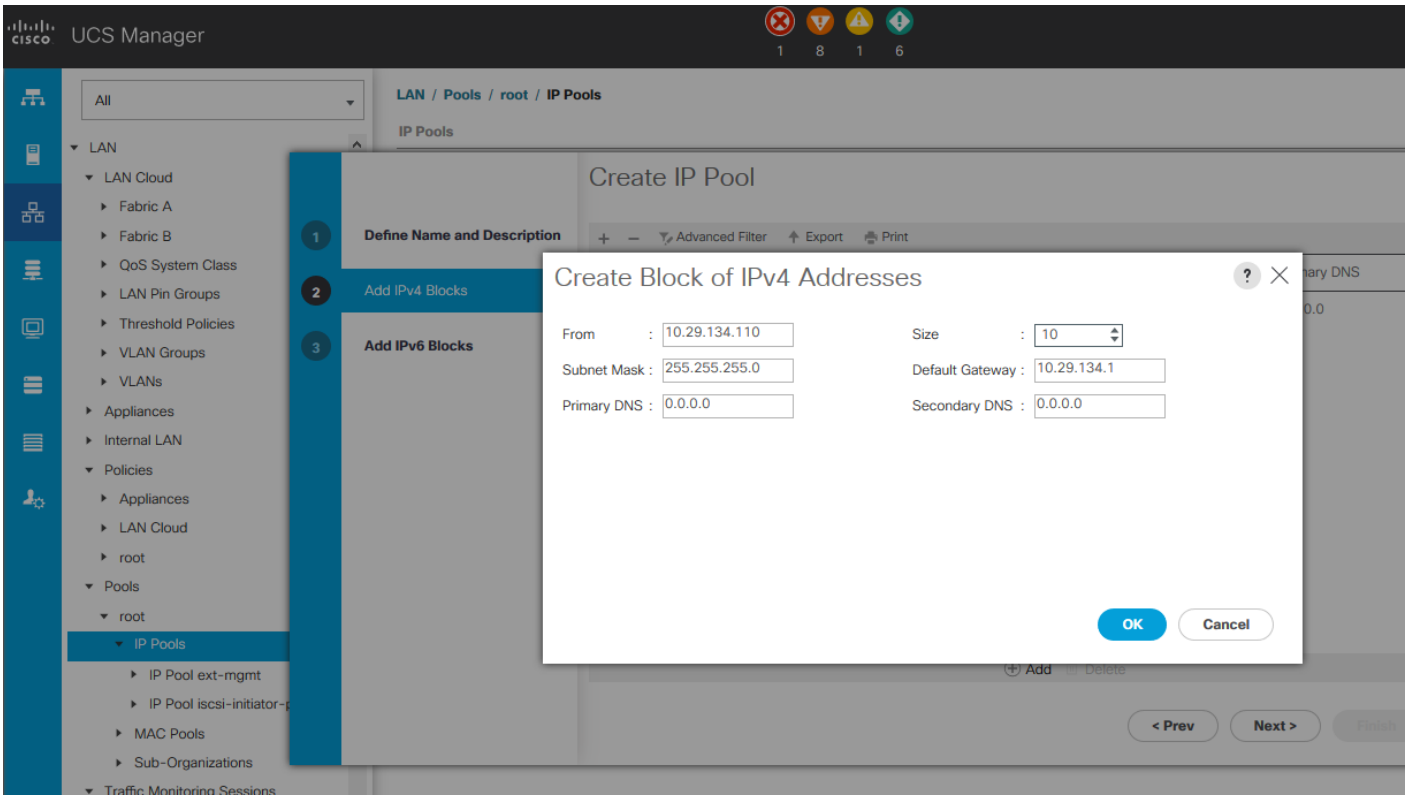

6. Click Next and then click Finish to create the IP block.

UUID Suffix Pool Creation

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools and then select Create UUID Suffix Pool.

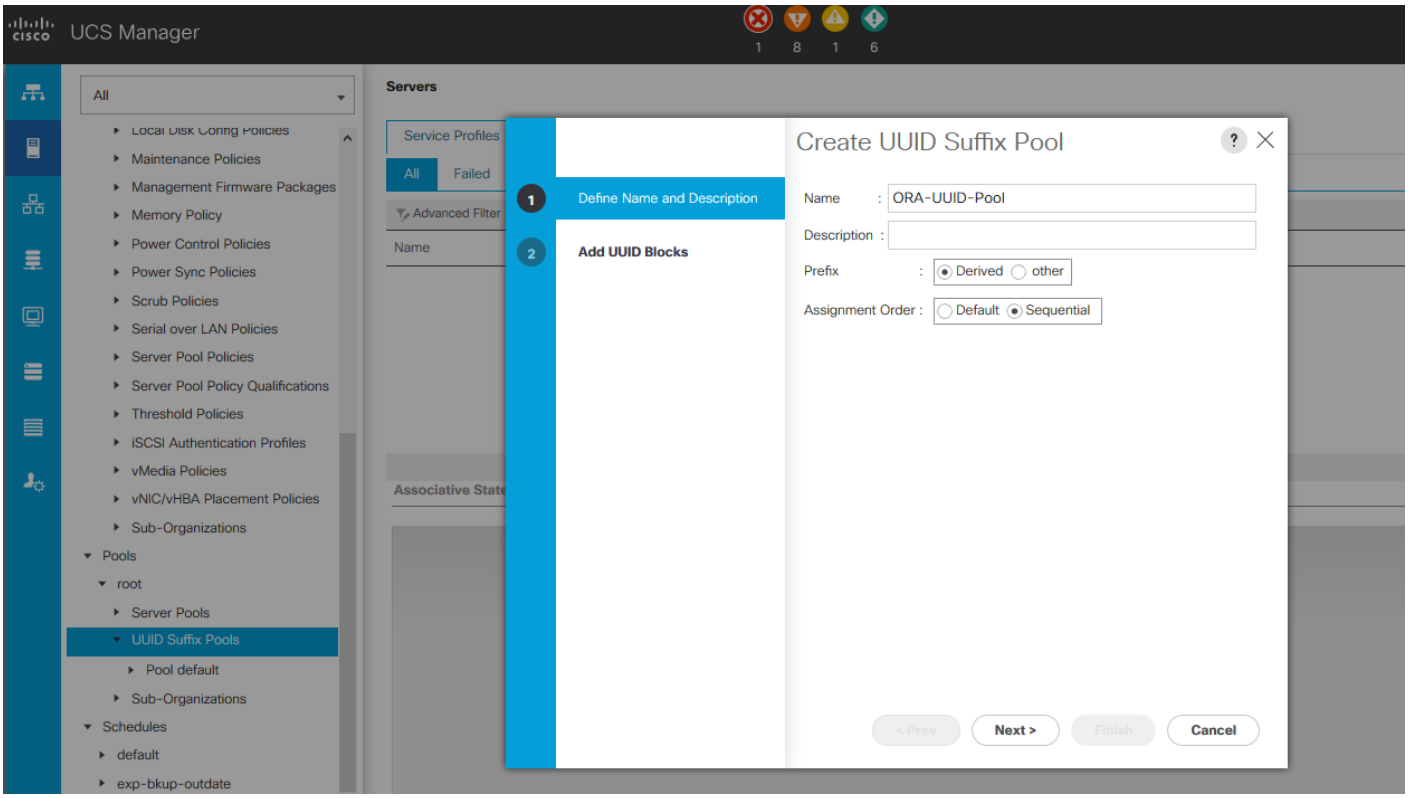

4. Enter ORA-UUID-Pool as the name of the UUID name.
5. Optional: Enter a description for the UUID pool.
6. Keep the prefix at the derived option and select Sequential in as Assignment Order then click Next.
7. Click Add to add a block of UUIDs.
8. Create a starting point UUID as per your environment.

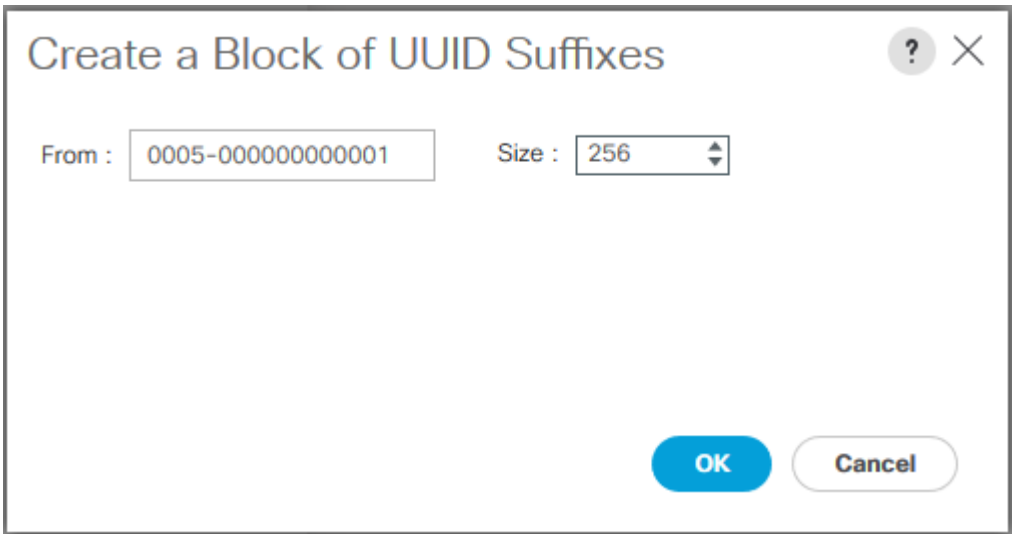

9. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

Server Pool Creation

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:

Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root > right-click Server Pools > Select Create Server Pool.
3. Enter ORA-Pool as the name of the server pool.
4. Optional: Enter a description for the server pool then click Next
5. Select all the eight servers to be used for the Oracle RAC management and click > to add them to the server pool.
6. Click Finish and then click OK.

MAC Pool Creation

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > right-click MAC Pools under the root organization.
3. Select Create MAC Pool to create the MAC address pool.
4. Enter ORA-MAC-A as the name for MAC pool.

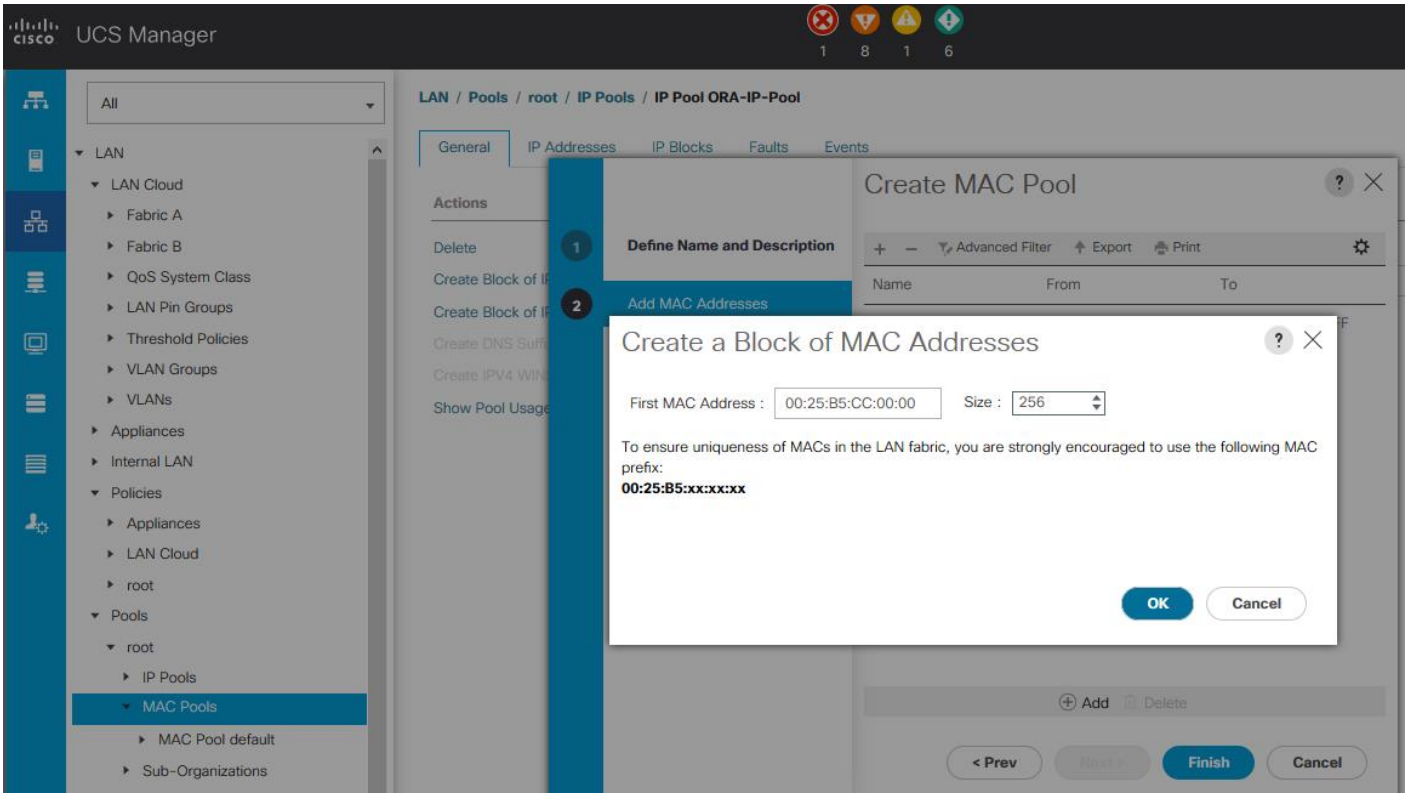

5. Enter the seed MAC address and provide the number of MAC addresses to be provisioned.
6. Click OK and then click Finish.
7. In the confirmation message, click OK.
8. Create MAC Pool B and assign unique MAC Addresses as shown below.

We created Oracle-MAC-A and Oracle-MAC-B as shown below for all the vNIC MAC Addresses.

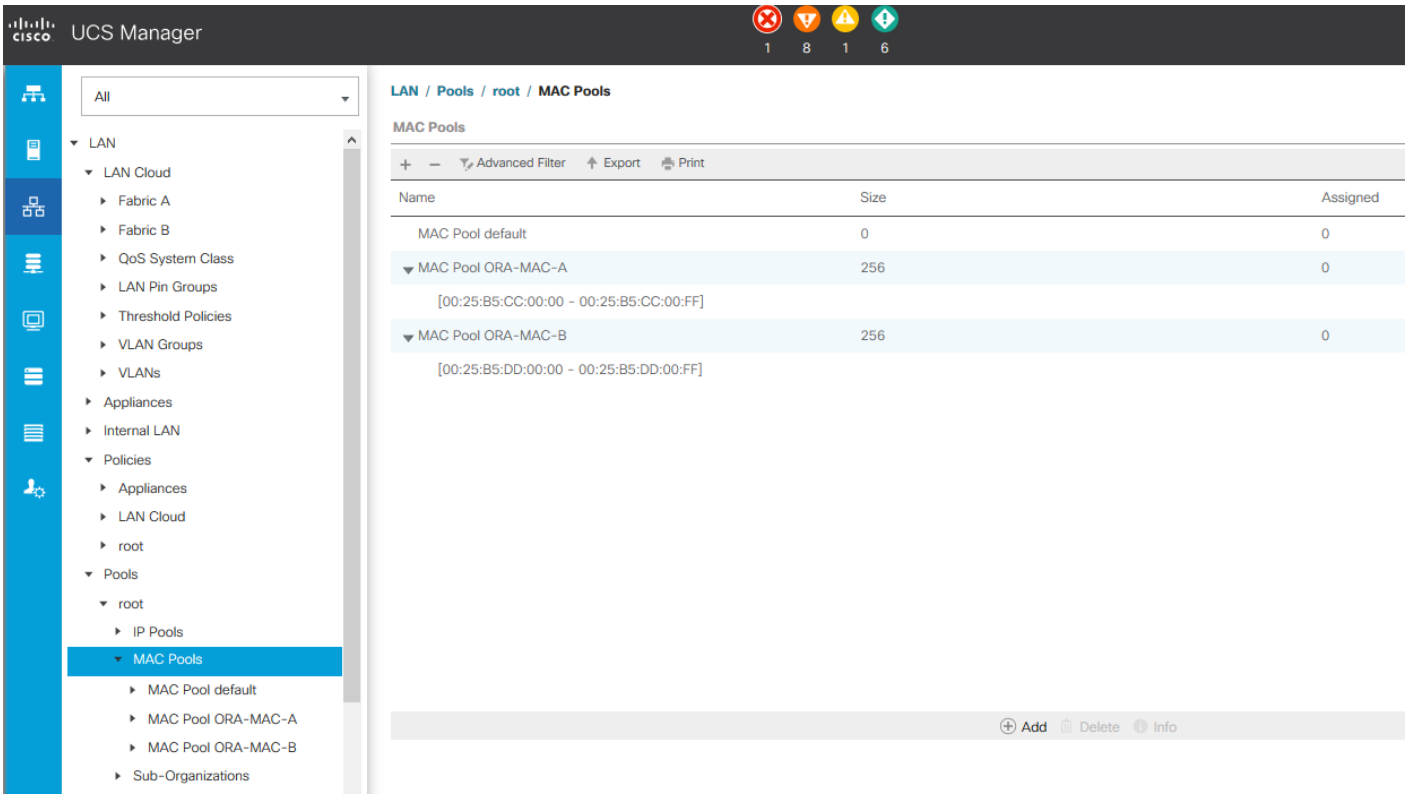

WWNN and WWPN Pool Creation

To configure the necessary WWNN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > Root > WWNN Pools > right-click WWNN Pools > select Create WWNN Pool.
3. Assign name and Assignment Order as sequential, as shown below.

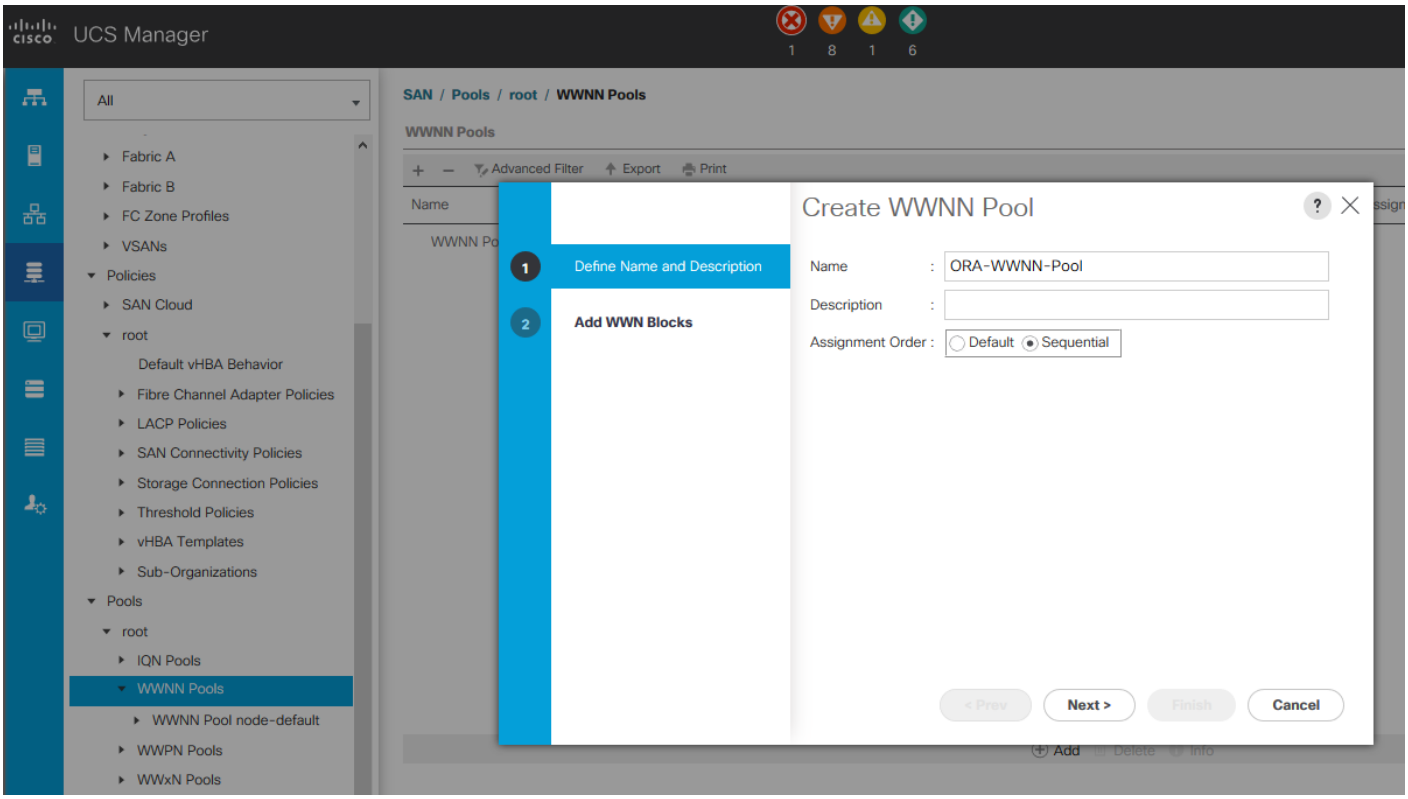

4. Click Next and then click Add to add block of Ports.

5. Enter Block for WWN and size of WWNN Pool as shown below.

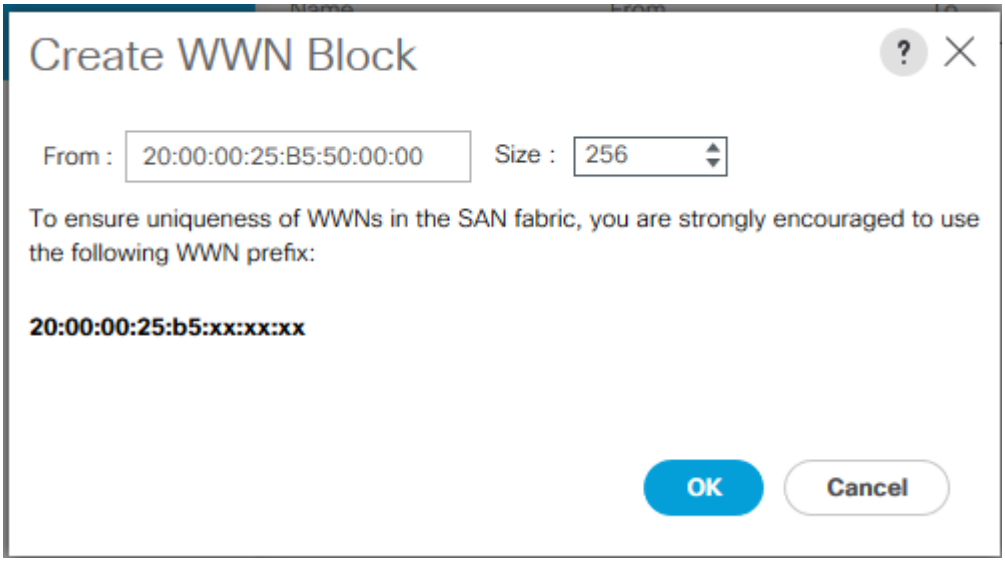

6. Click OK and then click Finish.

To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:

We created two WWPN as ORA-WWPN-A Pool and ORA-WWPN-B as World Wide Port Name as shown below. These WWNN and WWPN entries will be used to access storage through SAN configuration.

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > Root > WWPN Pools > right-click WWPN Pools > select Create WWPN Pool.
3. Assign name as ORA-WWPN-A and Assignment Order as sequential.

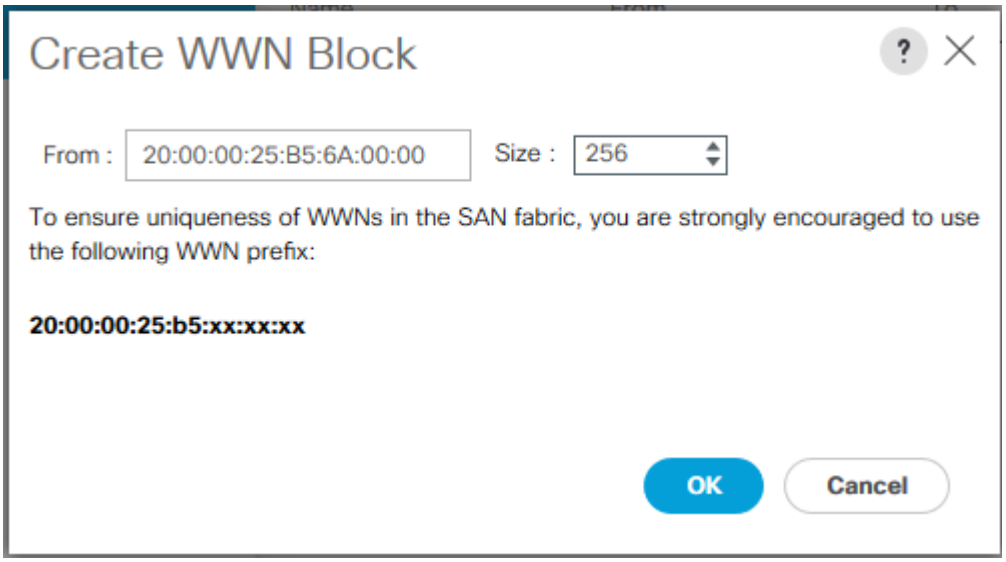

4. Click Next and then click Add to add block of Ports.
5. Enter Block for WWN and size.
6. Click OK and then click Finish.
7. Configure ORA-WWPN-Bs Pool as well and assign the unique block IDs as shown below.

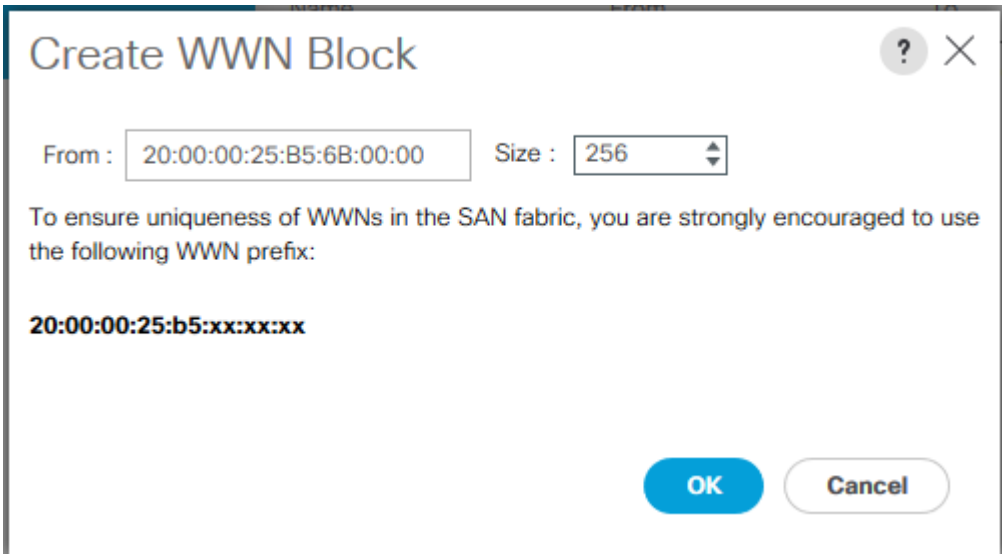

Set Jumbo Frames in both the Cisco Fabric Interconnect

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.

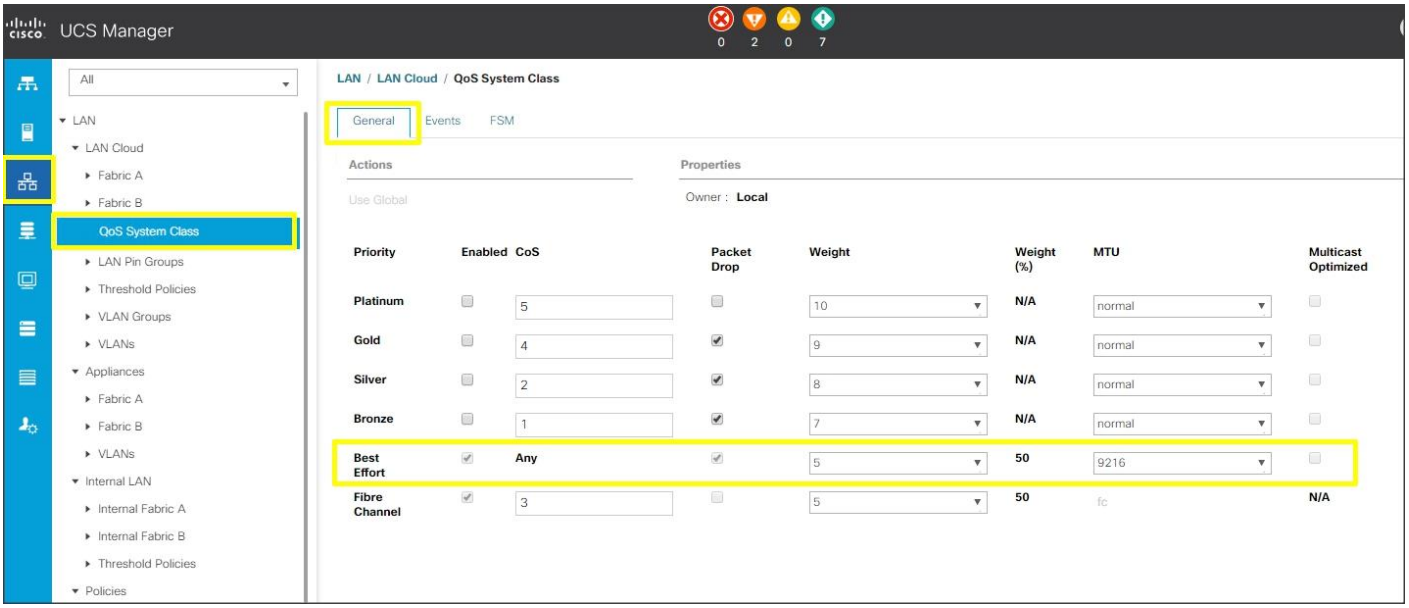

5. Click Save Changes.
6. Click OK.

Create Adapter Policy

To create an Adapter Policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > right-click Adapter Policies.
3. Select Create Ethernet Adapter Policy.
4. Provide a name for the Ethernet adapter policy. Change the following fields and click Save Changes when you are finished:
   - Resources
     - Transmit Queues: 8
     - Ring Size: 4096
     - Receive Queues: 8
     - Ring Size: 4096
     - Completion Queues: 16
     - Interrupts: 32
   - Options
     - Receive Side Scaling (RSS): Enabled

5. Configure adapter policy as shown below:

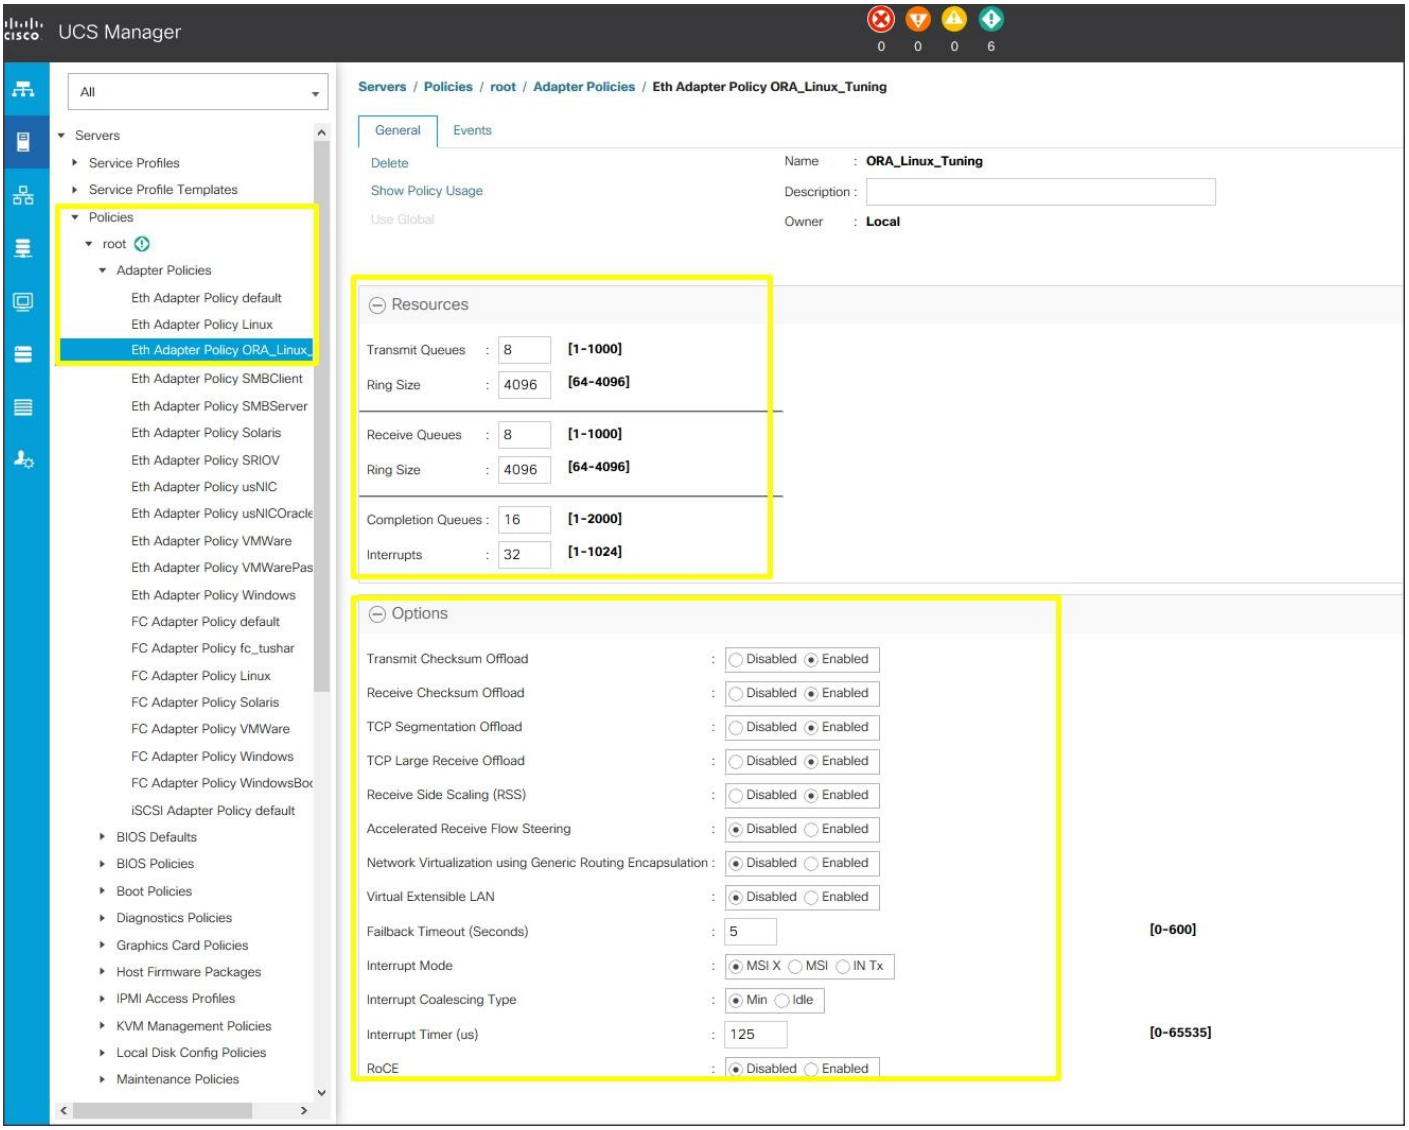

RSS distributes network receive processing across multiple CPUs in multiprocessor systems. This can be one of the following:

- Disabled—Network receive processing is always handled by a single processor even if additional processors are available.
- Enabled—Network receive processing is shared across processors whenever possible.

6. Click OK to finish.

Configure Update Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Maintenance Policies > Default.

3. Change the Reboot Policy to User Ack.
4. Click Save Changes.
5. Click OK to accept the changes.

Configure vNIC and vHBA Template

We created two vNIC template for Public Network and Private Network Traffic. We will use these vNIC Templates during the creation of the Service Profile later in this section.

Create Public vNIC Template

To create vNIC (virtual network interface card) template for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root > vNIC Templates > right-click to vNIC Template and Select "Create vNIC Template"
3. Enter ORA-vNIC-A as the vNIC template name and keep Fabric A selected.
4. Select the Enable Failover checkbox for high availability of the vNIC.
5. Select Template Type as Updating Template.
6. Under VLANs, select the checkboxes default and Public_Traffic and set Native-VLAN as the Public_Traffic.
7. Keep MTU value 1500 for Public Network Traffic.
8. In the MAC Pool list, select ORA-MAC-A.
9. Click OK to create the vNIC template as shown below:

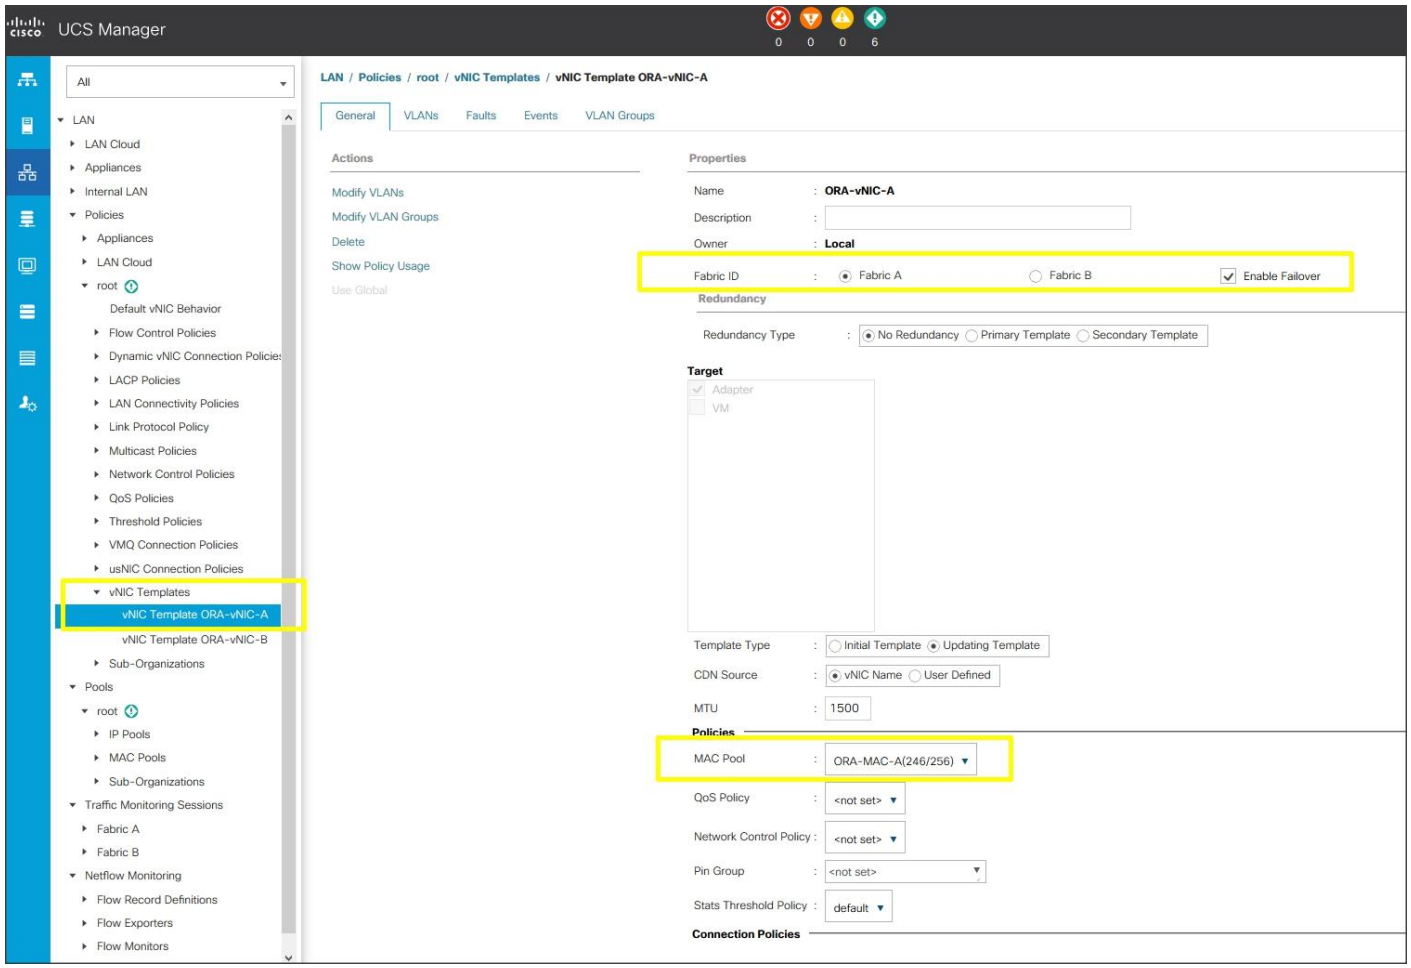

10. Click OK to finish.
11. Create another vNIC template for Private Network Traffic.
12. Enter ORA-vNIC-B as the vNIC template name for Private Network Traffic.
13. Select the Fabric B and Enable Failover for Fabric ID options.
14. Select Template Type as Updating Template.
15. Under VLANs, select the checkboxes default and Private_Traffic and set Native-VLAN as the Private_Traffic.
16. Set MTU value to 9000 and MAC Pool as ORA-MAC-B.
17. Click OK to create the vNIC template as shown below:

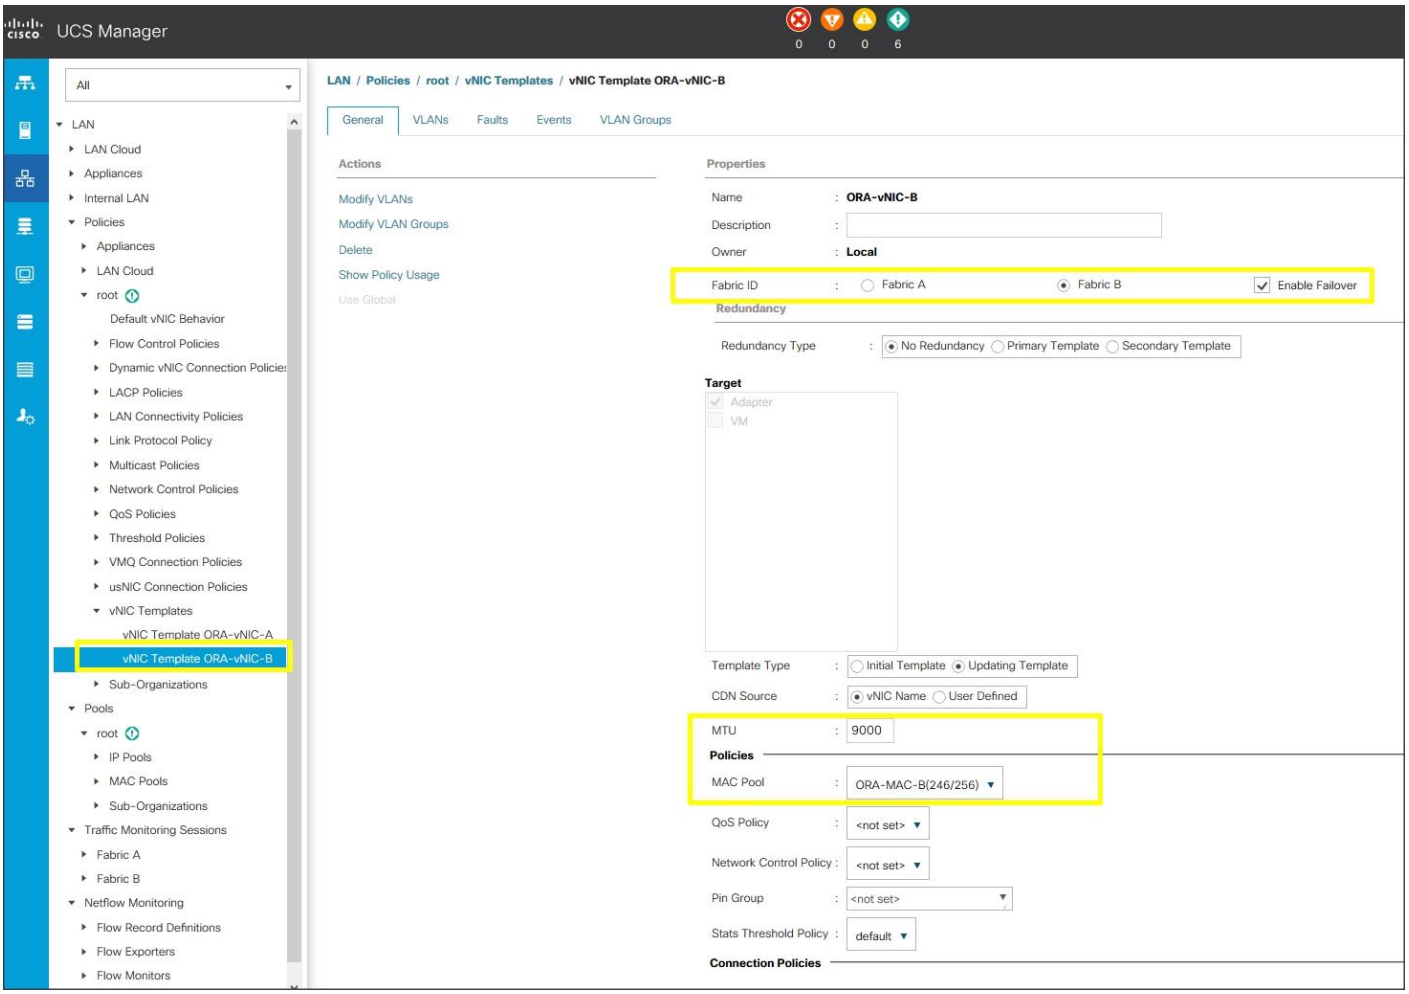

Create Storage vHBA Template

To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root > right-click vHBA Templates > Select “Create vHBA Template” to create vHBAs.
3. Enter name as ORA-HBA-A and keep Fabric A selected.
4. Select VSAN as ORA-VSAN-A and template type to Updating Template.
5. Select WWPN Pool as Oracle-WWPN-A from the drop-down list as shown below.

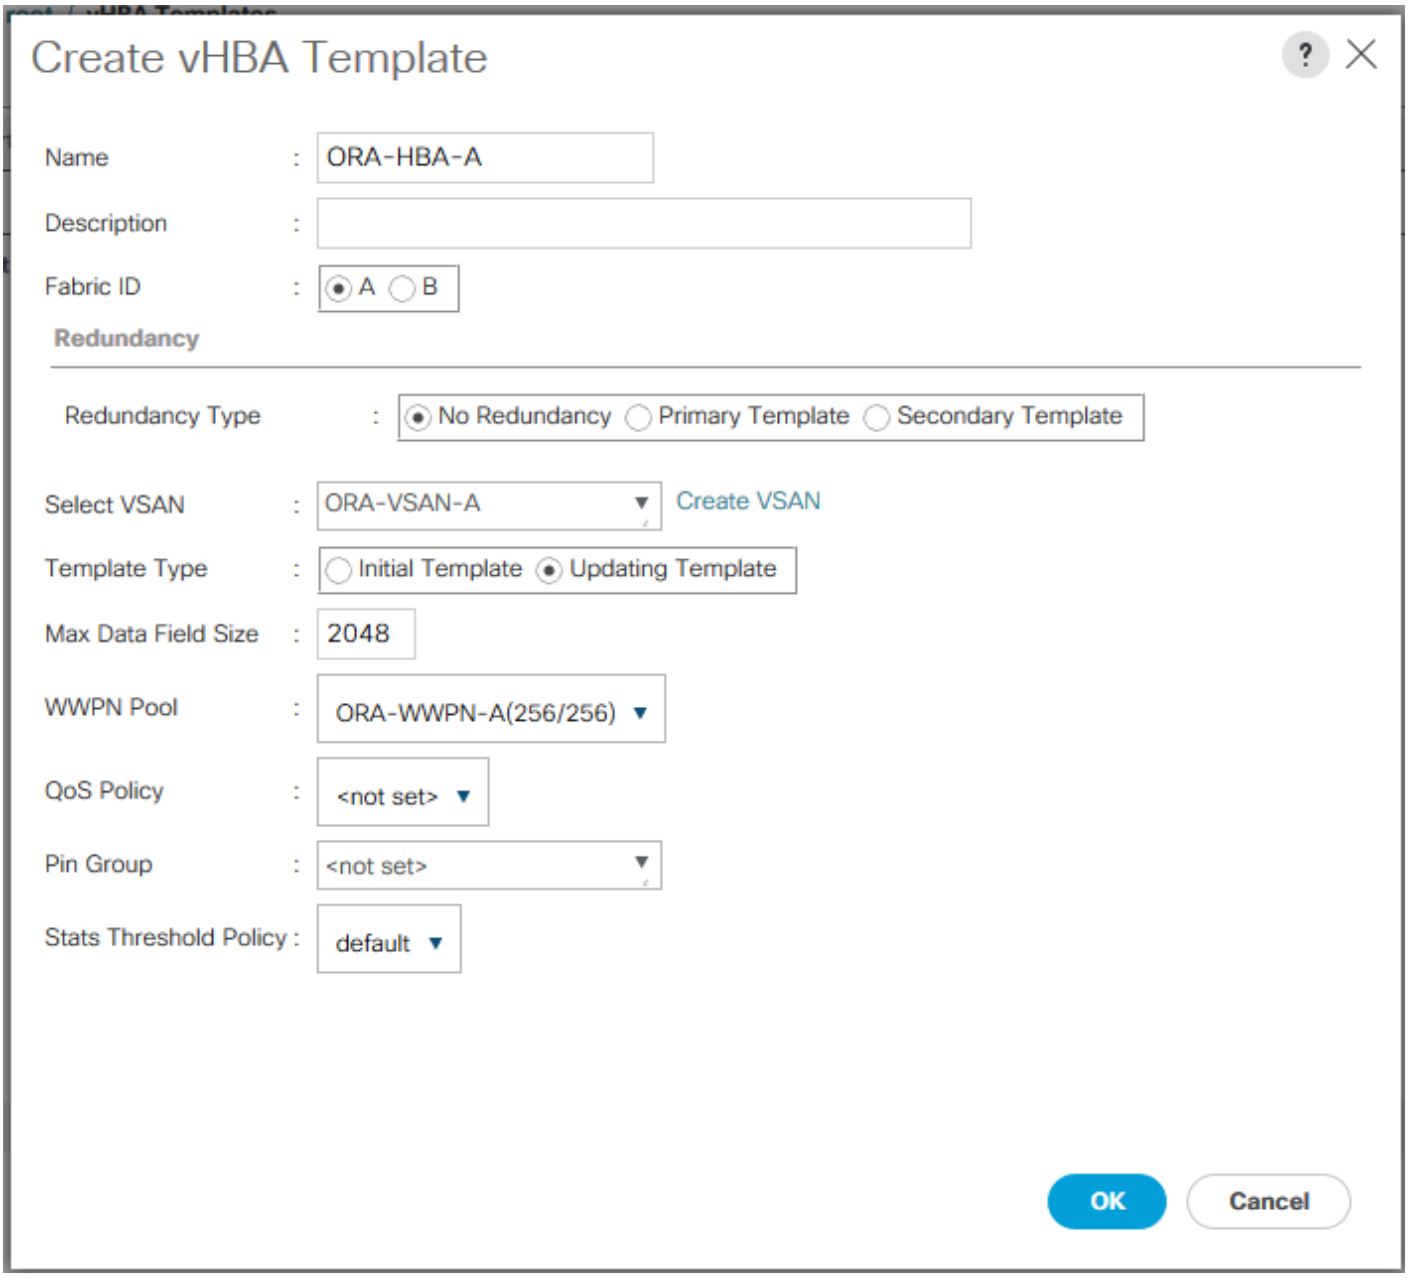

For this solution, we created two vHBA as Oracle-HBA-A and Oracle-HBA-B.

6. Enter name as ORA-HBA-B and select Fabric B. Select WWPN Pool for Oracle-HBA-B as “Oracle-WWPN-B” as shown below:

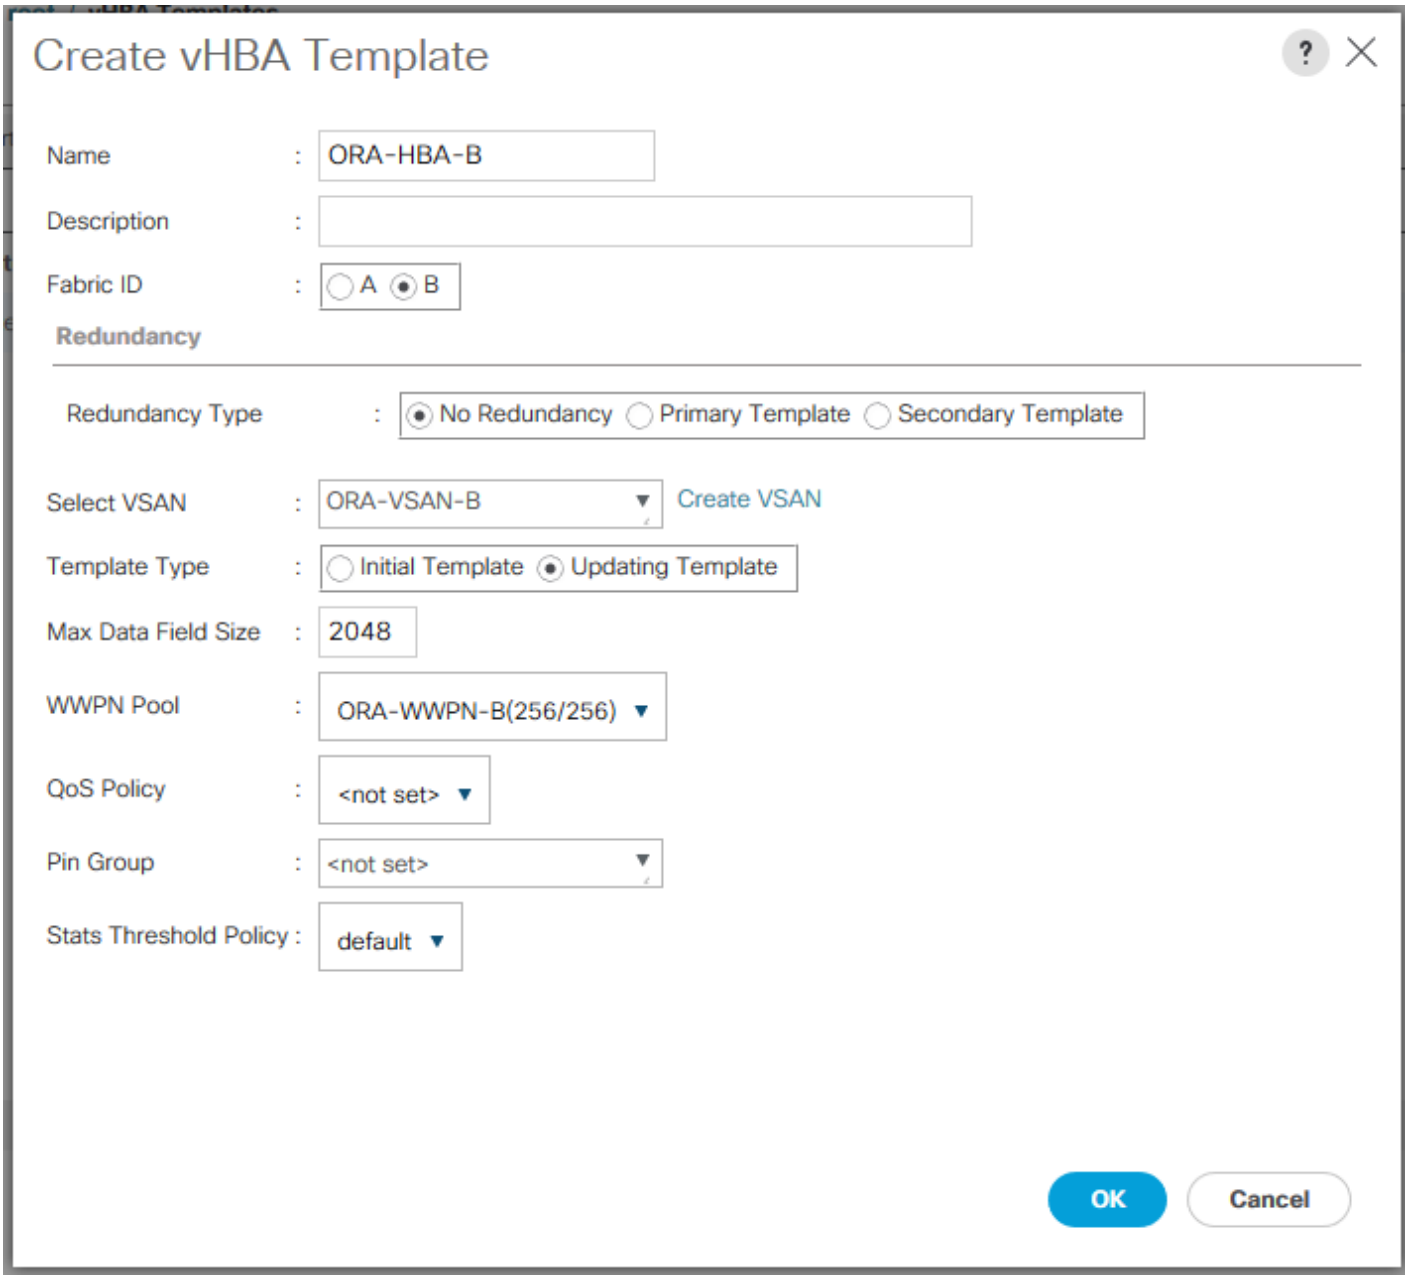

Create Server Boot Policy for SAN Boot

All Oracle nodes were set to boot from SAN for this Cisco Validated Design as part of the Service Profile template. The benefits of booting from SAN are numerous; disaster recovery, lower cooling and power requirements for each server since a local drive is not required, and better performance, to name just a few.

We strongly recommend to use “Boot from SAN” to realize the full benefits of Cisco UCS stateless computing feature, such as service profile mobility.

This process applies to a Cisco UCS environment in which the storage SAN ports are configured in the following section.

A Local disk configuration for the Cisco UCS is necessary if the servers in the environments have a local disk.

To configure Local disk policy, complete the following steps:

1. Go to tab Servers > Policies > root > right-click Local Disk Configuration Policy > Enter “SAN-Boot” as the local disk configuration policy name and change the mode to “No Local Storage.”

2. Click OK to create the policy as shown in the screenshot below:

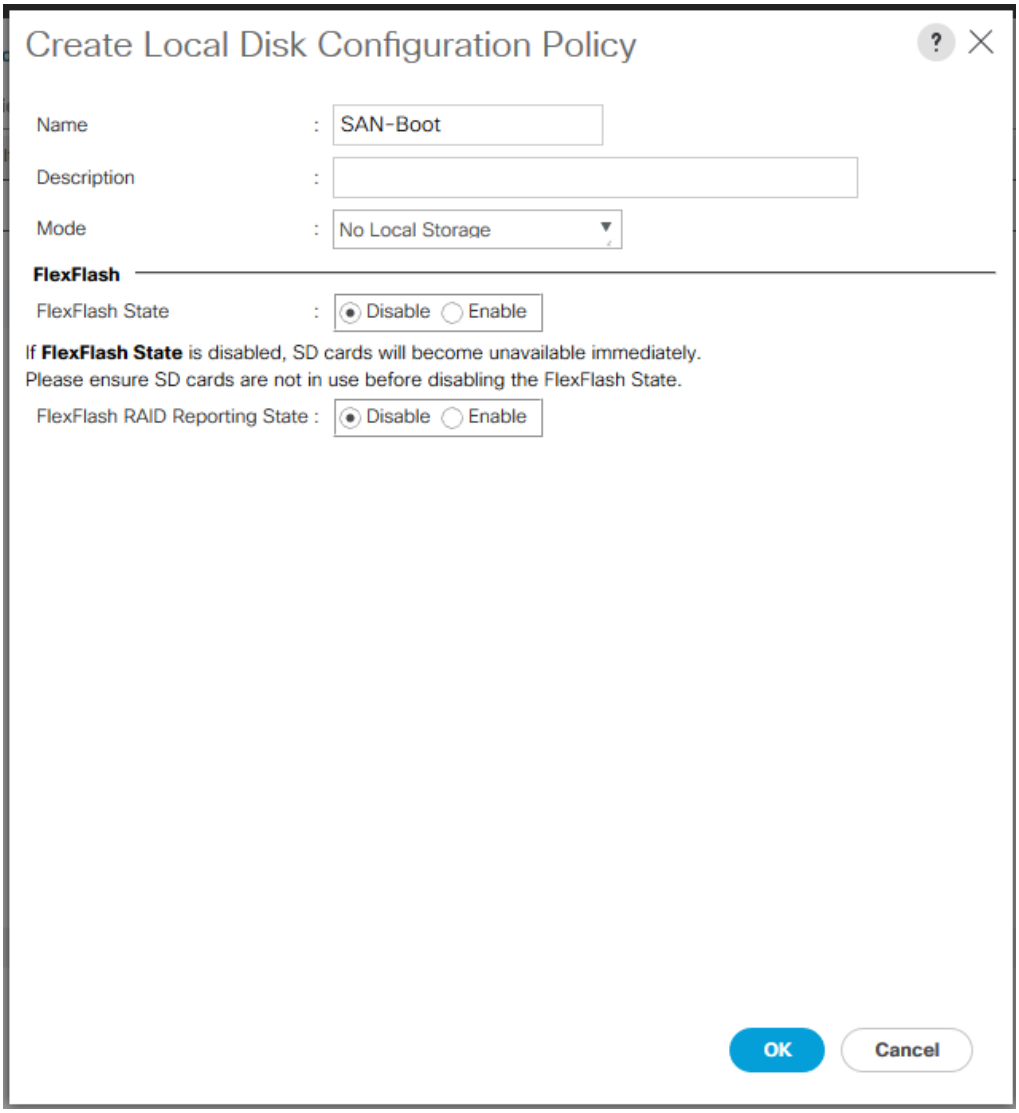

As shown in the screenshot below, the Pure Storage FlashArray have eight active FC connections that go to the Cisco MDS switches. Four FC ports are connected to Cisco MDS-A and the other four FC ports are connected to Cisco MDS-B Switches. All FC ports are 16 Gb/s. The SAN Ports CT0.FC0, CT0.FC6, of Pure Storage FlashArray Controller 0 are connected to Cisco MDS Switch A and CT0.FC1, CT0.FC7 are connected to Cisco MDS Switch B. Similarly, the SAN Ports CT1.FC0, CT1.FC6, of Pure Storage FlashArray Controller 1 are connected to Cisco MDS Switch A and CT1.FC1, CT1.FC7 are connected to Cisco MDS Switch B.

Figure 6 Pure Storage FC Ports

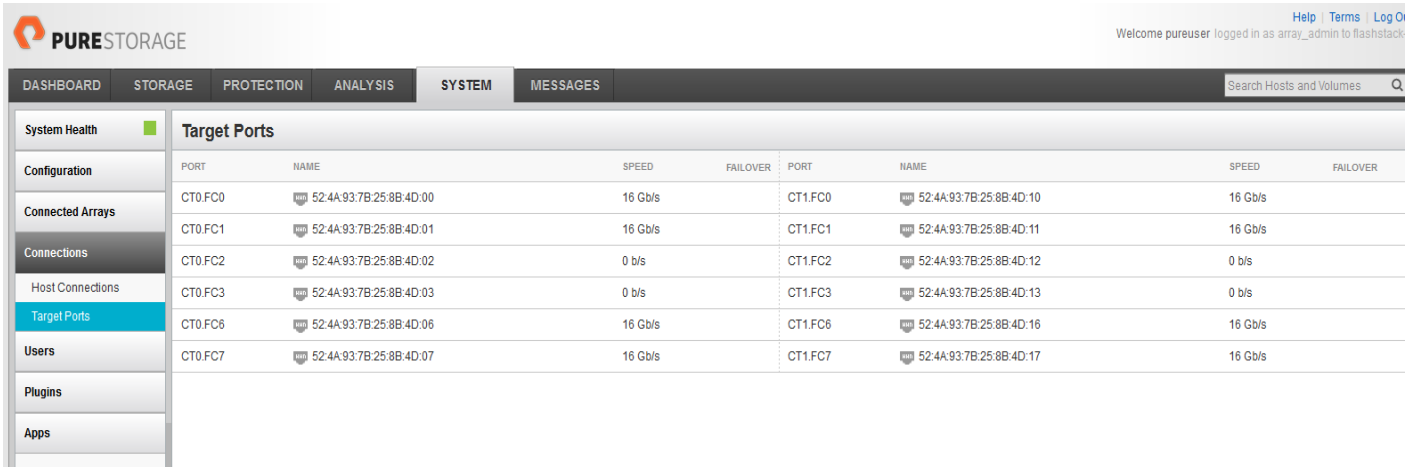

Create SAN Policy A

The SAN-A boot policy configures the SAN Primary's primary-target to be port CT0.FC0 on the Pure Storage cluster and SAN Primary's secondary-target to be port CT1.FC0 on the Pure Storage cluster. Similarly, the SAN Secondary's primary-target should be port CT1.FC1 on the Pure Storage cluster and SAN Secondary's secondary-target should be port CT0.FC1 on the Pure Storage cluster.

Log into the storage controller and verify all the port information is correct. This information can be found in the Pure Storage GUI under System > Connections > Target Ports.

You have to create SAN Primary (hba0) and SAN Secondary (hba1) in SAN-A Boot Policy by entering WWPN of Pure Storage FC Ports as detailed in the following section.

To create Boot Policies for the Cisco UCS environments, complete the following steps:

1. Go to Cisco UCS Manager and then go to Servers > Policies > root > Boot Policies.

2. Right-click and select Create Boot Policy. Enter SAN-A as the name of the boot policy as shown below:

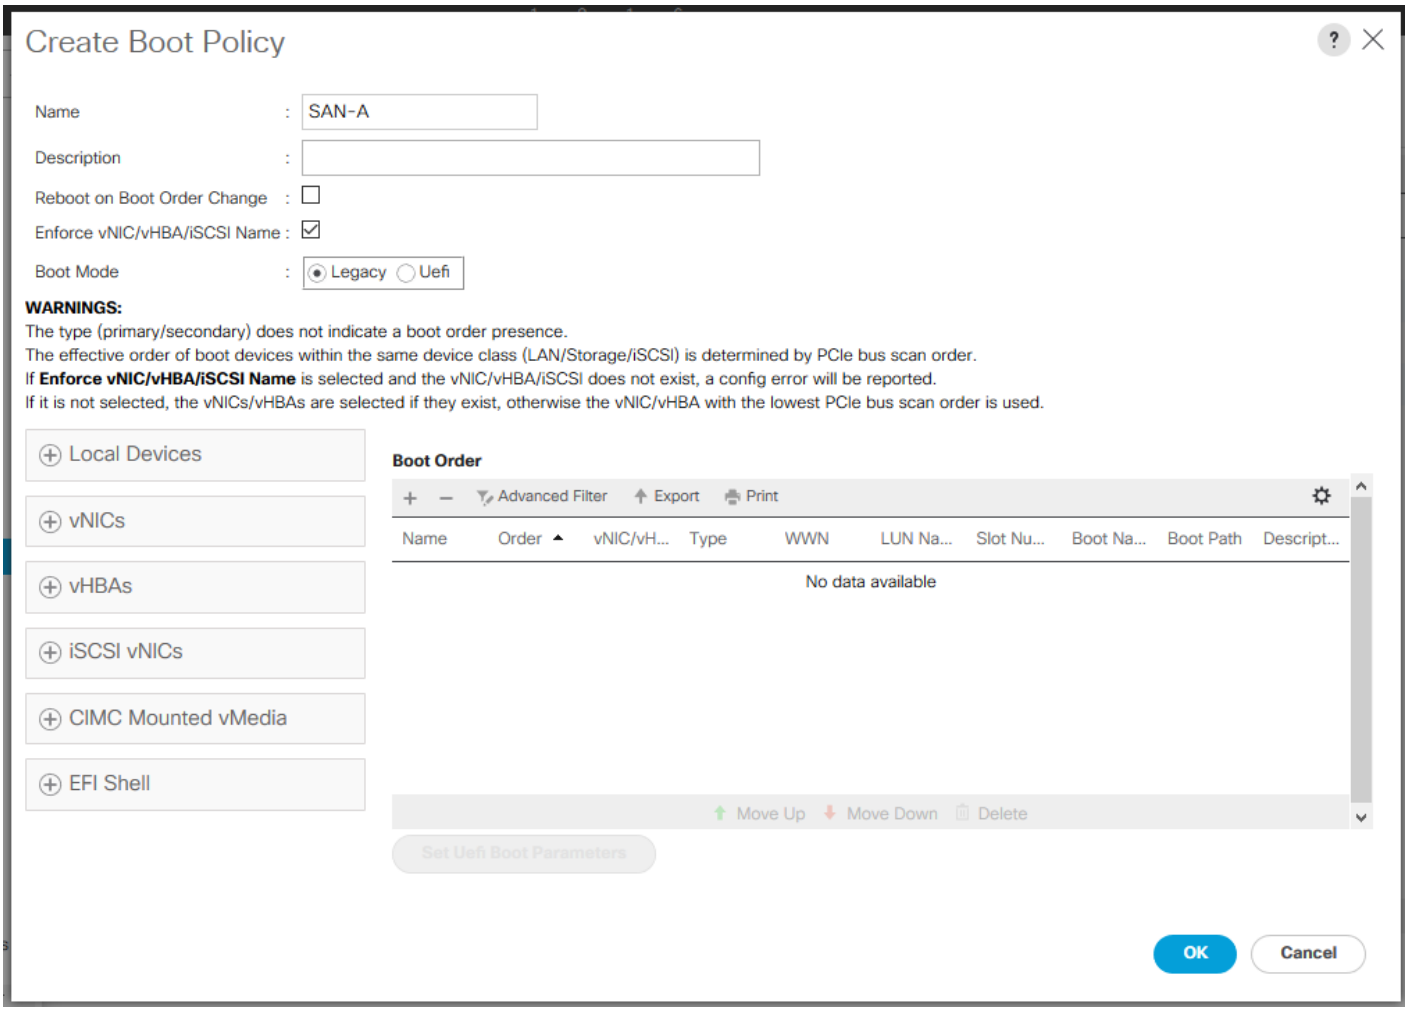

3. Expand the Local Devices drop-down menu and Choose Add CD/DVD. Expand the vHBAs drop-down list and Choose Add SAN Boot.

The SAN boot paths and targets will include primary and secondary options in order to maximize resiliency and number of paths.

---

4. In the Add SAN Boot dialog box, select Type as “Primary” and name vHBA as “hba0”. Click OK to add SAN Boot.

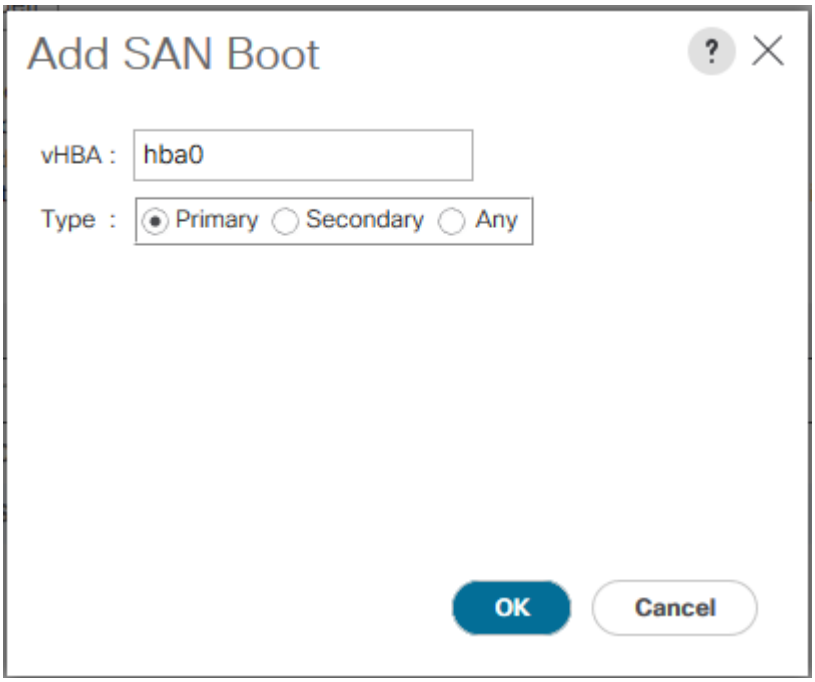

5. Select add SAN Boot Target to enter WWPN address of storage port. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC0 of Pure Storage and add SAN Boot Primary Target.

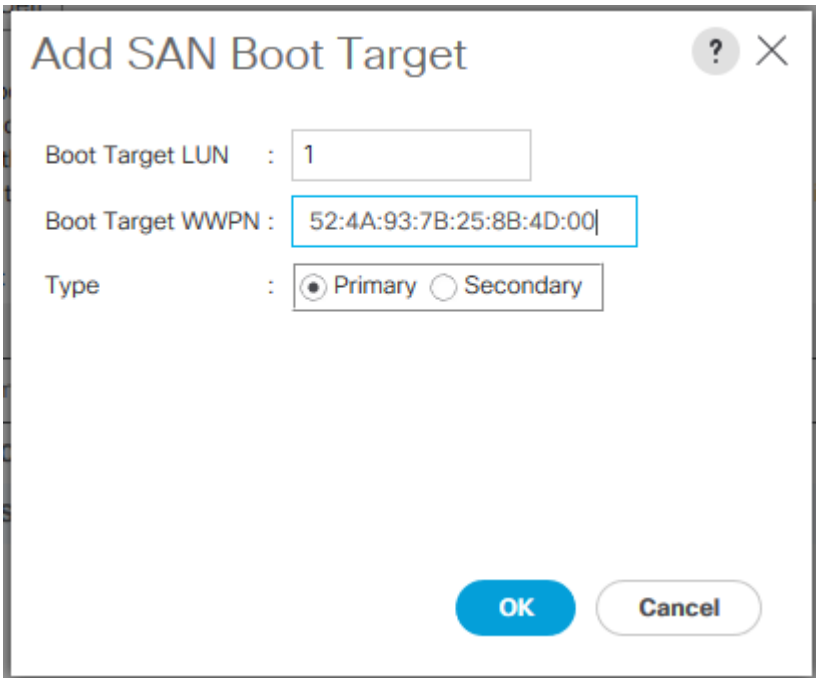

6. Add secondary SAN Boot target into same hba0, enter the boot target LUN as 1 and WWPN for FC port CT1.FC0 of Pure Storage, and add SAN Boot Secondary Target.

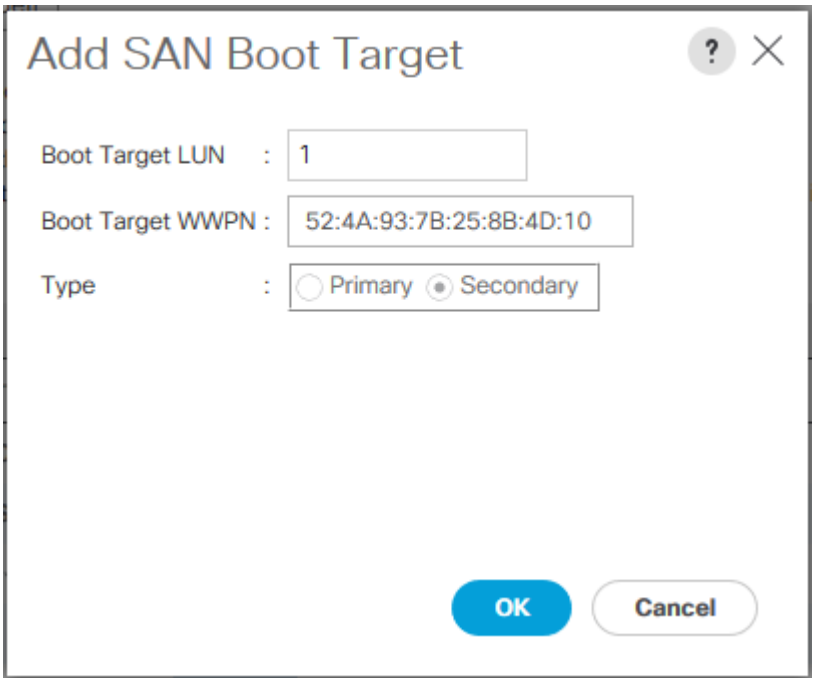

7. From the vHBA drop-down menu and choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba1" in the vHBA field. Click OK to SAN Boot, then choose Add SAN Boot Target.

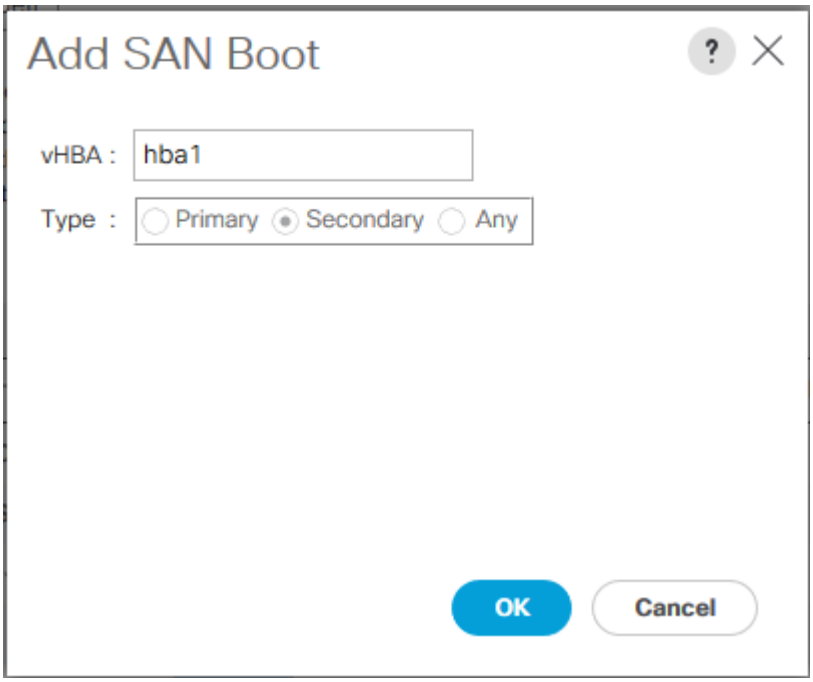

8. Keep 1 as the value for the Boot Target LUN. Enter the WWPN for FC port CT1.FC1 of Pure Storage and add SAN Boot Primary Target.

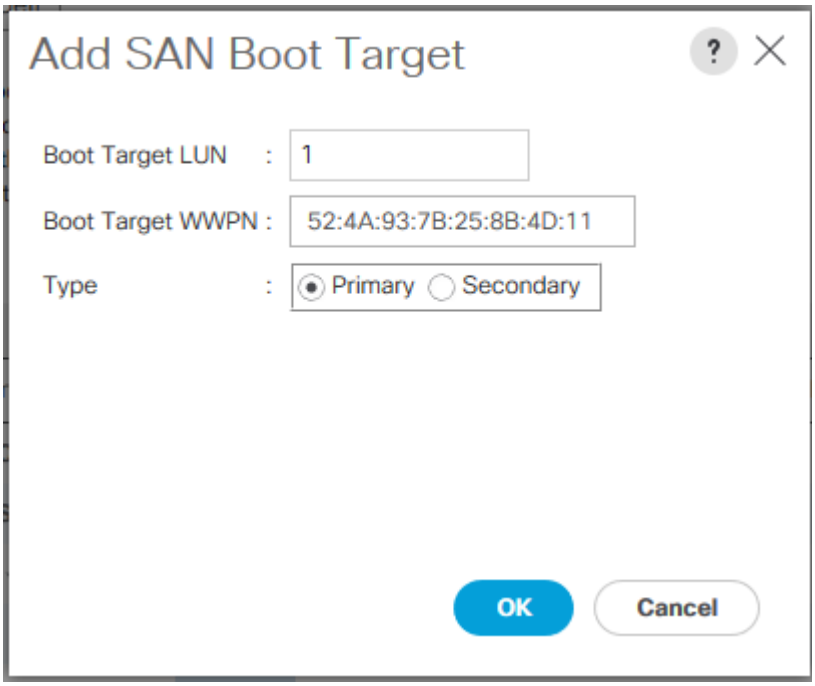

9. Add a secondary SAN Boot target into same hba1 and enter the boot target LUN as 1 and WWPN for FC port CT0.FC1 of Pure Storage and add SAN Boot Secondary Target.

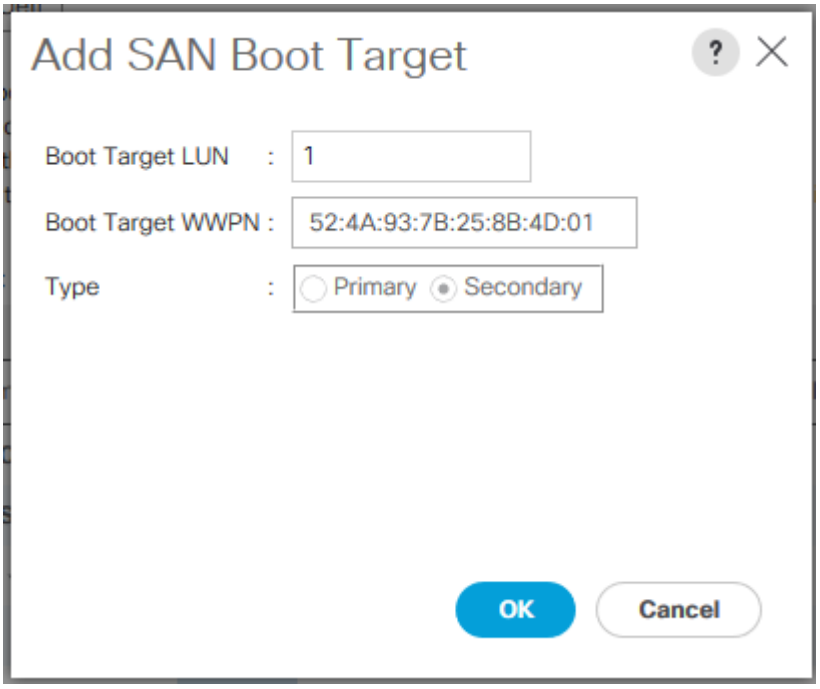

10. After creating the FC boot policies, you can view the boot order in the Cisco UCS Manager GUI. To view the boot order, navigate to Servers > Policies > Boot Policies. Click Boot Policy SAN-Boot-A to view the boot order in the right pane of the Cisco UCS Manager as shown below:

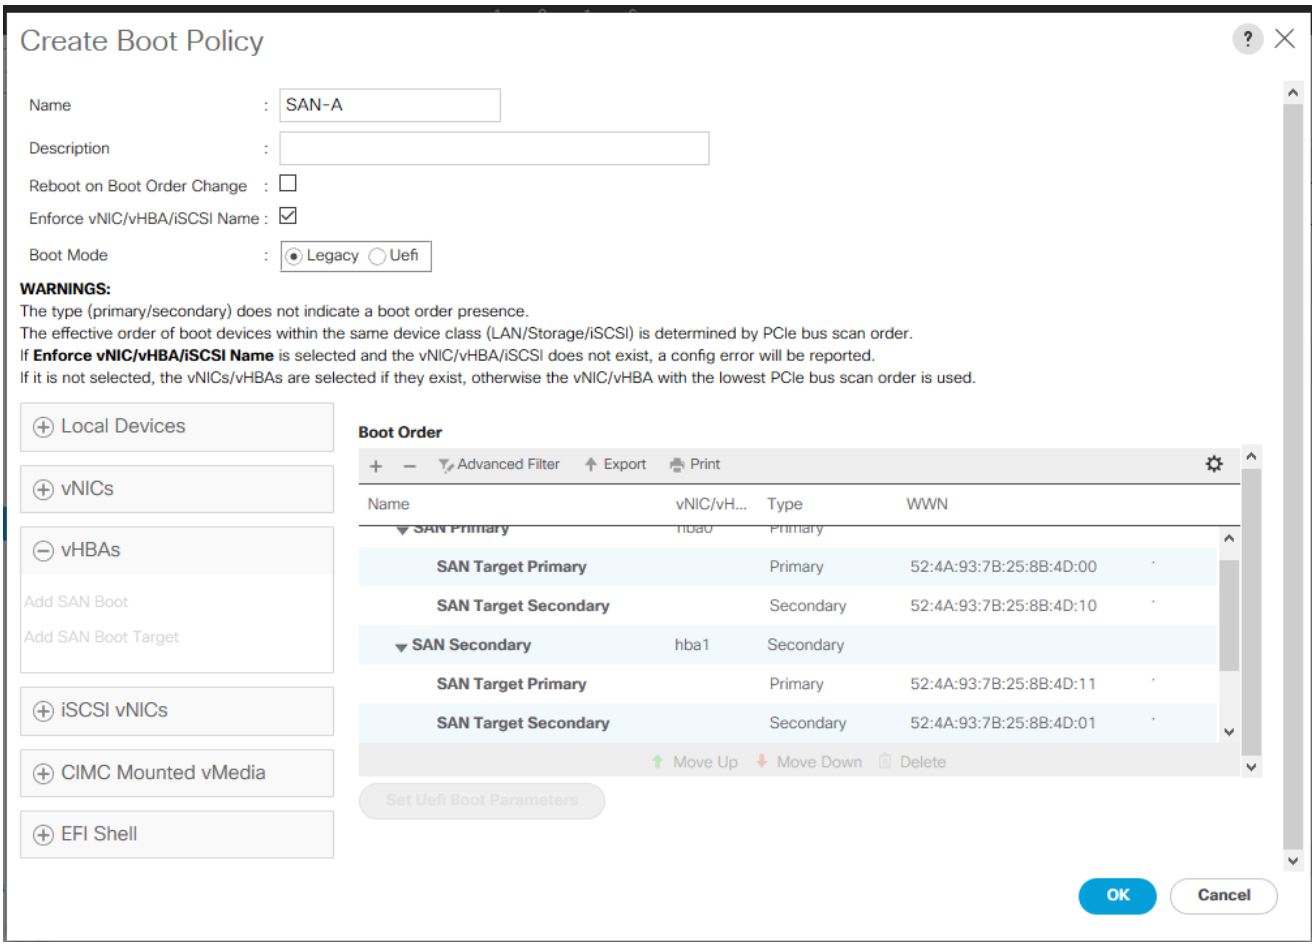

Create SAN Policy B

The SAN-B boot policy configures the SAN Primary's primary-target to be port CT0.FC6 on the Pure Storage cluster and SAN Primary's secondary-target to be port CT1.FC6 on the Pure Storage cluster. Similarly, the SAN Secondary's primary-target should be port CT1.FC7 on the Pure Storage cluster and SAN Secondary's secondary-target should be port CT0.FC7 on the Pure Storage cluster.

Log into the storage controller and verify all the port information is correct. This information can be found in the Pure Storage GUI under System > Connections > Target Ports.

You have to create SAN Primary (hba0) and SAN Secondary (hba1) in SAN-B Boot Policy by entering WWPN of Pure Storage FC Ports as explained in the following section.

To create boot policies for the Cisco UCS environments, complete the following steps:

1. Go to UCS Manager and then go to tab Servers > Policies > root > Boot Policies.

2. Right-click and select Create Boot Policy. Enter SAN-B as the name of the boot policy as shown in the figure below:

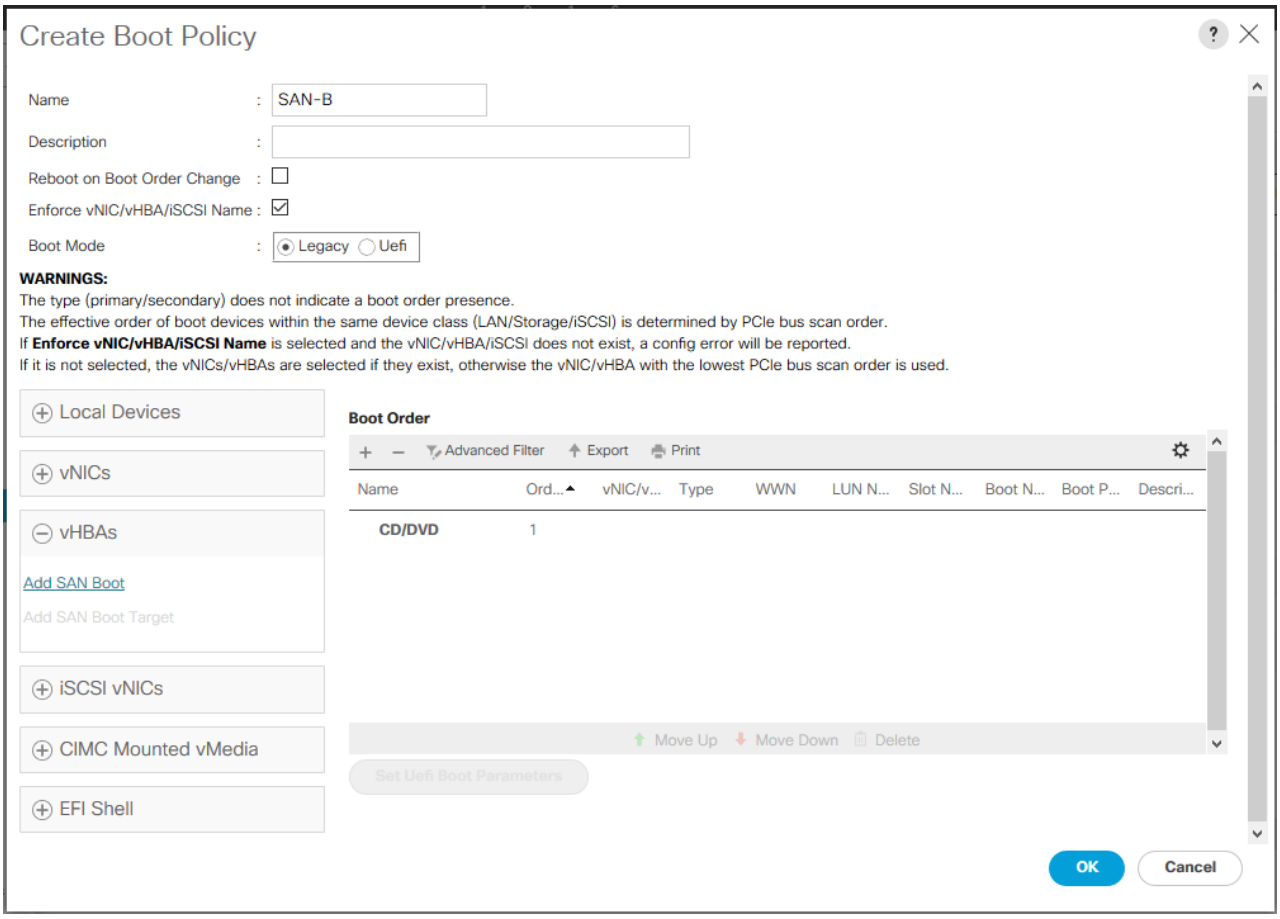

3. Expand the Local Devices drop-down list and Choose Add CD/DVD. Expand the vHBAs drop-down list and choose Add SAN Boot.

The SAN boot paths and targets will include primary and secondary options in order to maximize resiliency and number of paths.

4. In the Add SAN Boot dialog box, select Type as “Primary” and name vHBA as “hba0”. Click OK to add SAN Boot.

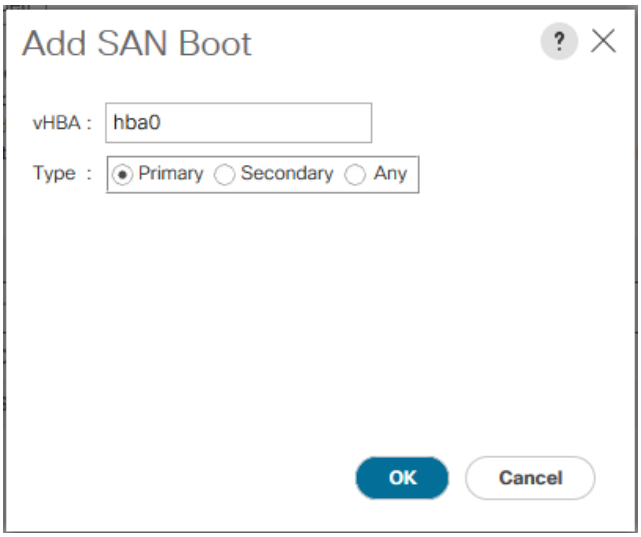

5. Select add SAN Boot Target to enter WWPN address of storage port. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC6 of Pure Storage and add SAN Boot Primary Target.

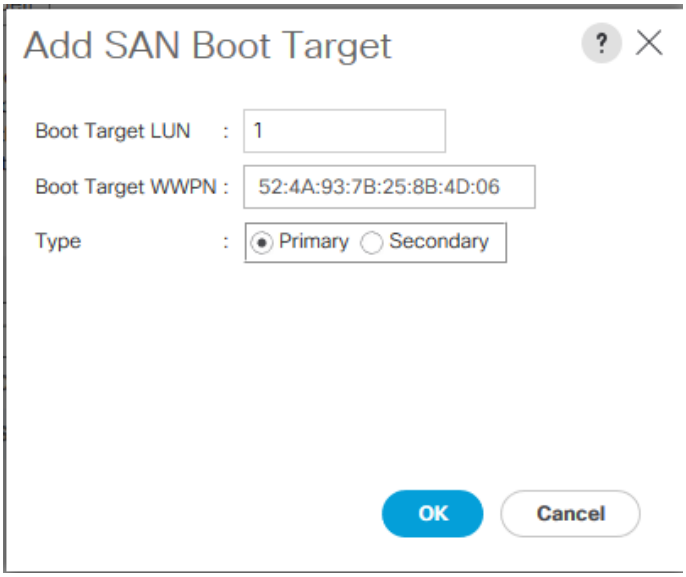

6. Add the secondary SAN Boot target into the same hba0; enter boot target LUN as 1 and WWPN for FC port CT1.FC6 of Pure Storage, and add SAN Boot Secondary Target.

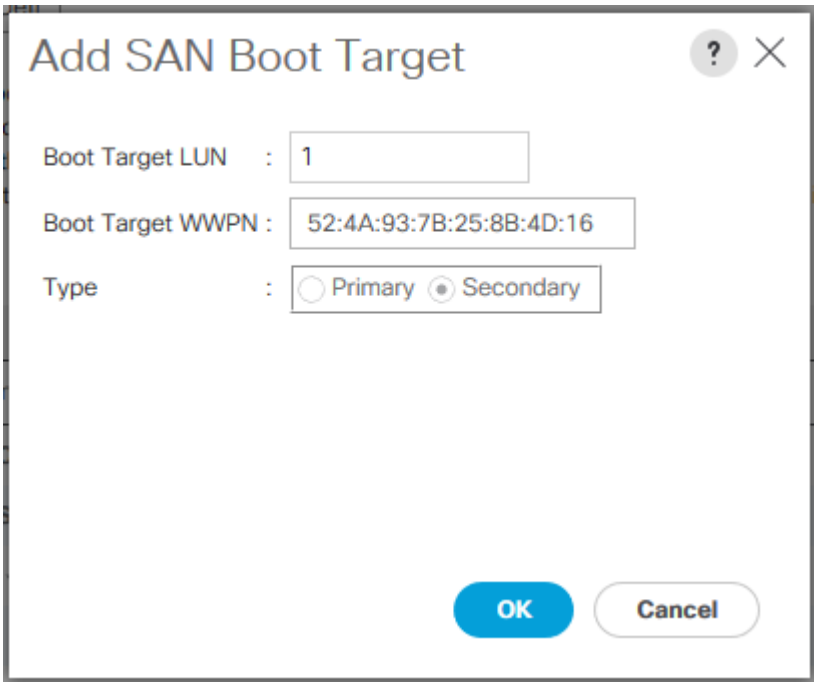

7. From the vHBA drop-down list, choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba1" in the vHBA field. Click OK to SAN Boot, then choose Add SAN Boot Target.

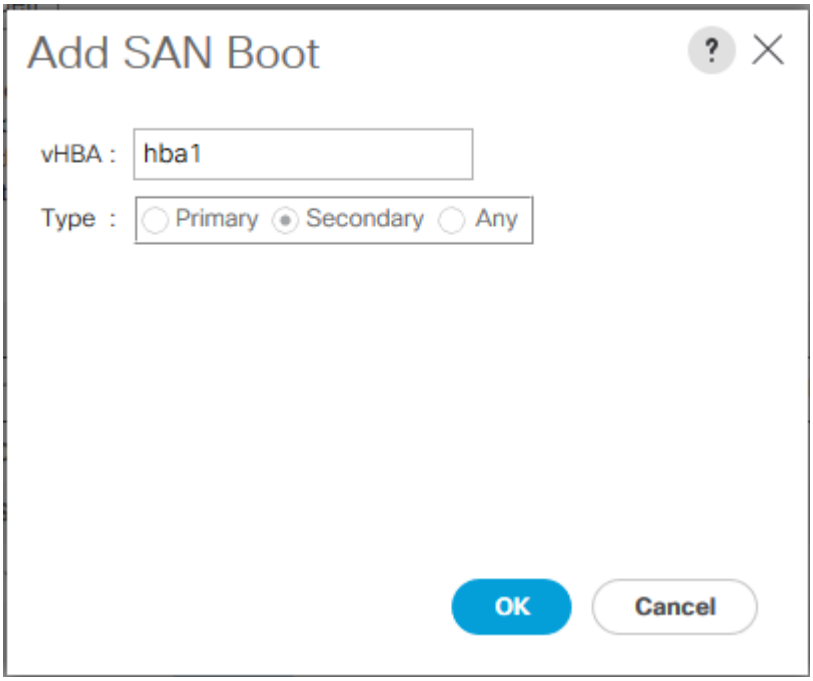

8. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT1.FC7 of Pure Storage and Add SAN Boot Primary Target.

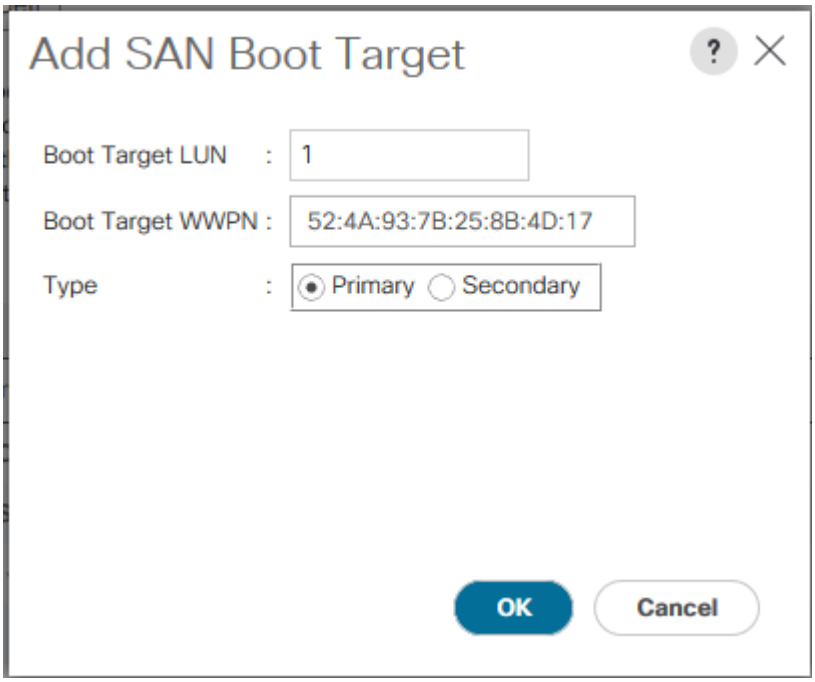

9. Add secondary SAN Boot target into same hba1 and enter boot target LUN as 1 and WWPN for FC port CT0.FC7 of Pure Storage and add SAN Boot Secondary Target.

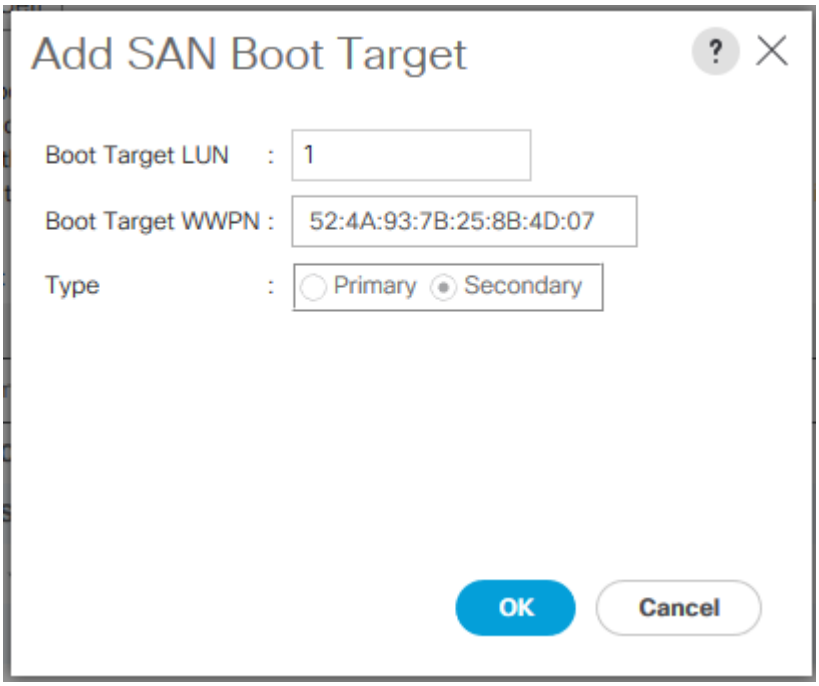

10. After creating the FC boot policies, you can view the boot order in the Cisco UCS Manager GUI. To view the boot order, navigate to Servers > Policies > Boot Policies. Click Boot Policy SAN-Boot-A to view the boot order in the right pane of the Cisco UCS Manager as shown below:

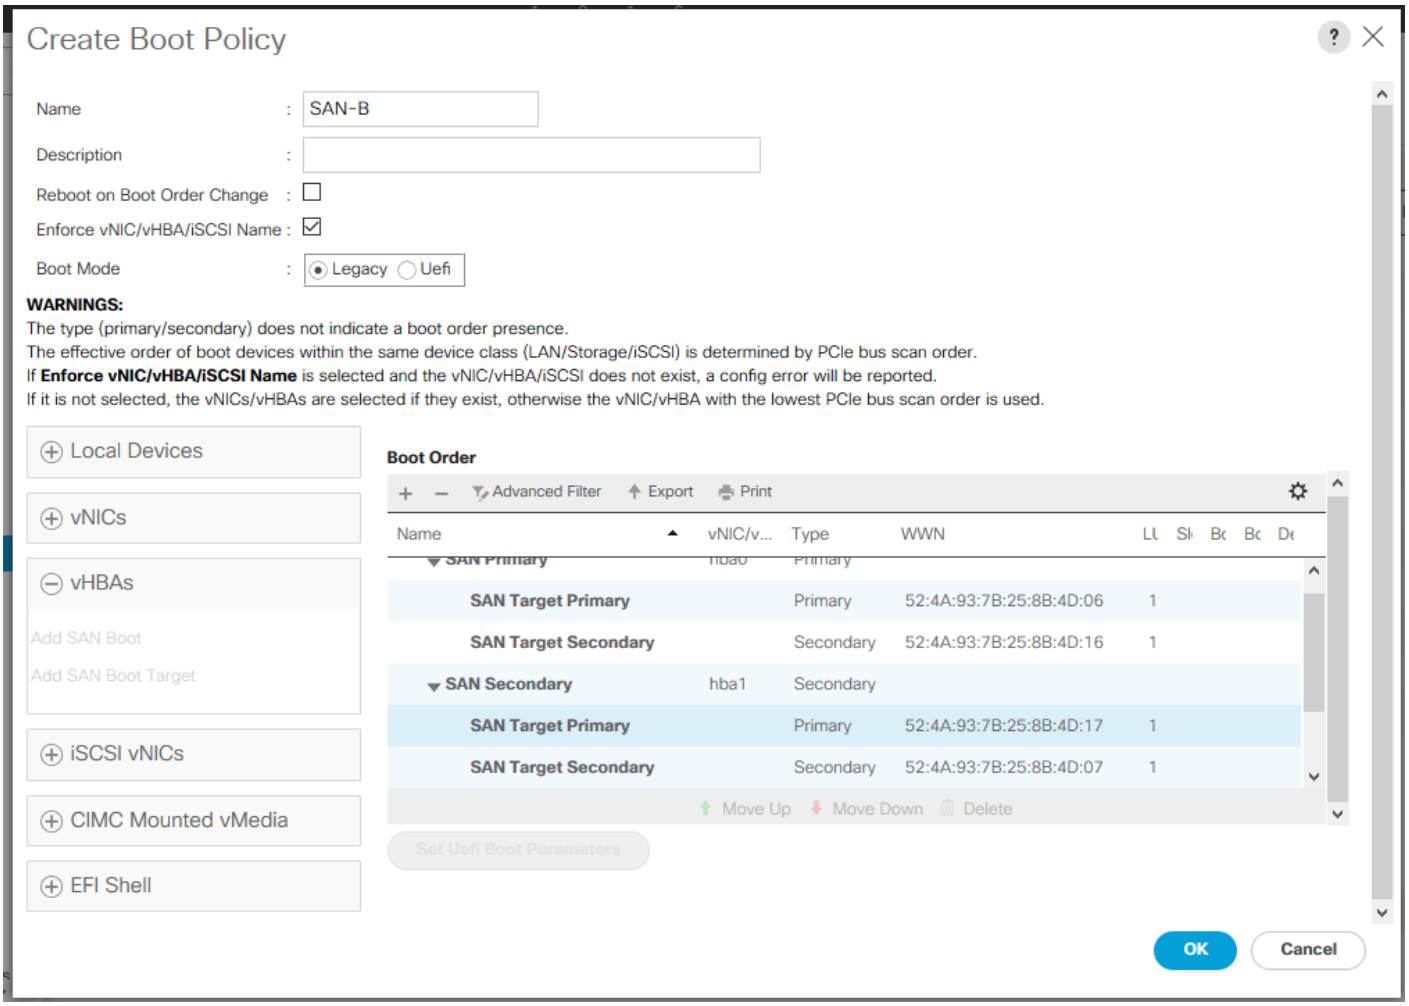

For this solution, we created two Boot Policy as “SAN-A” and “SAN-B”. For 8 Oracle RAC Nodes, you will assign first 4 Service Profiles with SAN-A to first 4 RAC nodes (oraracx1, oraracx2, oraracx3 and oraracx4) and remaining 4 Service Profiles with SAN-B to remaining 4 Oracle RAC nodes (oraracx5, oraracx6, oraracx7 and oraracx8) as explained in the following section.

Configure and Create a Service Profile Template

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

You will create two Service Profile Template. First Service profile template “ORAX-1” using boot policy as “SAN-A” and second Service profile template “ORAX-2” using boot policy as “SAN-B” to utilize all the FC ports from Pure Storage for high-availability in case of any FC links go down.

You will create the first ORAX-1 as explained in the following section.

Create Service Profile Template

To create a service profile template, complete the following steps:

1. In the Cisco UCS Manager, go to Servers > Service Profile Templates > root and right-click to “Create Service Profile Template” as shown below.

2. Enter the Service Profile Template name, select the UUID Pool that was created earlier, and click Next.

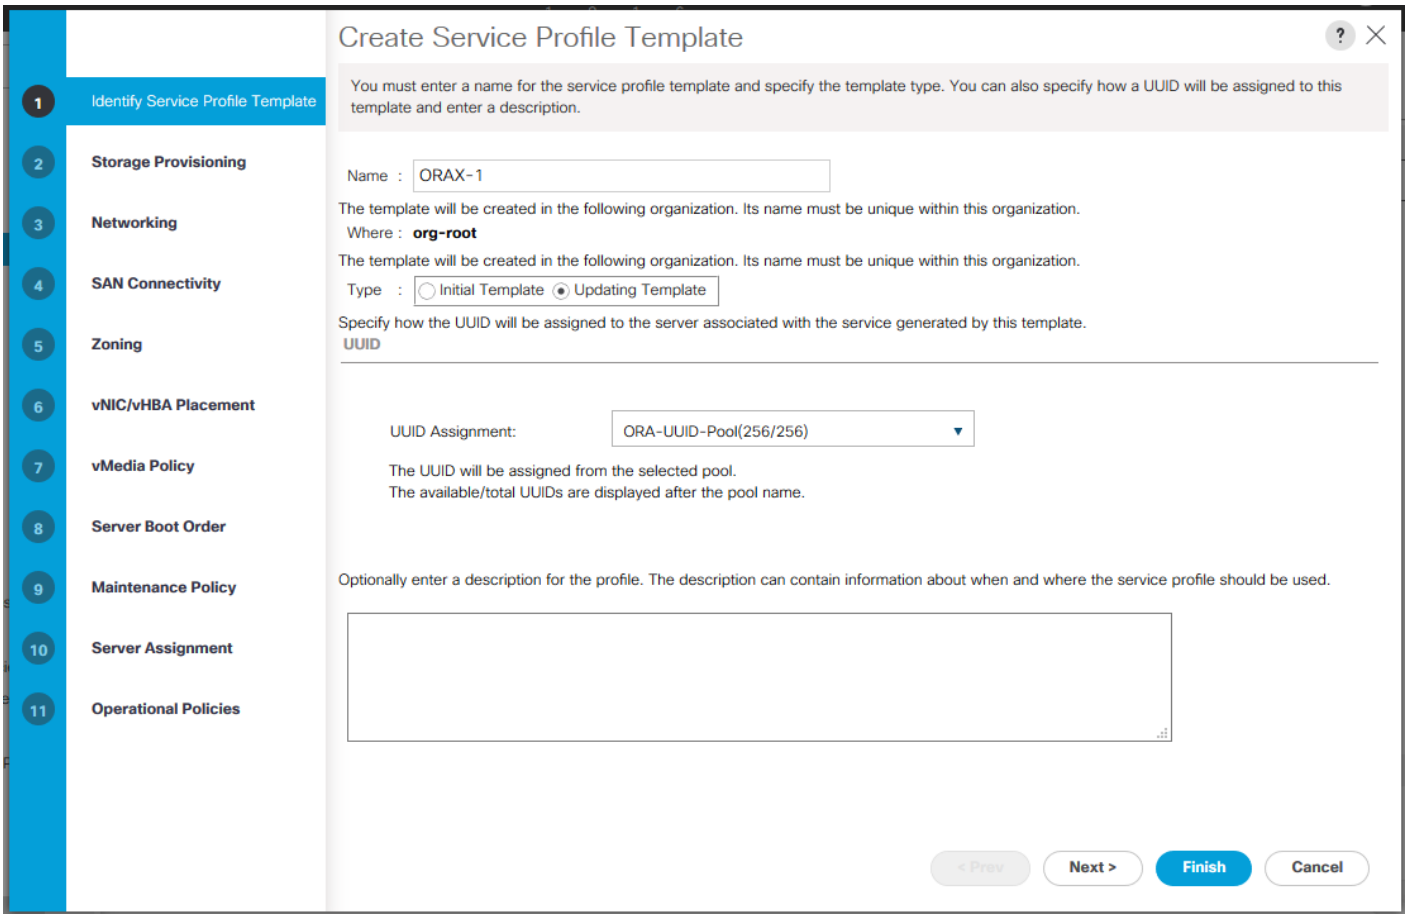

3. Select Local Disk Configuration Policy to SAN-Boot as No Local Storage.

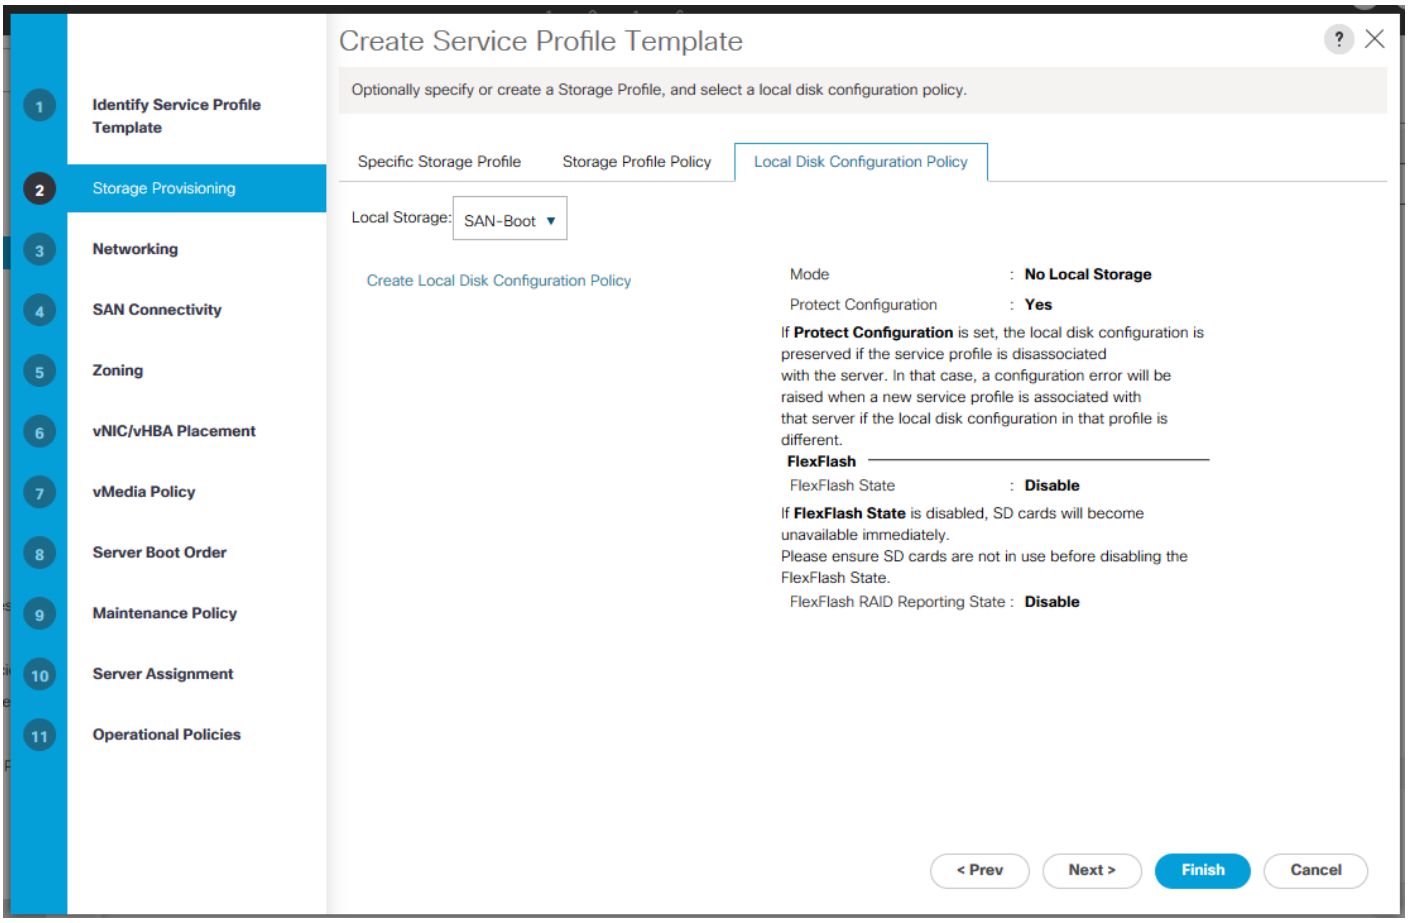

4. In the networking window, select “Expert” and click “Add” to create vNICs. Add one or more vNICs that the server should use to connect to the LAN.

5. Now there are two vNIC in the create vNIC menu. You have given name to first vNIC as “eth0” and second vNIC as “eth1.”

6. As shown below, select vNIC Template as Oracle-vNIC-A and Adapter Policy as ORA_Linux_Tuning which was created earlier for vNIC “eth0”.

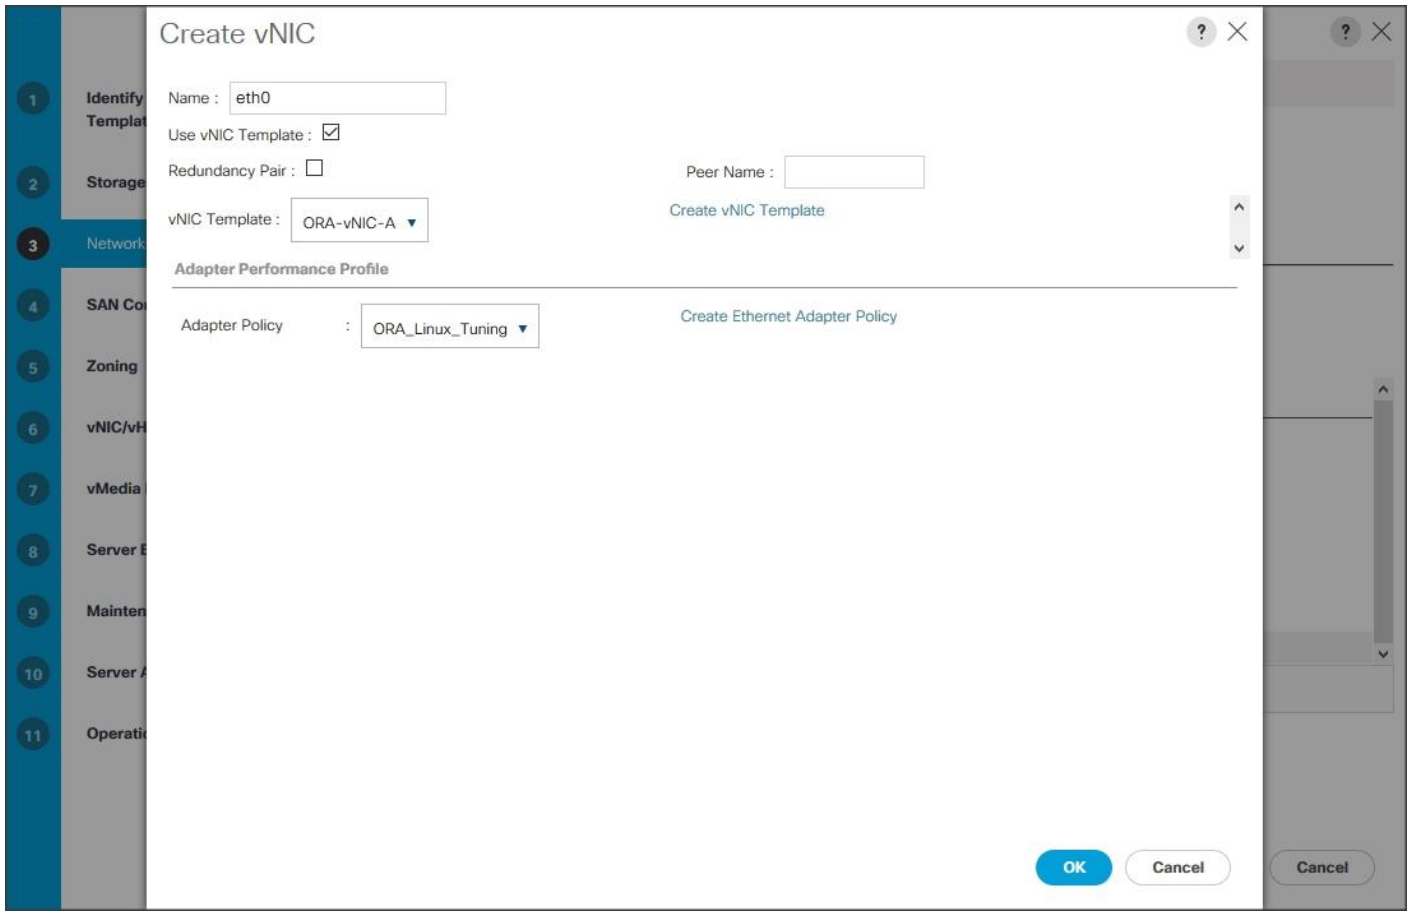

7. Similarly, as shown below, select vNIC Template as Oracle-vNIC-B and Adapter Policy as ORA_Linux_Tuning for vNIC “eth1”.

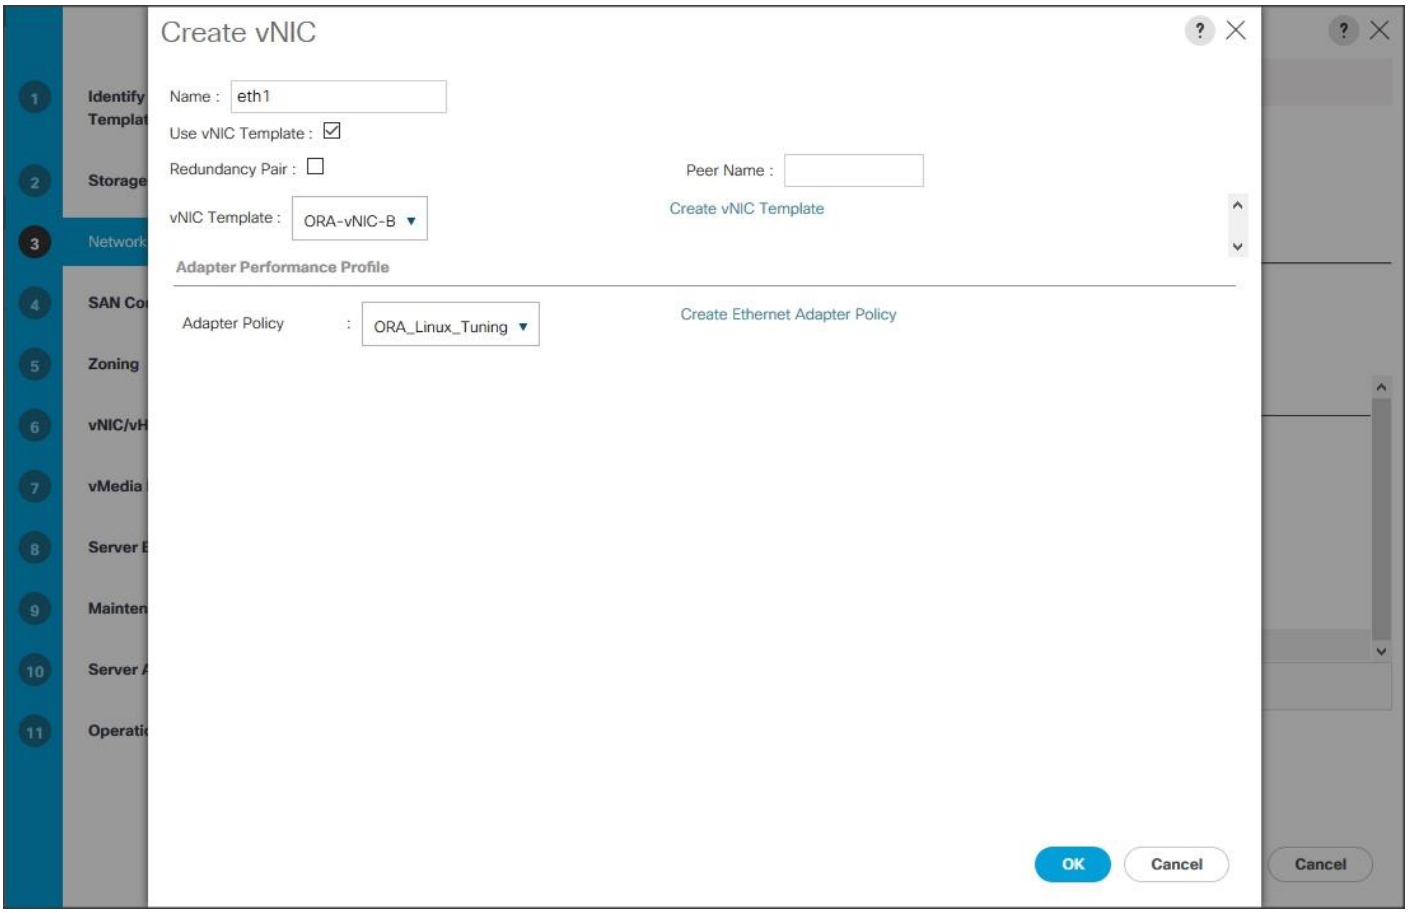

As shown above, eth0 and eth1 vNICs are created so that Servers can connect to the LAN.

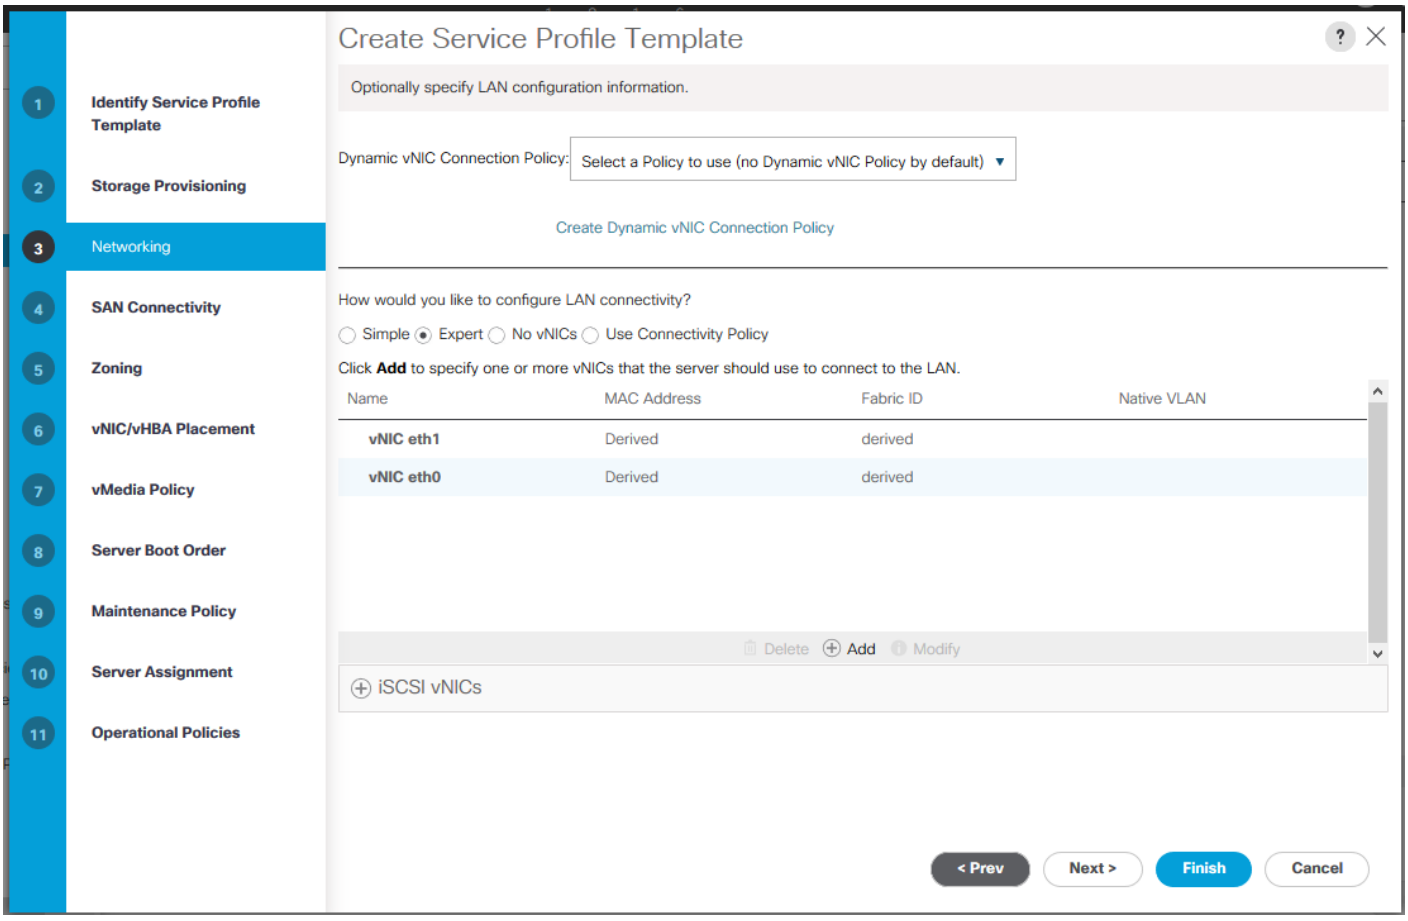

8. Once vNICs are created, you need to create vHBAs. Click Next.

9. In the SAN Connectivity menu, select “Expert” to configure as SAN connectivity. Select WWNN (World Wide Node Name) pool, which we created earlier. Click on “Add” to add vHBAs as shown below. The following four HBA were created:

   - Hba0 using vHBA Template Oracle-HBA-A
   - Hba1 using vHBA Template Oracle-HBA-B
   - Hba2 using vHBA Template Oracle-HBA-A
   - Hba3 using vHBA Template Oracle-HBA-B

Figure 7 vHBA0

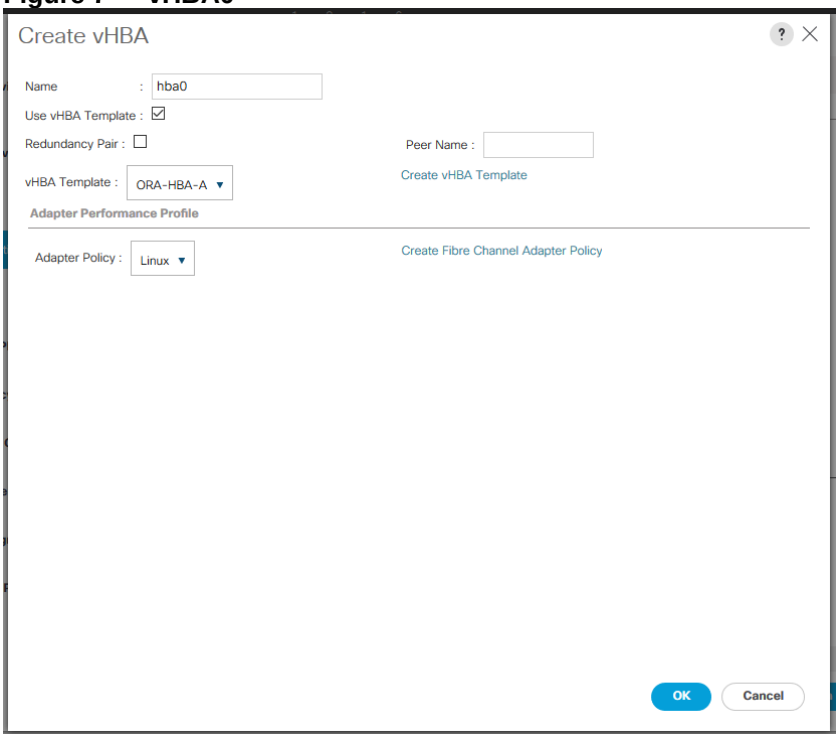

Figure 8 vHBA1

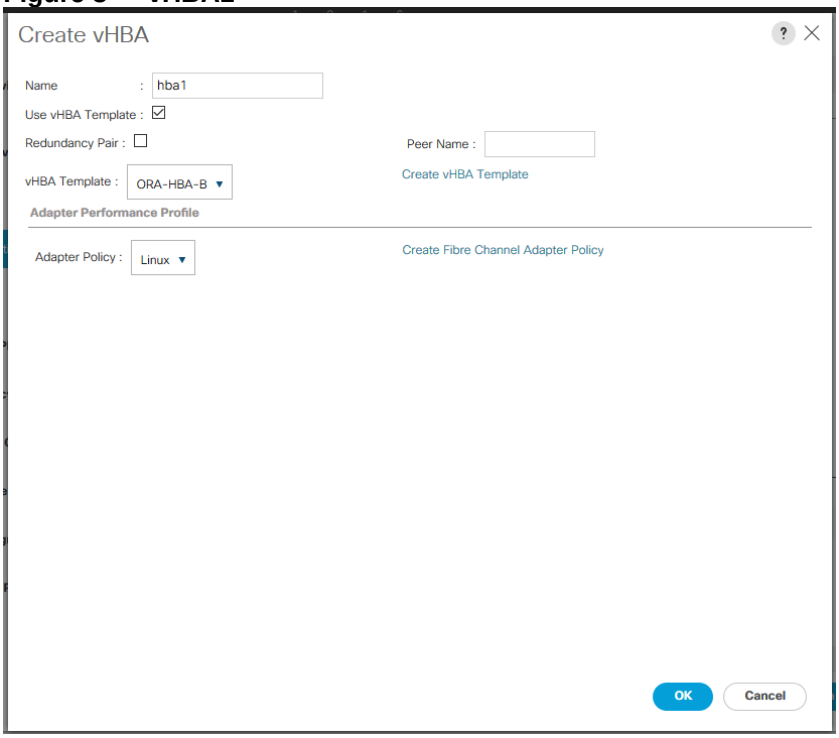

Figure 9 All vHBAs

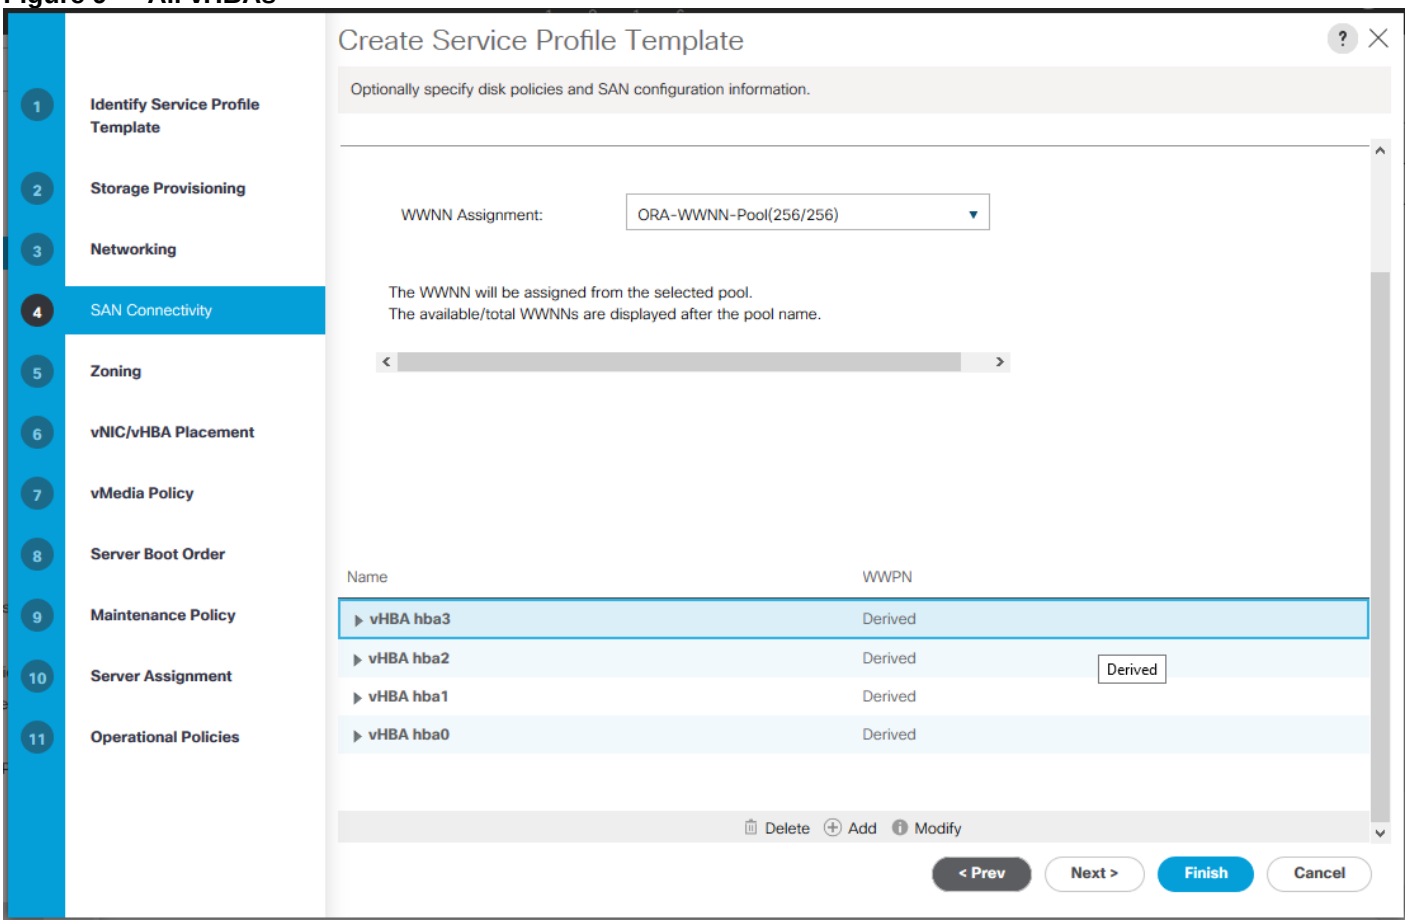

Skip zoning; for this Oracle RAC Configuration, the Cisco MDS 9148S is used for zoning.

10. Select default option as Let System Perform Placement in the Placement Selection menu.
11. For the Server Boot Policy, select “SAN-A” as Boot Policy which you created earlier.

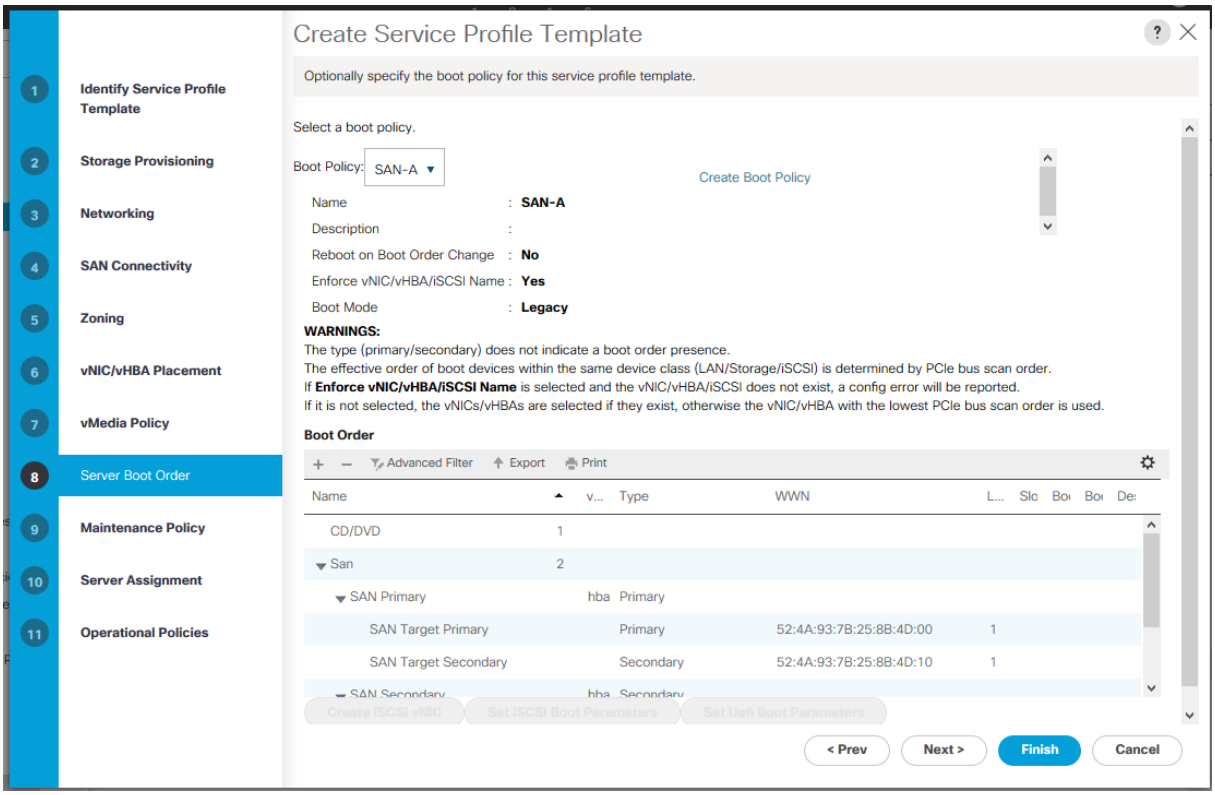

12. The remaining maintenance and assignment policies were left as default in the configuration. However, they may vary from site-to-site depending on workloads, best practices, and policies.

13. Click Next and then click Finish to create service profile template as “ORAX-1.” This service profile template is be used to create first 4 service profiles for oracle RAC node 1 to 4.

14. Create another service profile template as “ORAX-2”. This ORAX-2 service profile template will be used to create remaining 4 service profiles for oracle RAC node 5 to 8.

You can achieve this quickly by cloning the template. For this solution, we have cloned service profile template ORAX-1 to ORAX-2.

15. In this service profile template ORAX-2, you have modified Boot Policy as “SAN-B” to use all the remaining FC paths of storage for high availability.

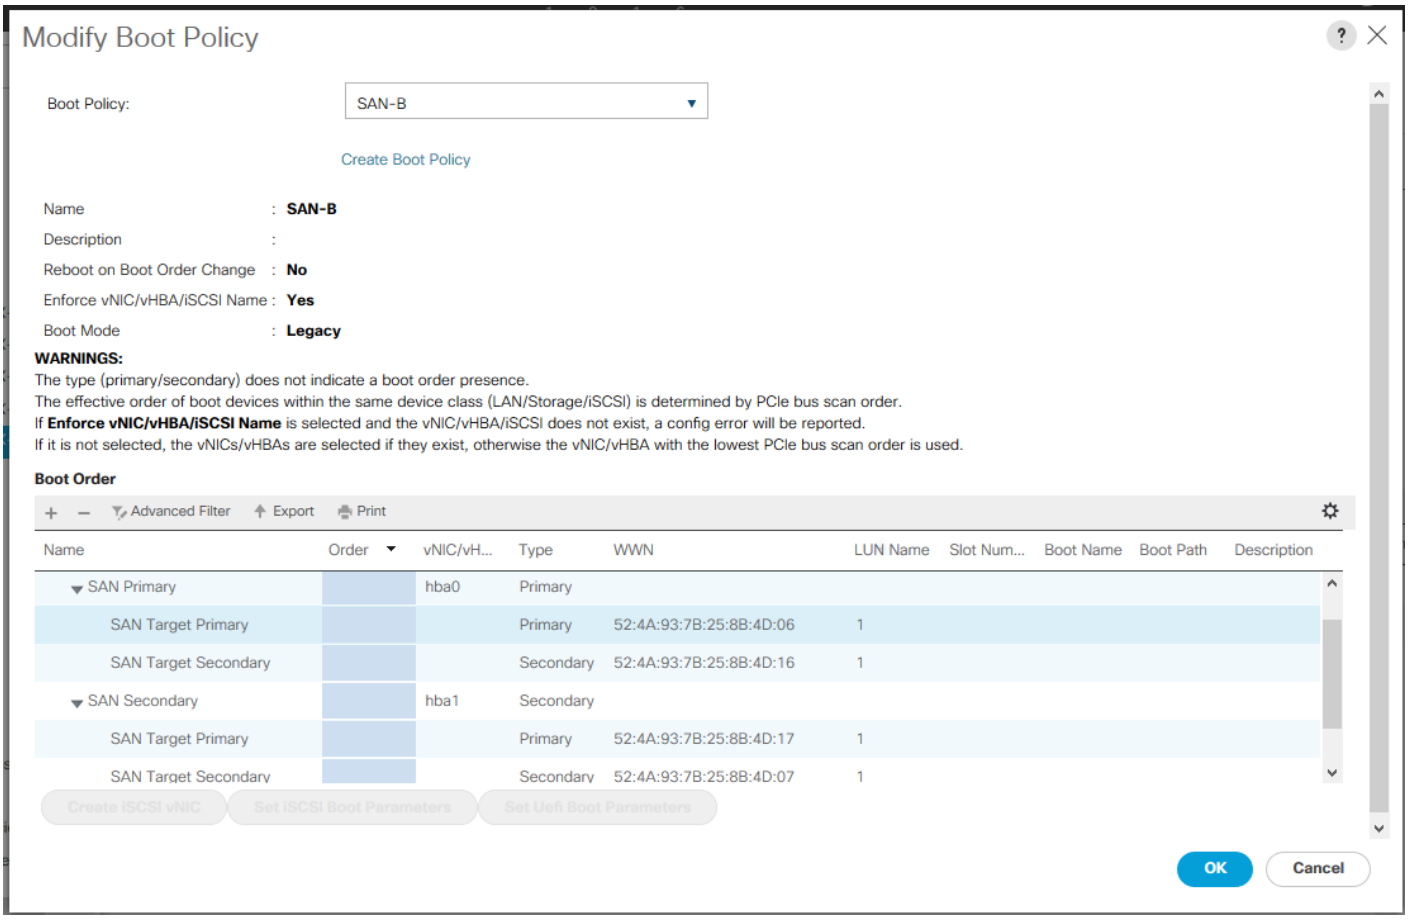

You have now created Service profile template as “ORAX-1” and “ORAX-2” with each having four vHBAs and two vNICs.

Create Service Profiles from Template and Associate to Servers

Create Service Profiles from Template

You will create eight Service profiles for eight Oracle RAC nodes as explained in the following sections.

For the first four Oracle RAC Nodes (oraracx1, oraracx2, oraracx3 and oraracx4), you will create four Service Profiles from Template “ORAX-1.” The remaining four Oracle RAC Nodes (oraracx5, oraracx6, oraracx7 and oraracx8), will require creating another four Service Profiles from Template “ORAX-2”.

To create first four Service Profiles from Template, complete the following steps:

1. Go to tab Servers > Service Profiles > root > and right-click “Create Service Profiles from Template.”

2. Select the Service profile template as “ORAX-1” which you created earlier and name the service profile as “ORARACX.”

3. To create four service profiles, enter “Number of Instances” as 4 as shown below. This process will create service profiles as “ORARACX1”, “ORARACX2”, “ORARACX3” and “ORARACX4.”

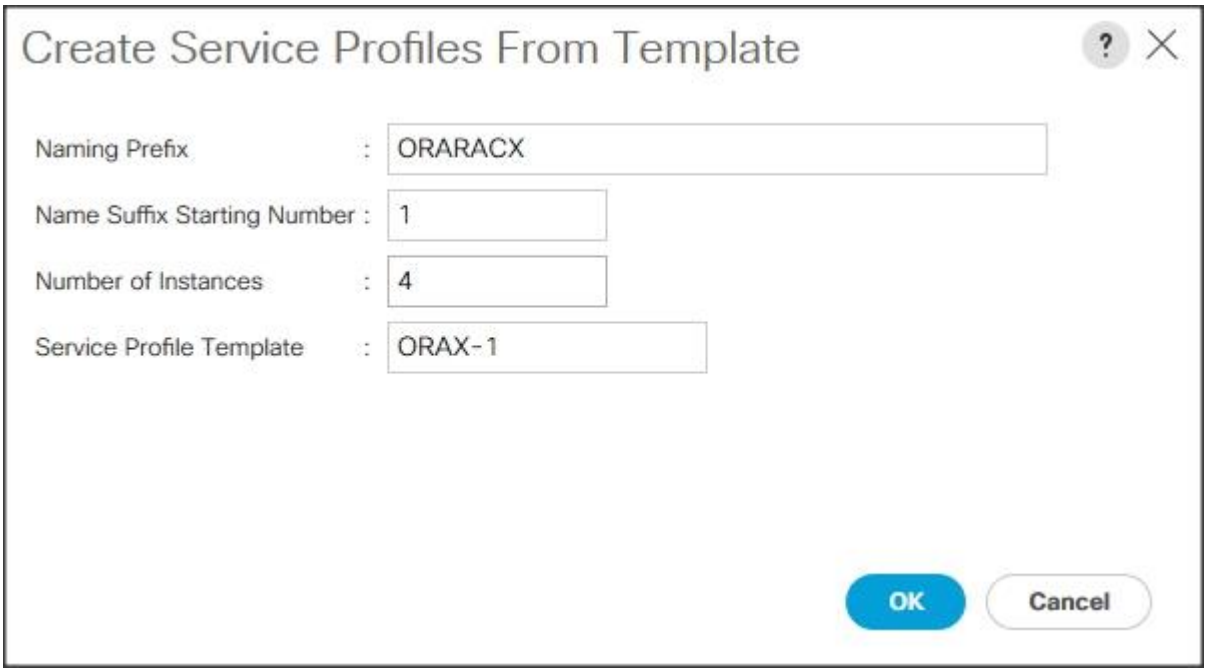

4. Create remaining four Service Profiles “ORARACX5”, “ORARACX6”, “ORARACX7” and “ORARACX8” from Template “ORAX-2.”

When the service profiles are created, associate them to the servers as described in the following section.

Associate Service Profiles to the Servers

To associate service profiles to the servers, complete the following steps.

1. Under the servers tab, select the desired service profile, and select Change Service Profile Association.
2. Right-click the name of service profile you want to associate with the server and select the option "Change Service Profile Association."
3. In the Change Service Profile Association page, from the Server Assignment drop-down list, select existing server that you would like to assign and click OK.

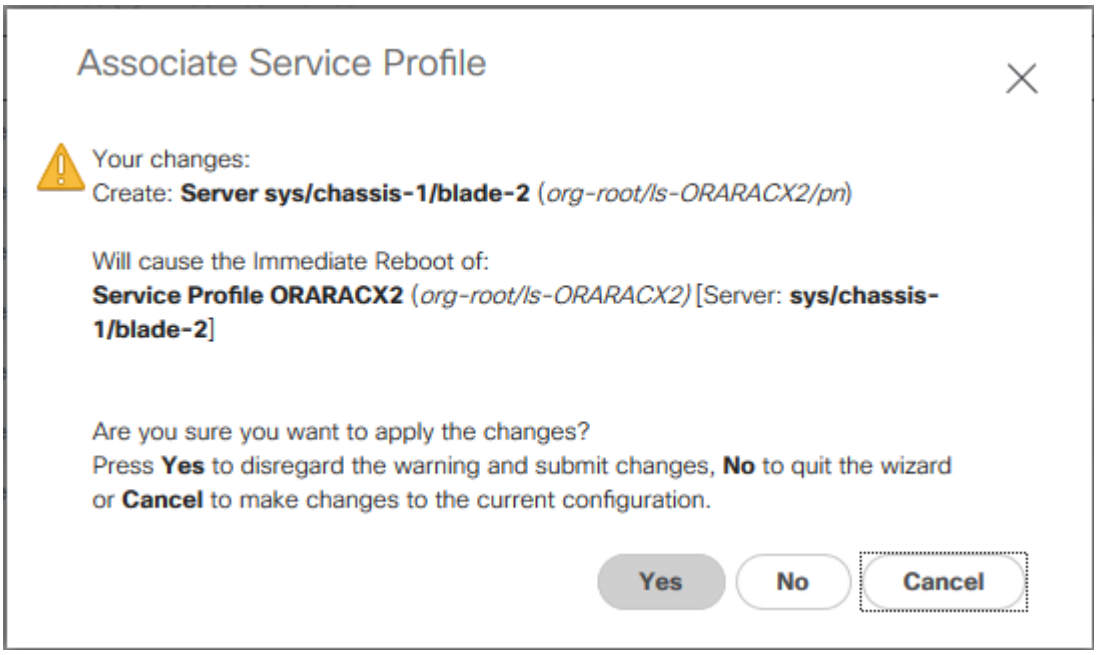

4. You will assign service profiles ORARAXC1 to ORARACX4 to Chassis 1 Servers and ORARACX5 to ORARACX8 to Chassis 2 Servers.

5. Repeat the same steps to associate remaining seven service profiles for the blade servers.

   You have assigned “ORARACX1” to Chassis 1 Server 1, Service Profile “ORARACX2” to Chassis 1 Server 2, Service Profile “ORARACX3” to Chassis 1 Server 3 and, Service Profile “ORARACX4” to Chassis 1 Server 4.

   You have assigned Service Profile “ORARACX5” to Chassis 2 Server 1, Service Profile “ORARACX6” to Chassis 2 Server 2, Service Profile “ORARACX7” to Chassis 2 Server 3 and Service Profile “ORARACX8” to Chassis 2 Server 4.

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

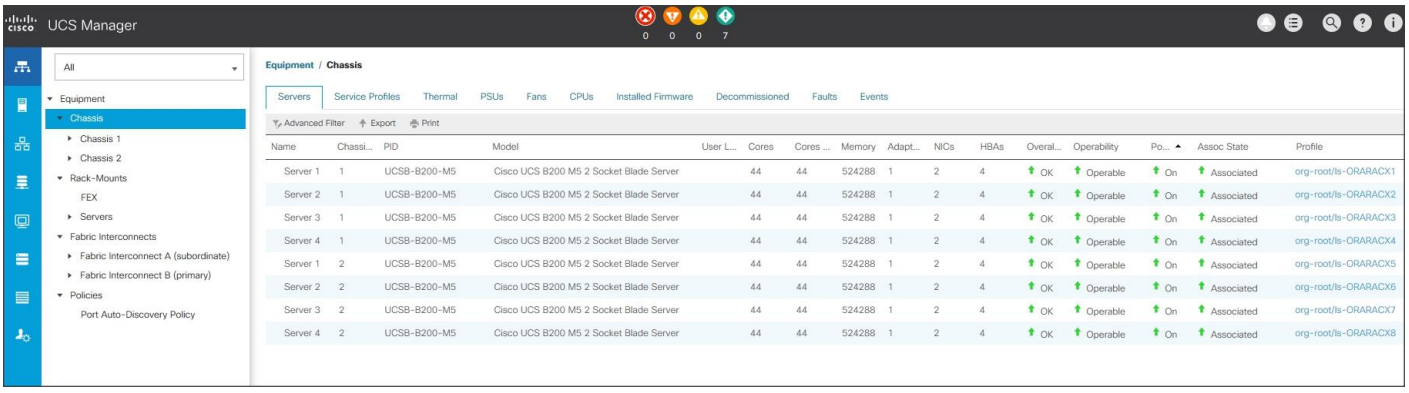

7. As shown above, make sure all the server nodes has no major or critical fault and all are in operable state.

This completes the configuration required for Cisco UCS Manager Setup.

Configure Cisco Nexus 9372PX-E Switches

The following sections detail the steps for the Nexus 9372PX-E switch configuration. The details of “show run” output is listed in the Appendix.

Configure Global Settings for Cisco Nexus A and Cisco Nexus B

To set global configuration, complete the following steps on both the Nexus switches

1. Login as admin user into the Nexus Switch A and run the following commands to set global configurations and jumbo frames in QoS:

conf terminal

spanning-tree port type network default

spanning-tree port type edge bpduguard default

port-channel load-balance ethernet source-dest-port

policy-map type network-qos jumbo

class type network-qos class-default

mtu 9216

exit

class type network-qos class-fcoe

pause no-drop

mtu 2158

exit

exit

system qos

service-policy type network-qos jumbo

exit

copy run start

2. Login as admin user into the Nexus Switch B and run the same above commands to set global configurations and jumbo frames in QoS.

Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

To create the necessary virtual local area networks (VLANs), complete the following steps on both Nexus switches.

1. Login as admin user into the Nexus Switch A.

2. Create VLAN 134 for Public Network Traffic:

PURESTG-NEXUS-A# config terminal

PURESTG-NEXUS-A(config)# VLAN 134

PURESTG-NEXUS-A(config-VLAN)# name Oracle_Public_Traffic

PURESTG-NEXUS-A(config-VLAN)# no shutdown

PURESTG-NEXUS-A(config-VLAN)# exit

PURESTG-NEXUS-A(config)# copy running-config startup-config

PURESTG-NEXUS-A(config)# exit

3. Create VLAN 10 for Private Network Traffic:

PURESTG-NEXUS-A# config terminal

PURESTG-NEXUS-A(config)# VLAN 10

PURESTG-NEXUS-A(config-VLAN)# name Oracle_Private_Traffic

PURESTG-NEXUS-A(config-VLAN)# no shutdown

PURESTG-NEXUS-A(config-VLAN)# exit

PURESTG-NEXUS-A(config)# copy running-config startup-config

PURESTG-NEXUS-A(config)# exit

4. Login as admin user into the Nexus Switch B and create VLAN 134 for Public Network Traffic and VLAN 10 for Private Network Traffic.

Virtual Port Channel (vPC) Summary for Data and Storage Network

In the Cisco Nexus 9372PX-E switch topology, a single vPC feature is enabled to provide HA, faster convergence in the event of a failure, and greater throughput. Cisco Nexus 9372PX-E vPC configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers is shown below:

Table 6 vPC Summary

vPC Domain	vPC Name	vPC ID
1	Peer-Link	1
1	vPC Public	21
1	vPC Private	22

As listed in the table above, a single vPC domain with Domain ID 1 is created across two Cisco Nexus 9372PX-E member switches to define vPC members to carry specific VLAN network traffic. In this topology, we defined a total number of 3 vPCs.

vPC ID 1 is defined as Peer link communication between two Nexus switches in Fabric A and B.

vPC IDs 21 and 22 are defined for public and private network traffic from Cisco UCS fabric interconnects.

Create vPC Peer-Link Between the Two Nexus Switches

To create the vPC Peer-Link, complete the following steps:

Figure 10 Nexus Switch Peer-Link

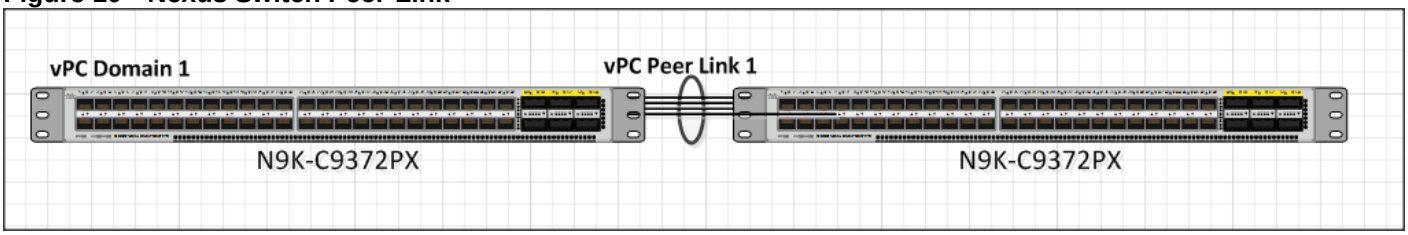

1. Login as “admin” user into the Nexus Switch A.

For vPC 1 as Peer-link, we used interfaces 1-2 for Peer-Link. You may choose the appropriate number of ports for your needs.

To create the necessary port channels between devices, complete the following on both the Nexus Switches:

PURESTG-NEXUS-A# config terminal

PURESTG-NEXUS-A(config)#feature vpc

PURESTG-NEXUS-A(config)#feature lacp

PURESTG-NEXUS-A(config)#vpc domain 1

PURESTG-NEXUS-A(config-vpc-domain)# peer-keepalive destination 10.29.134.154 source 10.29.134.153

PURESTG-NEXUS-A(config-vpc-domain)# exit

PURESTG-NEXUS-A(config)# interface port-channel 1

PURESTG-NEXUS-A(config-if)# description VPC peer-link

PURESTG-NEXUS-A(config-if)# switchport mode trunk

PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134

PURESTG-NEXUS-A(config-if)# spanning-tree port type network

PURESTG-NEXUS-A(config-if)# vpc peer-link

PURESTG-NEXUS-A(config-if)# exit

PURESTG-NEXUS-A(config)# interface Ethernet1/1

PURESTG-NEXUS-A(config-if)# description Nexus5k-B-Cluster-Interconnect

PURESTG-NEXUS-A(config-if)# switchport mode trunk

PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134

PURESTG-NEXUS-A(config-if)# channel-group 1 mode active

PURESTG-NEXUS-A(config-if)# no shutdown

PURESTG-NEXUS-A(config-if)# exit

PURESTG-NEXUS-A(config)# interface Ethernet1/2

PURESTG-NEXUS-A(config-if)# description Nexus5k-B-Cluster-Interconnect

PURESTG-NEXUS-A(config-if)# switchport mode trunk

PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134

PURESTG-NEXUS-A(config-if)# channel-group 1 mode active

```
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit


PURESTG-NEXUS-A(config)# interface Ethernet1/15
PURESTG-NEXUS-A(config-if)# description connect to uplink switch
PURESTG-NEXUS-A(config-if)# switchport access vlan 134
PURESTG-NEXUS-A(config-if)# speed 1000
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# copy running-config startup-config
```

2. Login as admin user into the Nexus Switch B and repeat the above steps to configure second nexus switch. (Note: Make sure to change peer-keepalive destination and source IP address appropriately for Nexus Switch B)

Create vPC Configuration Between Nexus 9372PX-E and Fabric Interconnects

Create and configure vPC 21 and 22 for Data network between Nexus switches and Fabric Interconnects.

Figure 11 vPC Configuration Between Nexus Switches and Fabric Interconnects

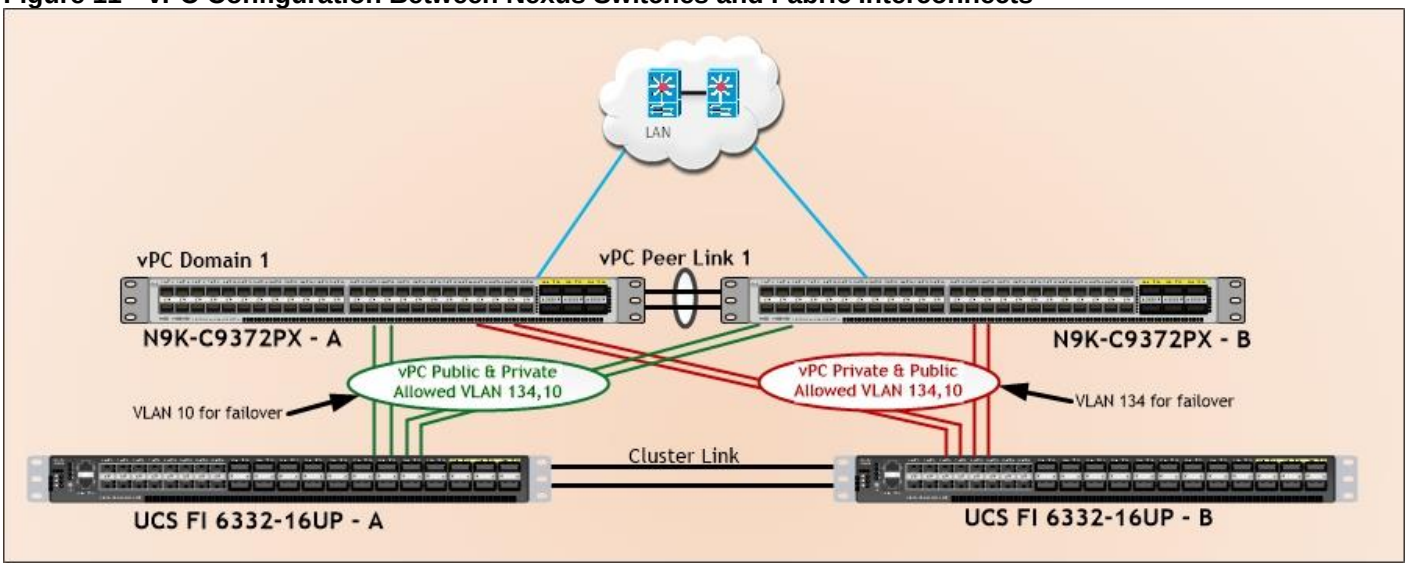

The table below lists the vPC IDs, allowed VLAN IDs, and Ethernet uplink ports.

Table 7 vPC IDs & VLAN IDs

vPC Description	vPC ID	Fabric Interconnects Ports	Nexus Ports	Allowed VLANs
Port Channel FI-A	21	FI-A Port 1/11	N9K-A Port 11	134,10 Note: VLAN 10 needed for failover
		FI-A Port 1/12	N9K-A Port 12	
		FI-A Port 1/13	N9K-B Port 11	
		FI-A Port 1/14	N9K-B Port 12	
Port-Channel FI-B	22	FI-B Port 1/11	N9K-A Port 13	10,134 Note: VLAN 134 needed for failover
		FI-B Port 1/12	N9K-A Port 14	
		FI-B Port 1/13	N9K-B Port 13	
		FI-B Port 1/14	N9K-B Port 14	

To create the necessary port channels between devices, complete the following steps on both Nexus Switches:

1. Login as admin user into Nexus Switch A and perform the following:

PURESTG-NEXUS-A# config Terminal

PURESTG-NEXUS-A(config)# interface port-channel21

PURESTG-NEXUS-A(config-if)# description connect to Fabric Interconnect A

PURESTG-NEXUS-A(config-if)# switchport mode trunk

PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134

PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk

PURESTG-NEXUS-A(config-if)# vpc 21

```
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface port-channel22
PURESTG-NEXUS-A(config-if)# description connect to Fabric Interconnect B
PURESTG-NEXUS-A(config-if)# switchport mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# vpc 22
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit

PURESTG-NEXUS-A(config)# interface Ethernet1/11
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-A:1/11
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 21 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface Ethernet1/12
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-A:1/12
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 21 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface Ethernet1/13
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-B:1/11
PURESTG-NEXUS-A(config-if)# switch mode trunk
```

```
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 22 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface Ethernet1/14
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-B:1/12
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 22 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# copy running-config startup-config
```

2. Login as admin user into the Nexus Switch B and complete the following for the second switch configuration:

```
PURESTG-NEXUS-B# config Terminal
PURESTG-NEXUS-B(config)# interface port-channel21
PURESTG-NEXUS-B(config-if)# description connect to Fabric Interconnect A
PURESTG-NEXUS-B(config-if)# switchport mode trunk
PURESTG-NEXUS-B(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-B(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-B(config-if)# vpc 21
PURESTG-NEXUS-B(config-if)# no shutdown
PURESTG-NEXUS-B(config-if)# exit
PURESTG-NEXUS-B(config)# interface port-channel22
PURESTG-NEXUS-B(config-if)# description connect to Fabric Interconnect B
PURESTG-NEXUS-B(config-if)# switchport mode trunk
PURESTG-NEXUS-B(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-B(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-B(config-if)# vpc 22
```

```
PURESTG-NEXUS-B(config-if)# no shutdown
PURESTG-NEXUS-B(config-if)# exit

PURESTG-NEXUS-B(config)# interface Ethernet1/11
PURESTG-NEXUS-B(config-if)# description Fabric-Interconnect-A:1/13
PURESTG-NEXUS-B(config-if)# switch mode trunk
PURESTG-NEXUS-B(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-B(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-B(config-if)# mtu 9216
PURESTG-NEXUS-B(config-if)# channel-group 21 mode active
PURESTG-NEXUS-B(config-if)# no shutdown
PURESTG-NEXUS-B(config-if)# exit
PURESTG-NEXUS-B(config)# interface Ethernet1/12
PURESTG-NEXUS-B(config-if)# description Fabric-Interconnect-A:1/14
PURESTG-NEXUS-B(config-if)# switch mode trunk
PURESTG-NEXUS-B(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-B(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-B(config-if)# mtu 9216
PURESTG-NEXUS-B(config-if)# channel-group 21 mode active
PURESTG-NEXUS-B(config-if)# no shutdown
PURESTG-NEXUS-B(config-if)# exit
PURESTG-NEXUS-B(config)# interface Ethernet1/13
PURESTG-NEXUS-B(config-if)# description Fabric-Interconnect-B:1/13
PURESTG-NEXUS-B(config-if)# switch mode trunk
PURESTG-NEXUS-B(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-B(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-B(config-if)# mtu 9216
PURESTG-NEXUS-B(config-if)# channel-group 22 mode active
PURESTG-NEXUS-B(config-if)# no shutdown
PURESTG-NEXUS-B(config-if)# exit
PURESTG-NEXUS-B(config)# interface Ethernet1/14
PURESTG-NEXUS-B(config-if)# description Fabric-Interconnect-B:1/14
```

```
PURESTG-NEXUS-B(config-if)# switch mode trunk
PURESTG-NEXUS-B(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-B(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-B(config-if)# mtu 9216
PURESTG-NEXUS-B(config-if)# channel-group 22 mode active
PURESTG-NEXUS-B(config-if)# no shutdown
PURESTG-NEXUS-B(config-if)# exit
PURESTG-NEXUS-B(config)# copy running-config startup-config
```

Verify All vPC Status Is Up on Both Cisco Nexus Switches

Figure 12 Cisco Nexus Switch A Port-Channel Summary

```
PURESTG-NEXUS-A#
PURESTG-NEXUS-A# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
19    Po19(SU)    Eth      LACP      Eth1/17(P)   Eth1/18(P)
20    Po20(SU)    Eth      LACP      Eth1/19(P)   Eth1/20(P)
21    Po21(SU)    Eth      LACP      Eth1/11(P)   Eth1/12(P)
22    Po22(SU)    Eth      LACP      Eth1/13(P)   Eth1/14(P)
PURESTG-NEXUS-A#
```

Figure 13 Cisco Nexus Switch B Port-Channel Summary

```
PURESTG-NEXUS-B#
PURESTG-NEXUS-B# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
19    Po19(SU)    Eth      LACP      Eth1/17(P)   Eth1/18(P)
20    Po20(SU)    Eth      LACP      Eth1/19(P)   Eth1/20(P)
21    Po21(SU)    Eth      LACP      Eth1/11(P)   Eth1/12(P)
22    Po22(SU)    Eth      LACP      Eth1/13(P)   Eth1/14(P)
PURESTG-NEXUS-B#
```

Figure 14 vPC Description for Cisco Nexus Switch A

```
PURESTG-NEXUS-A#
PURESTG-NEXUS-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 inconsistency reason       : Consistency Check Not Performed
vPC role                          : primary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

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

vPC status
----------------------------------------------------------------------------
id   Port   Status Consistency Reason                     Active vlans
--   ----   ------ ----------- ------                     ------------
19   Po19   up     success     success                    1,10,134

20   Po20   up     success     success                    1,10,134

21   Po21   up     success     success                    1,10,134

22   Po22   up     success     success                    1,10,134

PURESTG-NEXUS-A#
```

Figure 15 vPC Description for Cisco Nexus Switch B

```
PURESTG-NEXUS-B#
PURESTG-NEXUS-B# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 inconsistency reason       : Consistency Check Not Performed
vPC role                          : secondary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled (timeout = 240 seconds)

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

vPC status
----------------------------------------------------------------------------
id   Port   Status Consistency Reason                Active vlans
--   ----   ------ ----------- ------                ------------
19   Po19   up     success     success               1,10,134

20   Po20   up     success     success               1,10,134

21   Po21   up     success     success               1,10,134

22   Po22   up     success     success               1,10,134

PURESTG-NEXUS-B#
```

Configure Cisco MDS 9148S Switches

Connect MDS Switches to Fabric Interconnects and Pure Storage System as shown in the figure below:

Figure 16 MDS, FI, and Pure Storage Layout

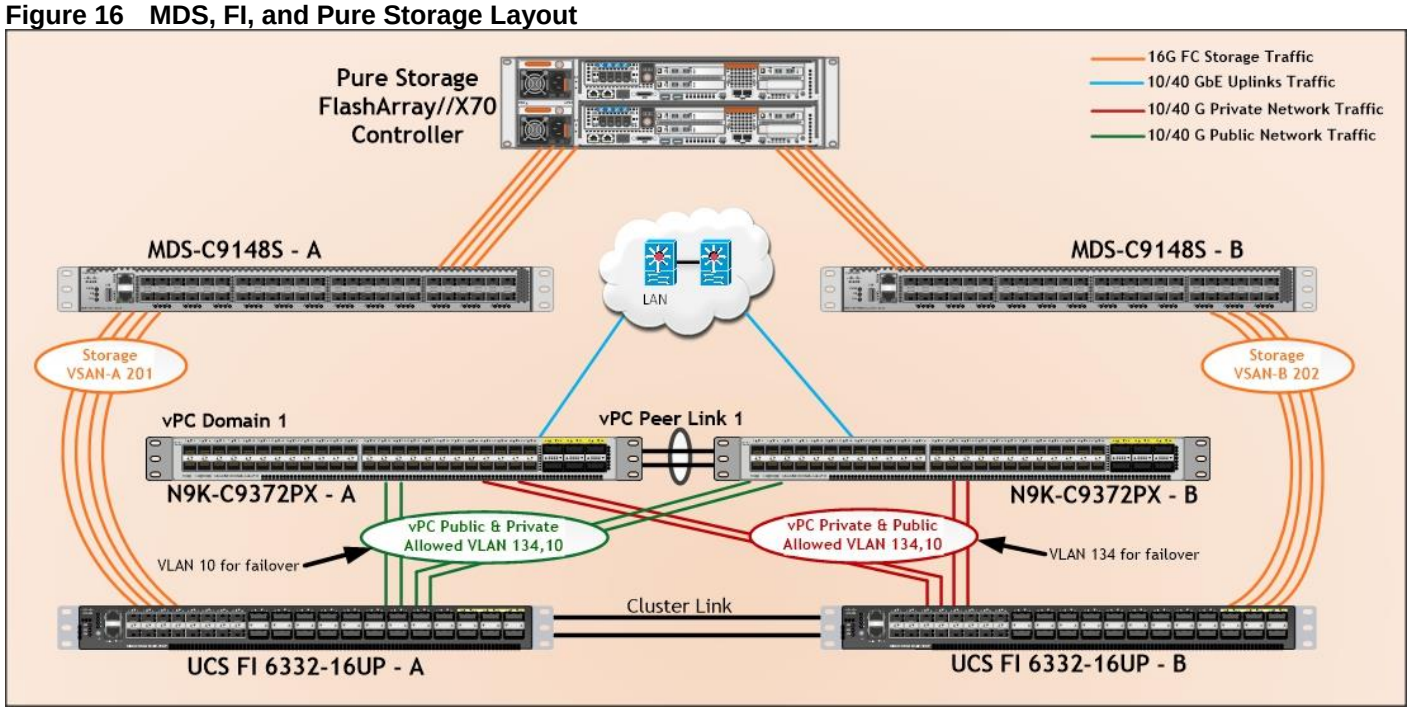

For this solution, we have connected four ports (ports 33 to 36) of MDS Switch A to Fabric Interconnect A (ports 1-4). Similarly, we have connected four ports (ports 33 to 36) of MDS Switch B to Fabric Interconnect B (ports 1-4) as shown in the table below. All ports carry 16 Gb/s FC Traffic.

Table 8 MDS 9148S Port Connection to Fabric Interconnects

MDS Switch	MDS Switch Port	FI Ports	Fabric Interconnects
MDS Switch A	FC Port 1/33	FI-A Port 1/1	Fabric Interconnect A (FI-A)
	FC Port 1/34	FI-A Port 1/2	
	FC Port 1/35	FI-A Port 1/3	
	FC Port 1/36	FI-A Port 1/4	
MDS Switch B	FC Port 1/33	FI-B Port 1/1	Fabric Interconnect B (FI-B)
	FC Port 1/34	FI-B Port 1/2	
	FC Port 1/35	FI-B Port 1/3	
	FC Port 1/36	FI-B Port 1/4	

For this solution, we connected four ports (ports 25 to 28) of MDS Switch A to Pure Storage System. Similarly, we connected four ports (ports 25 to 28) of MDS Switch B to Pure Storage System as shown in the table below. All ports carry 16 Gb/s FC Traffic.

Table 9 MDS 9148S Port Connection to Pure Storage System

MDS Switch	MDS Switch Port	Pure Storage	Storage Port
MDS Switch A	FC Port 1/25	Storage Controller-0	CT0-FC0
	FC Port 1/26	Storage Controller-0	CT0-FC6
	FC Port 1/27	Storage Controller-1	CT1-FC0
	FC Port 1/28	Storage Controller-1	CT1-FC6
MDS Switch B	FC Port 1/25	Storage Controller-0	CT0-FC1
	FC Port 1/26	Storage Controller-0	CT0-FC7
	FC Port 1/27	Storage Controller-1	CT1-FC1
	FC Port 1/28	Storage Controller-1	CT1-FC7

Configure Feature for MDS Switch A and MDS Switch B

To set feature on MDS Switches, complete the following steps on both MDS switches:

1. Login as admin user into MDS Switch A.

```
PURESTG-MDS-A# config terminal
PURESTG-MDS-A(config)# feature npiv
PURESTG-MDS-A(config)# feature telnet
PURESTG-MDS-A(config)# switchname PURESTG-MDS-A
PURESTG-MDS-A(config)# copy running-config startup-config
```

(1) Login as admin user into MDS Switch B.

```
PURESTG-MDS-B# config terminal
PURESTG-MDS-B(config)# feature npiv
PURESTG-MDS-B(config)# feature telnet
PURESTG-MDS-B(config)# switchname PURESTG-MDS-B
PURESTG-MDS-B(config)# copy running-config startup-config
```

Configure VSANs for MDS Switch A and MDS Switch B

To create VSANs, complete the following steps on both MDS switches:

1. Login as admin user into MDS Switch A.
2. Create VSAN 201 for Storage Traffic:

```
PURESTG-MDS-A # config terminal
PURESTG-MDS-A(config)# VSAN database
PURESTG-MDS-A(config-vsan-db)# vsan 201
PURESTG-MDS-A(config-vsan-db)# vsan 201 interface fc 1/25-36
PURESTG-MDS-A(config-vsan-db)# exit

PURESTG-MDS-A(config)# interface fc 1/25-36
PURESTG-MDS-A(config-if)# switchport trunk allowed vsan 201
PURESTG-MDS-A(config-if)# switchport trunk mode off
PURESTG-MDS-A(config-if)# port-license acquire
PURESTG-MDS-A(config-if)# no shutdown
PURESTG-MDS-A(config-if)# exit
PURESTG-MDS-A(config)# copy running-config startup-config
```

3. Login as admin user into MDS Switch B.

4. Create VSAN 202 for Storage Traffic:

```
PURESTG-MDS-B # config terminal
PURESTG-MDS-B(config)# VSAN database
PURESTG-MDS-B(config-vsan-db)# vsan 202
PURESTG-MDS-B(config-vsan-db)# vsan 202 interface fc 1/25-36
PURESTG-MDS-B(config-vsan-db)# exit

PURESTG-MDS-B(config)# interface fc 1/25-36
PURESTG-MDS-B(config-if)# switchport trunk allowed vsan 202
PURESTG-MDS-B(config-if)# switchport trunk mode off
PURESTG-MDS-B(config-if)# port-license acquire
PURESTG-MDS-B(config-if)# no shutdown
PURESTG-MDS-B(config-if)# exit
PURESTG-MDS-B(config)# copy running-config startup-config
```

Create and Configure Fiber Channel Zoning

This procedure sets up the Fibre Channel connections between the Cisco MDS 9148S switches, the Cisco UCS Fabric Interconnects, and the Pure Storage FlashArray systems.

Before you configure the zoning details, decide how many paths are needed for each LUN and extract the WWPN numbers for each of the HBAs from each server. We used 4 HBAs for each Server. Two HBAs (HBA0 and HBA2) are connected to MDS Switch-A and other two HBAs (HBA1 and HBA3) are connected to MDS Switch-B.

To create and configure the fiber channel zoning, complete the following steps:

1. Log in to the Cisco UCS Manager > Equipment > Chassis > Servers and select the desired server. On the right hand menu, click the Inventory tab and HBA's sub-tab to get the WWPN of HBA's as shown in the screenshot below:

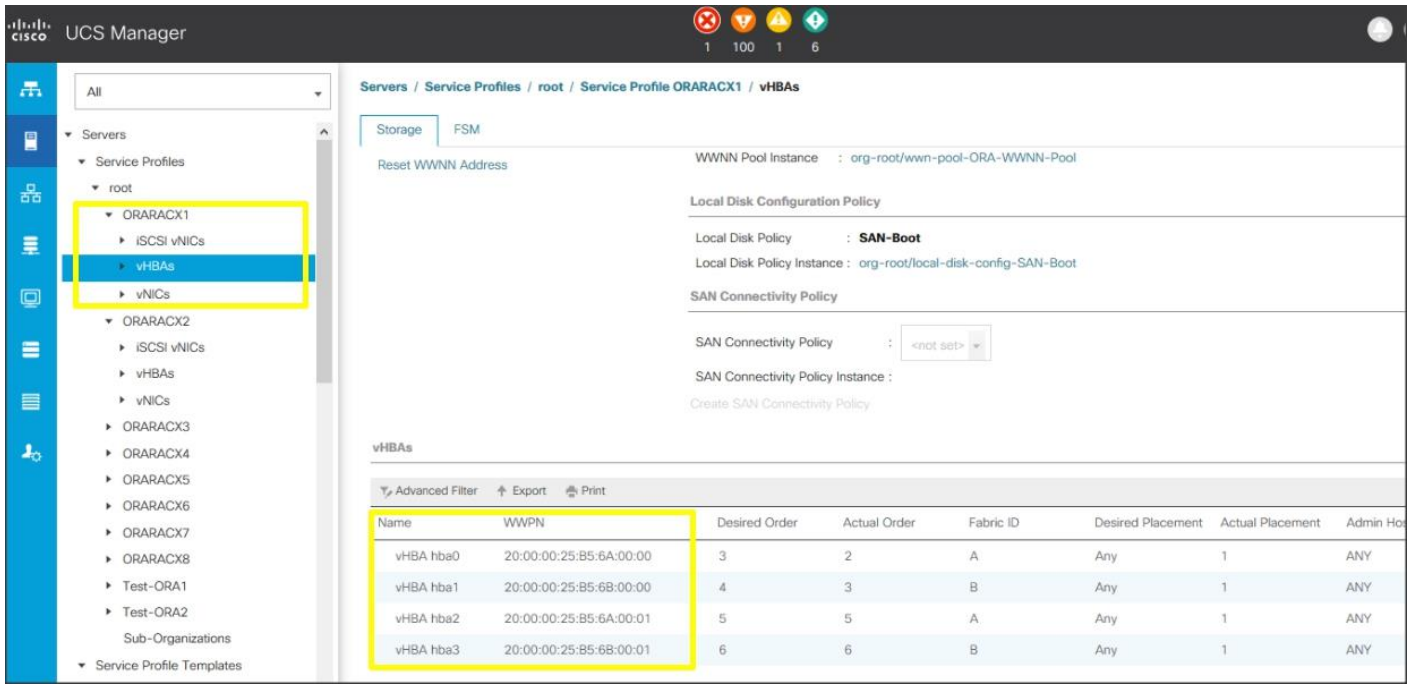

2. Connect to the Pure Storage System and extract the WWPN of FC Ports connected to the Cisco MDS Switches. We have connected 8 FC ports from Pure Storage System to Cisco MDS Switches. FC ports CT0.FC0, CT1.FC0, CT0.FC6, CT1.FC6 are connected to MDS Switch-A and similarly FC ports CT0.FC1, CT1.FC1, CT0.FC7, CT1.FC7 are connected to MDS Switch-B.

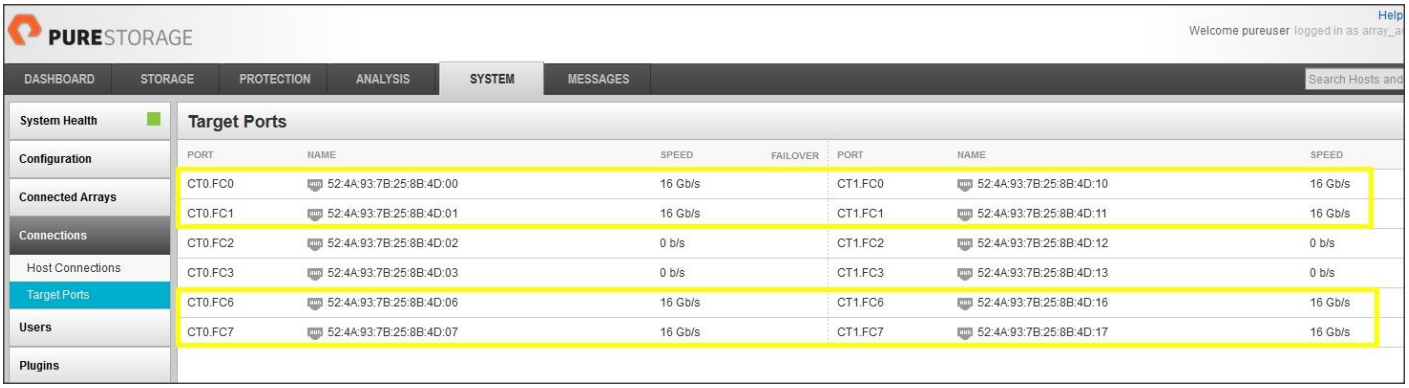

Create Device Aliases for Fiber Channel Zoning

Cisco MDS Switch A

To configure device aliases and zones for the SAN boot paths as well as datapaths of MDS switch A, complete the following steps:

1. Login as admin user and run the following commands:

conf t

device-alias database

device-alias name oraracx1-hba0 pwwn 20:00:00:25:b5:6a:00:00

device-alias name oraracx1-hba2 pwwn 20:00:00:25:b5:6a:00:01

device-alias name oraracx2-hba0 pwwn 20:00:00:25:b5:6a:00:02

device-alias name oraracx2-hba2 pwwn 20:00:00:25:b5:6a:00:03

device-alias name oraracx3-hba0 pwwn 20:00:00:25:b5:6a:00:04

device-alias name oraracx3-hba2 pwwn 20:00:00:25:b5:6a:00:05

device-alias name oraracx4-hba0 pwwn 20:00:00:25:b5:6a:00:06

device-alias name oraracx4-hba2 pwwn 20:00:00:25:b5:6a:00:07

device-alias name oraracx5-hba0 pwwn 20:00:00:25:b5:6a:00:08

device-alias name oraracx5-hba2 pwwn 20:00:00:25:b5:6a:00:09

device-alias name oraracx6-hba0 pwwn 20:00:00:25:b5:6a:00:0a

device-alias name oraracx6-hba2 pwwn 20:00:00:25:b5:6a:00:0b

device-alias name oraracx7-hba0 pwwn 20:00:00:25:b5:6a:00:0c

device-alias name oraracx7-hba2 pwwn 20:00:00:25:b5:6a:00:0d

device-alias name oraracx8-hba0 pwwn 20:00:00:25:b5:6a:00:0e

device-alias name oraracx8-hba2 pwwn 20:00:00:25:b5:6a:00:0f

device-alias name FLASHSTACK-X-CT0-FC0 pwwn 52:4a:93:7b:25:8b:4d:00

device-alias name FLASHSTACK-X-CT0-FC6 pwwn 52:4a:93:7b:25:8b:4d:06

device-alias name FLASHSTACK-X-CT1-FC0 pwwn 52:4a:93:7b:25:8b:4d:10

device-alias name FLASHSTACK-X-CT1-FC6 pwwn 52:4a:93:7b:25:8b:4d:16

device-alias commit

Cisco MDS Switch B

To configure device aliases and zones for the SAN boot paths as well as datapaths of MDS switch B, complete the following steps:

1. Login as admin user and run the following commands:

conf t

device-alias database

device-alias name oraracx1-hba1 pwwn 20:00:00:25:b5:6b:00:00

device-alias name oraracx1-hba3 pwwn 20:00:00:25:b5:6b:00:01

device-alias name oraracx2-hba1 pwwn 20:00:00:25:b5:6b:00:02

device-alias name oraracx2-hba3 pwwn 20:00:00:25:b5:6b:00:03

device-alias name oraracx3-hba1 pwwn 20:00:00:25:b5:6b:00:04

device-alias name oraracx3-hba3 pwwn 20:00:00:25:b5:6b:00:05

device-alias name oraracx4-hba1 pwwn 20:00:00:25:b5:6b:00:06

device-alias name oraracx4-hba3 pwwn 20:00:00:25:b5:6b:00:07

device-alias name oraracx5-hba1 pwwn 20:00:00:25:b5:6b:00:08

device-alias name oraracx5-hba3 pwwn 20:00:00:25:b5:6b:00:09

device-alias name oraracx6-hba1 pwwn 20:00:00:25:b5:6b:00:0a

device-alias name oraracx6-hba3 pwwn 20:00:00:25:b5:6b:00:0b

device-alias name oraracx7-hba1 pwwn 20:00:00:25:b5:6b:00:0c

device-alias name oraracx7-hba3 pwwn 20:00:00:25:b5:6b:00:0d

device-alias name oraracx8-hba1 pwwn 20:00:00:25:b5:6b:00:0e

device-alias name oraracx8-hba3 pwwn 20:00:00:25:b5:6b:00:0f

device-alias name FLASHSTACK-X-CT0-FC1 pwwn 52:4a:93:7b:25:8b:4d:01

device-alias name FLASHSTACK-X-CT0-FC7 pwwn 52:4a:93:7b:25:8b:4d:07

device-alias name FLASHSTACK-X-CT1-FC1 pwwn 52:4a:93:7b:25:8b:4d:11

device-alias name FLASHSTACK-X-CT1-FC7 pwwn 52:4a:93:7b:25:8b:4d:17

device-alias commit

Create Zoning

Cisco MDS Switch A

To configure zones for the MDS switch A, complete the following steps:

1. Create a zone for each service profile.
2. Login as admin user and create the zone as shown below:

conf t

zone name oraracx1 vsan 201

member pwwn 52:4a:93:7b:25:8b:4d:00

member pwwn 52:4a:93:7b:25:8b:4d:06

member pwwn 52:4a:93:7b:25:8b:4d:10

member pwwn 52:4a:93:7b:25:8b:4d:16

member pwwn 20:00:00:25:b5:6a:00:00

member pwwn 20:00:00:25:b5:6a:00:01

zone name oraracx2 vsan 201

member pwwn 52:4a:93:7b:25:8b:4d:00

member pwwn 52:4a:93:7b:25:8b:4d:06

member pwwn 52:4a:93:7b:25:8b:4d:10

member pwwn 52:4a:93:7b:25:8b:4d:16

member pwwn 20:00:00:25:b5:6a:00:02

```
    member pwwn 20:00:00:25:b5:6a:00:03
  zone name oraracx3 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
    member pwwn 52:4a:93:7b:25:8b:4d:06
    member pwwn 52:4a:93:7b:25:8b:4d:10
    member pwwn 52:4a:93:7b:25:8b:4d:16
    member pwwn 20:00:00:25:b5:6a:00:04
    member pwwn 20:00:00:25:b5:6a:00:05
  zone name oraracx4 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
    member pwwn 52:4a:93:7b:25:8b:4d:06
    member pwwn 52:4a:93:7b:25:8b:4d:10
    member pwwn 52:4a:93:7b:25:8b:4d:16
    member pwwn 20:00:00:25:b5:6a:00:06
    member pwwn 20:00:00:25:b5:6a:00:07
  zone name oraracx5 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
    member pwwn 52:4a:93:7b:25:8b:4d:06
    member pwwn 52:4a:93:7b:25:8b:4d:10
    member pwwn 52:4a:93:7b:25:8b:4d:16
    member pwwn 20:00:00:25:b5:6a:00:08
    member pwwn 20:00:00:25:b5:6a:00:09
  zone name oraracx6 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
    member pwwn 52:4a:93:7b:25:8b:4d:06
    member pwwn 52:4a:93:7b:25:8b:4d:10
    member pwwn 52:4a:93:7b:25:8b:4d:16
    member pwwn 20:00:00:25:b5:6a:00:0a
    member pwwn 20:00:00:25:b5:6a:00:0b
  zone name oraracx7 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
    member pwwn 52:4a:93:7b:25:8b:4d:06
```

```
    member pwwn 52:4a:93:7b:25:8b:4d:10
    member pwwn 52:4a:93:7b:25:8b:4d:16
    member pwwn 20:00:00:25:b5:6a:00:0c
    member pwwn 20:00:00:25:b5:6a:00:0d
zone name oraracx8 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
    member pwwn 52:4a:93:7b:25:8b:4d:06
    member pwwn 52:4a:93:7b:25:8b:4d:10
    member pwwn 52:4a:93:7b:25:8b:4d:16
    member pwwn 20:00:00:25:b5:6a:00:0e
    member pwwn 20:00:00:25:b5:6a:00:0f
exit
```

3. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members.

```
conf t
zoneset name oraracx vsan 201
    member oraracx1
    member oraracx2
    member oraracx3
    member oraracx4
    member oraracx5
    member oraracx6
    member oraracx7
    member oraracx8
exit
```

4. Activate the zone set by running following commands.

```
zoneset activate name oraracx vsan 201
exit
```

copy run start

Cisco MDS Switch B

To configure zones for the MDS switch B, complete the following steps:

1. Create a zone for each service profile.
2. Login as admin user and run the following commands:

```
conf t
zone name oraracx1 vsan 202
   member pwwn 52:4a:93:7b:25:8b:4d:01
   member pwwn 52:4a:93:7b:25:8b:4d:07
   member pwwn 52:4a:93:7b:25:8b:4d:11
   member pwwn 52:4a:93:7b:25:8b:4d:17
   member pwwn 20:00:00:25:b5:6b:00:00
   member pwwn 20:00:00:25:b5:6b:00:01
zone name oraracx2 vsan 202
   member pwwn 52:4a:93:7b:25:8b:4d:01
   member pwwn 52:4a:93:7b:25:8b:4d:07
   member pwwn 52:4a:93:7b:25:8b:4d:11
   member pwwn 52:4a:93:7b:25:8b:4d:17
   member pwwn 20:00:00:25:b5:6b:00:02
   member pwwn 20:00:00:25:b5:6b:00:03
zone name oraracx3 vsan 202
   member pwwn 52:4a:93:7b:25:8b:4d:01
   member pwwn 52:4a:93:7b:25:8b:4d:07
   member pwwn 52:4a:93:7b:25:8b:4d:11
   member pwwn 52:4a:93:7b:25:8b:4d:17
   member pwwn 20:00:00:25:b5:6b:00:04
   member pwwn 20:00:00:25:b5:6b:00:05
zone name oraracx4 vsan 202
   member pwwn 52:4a:93:7b:25:8b:4d:01
   member pwwn 52:4a:93:7b:25:8b:4d:07
   member pwwn 52:4a:93:7b:25:8b:4d:11
   member pwwn 52:4a:93:7b:25:8b:4d:17
```

```
    member pwwn 20:00:00:25:b5:6b:00:06
    member pwwn 20:00:00:25:b5:6b:00:07
  zone name oraracx5 vsan 202
    member pwwn 52:4a:93:7b:25:8b:4d:01
    member pwwn 52:4a:93:7b:25:8b:4d:07
    member pwwn 52:4a:93:7b:25:8b:4d:11
    member pwwn 52:4a:93:7b:25:8b:4d:17
    member pwwn 20:00:00:25:b5:6b:00:08
    member pwwn 20:00:00:25:b5:6b:00:09
  zone name oraracx6 vsan 202
    member pwwn 52:4a:93:7b:25:8b:4d:01
    member pwwn 52:4a:93:7b:25:8b:4d:07
    member pwwn 52:4a:93:7b:25:8b:4d:11
    member pwwn 52:4a:93:7b:25:8b:4d:17
    member pwwn 20:00:00:25:b5:6b:00:0a
    member pwwn 20:00:00:25:b5:6b:00:0b
  zone name oraracx7 vsan 202
    member pwwn 52:4a:93:7b:25:8b:4d:01
    member pwwn 52:4a:93:7b:25:8b:4d:07
    member pwwn 52:4a:93:7b:25:8b:4d:11
    member pwwn 52:4a:93:7b:25:8b:4d:17
    member pwwn 20:00:00:25:b5:6b:00:0c
    member pwwn 20:00:00:25:b5:6b:00:0d
  zone name oraracx8 vsan 202
    member pwwn 52:4a:93:7b:25:8b:4d:01
    member pwwn 52:4a:93:7b:25:8b:4d:07
    member pwwn 52:4a:93:7b:25:8b:4d:11
    member pwwn 52:4a:93:7b:25:8b:4d:17
    member pwwn 20:00:00:25:b5:6b:00:0e
    member pwwn 20:00:00:25:b5:6b:00:0f
  exit
```

3. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members:

```
conf t
zoneset name oraracx vsan 202
   member oraracx1
   member oraracx2
   member oraracx3
   member oraracx4
   member oraracx5
   member oraracx6
   member oraracx7
   member oraracx8
exit
```

4. Activate the zone set by running following commands:

```
zoneset activate name oraracx vsan 202
exit
copy run start
```

Configure Pure Storage

The design goal of the reference architecture was to best represent a real-world environment as closely as possible. The approach included features of Cisco UCS to rapidly deploy stateless servers and use Pure Storage FlashArray's boot LUNs to provision the O.S on top it. Zoning was performed on the Cisco MDS 9148S switches to enable the initiators discover the targets during boot process.

A Service Profile was created within Cisco UCS Manager to deploy the 8 servers quickly with a standard configuration. SAN boot volumes for these servers were hosted on the same Pure Storage FlashArray //X70. Once the stateless servers were provisioned, following process was performed to enable Rapid deployment of 8 RAC nodes.

Each Server node has dedicated single LUN to install operating system and all the eight server node was booted off SAN. For this solution, we have installed Oracle Linux 7.4 on this LUNs and performed all the pre-requisite packages for Oracle Database 12cR2 to create eight node Oracle RAC database solution.

Using logical servers that are disassociated from the physical hardware removes many limiting constraints around how servers are provisioned. Cisco UCS Service Profiles contain values for a server's property settings, including virtual network interface cards (vNICs), MAC addresses, boot policies, firmware policies, fabric connectivity, external management, and HA information. The service profiles represent all the attributes of a logical server in Cisco UCS model. By abstracting these settings from the physical server into a Cisco Service Profile, the Service Profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, Service Profiles can, at any time, be migrated from one physical server to another. Furthermore, Cisco is the only hardware provider to offer a truly unified management platform, with Cisco UCS Service Profiles and hardware abstraction capabilities extending to both blade and rack servers.

In addition to the service profiles, the use of Pure Storage's FlashArray's with SAN boot policy provides the following benefits:

- Scalability - Rapid deployment of new servers to the environment in a very few steps
- Manageability - Enables seamless hardware maintenance and upgrades without any restrictions. This is a huge benefit in comparison to other appliance model like Exadata
- Flexibility - Easy to repurpose physical servers for different applications and services as needed
- Availability - Hardware failures are not impactful and critical. In rare case of a server failure, it is easier to associate the logical service profile to another healthy physical server to reduce the impact

Configure Host

Before using a volume (LUN) on a host, the host has to be defined on Pure FlashArray. To set up a host complete the following steps:

1. Log into Pure Storage dashboard.
2. In the PURE GUI, go to Storage tab.
3. Under Hosts option in the left frame, click the + sign to create a host.
4. Enter the name of the host and click Create. This will create a Host entry under the Hosts category.

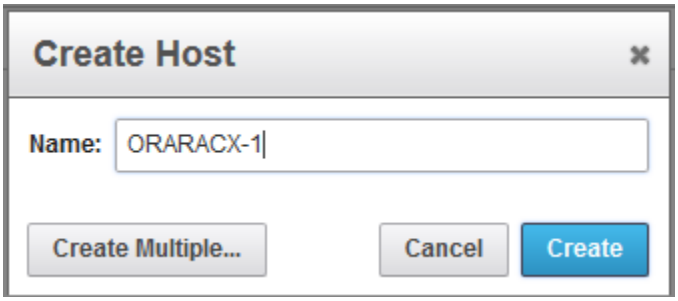

5. To update the host with the connectivity information by providing the Fibre Channel WWNs or iSCSI IQNs, click the Host that was created.

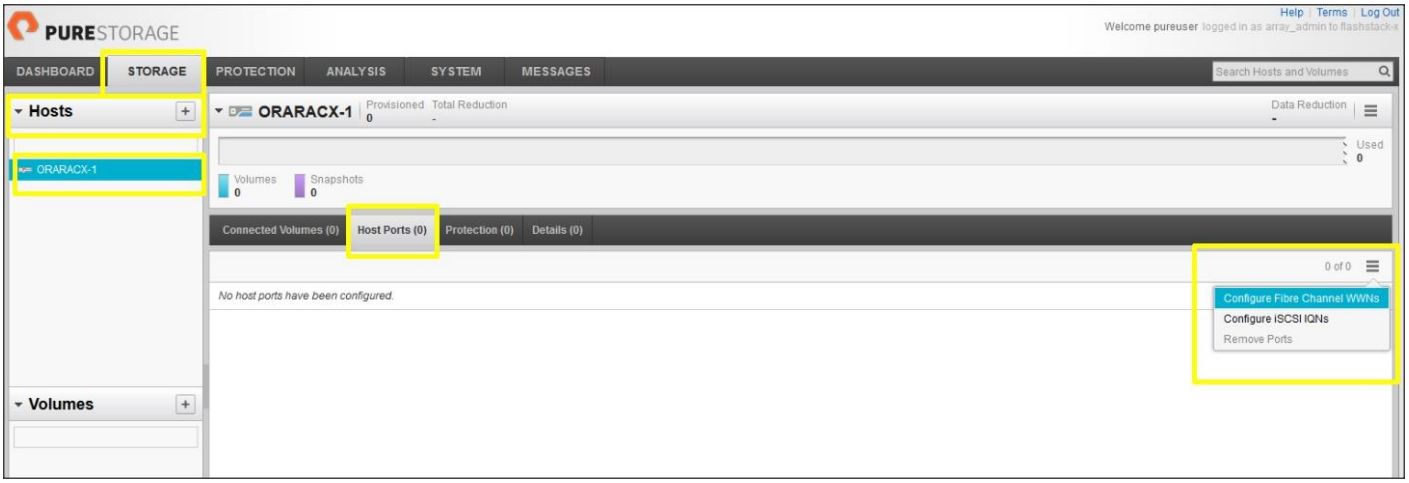

6. In the host context, click the Host Ports tab and click the settings button and select "Configure Fibre Channel WWNs" which will display a window with the available WWNs in the left side.

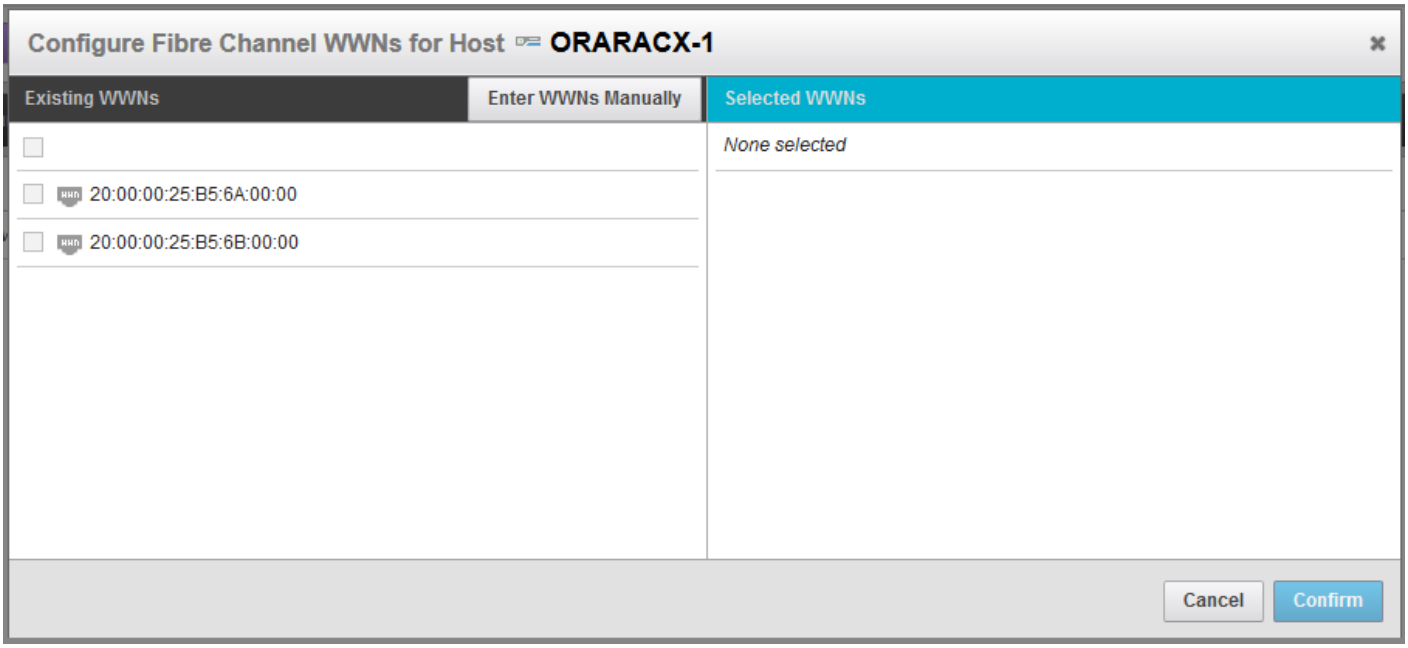

WWNs will appear only if the appropriate FC connections were made and the zones were setup on the underlying FC switch.

7. Select the list of WWNs that belongs to the host in the next window and click “Confirm.”

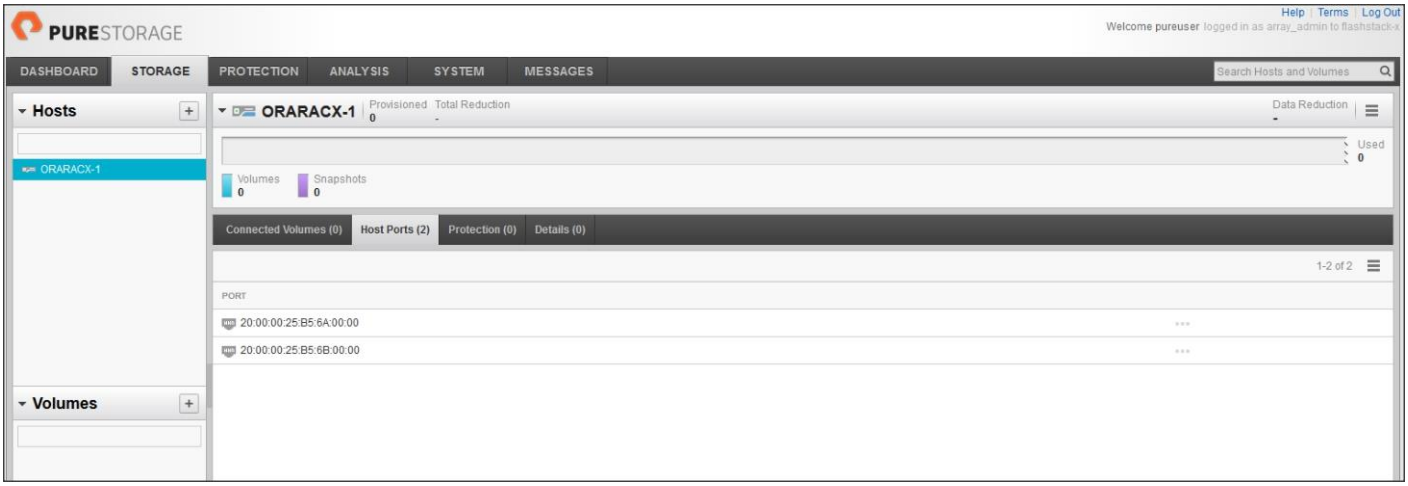

Make sure the zoning has been setup to include the WWNs details of the initiators along with the target, without which the SAN boot will not work.

Configure Volume

To configure a volume, complete the following steps:

1. Go to tab Storage > Volumes > and click on + sign to “Create Volume.”

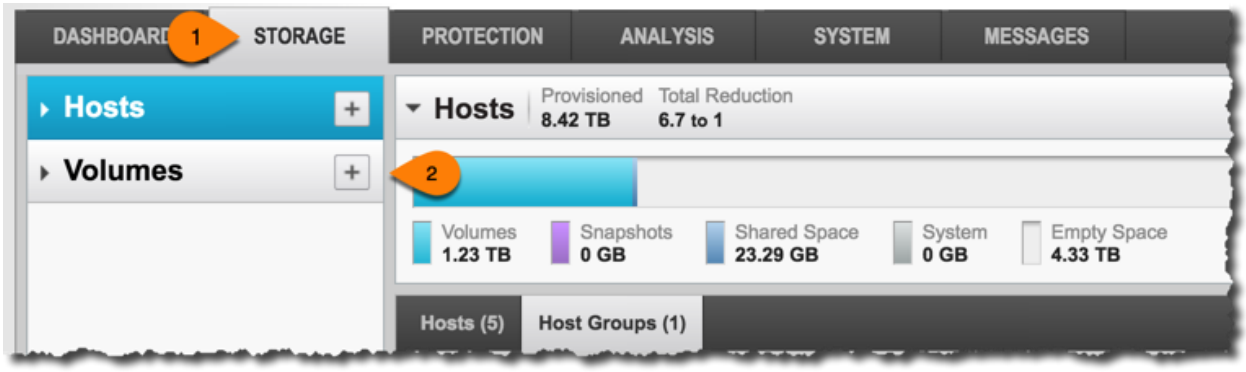

2. Provide the name of the volume, size, choose the size type (KB, MB, GB, TB, PB) and click Create to create the volume.

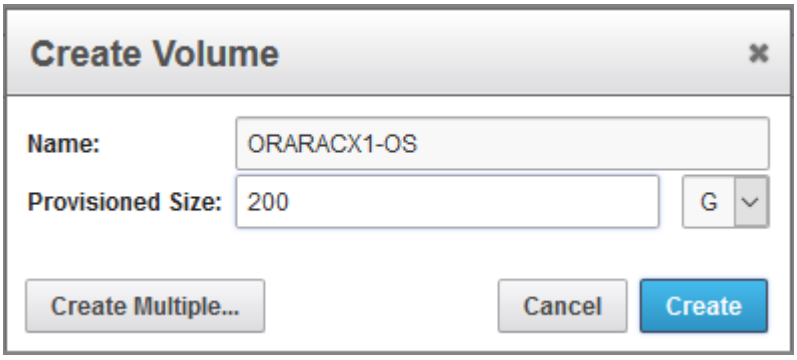

3. Attach the volume to a host by going to the “Connected Hosts and Host Groups” tab under the volume context menu.

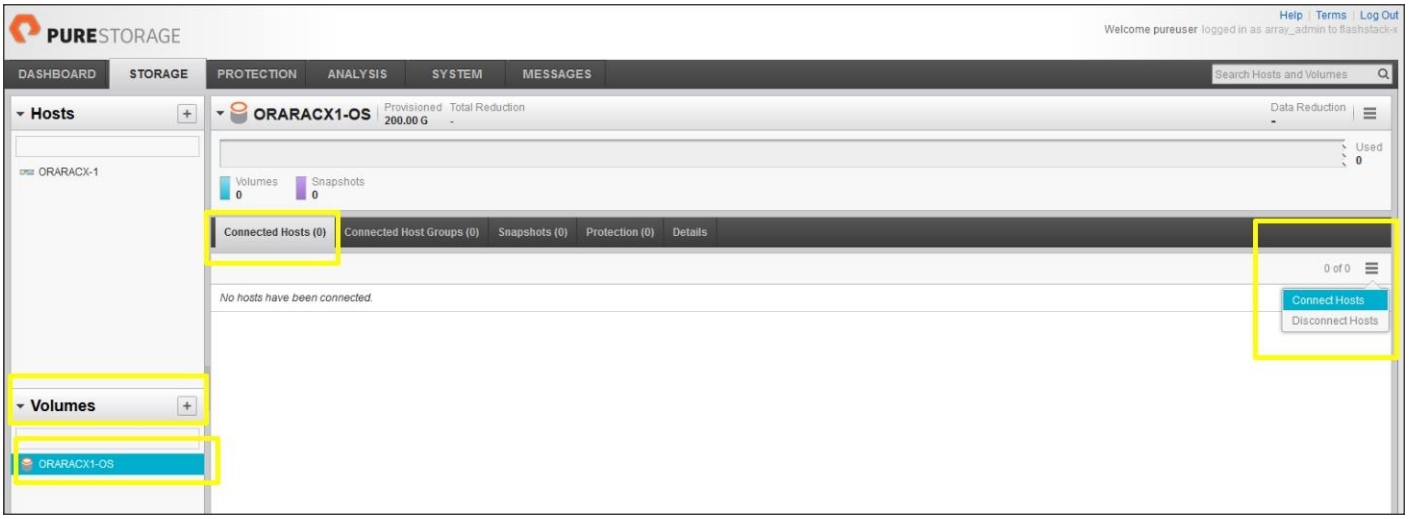

4. Click the Settings icon and select “Connect Hosts.” Select the host where the volume should be attached and click Confirm.

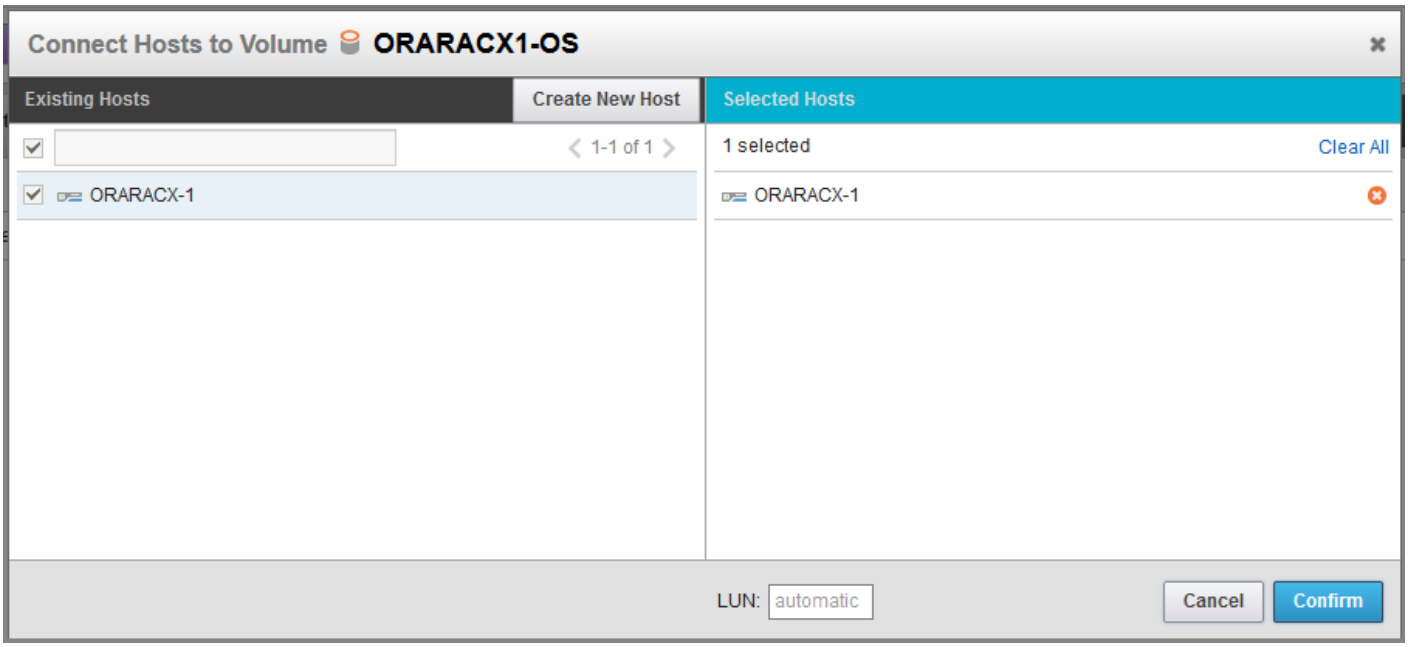

This completes the connectivity of Storage LUN to the Server Node. You created one boot LUN (ORARACX1-OS) of 200GB and assigned this LUN to the first oracle RAC node “ORARACX-1”. You will install OS and perform all prerequisites for Oracle RAC Database to this LUN. Similarly, repeat the steps to create 7 more Hosts and LUNs to get ready 8 nodes for OS installation in the next steps.

OS and Database Deployment

Operating System Configuration

Step-by-step OS installation details are not detailed in this document, but the following section describes the key steps for OS install.

1. Download Oracle Linux 7.4 OS image from https://edelivery.oracle.com/linux.

2. Launch KVM console on desired server by going to tab Equipment > Chassis > Chassis 1 > Servers > Server 1 > from right side windows General > and select KVM Console to open KVM.

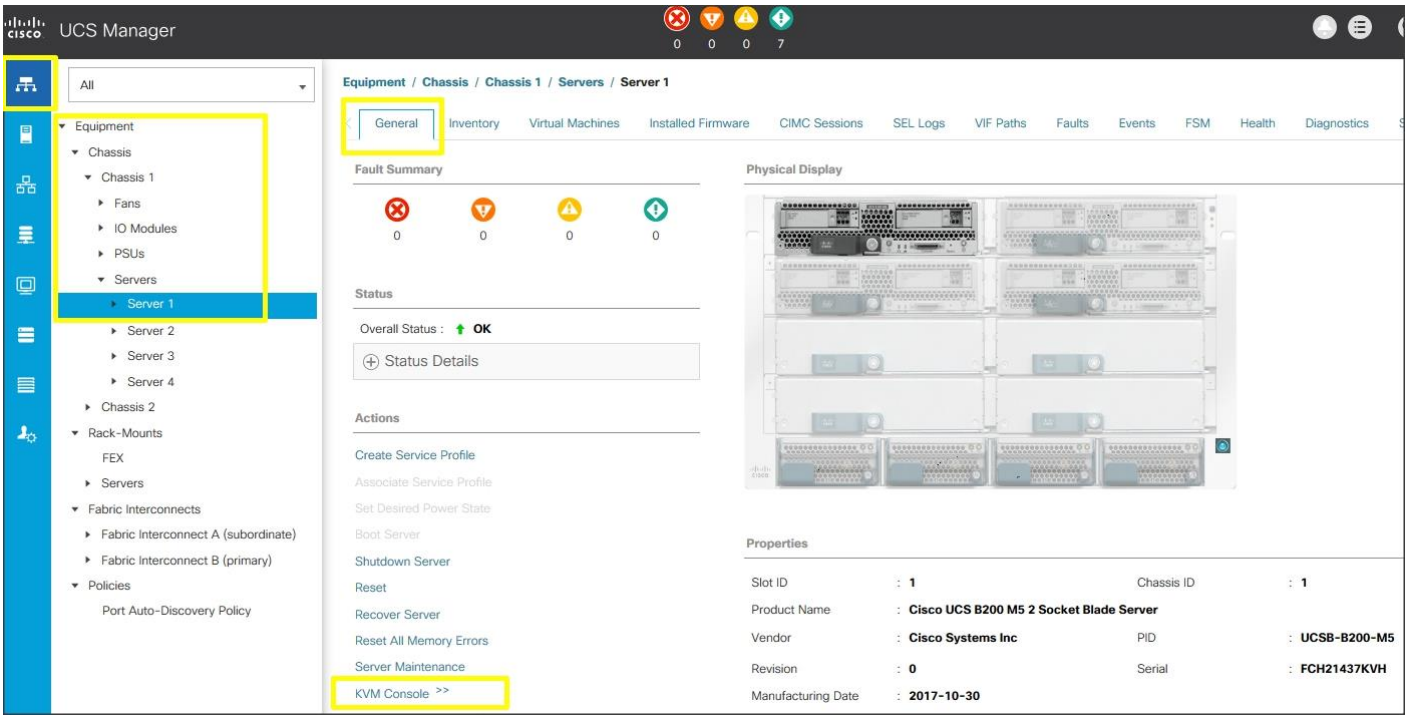

3. Click Accept security and open KVM. Enable virtual media, map the Oracle Linux ISO image and reset the server.

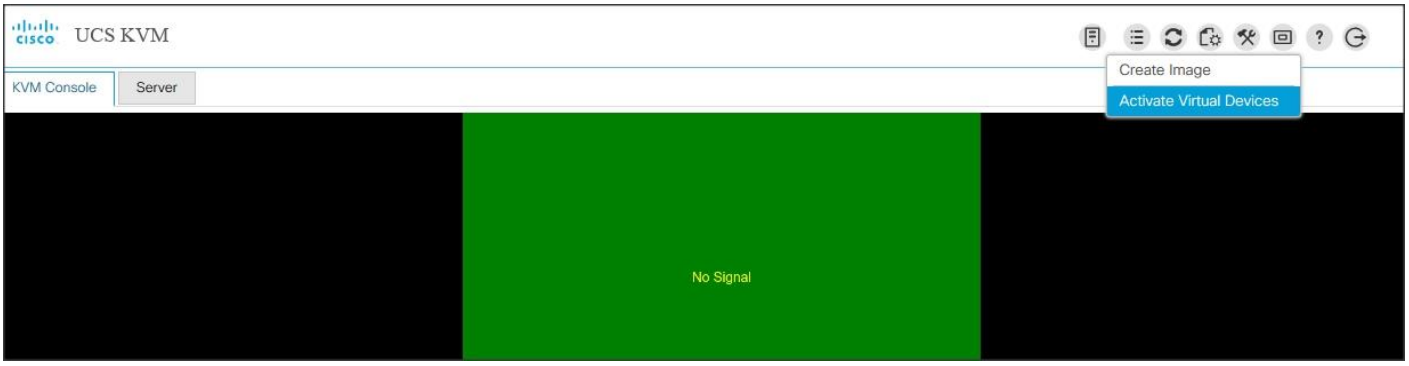

4. When the Server starts booting, it will detect the Pure Storage active FC paths as shown below. If you see the following message in the KVM console while the server is rebooting along with the target WWPNs, it confirms the setup is done correctly and boot from SAN will be successful.

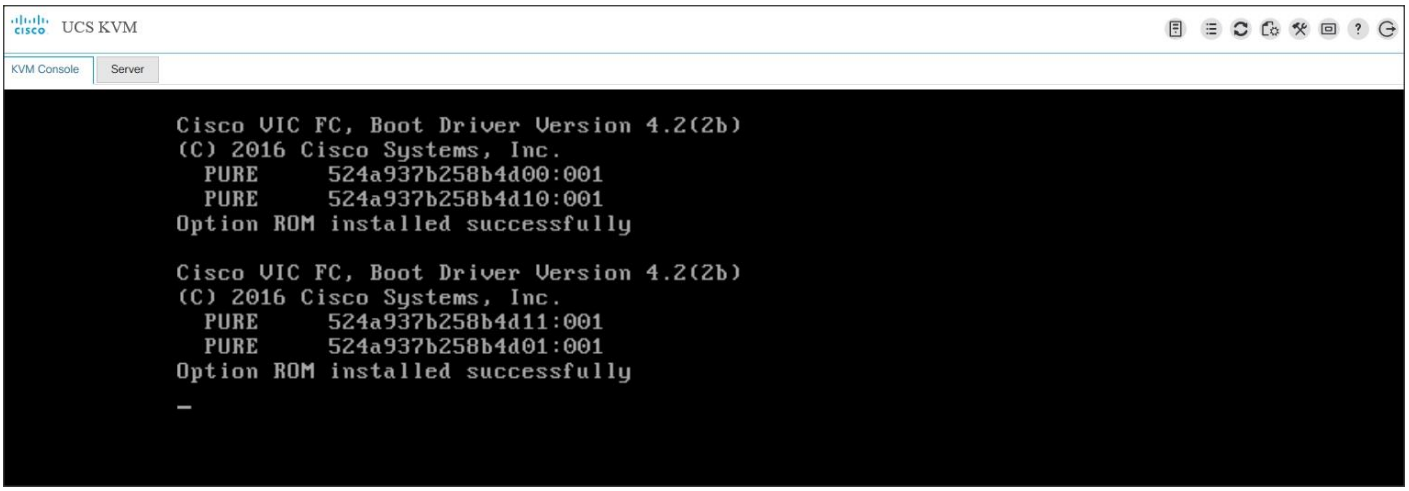

5. During server boot order, it will detect the virtual media connected as Oracle Linux cd. It should launch the Oracle Linux installer. Select language and assign the Installation destination as Pure Storage FlashArray LUN. Apply hostname and click “Configure Network” to configure all network interfaces. Alternatively, you can only configure “Public Network” in this step. You can configure additional interfaces as part of post install steps.

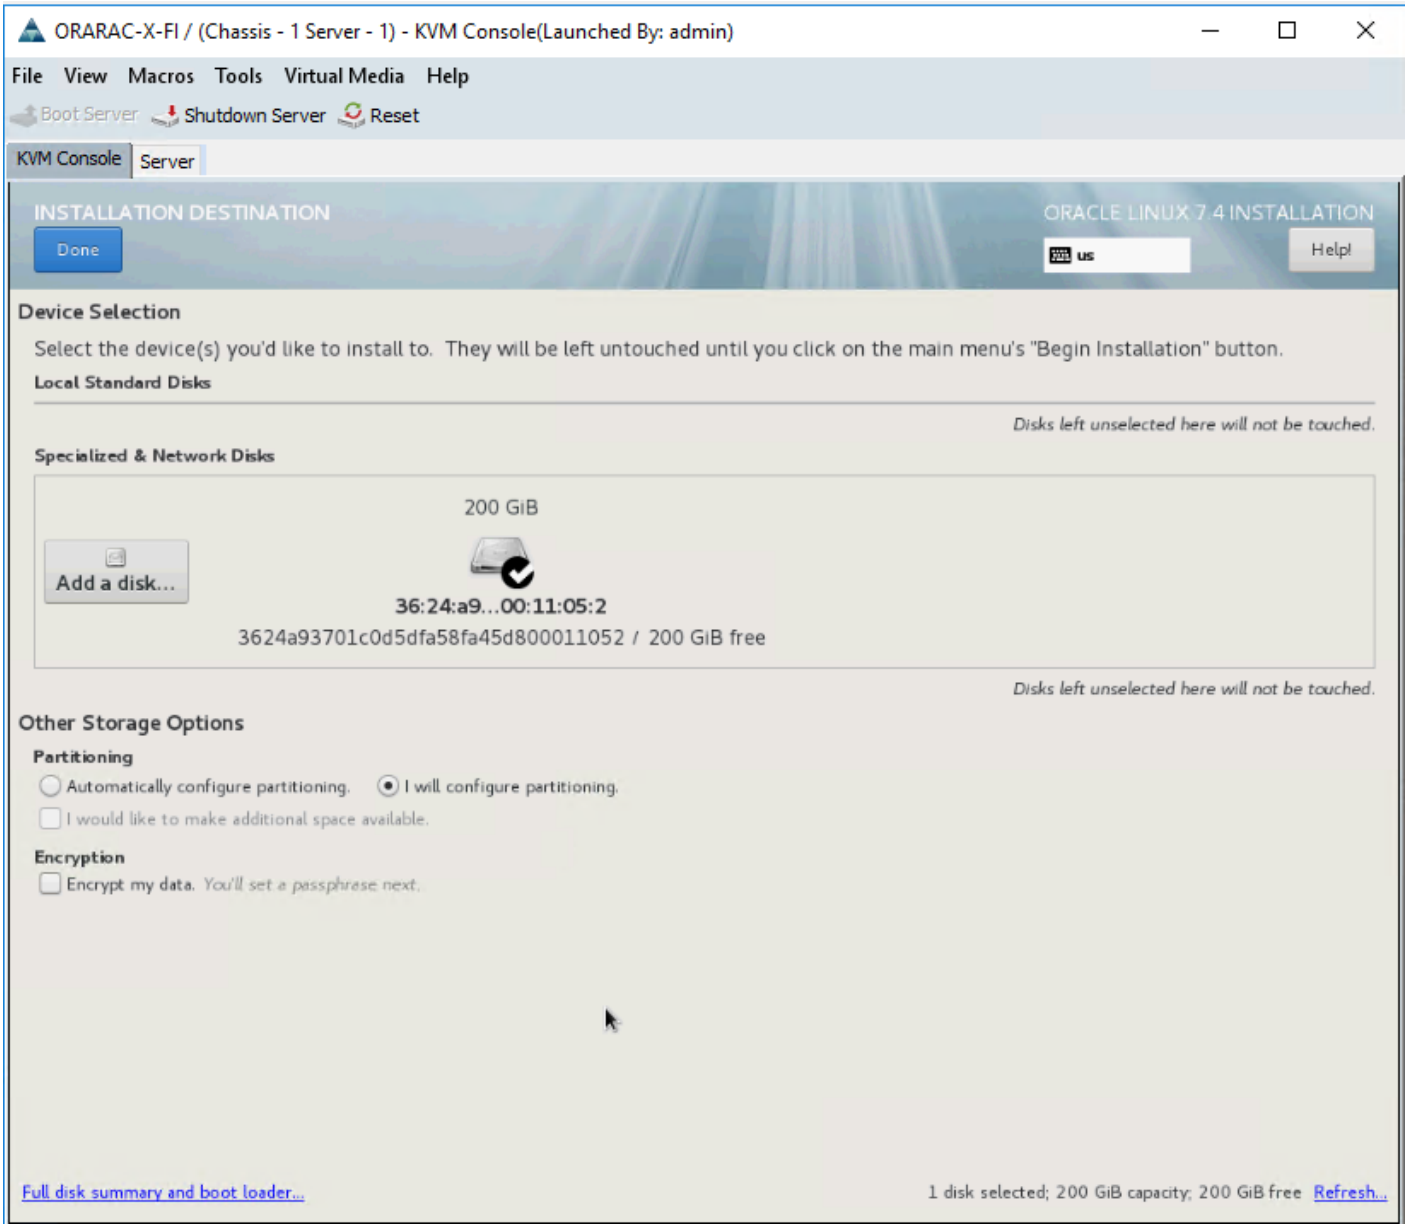

6. As a part of additional RPM package, we recommend to select “Customize Now” and configure “UEK kernel Repo.”

7. After the OS install, reboot the server, complete appropriate registration steps. You can choose to synchronize the time with ntp server. Alternatively, you can choose to use Oracle RAC cluster synchronization daemon (OCSSD). Both ntp and OCSSD are mutually exclusive and OCSSD will be setup during GRID install if ntp is not configured.

Operating System Prerequisites for Oracle Software Installation

Configure BIOS

This section describes how to optimize the BIOS settings to meet requirements for the best performance and energy efficiency for the Cisco UCS M5 generation of blade and rack servers.

Configure BIOS for OLTP Workloads

OLTP systems are often decentralized to avoid single points of failure. Spreading the work over multiple servers can also support greater transaction processing volume and reduce response time. Make sure to disable Intel IDLE driver in the OS configuration section. When Intel idle driver is disabled, the OS uses acpi_idle driver to control the c-states.

For latency sensitive workloads, it is recommended to always disable c-states in both OS and BIOS to ensure c-states are disabled.

The following are the recommended options for optimizing OLTP workloads on Cisco UCS M5 platforms managed by Cisco UCS Manager.

Figure 17 BIOS Options for OLTP Workloads

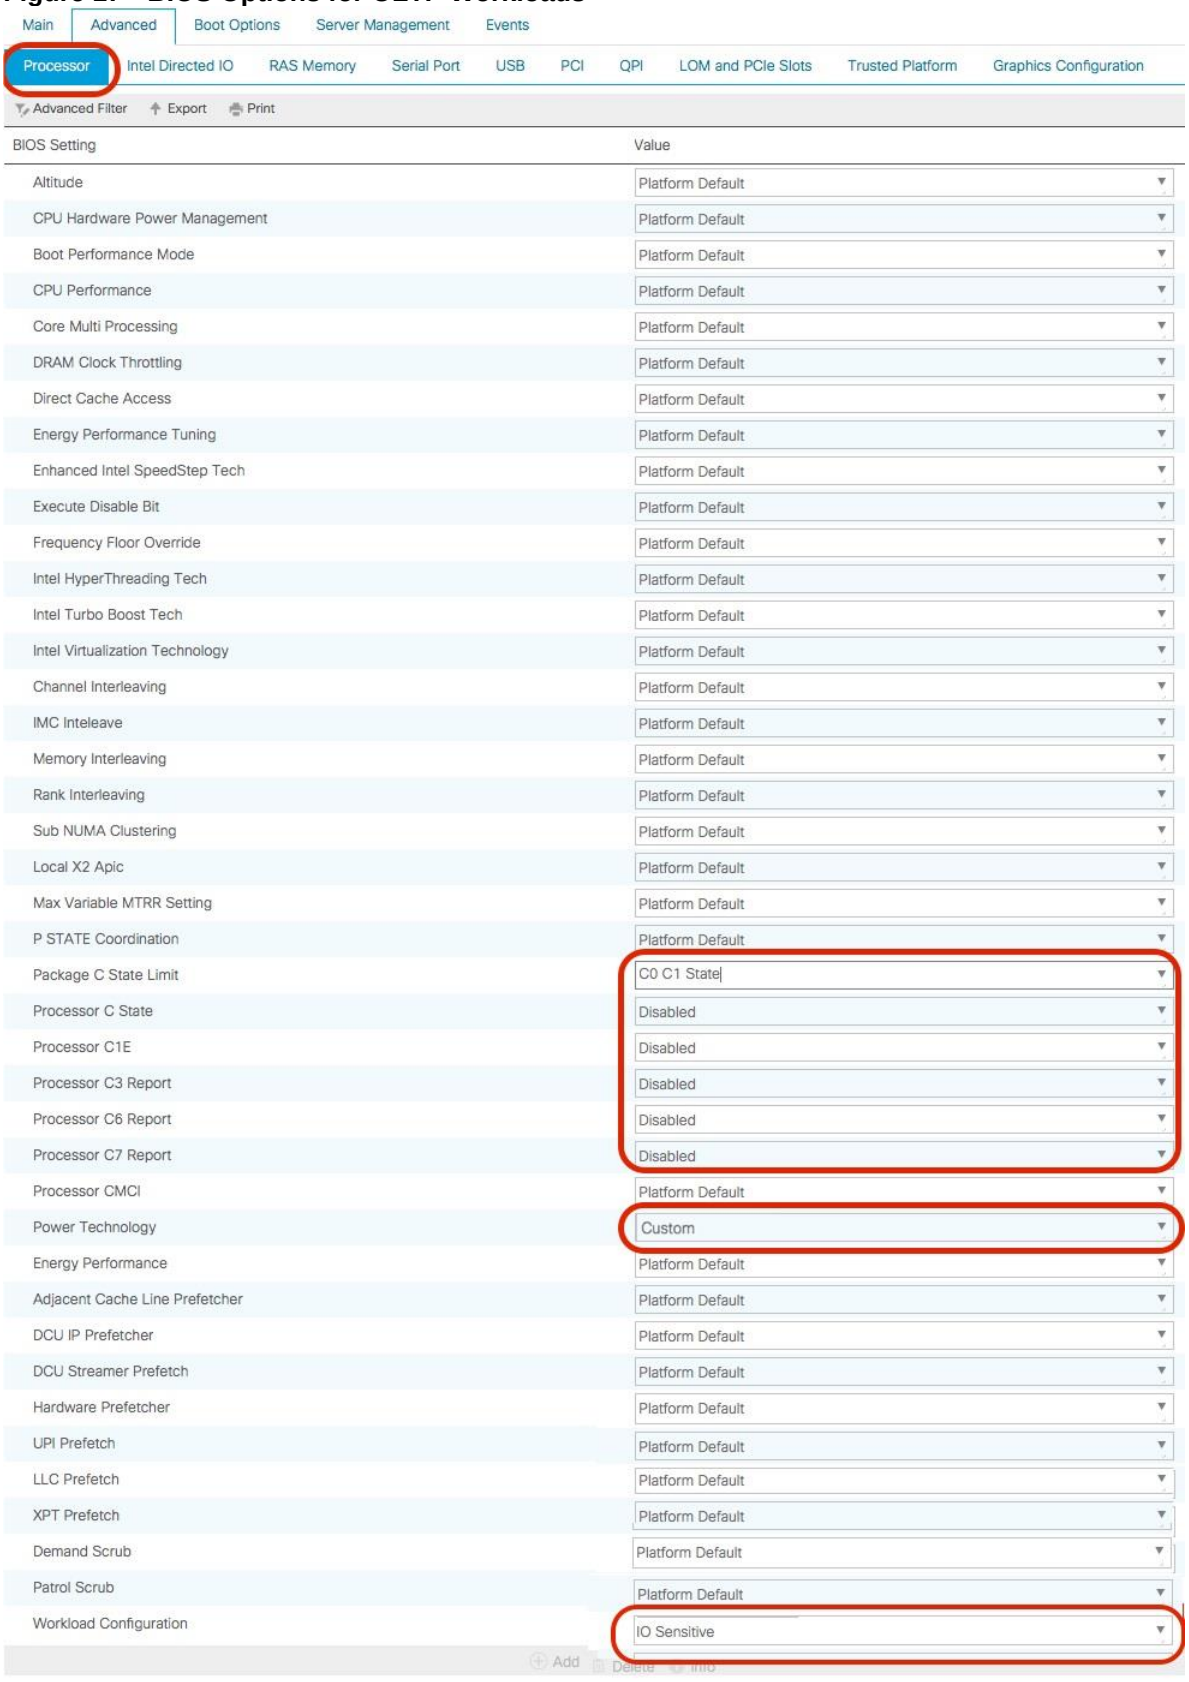

For more information about BIOS settings, refer to: https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/whitepaper_c11-740098.pdf

If the CPU gets into a deeper C-state and not able to get out to deliver full performance quickly. The result is unwanted latency spikes for workloads. To address this, it is recommended to disable C states in the BIOS and in addition, Oracle recommends disabling it from OS level as well by modifying grub entries. For this solution, we have configured BIOS options by modifying in /etc/default/grub file as shown below:

```
[root@oraracx1 ~]# cat /etc/default/grub

GRUB_TIMEOUT=5

GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"

GRUB_DEFAULT=saved

GRUB_DISABLE_SUBMENU=true

GRUB_TERMINAL_OUTPUT="console"

GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=ol/root rd.lvm.lv=ol/swap rhgb quiet numa=off transparent_hugepage=never intel_idle.max_cstate=0 processor.max_cstate=0"

GRUB_DISABLE_RECOVERY="true"
```

Prerequisites Automatic Installation

After installing Oracle Linux 7.4 on all the server nodes (oraracx1, oraracx2, oraracx3, oraracx4, oraracx5, oraracx6, oraracx7 and oraracx8), you have to configure operating system pre-requisites on all the eight nodes to successfully install Oracle RAC Database 12cR2.

To configure operating system pre-requisite for Oracle 12cR2 software on all eight nodes, complete the following steps:

Follow the steps according to your environment and requirements. Refer to the Install and Upgrade Guide for Linux for Oracle Database 12c R2: https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwlin/configuring-operating-systems-for-oracle-grid-infrastructure-on-linux.html#GUID-B8649E42-4918-49EA-A608-446F864EB7A0.

To configure the prerequisites on all the eight nodes, complete the following steps:

You can perform either the Automatic Setup or the Manual Setup to complete the basic prerequisites. The Additional Setup is required for all installations.

For this solution, we configured the prerequisites automatically by installing the “oracle-database-server-12cR2-preinstall" rpm package. You can also download the required packages from http://public-yum.oracle.com/oracle-linux-7.html. If you plan to use the “oracle-database-server-12cR2-preinstall" rpm package to perform all your prerequisite setup automatically, then login as root user and issue the following command.

```
[root@oraracx1 ~]# yum install oracle-database-server-12cR2-preinstall
```

If you have not used the "oracle-database-server-12cR2-preinstall" package, then you will have to manually perform the prerequisites tasks on all the eight nodes.

Additional Prerequisites Setup

After configuring automatic or manual prerequisites steps, you have to configure a few additional steps to complete the prerequisites for the Oracle Software installations on all the eight nodes as described below.

Disable SELinux

As most of the Organizations might already be running hardware-based firewalls to protect their corporate networks, we disabled Security Enhanced Linux (SELinux) and the firewalls at the server level for this reference architecture.

You can set secure linux to permissive by editing the "/etc/selinux/config" file, making sure the SELINUX flag is set as follows.

SELINUX= permissive

Disable Firewall

Check the status of the firewall by running following commands. (The status displays as active (running) or inactive (dead)). If the firewall is active / running, enter the following command to stop it:

systemctl status firewalld.service

systemctl stop firewalld.service

Also, to completely disable the firewalld service, so it does not reload when you restart the host machine, run the following command:

systemctl disable firewalld.service

Set the User Passwords

Run the following commands to change the password for Oracle and Grid Users:

passwd oracle

passwd grid

Configure Multipath Setup

 Follow the steps below on all the eight oracle RAC nodes.

Configure multipaths to access the LUNs presented from Pure Storage to the nodes. Device Mapper Multipath provides the ability to aggregate multiple IO paths to a newly created device mapper mapping to achieve high availability, I/O load balancing and persistent naming. We made sure the multipathing packages were installed and enabled for automatic restart across reboots.

Add or modify “/etc/multipath.conf" file accordingly to give the alias name of each LUN id presented from Pure Storage as given below into all eight nodes:

Run "multipath -ll" command to view all the LUN id.

[root@oraracx1 ~]# cat /etc/multipath.conf

```
defaults {
        polling_interval 10
}
```

```
devices {
        device {
                vendor                "PURE"
                path_grouping_policy    multibus
                path_checker            tur
                path_selector           "queue-length 0"
                fast_io_fail_tmo        10
                dev_loss_tmo            60
                no_path_retry           0
        }
}
multipaths {
        multipath {
                wwid          3624a93701c0d5dfa58fa45d800011066
                alias         orarax1_os
        }
}
```

Make sure the LUNs wwid address reflects the correct value for all eight nodes in /etc/multipath.conf

Configure /etc/multipath.conf (Appendix) as per Pure Storage's recommended multipath config for Oracle Linux as documented in the Pure Support Page:

https://support.purestorage.com/Solutions/Operating_Systems/Linux/Linux_Recommended_Settings

Configure UDEV rules

Follow the steps below on all the eight oracle RAC nodes.

As per Pure Storage FlashArray's best practice, setup the queue settings with udev rules. Refer to the updated Linux best practices for Pure Storage FlashArray on Pure's support site.

Create file named /etc/udev/rules.d/99-pure-storage.rules with the following entries:

```
# Recommended settings for PURE Storage FlashArray

# Use noop scheduler for high-performance solid-state storage

ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE", ATTR{queue/scheduler}="noop"


# Reduce CPU overhead due to entropy collection
```

ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE", ATTR{queue/add_random}="0"

Schedule I/O on the core that initiated the process

ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE", ATTR{queue/rq_affinity}="2"

These steps complete the prerequisite for Oracle Database 12cR2 Installation at OS level on Oracle RAC Node 1 (ORARACX1). For this FlashStack Solution, we used 8 identical Cisco UCS B-series B200 M5 blade servers for hosting the 8-node Oracle RAC database. All of the steps above were also performed on all eight nodes to create 8 node Oracle RAC solution.

This process will complete all 8 Oracle RAC Nodes with OS and all prerequisites to install Oracle Database Software 12cR2.

Configure Host Groups

After completion of all Oracle Nodes OS Boot LUN, create and configure Host Group. You will assign all the Oracle Nodes to the Host Group into Pure FlashArray.

To configure the Host Group, complete the following steps:

1. Log into Pure Storage dashboard.

2. In the PURE GUI, go to tab Storage > Hosts > click the + sign to “Create Host Group” as shown below:

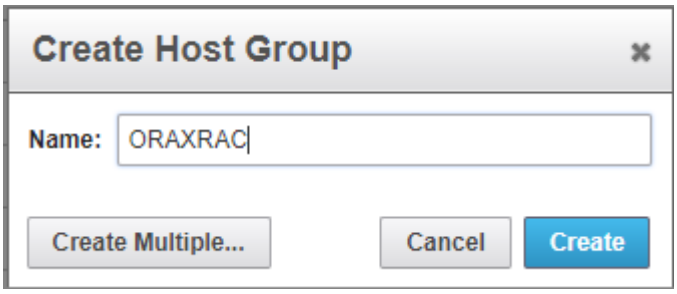

3. Select from Host Group “ORARACX” > Hosts > click Settings and “Add Hosts.”

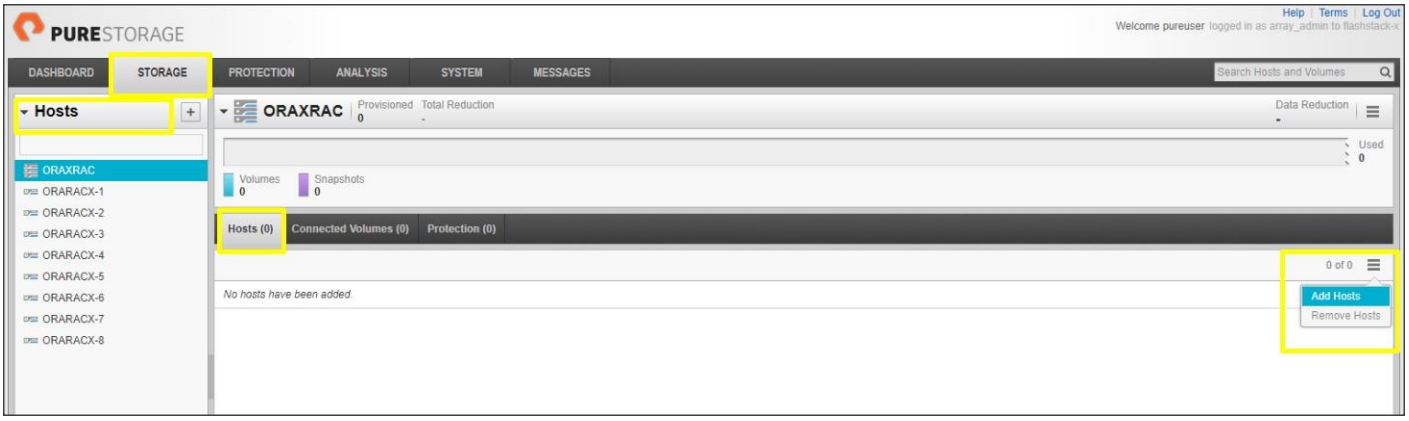

4. Add all nodes from “Existing Hosts” to “Selected Hosts” and click Confirm.

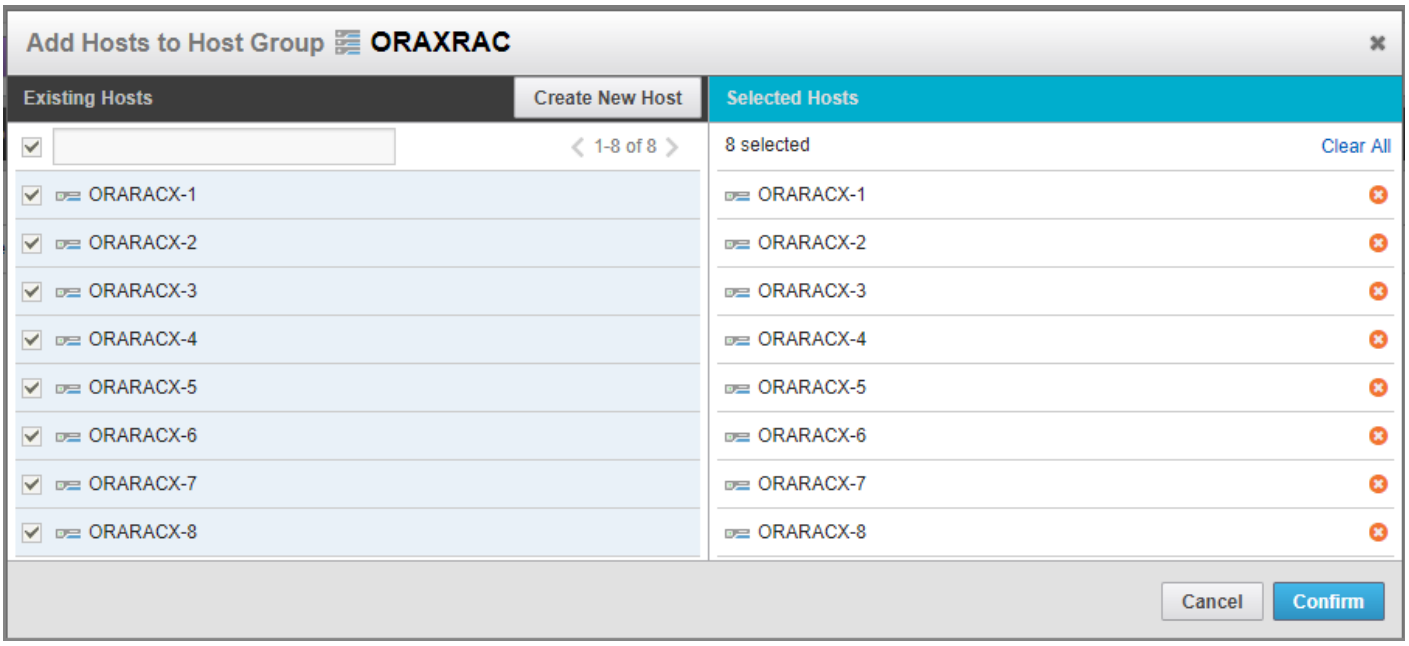

5. Verify that all nodes are added into Host Group as shown below:

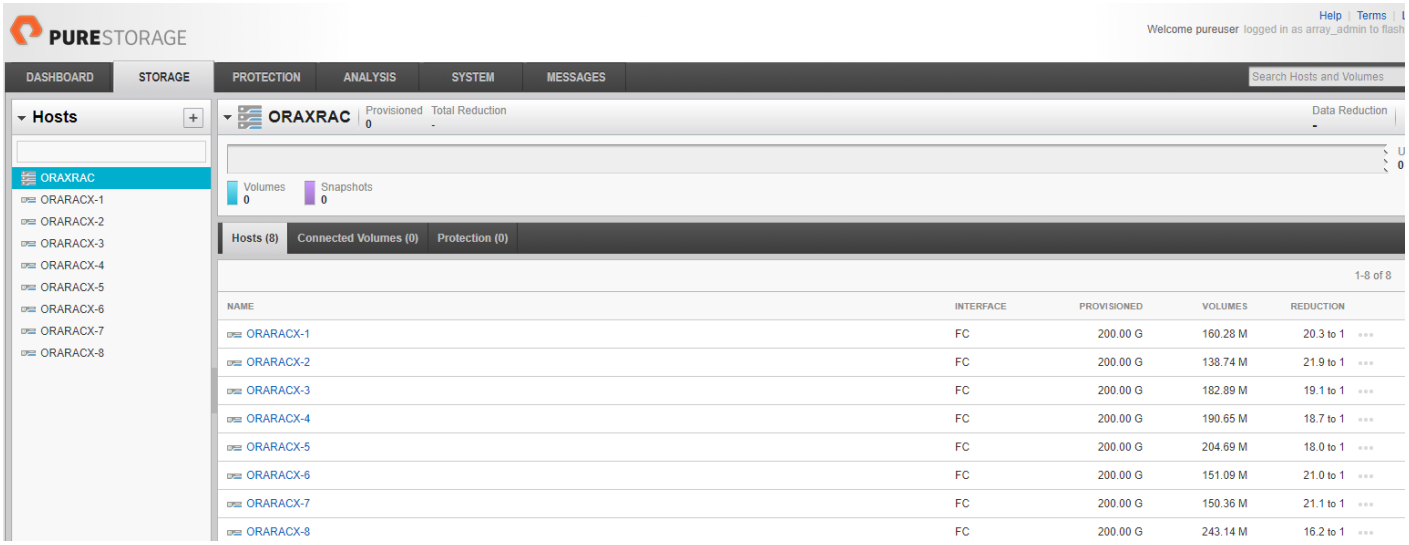

Configure CRS, Data and Redo log Volumes

You will create and assign CRS, Data and Redo Log Volumes to Host Group, which was created earlier. By doing this, all the nodes into Host Group can able to read/write data from/to the Volume.

For this FlashStack solution, we created two OLTP Database (OLTPDB1 and SOEDB1) and one DSS Database (DSSDB1). The following table shows the LUNs are created and the description of the LUNs.

Table 10 LUN Description

Database	Volume Name	Size	Purpose
	dg_orarac_crs	200 GB	Store OCR and Voting Disk information for All Database
OLTPDB1	dg_oradata_oltp1	5 TB	Store Datafiles for OLTPDB1 Database
	dg_redolog_oltp1	500 GB	Store Redolog Files for OLTPDB1 Database
SOEDB1	dg_oradata_soe1	10 TB	Store Datafiles for SOEDB1 Database

	dg_redolog_soe1	500 GB	Store Redolog Files for SOEDB1 Database
DSSDB1	dg_oradata_dss1	13 TB	Store Datafiles for DSSDB1 Database
	dg_redolog_dss1	500 GB	Store Redolog Files for DSSDB1 Database
SLOBDB	dg_oradata_slob01	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob02	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob03	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob04	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob05	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob06	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob07	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob08	600 GB	Store Datafiles for SLOBDB Database
	dg_oradata_slob09	600 GB	Store Datafiles for SLOBDB Database
	dg_oraredo_slob	100 GB	Store Redolog Files for SLOBDB Database

1. Create CRS Volume of 200 GB for storing OCR and Voting Disk files for all the database.

2. Create Data Volumes for each database to store database files.

3. Create Redo Volumes for each database to store redo log files.

4. After you create all appropriate volumes, assign these volumes to Host Group. Attach all the volumes to a host group by going to the “Connected Hosts and Host Groups” tab under the volume context menu, click the Settings icon and select “Connect Host Group.” Select the host group where the volume should be attached and click Confirm.

5. Verify that all the volumes are visible into Host Group as shown below:

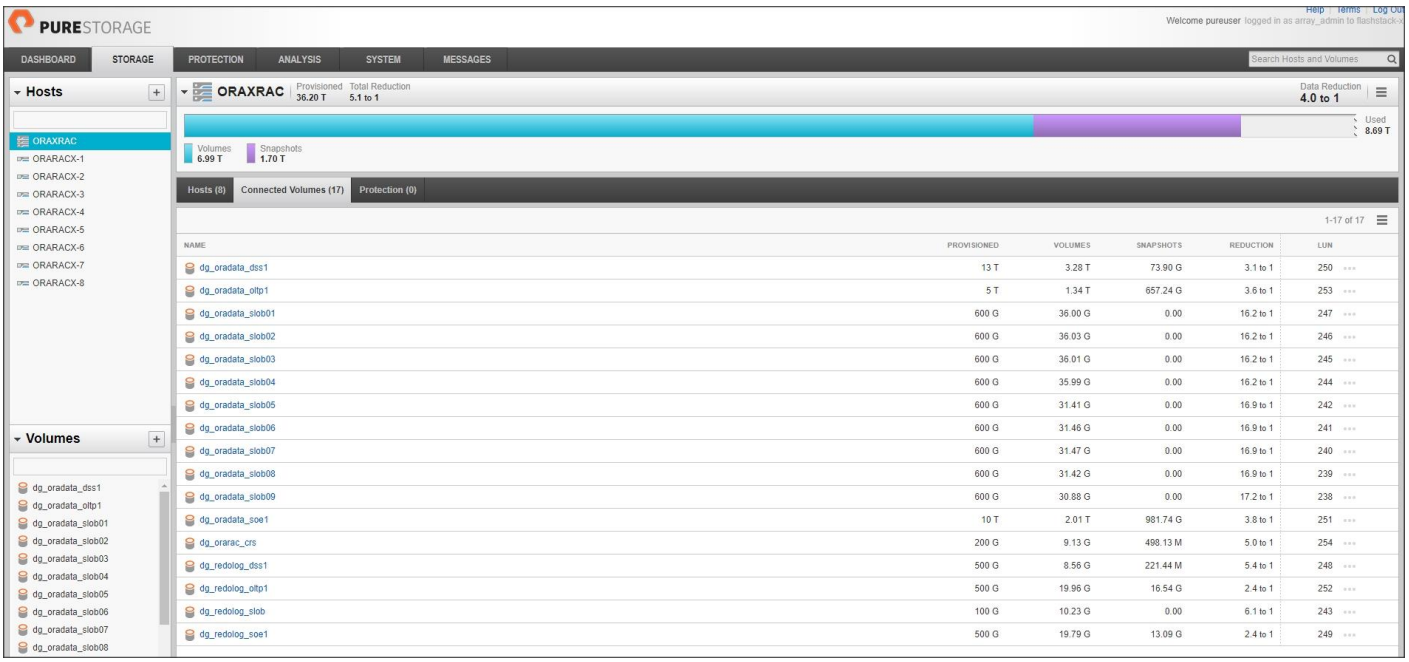

Configure UDEV rules

After creating all volumes into Pure Storage, you need to configure UDEV to assign permission in all Oracle RAC Nodes. This includes the device details along with required permissions to enable grid user to have read/write privileges on these devices. Configure UDEV rules on all Oracle Nodes as shown below:

IMPORTANT: The /etc/multipath.conf for the Oracle ASM devices and udev rules for these devices should be copied on to all the RAC nodes and verified to make sure the devices are visible and permissions are enabled for grid user on all the nodes

Create a new file named /etc/udev/rules.d/99-oracleasm.rules with the following entries:

#All volumes which starts with dg_orarac_* #

ENV{DM_NAME}=="dg_orarac_crs", OWNER:="grid", GROUP:="oinstall", MODE:="660"

#All volumes which starts with dg_oradata_* #

ENV{DM_NAME}=="dg_oradata_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"

#All volumes which starts with dg_oraredo_* #

ENV{DM_NAME}=="dg_oraredo_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"

Configure Public and Private NICs on Each RAC Node

If you have not configured network settings during OS installation, then configure it now. Each node must have at least two network interface cards (NIC), or network adapters. One adapter is for the public network interface and the other adapter is for the private network interface (the interconnect).

Login as a root user into each node and go to “/etc/sysconfig/network-scripts” and configure Public network and Private network IP Address. Configure the private and public NICs with the appropriate IP addresses across all the Oracle RAC nodes.

Configure “/etc/hosts” on Each RAC Node

Login as a root user into node and edit “/etc/hosts” file. Provide details for Public IP Address, Private IP Address, SCAN IP Address and Virtual IP Address for all nodes. Configure these settings in each Oracle RAC Nodes as shown below:

cat /etc/hosts

127.0.0.1 localhost localhost.localdomain

#Public IP

10.29.134.121 oraracx1 oraracx1.cisco.com

10.29.134.122 oraracx2 oraracx2.cisco.com

10.29.134.123 oraracx3 oraracx3.cisco.com

10.29.134.124 oraracx4 oraracx4.cisco.com

10.29.134.125 oraracx5 oraracx5.cisco.com

10.29.134.126 oraracx6 oraracx6.cisco.com

10.29.134.127 oraracx7 oraracx7.cisco.com

10.29.134.128 oraracx8 oraracx8.cisco.com

#Virtual IP

10.29.134.129 oraracx1-vip oraracx1-vip.cisco.com
10.29.134.130 oraracx2-vip oraracx2-vip.cisco.com
10.29.134.131 oraracx3-vip oraracx3-vip.cisco.com
10.29.134.132 oraracx4-vip oraracx4-vip.cisco.com
10.29.134.133 oraracx5-vip oraracx5-vip.cisco.com
10.29.134.134 oraracx6-vip oraracx6-vip.cisco.com
10.29.134.135 oraracx7-vip oraracx7-vip.cisco.com
10.29.134.136 oraracx8-vip oraracx8-vip.cisco.com

#Private IP

192.168.10.121 oraracx1-priv oraracx1-priv.cisco.com
192.168.10.122 oraracx2-priv oraracx2-priv.cisco.com
192.168.10.123 oraracx3-priv oraracx3-priv.cisco.com
192.168.10.124 oraracx4-priv oraracx4-priv.cisco.com
192.168.10.125 oraracx5-priv oraracx5-priv.cisco.com
192.168.10.126 oraracx6-priv oraracx6-priv.cisco.com
192.168.10.127 oraracx7-priv oraracx7-priv.cisco.com
192.168.10.128 oraracx8-priv oraracx8-priv.cisco.com

#SCAN IP

10.29.134.137 oraracx-cluster-scan oraracx-cluster-scan.cisco.com
10.29.134.138 oraracx-cluster-scan oraracx-cluster-scan.cisco.com
10.29.134.139 oraracx-cluster-scan oraracx-cluster-scan.cisco.com

#Oracle Client

10.29.134.196 linuxclient1 linuxclient1.cisco.com
10.29.134.197 linuxclient2 linuxclient2.cisco.com

You must configure the following addresses manually in your corporate setup.

- A Public IP Address for each node
- A Virtual IP address for each node
- Three single client access name (SCAN) address for the oracle cluster

When all the LUNs are created and O.S level prerequisites are completed, you are ready to install Oracle Grid Infrastructure as grid user. Download Oracle Database 12c Release 2 (12.2.0.1.0) for Linux x86-64 and Oracle Database 12c Release 2 Grid Infrastructure (12.2.0.1.0) for Linux x86-64 software from Oracle Software site. Copy these software binaries to Oracle RAC Node 1 and Unzip all files into appropriate directories.

Oracle Database 12c GRID Infrastructure Deployment

For this FlashStack Solution, you will install Oracle Grid and Database software on all eight nodes (oraracx1, oraracx2, oraracx3, oraracx4, oraracx5, oraracx6, oraracx7 and oraracx8). The installation guides you through gathering all node information and configuring ASM devices and all the prerequisite validations for GI.

It is not within the scope of this document to include the specifics of an Oracle RAC installation; you should refer to the Oracle installation documentation for specific installation instructions for your environment. We will provide a partial summary of details that might be relevant.

This section describes the high-level steps for Oracle Database 12c R2 RAC install. Prior to GRID and database install, verify all the prerequisites are completed. As an alternate, you can install Oracle validated RPM that will make sure all prerequisites are meet before Oracle grid install.

Use the following link for Oracle Database 12c Release 2 Install and Upgrade guide:
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/install-and-upgrade.html

Run Cluster Verification Utility

This step will verify that all prerequisites are meet to install Oracle Grid Infrastructure Software. Oracle Grid Infrastructure ships with the Cluster Verification Utility (CVU) that can run to validate pre and post installation configurations. To run this utility, login as Grid User in Oracle RAC Node 1 and go to the directory where oracle grid software binaries are located. Run script named as “runcluvfy.sh” as follows:

./runcluvfy.sh stage -pre crsinst -n oraracx1,oraracx2,oraracx3,oraracx4,oraracx5,oraracx6,oraracx7,oraracx8 –verbose

Configure HugePages

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

Advantages of HugePages:

- HugePages are not swappable so there is no page-in/page-out mechanism overhead.
- HugePages uses fewer pages to cover the physical address space, so the size of "book keeping"(mapping from the virtual to the physical address) decreases, so it requiring fewer entries in the TLB and so TLB hit ratio improves.
- HugePages reduces page table overhead. Also, HugePages eliminated page table lookup overhead: Since the pages are not subject to replacement, page table lookups are not required.

- Faster overall memory performance: On virtual memory systems each memory operation is actually two abstract memory operations. Since there are fewer pages to work on, the possible bottleneck on page table access is clearly avoided.

For our configuration, we used HugePages for all the OLTP and DSS workloads.

Please refer to the Oracle support for HugePages configuration details:
https://docs.oracle.com/en/database/oracle/oracle-database/12.2/unxar/administering-oracle-database-on-linux.html#GUID-CC72CEDC-58AA-4065-AC7D-FD4735E14416

Create Directory Structure

The directory structure should be create on all the RAC nodes but unzipping grid software happens on the first node only.

As the grid user, download the Oracle Grid Infrastructure image files and extract the files into the Grid home.

You must extract the zip image software into the directory where you want your Grid home to be located. Also, Download and copy the Oracle Grid Infrastructure image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.

mkdir -p /u01/app/12.2.0/grid

chown grid:oinstall /u01/app/12.2.0/grid

cd /u01/app/12.2.0/grid

unzip -q download_location/linuxx64_12201_grid_home

mkdir –p /u01/app/oracle/product/12.2.0/dbhome_1

chown –R oracle:oinstall /u01/app/oracle

Configure Oracle ASM Driver and Shared Disks

This step has to be done only on the first node.

1. Log in as the root user and set the environment variable ORACLE_HOME to the location of the Grid home:

   export ORACLE_HOME=/u01/app/12.2.0/grid

2. Use Oracle ASM command line tool (ASMCMD) to provision the disk devices for use with Oracle ASM Filter Driver:

   /u01/app/12.2.0/grid/bin/asmcmd afd_label OCRVOTE /dev/mapper/dg_orarac_crs –init

3. Verify the device has been marked for use with Oracle ASMFD:

   /u01/app/12.2.0/grid/bin/asmcmd afd_lslbl /dev/mapper/dg_orarac_crs

---

```
Label                Duplicate    Path
================================================================
OCRVOTE                           /dev/mapper/dg_orarac_crs
```

After configuring the Oracle ASM disk group, you will install Oracle Grid Infrastructure and Oracle Database 12c R2 standalone software. For this solution, we installed Oracle binaries on the boot LUN of the nodes. The OCR, Data, and redo log files reside in the Oracle ASM disk group created from CRS, Data and Redolog volume.

Log in as the grid user, and start the Oracle Grid Infrastructure installer as detailed in the next step.

Install and Configure Oracle Database Grid Infrastructure Software

It is not within the scope of this document to include the specifics of an Oracle RAC installation. However, we will provide partial summary of details that might be relevant. Please refer to the Oracle installation documentation for specific installation instructions for your environment.

To install Oracle Database Grid Infrastructure Software, complete the following steps:

1. Go to grid home where the Oracle 12c R2 Grid Infrastructure software binaries are located and launch the installer as the "grid" user.

2. Start the Oracle Grid Infrastructure installer by running the following command:

   ./gridSetup.sh

3. Select option “Configure Oracle Grid Infrastructure for a New Cluster” as shown below, then click Next:

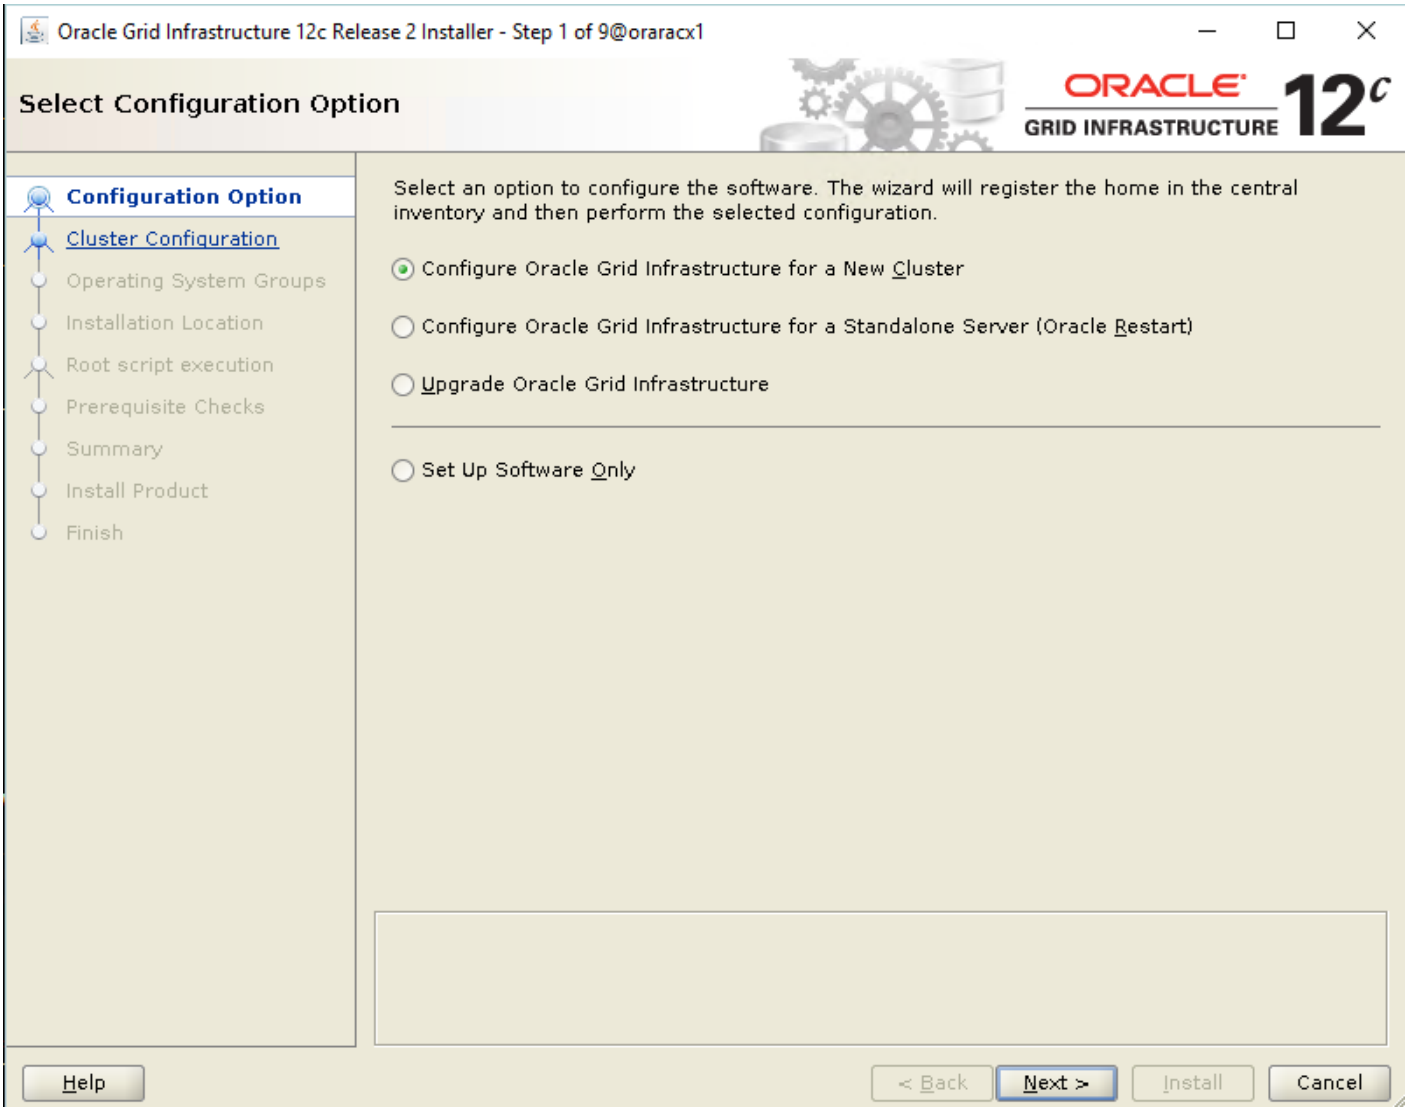

4. Select cluster configuration options “Configure an Oracle Standalone Cluster”, then click Next.

5. In next window, enter the Cluster Name and SCAN Name fields.

Enter the names for your cluster and cluster scan that are unique throughout your entire enterprise network. You can select Configure GNS if you have configured your domain name server (DNS) to send to the GNS virtual IP address name resolution requests

6. In next Cluster node information window, click the "Add" button to add all eight nodes Public Hostname and Virtual Hostname as shown below:

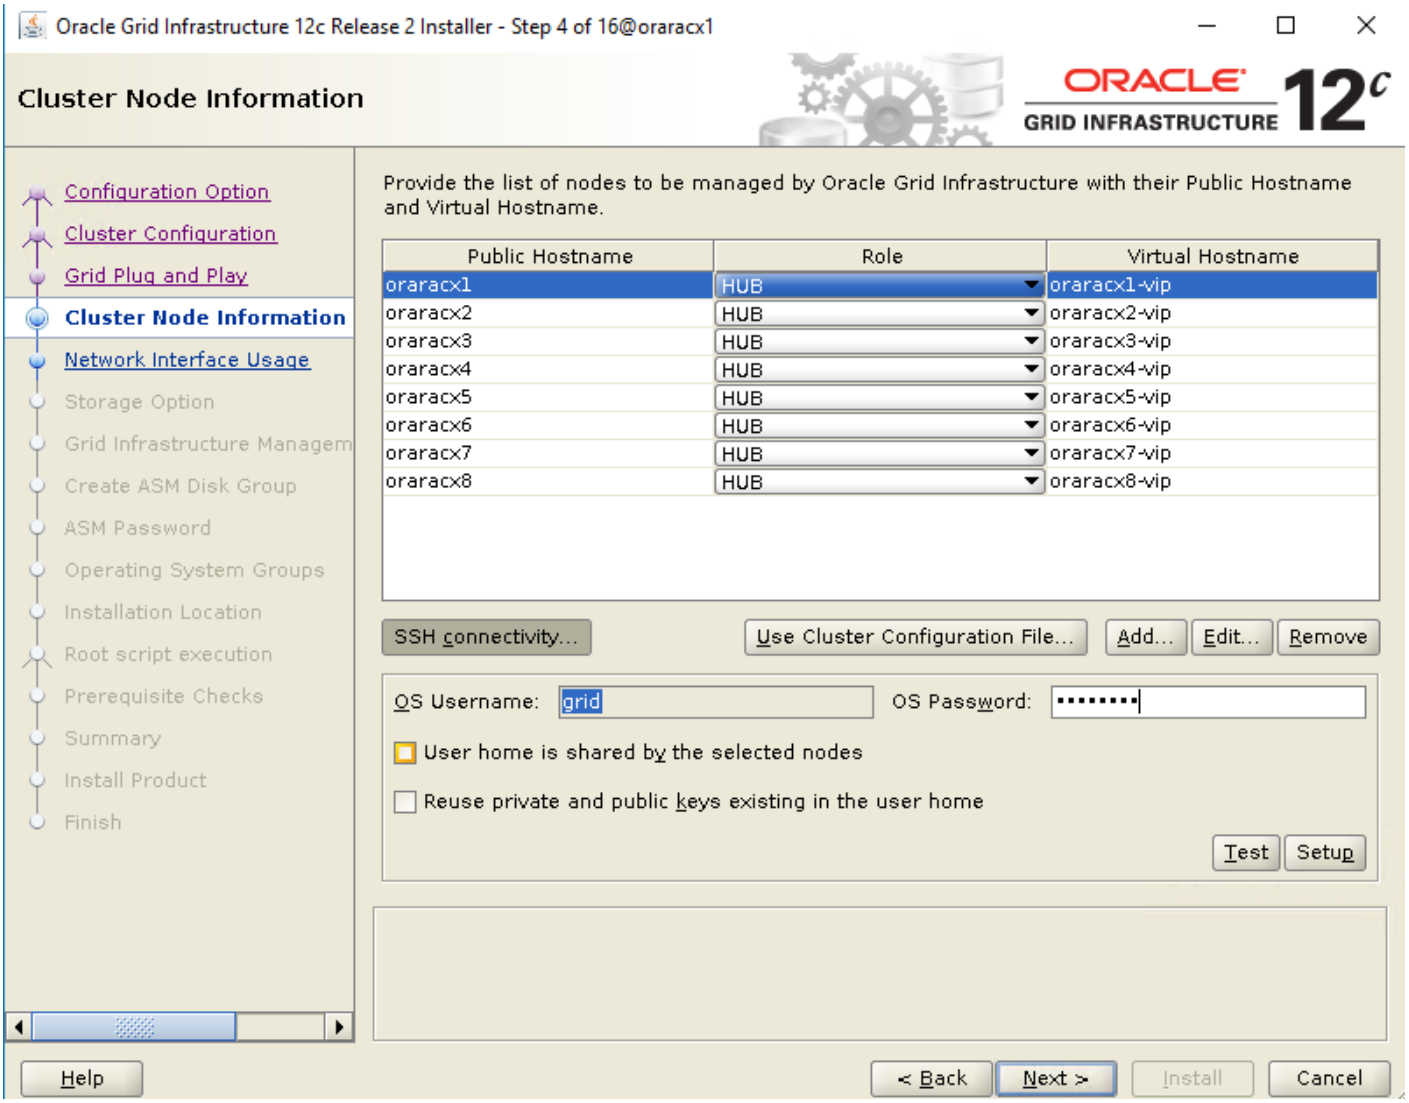

7. As shown above, you will see all nodes listed in the table of cluster nodes. Make sure the Role column is set to HUB for all eight nodes. Click the SSH Connectivity button at the bottom of the window. Enter the operating system user name and password for the Oracle software owner (grid). Click Setup.

8. A message window appears, indicating that it might take several minutes to configure SSH connectivity between the nodes. After sometime, another message window appears indicating that password-less SSH connectivity has been established between the cluster nodes. Click OK to continue.

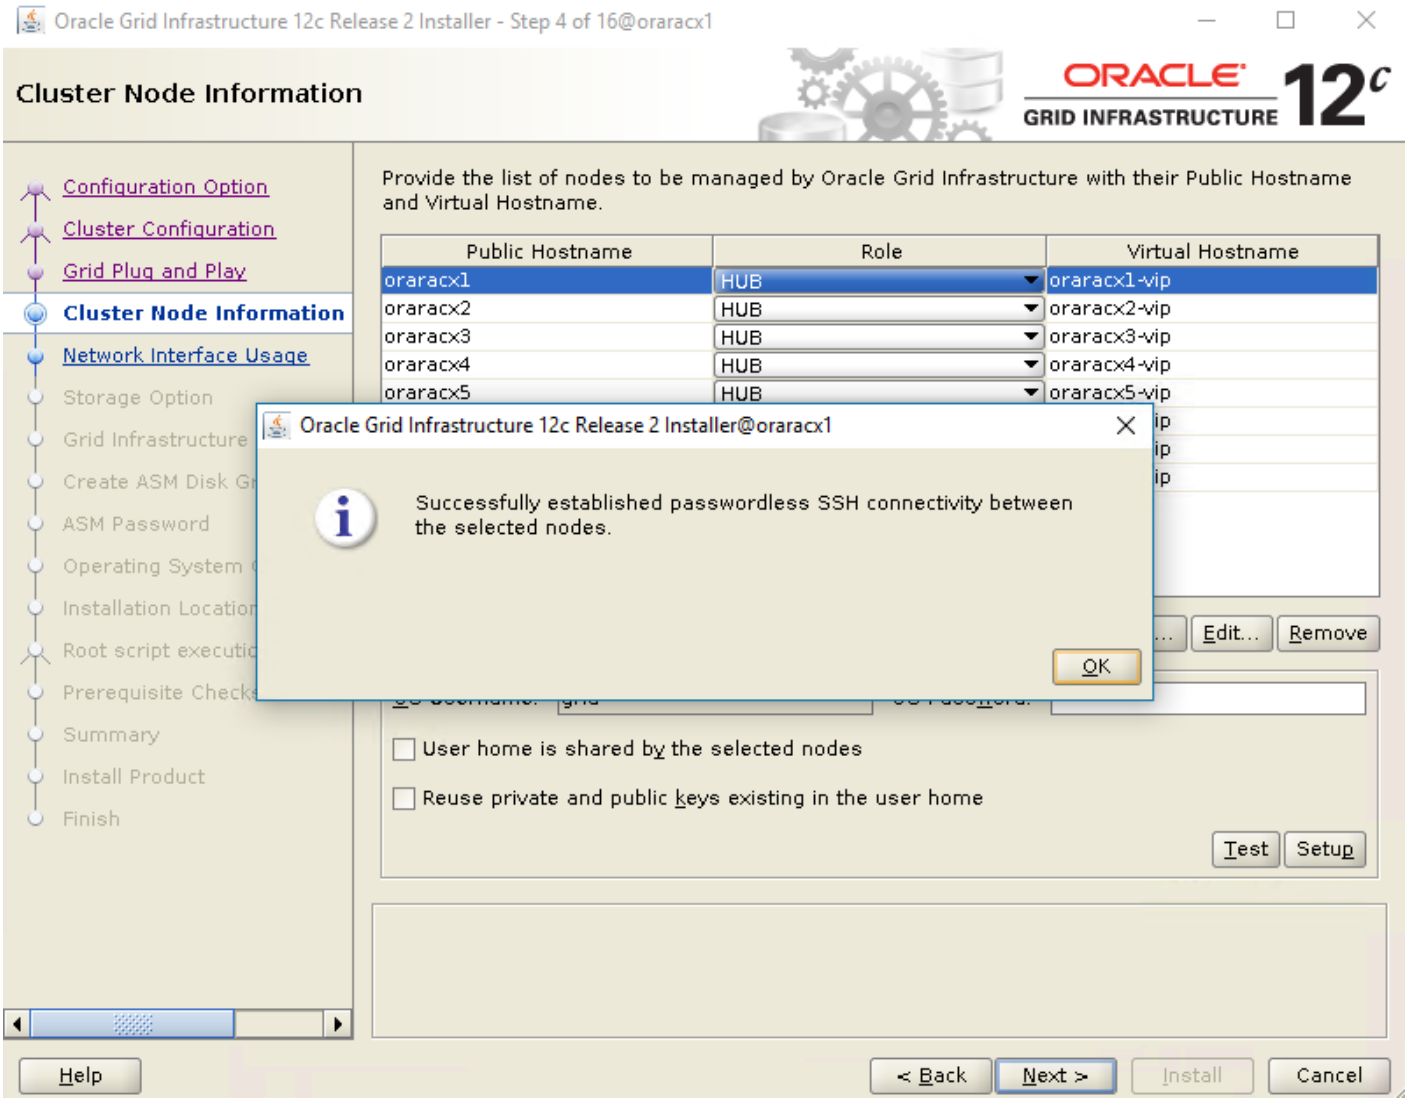

9. In Network Interface Usage screen, select the usage type for each network interface displayed as shown below:

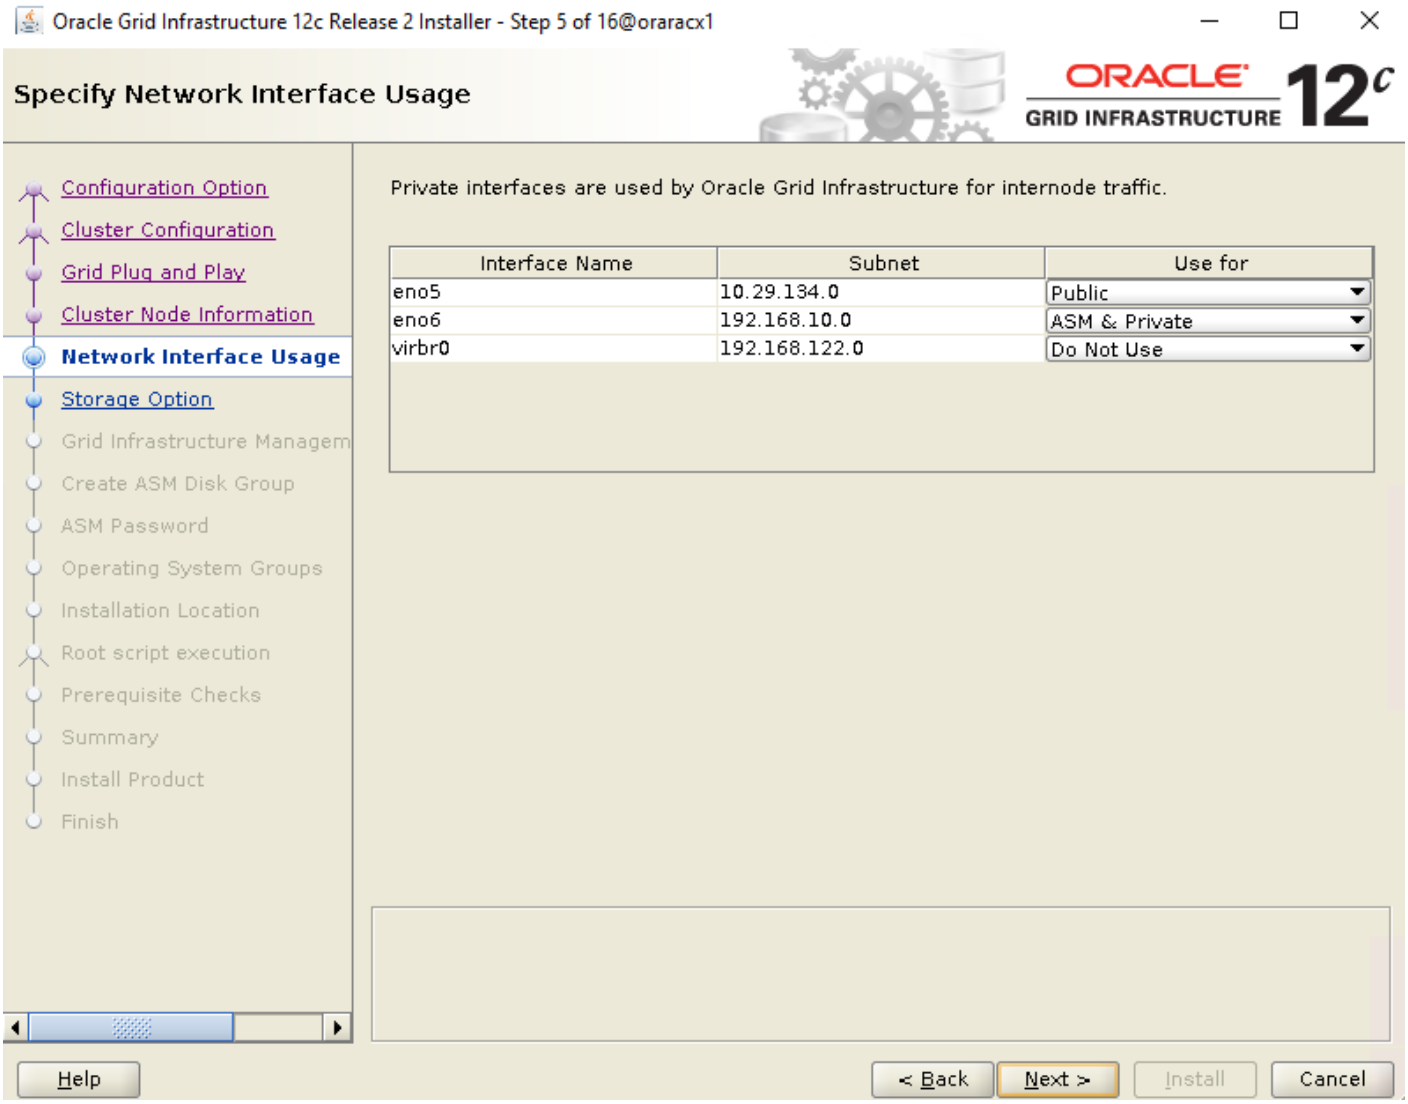

10. Select the Oracle ASM storage configuration option as “Configure ASM using block devices.” Choose whether you want to store the Grid Infrastructure Management Repository in a separate Oracle ASM disk group as “No”, and then click Next.

11. In the Create ASM Disk Group window, enter the name of disk group and select appropriate redundancy options as show below. We selected Oracle ASM Filter Driver (Oracle ASMFD) to manage Oracle ASM disk devices, so select the option Configure Oracle ASM Filter Driver. Select the OCRVOTE LUN assigned from Pure Storage to store OCR and Voting disk files.

During installation, disks labelled as Oracle ASMFD disks or Oracle ASMLIB disks are listed as candidate disks when using the default discovery string. However, if the disk has a header status of MEMBER, then it is not a candidate disk.

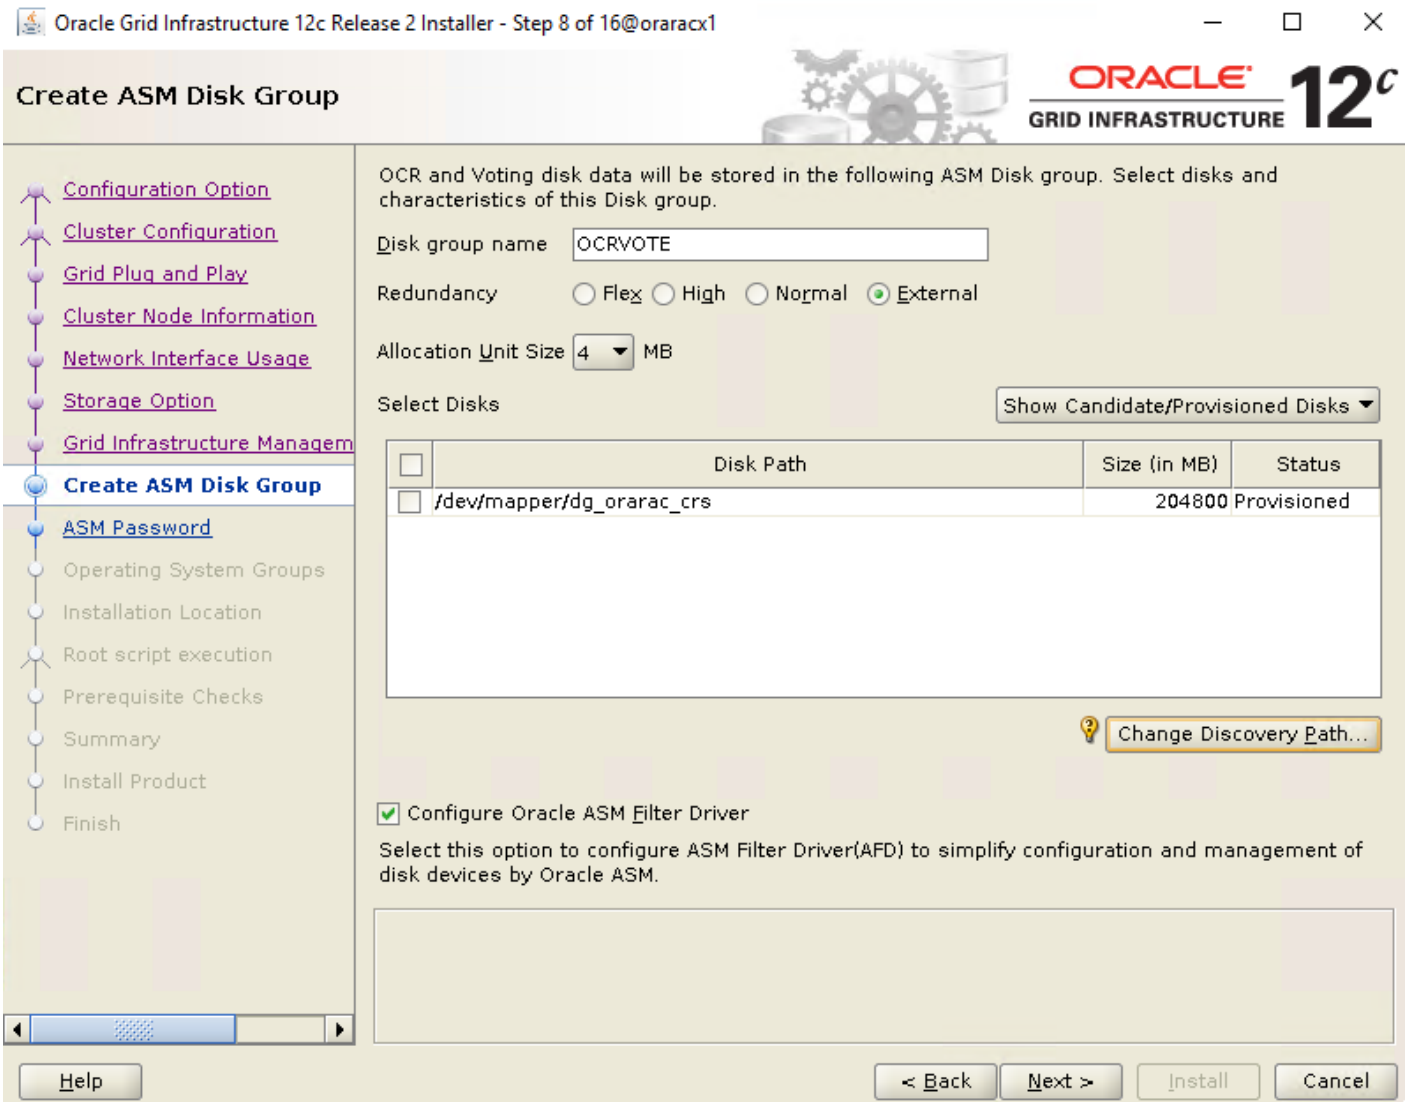

12. Choose the same password for the Oracle ASM SYS and ASMSNMP account, or specify different passwords for each account, then click Next.

13. Select the option “Do not use Intelligent Platform Management Interface (IPMI)”, then click Next.

You can choose to set it up according to your requirements.

14. Select the appropriate operating system group names for Oracle ASM according to your environments.

15. Specify the directory to use for the Oracle base for the Oracle Grid Infrastructure installation and then click Next. The Oracle base directory must be different from the Oracle home directory.

If you copied the Oracle Grid Infrastructure installation files into the Oracle Grid home directory as directed above, then the default location for the Oracle base directory should display as /u01/app/grid.

16. Click Automatically run configuration scripts to run scripts automatically and enter the relevant root user credentials. Click Next.

17. Wait while the prerequisite checks complete. If you have any issues, use the "Fix & Check Again" button.

If any of the checks have a status of Failed and are not fixable, then you must manually correct these issues. After you have fixed the issue, you can click the Check Again button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click Next.

18. Review the contents of the Summary window and then click Install. The installer displays a progress indicator enabling you to monitor the installation process.

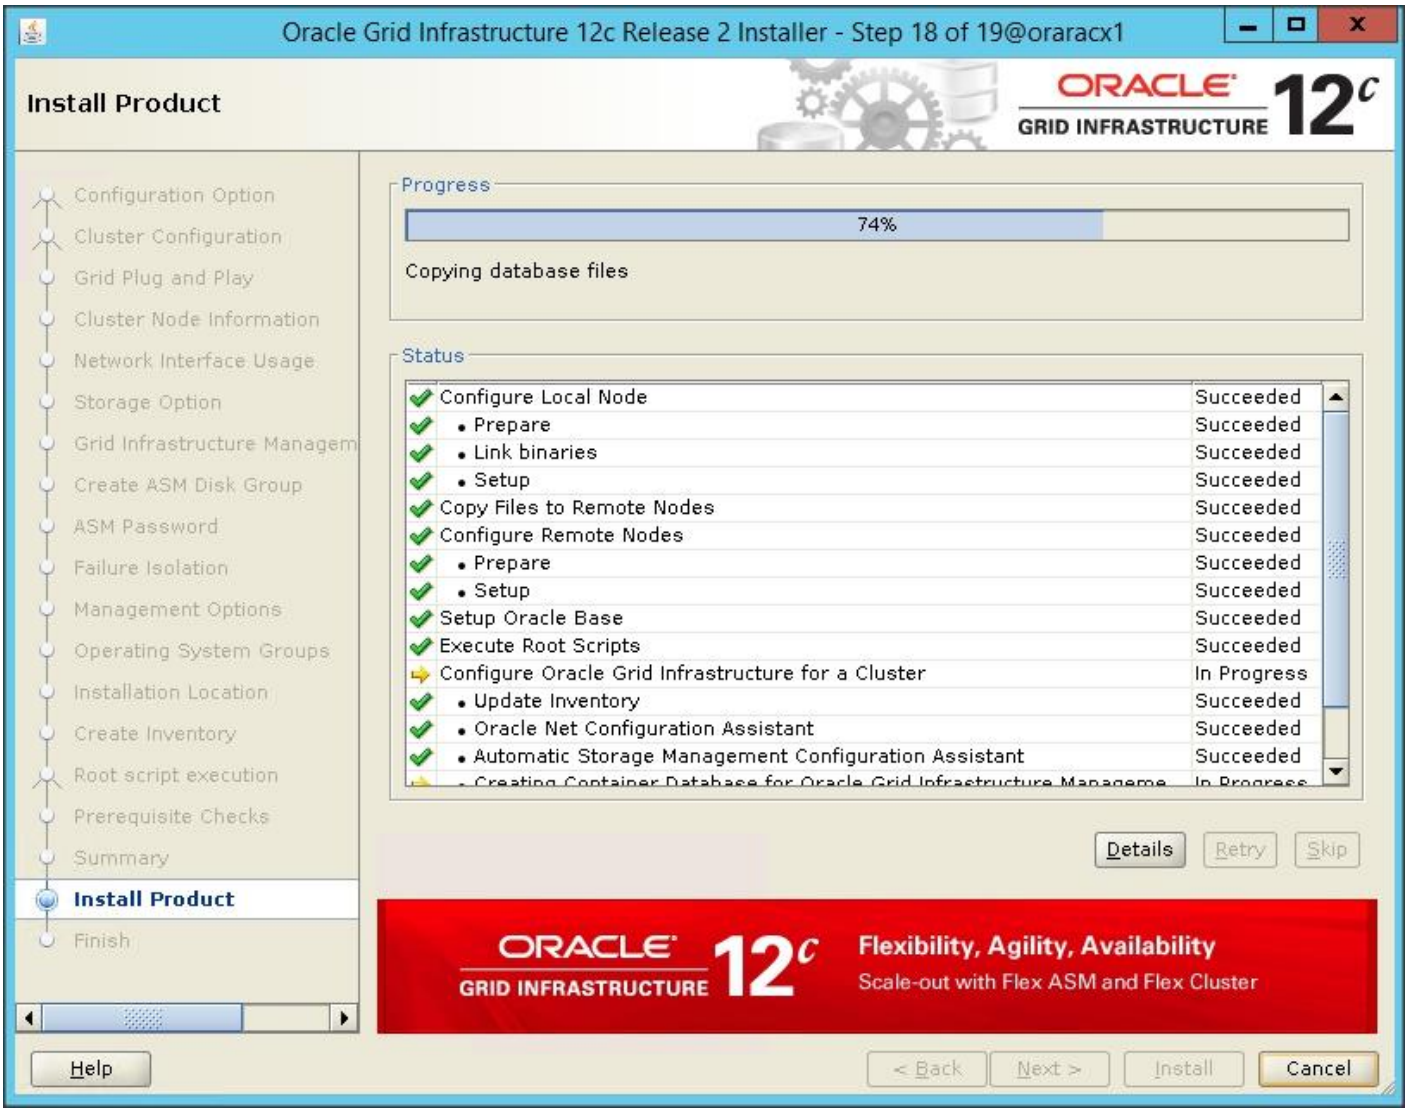

19. Wait for the grid installer configuration assistants to complete.

20. When the configuration complete successfully, click "Close" button to finish and exit the grid installer.

21. When GRID install is successful, login to each of the nodes and perform minimum health checks to make sure that Cluster state is healthy. After your Oracle Grid Infrastructure installation is complete, you can install Oracle Database on a cluster node for high availability, or install Oracle RAC.

Install Oracle Database Software

After successful GRID install, we recommend to install Oracle Database 12c software only. You can create databases using DBCA or database creation scripts at later stage.

It is not within the scope of this document to include the specifics of an Oracle RAC database installation. However, we will provide partial summary of details that might be relevant. Please refer to the Oracle database installation documentation for specific installation instructions for your environment.

To install Oracle Database Software, complete the following steps:

1. Start the runInstaller command from the Oracle Database 12*c* Release 2 (12.2) installation media where Oracle database software is located.

2. Select option “Install database software only” into Select Installation Option.

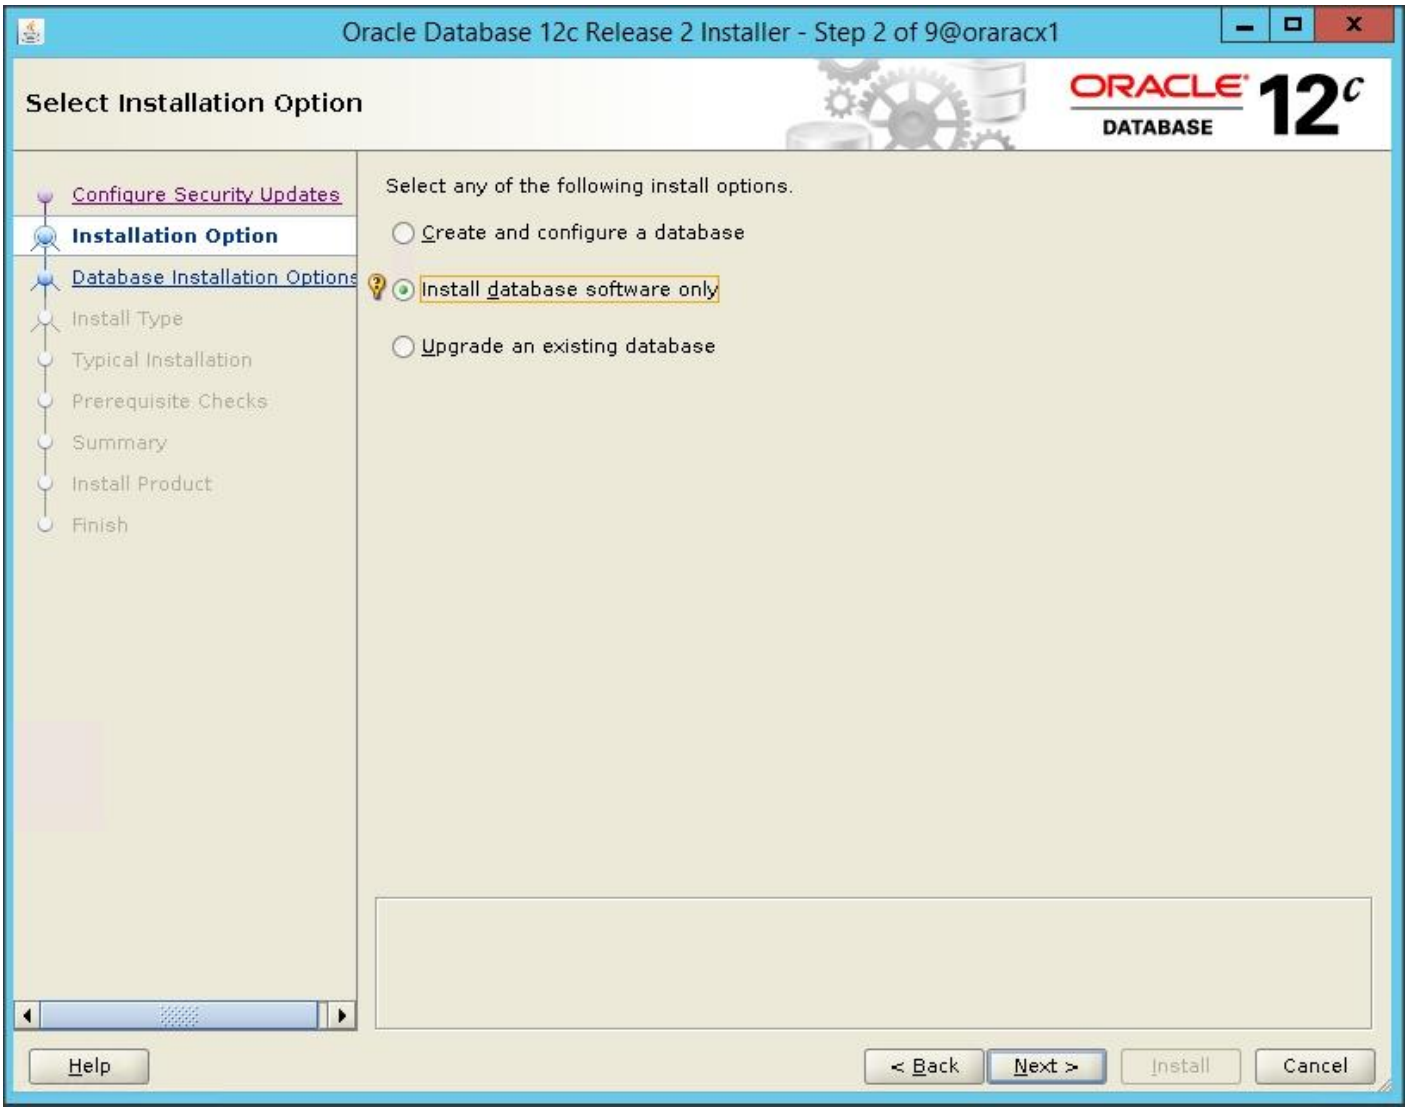

3. Select option "Oracle Real Application Clusters database installation" and click Next.

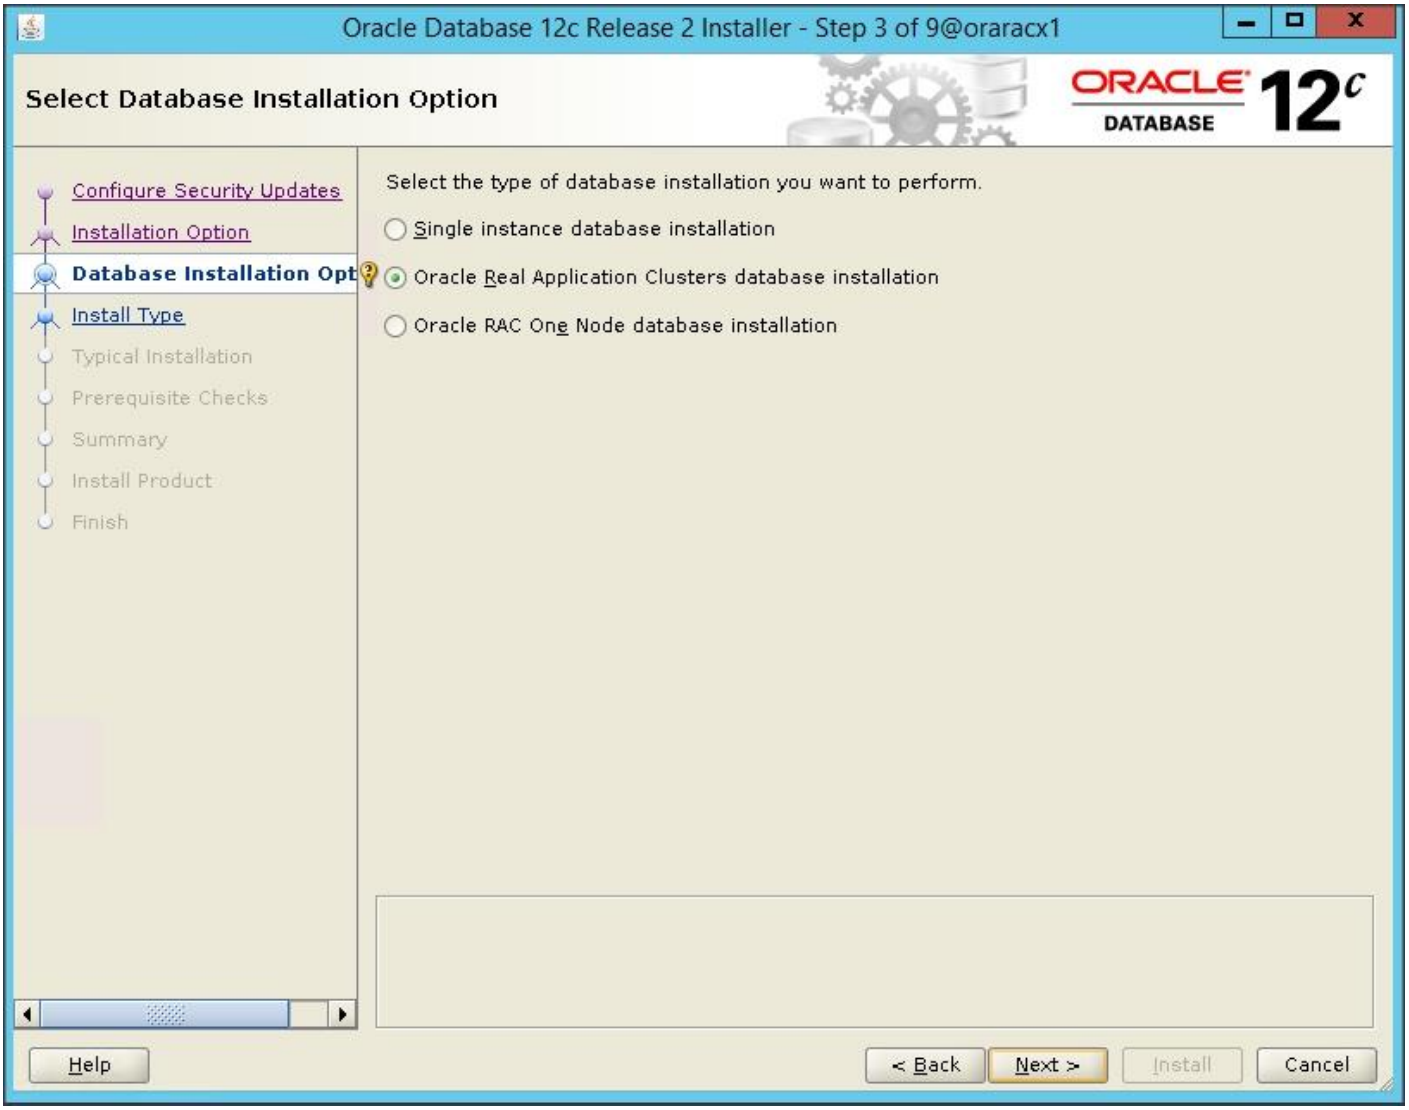

4. Select nodes in the cluster where installer should install Oracle RAC. For this setup, you will install software on all nodes as shown below.

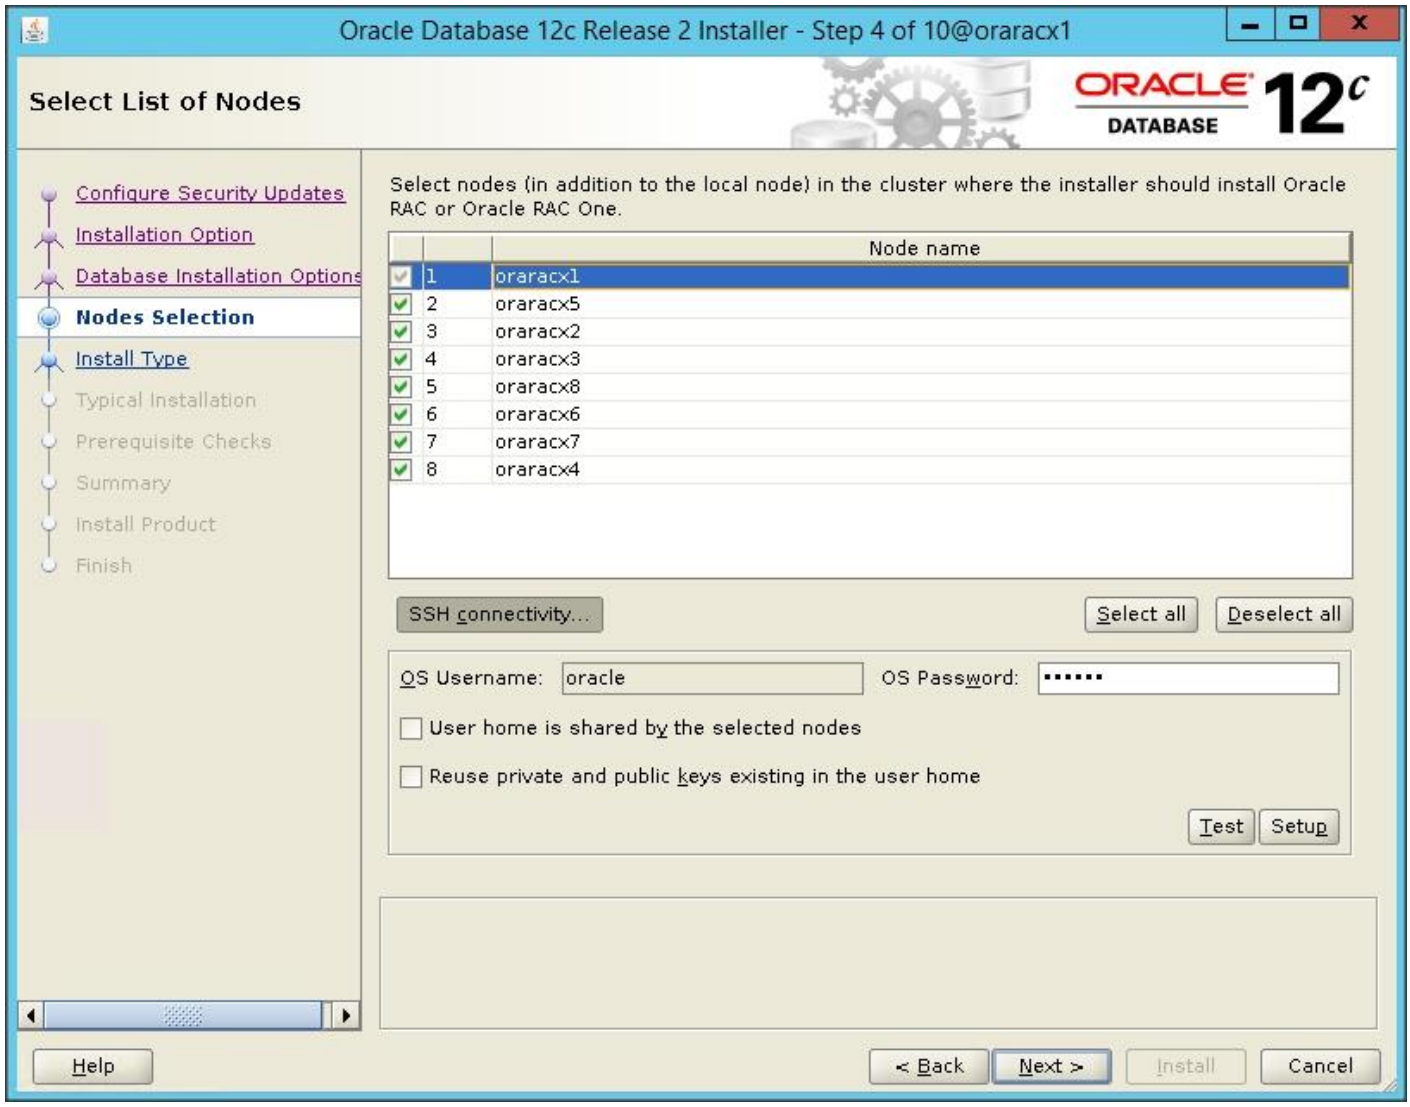

5. Click the "SSH Connectivity..." button and enter the password for the "oracle" user. Click the "Setup" button to configure passwordless SSH connectivity, and the "Test" button to test it once it is complete. When the test is complete, click Next.

6. Select Database Edition Options according to your environments and then click Next.

7. Enter Oracle Base as "/u01/app/oracle" and "/u01/app/oracle/product/12.2.0/dbhome_1" as the software location, then click Next.

8. Select the desired operating system groups and then click Next.

9. Wait for the prerequisite check to complete. If there are, any problems either click the "Fix & Check Again" button, or try to fix those by checking and manually installing required packages. Click Next.

10. Verify the Oracle Database summary information, click Install.

11. When prompted, run the configuration script on each node. When the scripts run successfully on each node, click OK.

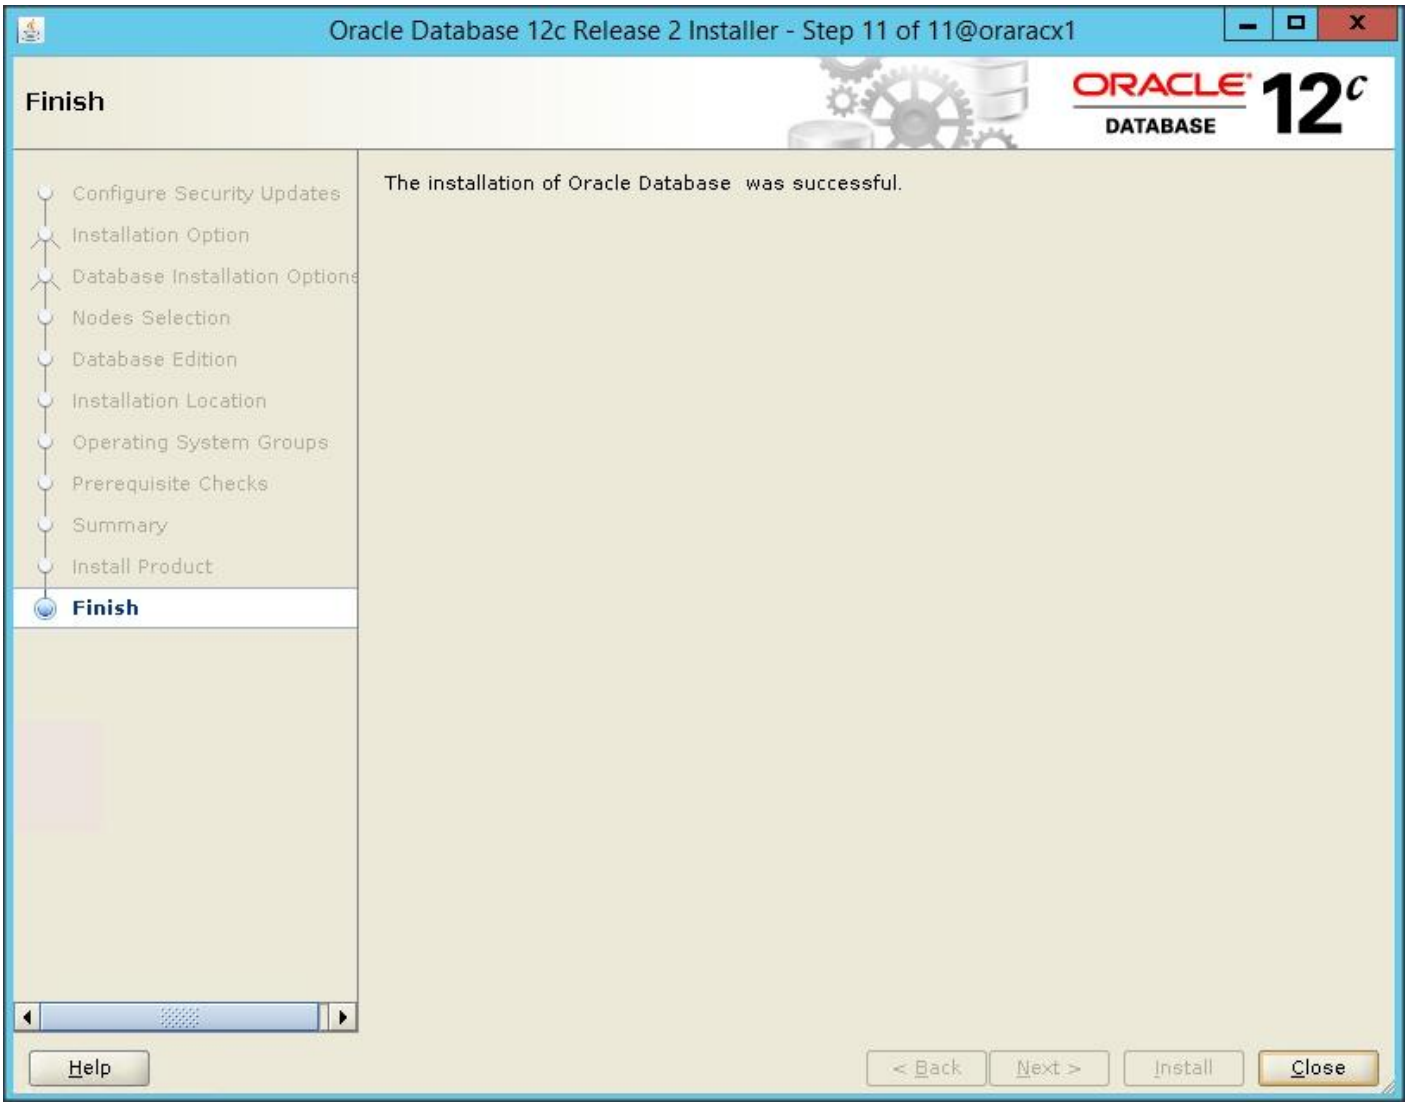

12. Click Close to exit the installer.

Configure Database LUNs

Use the Oracle ASM command line tool (ASMCMD) or ASM Configuration Assistance GUI to provision devices for use with Oracle ASM Filter Driver. You can label the disks by running asmcmd command line utility as shown below:

/u01/app/12.2.0/grid/bin/asmcmd afd_label DATAOLTP1 /dev/mapper/dg_oradata_oltp1 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label REDOOLTP1 /dev/mapper/dg_oraredo_oltp1 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DATASOE1 /dev/mapper/dg_oradata_soe1 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label REDOSOE1 /dev/mapper/dg_oraredo_soe1 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DATADSS1 /dev/mapper/dg_oradata_dss1 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label REDODSS1 /dev/mapper/dg_oraredo_dss1 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB1 /dev/mapper/dg_oradata_slob01 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB2 /dev/mapper/dg_oradata_slob02 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB3 /dev/mapper/dg_oradata_slob03 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB4 /dev/mapper/dg_oradata_slob04 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB5 /dev/mapper/dg_oradata_slob05 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB6 /dev/mapper/dg_oradata_slob06 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB7 /dev/mapper/dg_oradata_slob07 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB8 /dev/mapper/dg_oradata_slob08 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label DSLOB9 /dev/mapper/dg_oradata_slob09 –init

/u01/app/12.2.0/grid/bin/asmcmd afd_label REDOSLOB /dev/mapper/dg_oraredo_slob –init

Verify the device has been marked for use with Oracle ASMFD:

Figure 18 ASM Disk Groups

```
[root@oraracx1 ~]# export ORACLE_HOME=/u01/app/12.2.0/grid
[root@oraracx1 ~]# cd $ORACLE_HOME
[root@oraracx1 grid]# cd bin/
[root@oraracx1 bin]#
[root@oraracx1 bin]# ./asmcmd afd_lslbl
--------------------------------------------------------------------------------
Label                     Duplicate  Path
================================================================================
DATADSS1                             /dev/mapper/dg_oradata_dss1
DATAOLTP1                            /dev/mapper/dg_oradata_oltp1
DATASOE1                             /dev/mapper/dg_oradata_soe1
DSLOB1                               /dev/mapper/dg_oradata_slob01
DSLOB2                               /dev/mapper/dg_oradata_slob02
DSLOB3                               /dev/mapper/dg_oradata_slob03
DSLOB4                               /dev/mapper/dg_oradata_slob04
DSLOB5                               /dev/mapper/dg_oradata_slob05
DSLOB6                               /dev/mapper/dg_oradata_slob06
DSLOB7                               /dev/mapper/dg_oradata_slob07
DSLOB8                               /dev/mapper/dg_oradata_slob08
DSLOB9                               /dev/mapper/dg_oradata_slob09
OCRVOTE                              /dev/mapper/dg_orarac_crs
REDODSS1                             /dev/mapper/dg_oraredo_dss1
REDOOLTP1                            /dev/mapper/dg_oraredo_oltp1
REDOSLOB                             /dev/mapper/dg_oraredo_slob
REDOSOE1                             /dev/mapper/dg_oraredo_soe1
[root@oraracx1 bin]#
```

The figure below displays all the Disk Groups created to configure databases for this solution environment.

Figure 19 ASM Disk Groups

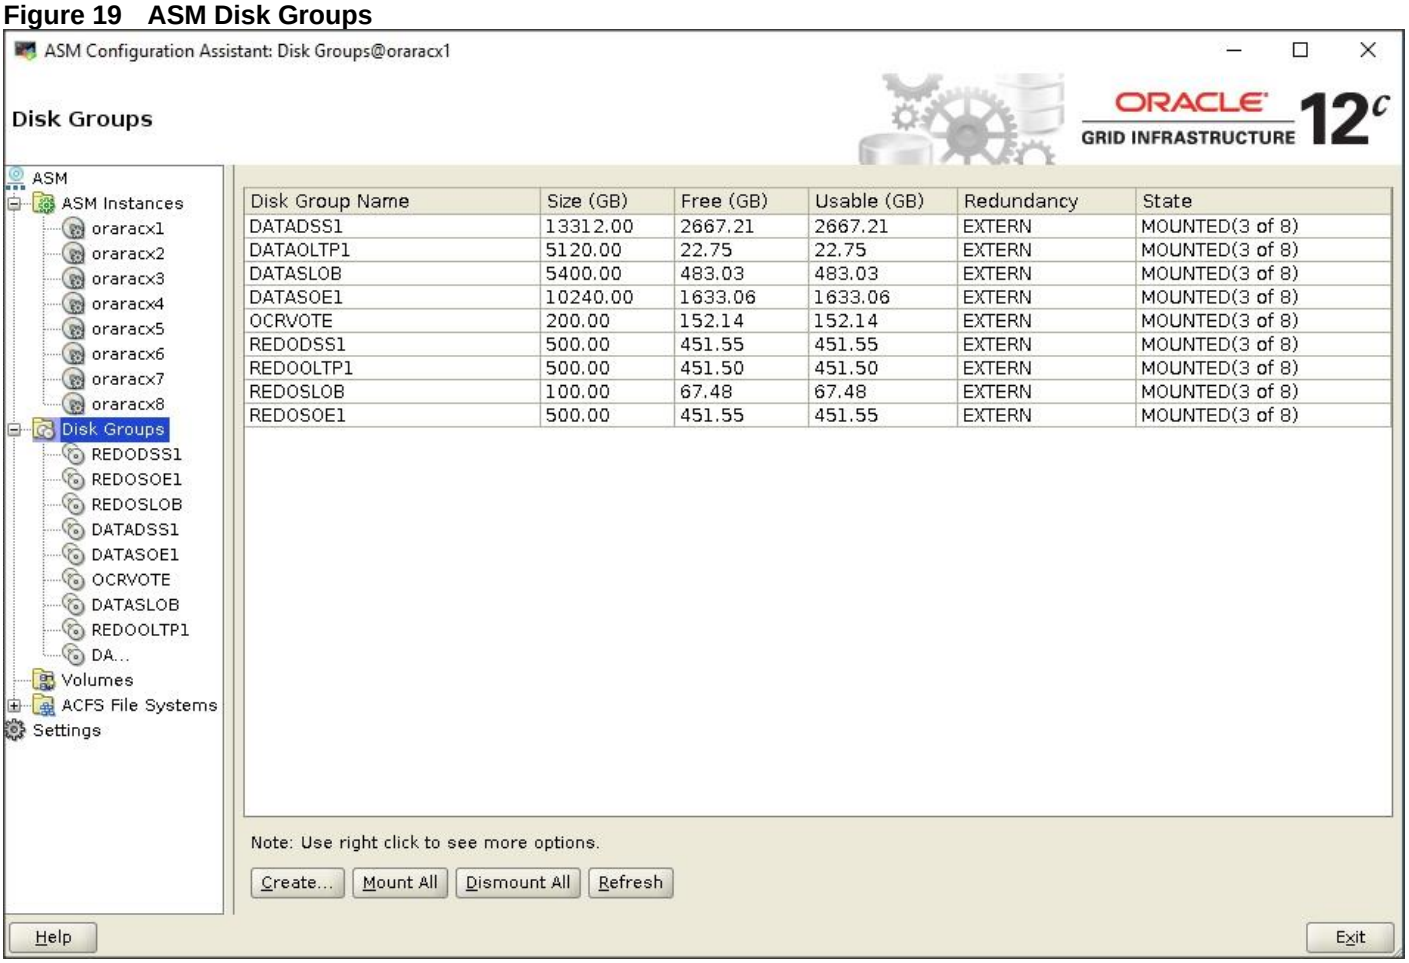

As shown above, we created the disk-group as DATADSS1 and REDODSS1 for DSS Database (DSSDB1) workload. Similarly, we have created disk-groups for OLTPDB1, SOEDB1 and SLOBDB1 for OLTP Databases workloads as explained in the following section.

Scalability Test and Results

Before configuring a database for workload tests, it is extremely important to validate that this is indeed a balanced configuration that is capable of delivering expected performance. In this FlashStack solution, we will test and validate node and user scalability on an 8 node Oracle RAC Databases with various database benchmarking tools. We used widely adopted database performance test tools to test and validate throughput, IOPS, and latency for various test scenarios on FlashArray //X70 system as follows:

- SLOB (Silly Little Oracle Benchmark)
- CalibrateIO
- Swingbench

Database Creation with DBCA

We used Oracle Database Configuration Assistant (DBCA) to create three OLTP (SLOBDB, OLTPDB1 and SOEDB1) and one DSS (DSSDB1) databases for SLOB and Swingbench calibration. Alternatively, you can use Database creation scripts to create the databases as well. Make sure to place the data files, redolog files and control files in appropriate directory paths discussed in the storage layout section.

SLOB Performance on FlashArray //X70

The Silly Little Oracle Benchmark (SLOB) is a toolkit for generating and testing I/O through an Oracle database. SLOB is very effective in testing the I/O subsystem with genuine Oracle SGA-buffered physical I/O. SLOB supports testing physical random single-block reads (db file sequential read) and random single block writes (DBWR flushing capability).

SLOB issues single block reads for the read workload that are generally 8K (as the database block size was 8K). For testing the SLOB workload, we created one OLTP database (SLOBDB) of 4 TB in Size. We created two disk-group to store the data and redolog files for the SLOBDB database. First disk-group “Data-SLOB” was created with 9 LUNs (600 GB each) while second disk-group “Redo-SLOB” was created with one LUN (100 GB). We loaded SLOB schema on “Data-SLOB” disk-group of up to 4 TB in size. The following tests were performed and various metrics like IOPS and latency were captured along with Oracle AWR reports for each test scenario.

User Scalability Performance on FlashArray //X70

SLOB was configured to run against all the 8 RAC nodes and the concurrent users were equally spread across all the nodes. For Pure Storage FlashArray //X70, we scale users from 32 to 256 for Oracle RAC 8 nodes and identify the maximum IOPS and latency as explained.

- User Scalability test with 32, 64, 128, 192 and 256 users on 8 Oracle RAC nodes
- Varying workloads
  - 100% read (0% update)
  - 90% read (10% update)
  - 70% read (30% update)
  - 50% read (50% update)

The following table illustrate user scalability for the total number of IOPS (both read and write) when run with 32, 64, 128, 192 and 256 Users. For all user scale, we recorded the following number of IOPS as shown in the table below:

Table 11 User Scalability

Users	Read/Write % (100/0)	Read/Write % (90/10)	Read/Write % (70/30)	Read/Write % (50/50)
32	74,363	87,205	106,367	105,738
64	154,943	169,370	198,618	204,928
128	302,211	316,118	329,088	322,619
192	414,616	397,346	367,135	327,229
256	496,892	412,491	364,626	342,852

The following graphs illustrate user scalability in terms of total IOPS while running SLOB workload for 32, 64, 128, 192 and 256 concurrent users for each test scenario.

Figure 20 SLOB User Scalability – IOPS

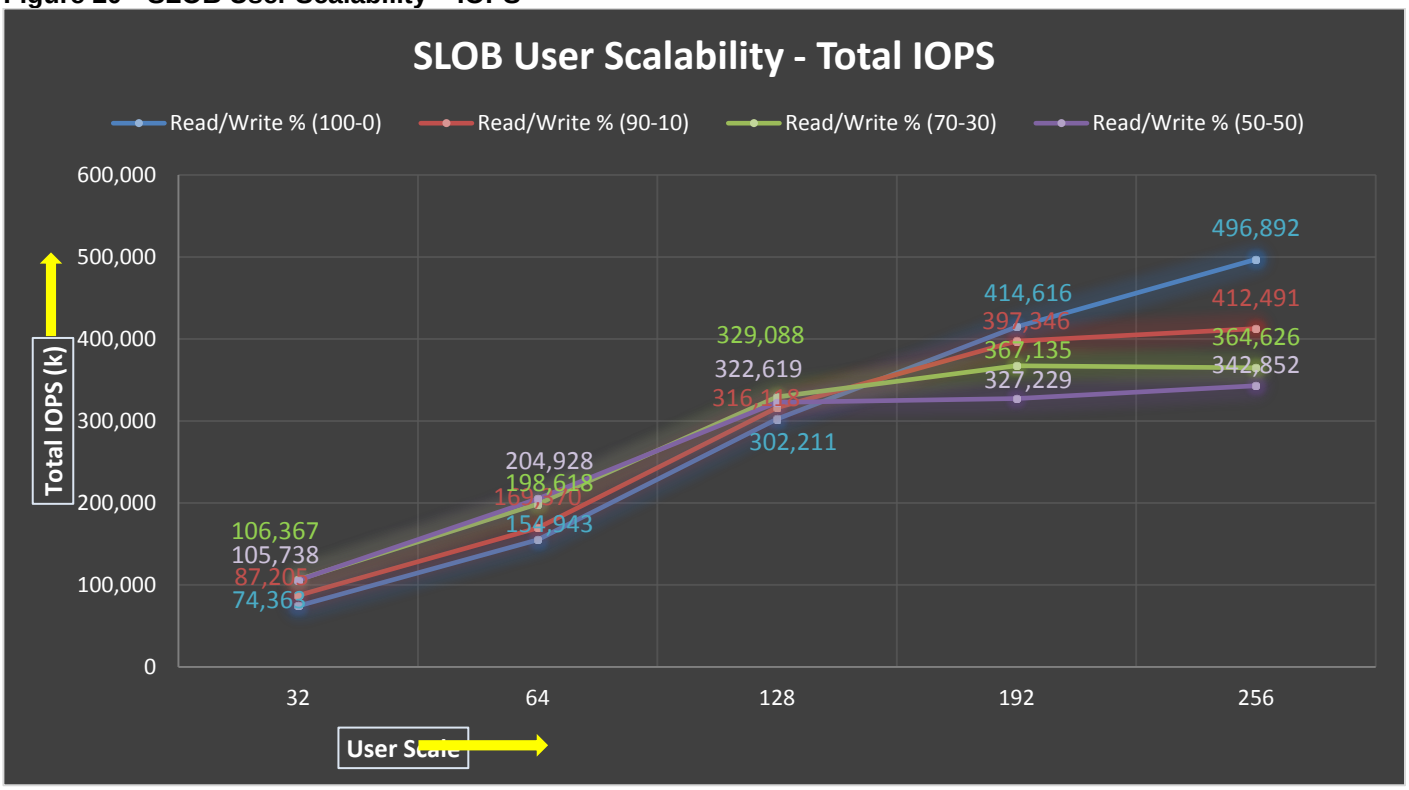

The graph illustrates the linear scalability with increased users and similar IOPS from 32 users to 256 users with 100% read, 90% read, 70% read and 50% read. The below snapshot was captured from a 100% Read (0% update) Test scenario while running SLOB test. The snapshot shows a section from 3-hour window of AWR report from the run that highlights Physical Reads/Sec and Physical Writes/Sec for each instance. This section highlights that IO load is distributed across all the cluster nodes performing workload operations. Due to variations in workload randomness, we conducted multiple runs to ensure consistency in behavior and test results.

Figure 21 SLOB – IOPS AWR Snapshot

System Statistics - Per Second DB/Inst: SLOB/slob1 Snaps: 46-47

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	67,666.78	61,597.7	1.2	2.8	8.8	0.5	1,029.7	2.3	0.12	0.0
2	68,027.41	61,921.7	1.1	2.9	8.8	0.5	1,035.5	2.3	0.12	0.0
3	68,204.55	62,089.6	1.0	2.7	8.5	0.7	1,037.9	2.4	0.12	0.0
4	72,683.00	62,249.4	1.1	2.9	8.8	0.5	1,040.4	2.4	0.12	0.0
5	68,223.68	62,098.3	1.1	2.8	8.9	0.5	1,039.1	2.4	0.12	0.0
6	68,461.38	62,338.5	1.1	2.7	8.3	0.5	1,041.7	2.3	0.11	0.1
7	68,004.02	61,936.0	1.1	2.8	8.5	0.5	1,035.1	2.4	0.12	0.0
8	72,693.44	62,651.9	1.2	2.6	36.6	0.5	1,045.1	2.8	0.11	0.0
Sum	553,964.26	496,883.0	8.9	22.1	97.1	4.3	8,304.4	19.2	0.93	0.3
Avg	69,245.53	62,110.4	1.1	2.8	12.1	0.5	1,038.0	2.4	0.12	0.0
Std	2,136.91	314.9	0.0	0.1	9.9	0.1	4.7	0.2	0.00	0.0

Even though the FlashArray //X70 can scale up to 9 GBps of reads, we were limited by the total number of IOPS and not on the bandwidth. The maximum bandwidth is validated with the DSS query as shown in next section.

The following graph illustrates the latency exhibited by the //X70 FlashArray across different workloads. All the workloads experienced less than 1 millisecond latency and it varies based on the workload. As expected, the 50% read (50% update) test exhibited higher latencies as the user counts increases. However, these are exceptional performance characteristics keeping the nature of the IO load.

Figure 22 SLOB User Scalability – Latency

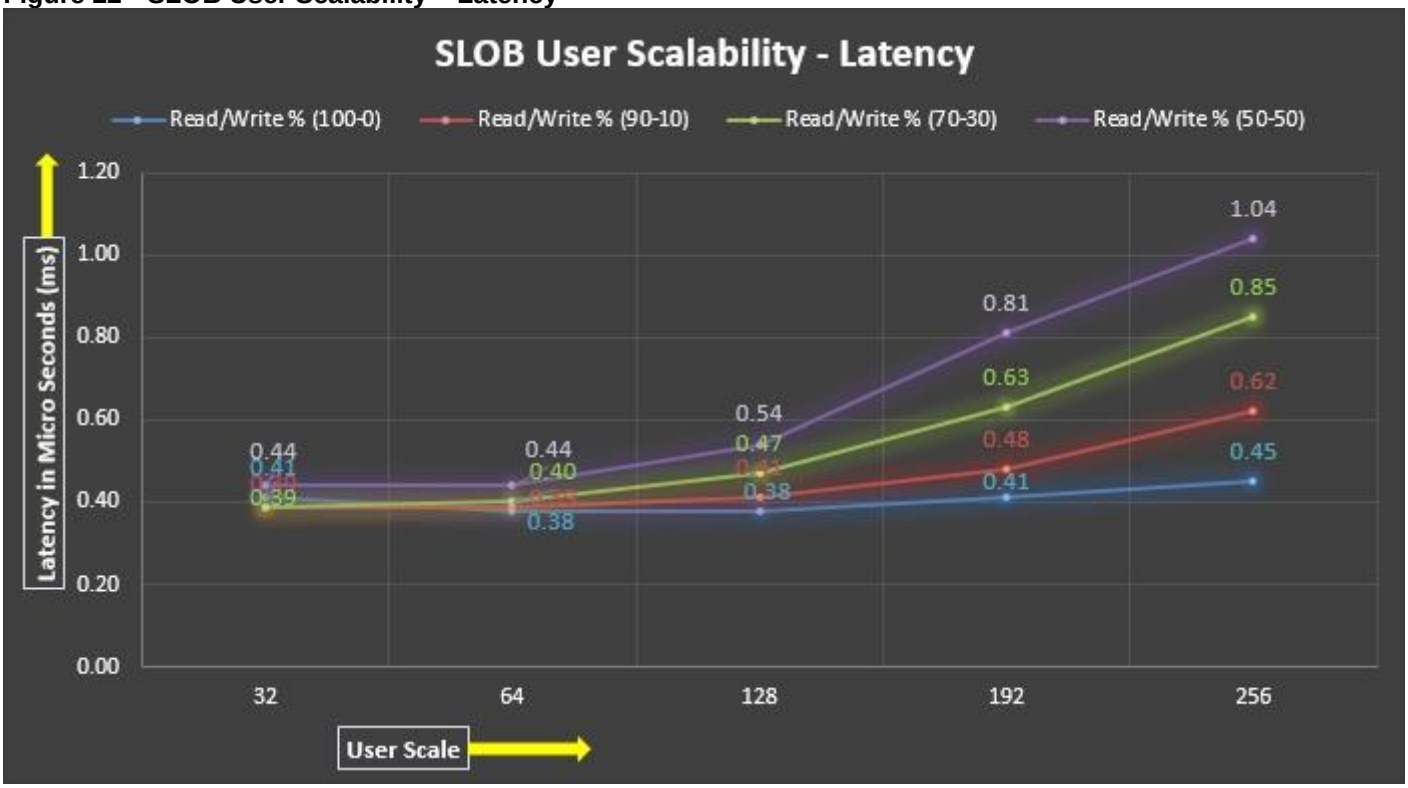

The following screenshot was captured from 50% Read (50% Update) Test scenario while running SLOB test. The snapshot shows a section from 3-hour window of AWR report from the run that highlights top timed Background Events.

Figure 23 SLOB User Scalability – Top Timed Background Events

Top Timed Background Events DB/Inst: SLOB/slob1 Snaps: 150-151
-> % of bg time: % of background elapsed time
-> Instance '*' - cluster wide summary
-> '*' Waits, %Timeouts, Wait Time Total(s) : Cluster-wide total for the wait event
-> '*' 'Wait Time Avg' : Cluster-wide average computed as (Wait Time Total / Event Waits)
-> '*' Summary 'Avg Wait Time' : Per-instance 'Wait Time Avg' used to compute the following statistics
-> '*' [Avg/Min/Max/Std Dev] : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
-> '*' Cnt : count of instances with wait times for the event

	Wait	Event			Wait Time			Summary Avg Wait Time				
I#	Class	Event	Waits	%Timeouts	Total(s)	Avg Wait	%DB time	Avg	Min	Max	Std Dev	Cnt
*		background cpu time	N/A	N/A	14,519.13		63.01					8
	System I/O	db file parallel write	13,242,015	0.0	6,555.66	495.07us	28.45	495.02us	490.23us	499.82us	3.46us	8
	System I/O	log file parallel write	3,123,071	0.0	4,747.55	1.52ms	20.60	1.52ms	1.50ms	1.55ms	20.96us	8
	Other	RMA: IPC0 completion sync	15,139	0.0	293.53	19.39ms	1.27	19.39ms	19.35ms	19.42ms	25.12us	8
	Cluster	service monitor: inst recovery completio	0	N/A	212.55		0.92					8
	System I/O	control file sequential read	101,438	0.0	78.96	778.41us	0.34	778.85us	763.67us	815.00us	17.31us	8
	Other	latch free	51,699	0.0	44.68	864.24us	0.19	885.27us	350.63us	1.86ms	672.41us	8
	Applicatio	enq: TX - row lock contention	17	0.0	40.10	2358.68m	0.17	2391.43ms	2022.23ms	2814.90ms	279.92ms	8
	Other	Sync ASM rebalance	3,504	0.0	32.06	9.15ms	0.14	9.03ms	14.68us	24.34ms	12.44ms	8
	System I/O	log file sequential read	13,156	0.0	25.89	1.97ms	0.11	1.97ms	1.91ms	2.02ms	36.83us	8

The following screenshot was captured from 90% Read (10% Update) Test scenario while running SLOB test. The snapshot shows a section from 3-hour window of AWR report from the run that highlights top timed Events.

Figure 24 SLOB User Scalability – Top Timed Events

Top Timed Events DB/Inst: SLOB/slob1 Snaps: 56-57
-> Instance '*' - cluster wide summary
-> '*' Waits, %Timeouts, Wait Time Total(s) : Cluster-wide total for the wait event
-> '*' 'Wait Time Avg' : Cluster-wide average computed as (Wait Time Total / Event Waits)
-> '*' Summary 'Avg Wait Time ' : Per-instance 'Wait Time Avg ' used to compute the following statistics
-> '*' [Avg/Min/Max/Std Dev] : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
-> '*' Cnt : count of instances with wait times for the event

	Wait	Event			Wait Time			Summary Avg Wait Time				
I#	Class	Event	Waits	%Timeouts	Total(s)	Avg Wait	%DB time	Avg	Min	Max	Std Dev	Cnt
*	User I/O	db file sequential read	670,680,974	0.0	416,854.63	621.54us	90.44	621.54us	620.33us	622.49us	660.67ns	8
		DB CPU	N/A	N/A	43,762.04		9.49					8
	Cluster	gc cr grant 2-way	281,631,284	0.0	18,822.33	66.83us	4.08	66.84us	65.21us	67.85us	826.36ns	8
	Cluster	gc current grant 2-way	32,504,757	0.0	2,179.90	67.06us	0.47	67.07us	65.84us	67.86us	656.23ns	8
	System I/O	db file parallel write	3,017,108	0.0	1,121.20	371.61us	0.24	371.63us	366.88us	377.31us	3.79us	8
	System I/O	log file parallel write	1,124,774	0.0	875.07	777.99us	0.19	778.00us	771.77us	781.56us	3.73us	8
	Other	gcs drm freeze in enter server mode	2,020	0.0	588.33	291.25ms	0.13	291.27ms	280.98ms	304.54ms	8.73ms	8
	Other	RMA: IPC0 completion sync	15,136	0.0	293.50	19.39ms	0.06	19.39ms	19.34ms	19.43ms	27.50us	8
	User I/O	db file scattered read	475,492	0.0	289.02	607.84us	0.06	607.84us	601.74us	614.25us	4.86us	8
	Other	enq: XL - fault extent map	486	0.0	234.84	483.20ms	0.05	354.61ms	6.31ms	871.36ms	332.16ms	8

SwingBench Performance on FlashArray //X70

Database Workload Configuration

We used Swingbench and Calibrate IO for workload testing. Swingbench is a simple to use, free, Java-based tool to generate database workload and perform stress testing using different benchmarks in Oracle database environments. Swingbench can be used to demonstrate and test technologies such as Real Application Clusters, Online table rebuilds, Standby databases, online backup and recovery, etc.

Swingbench provides four separate benchmarks, namely, Order Entry, Sales History, Calling Circle, and Stress Test. For the tests described in this solution, Swingbench Order Entry benchmark was used for OLTP workload testing and the Sales History benchmark was used for the DSS workload testing.

The Order Entry benchmark is based on SOE schema and is TPC-C like by types of transactions. The workload uses a very balanced read/write ratio around 60/40 and can be designed to run continuously and test the performance of a typical Order Entry workload against a small set of tables, producing contention for database resources.

The Sales History benchmark is based on the SH schema and is TPC-H like. The workload is query (read) centric and is designed to test the performance of queries against large tables.

Typically encountered in the real-world deployments, we tested a combination of scalability and stress related scenarios that ran on all the 8-node Oracle RAC cluster configuration.

- OLTP database user scalability and OLTP database node scalability representing small and random transactions

- DSS database workload representing larger transactions
- Mixed workload featuring OLTP and DSS database workloads running simultaneously for 24 hours

For Swingbench workload, we created two OLTP (Order Entry) and one DSS (Sales History) database to demonstrate database consolidation, multi-tenancy capability, performance and sustainability. We created approximately 4 TB of OLTPDB1, 5 TB of SOEDB1 and 9 TB of DSSDB1 database to perform swingbench testing.

After creating all the databases, we have to run the Calibrate IO tool to check the performance of the storage system as described below.

Oracle Calibrate IO Performance on FlashArray //X70

The I/O calibration feature of Oracle Database enables you to assess the performance of the storage subsystem, and determine whether I/O performance problems are caused by the database or the storage subsystem. Unlike other external I/O calibration tools that issue I/Os sequentially, the I/O calibration feature of Oracle Database issues I/Os randomly using Oracle datafiles to access the storage media, producing results that more closely match the actual performance of the database.

The I/O calibration feature of Oracle Database is accessed using the DBMS_RESOURCE_MANAGER.CALIBRATE_IO procedure. This procedure issues an I/O intensive read-only workload (made up of one megabytes of random of I/Os) to the database files to determine the maximum IOPS (I/O requests per second) and MBPS (megabytes of I/O per second) that can be sustained by the storage subsystem. Due to the overhead from running the I/O workload, I/O calibration should only be performed when the database is idle, or during off-peak hours, to minimize the impact of the I/O workload on the normal database workload.

To run I/O calibration and assess the I/O capability of the storage subsystem used by Oracle Database, use the DBMS_RESOURCE_MANAGER.CALIBRATE_IO procedure:

```
SET SERVEROUTPUT ON

DECLARE

  lat  INTEGER;

  iops INTEGER;

  mbps INTEGER;

BEGIN

   DBMS_RESOURCE_MANAGER.CALIBRATE_IO (2, 10, iops, mbps, lat);

  DBMS_OUTPUT.PUT_LINE ('max_iops = ' || iops);

  DBMS_OUTPUT.PUT_LINE ('latency  = ' || lat);

  dbms_output.put_line('max_mbps = ' || mbps);

end;

/
```

For Oracle Real Application Clusters (RAC) configurations, make sure that all instances are opened to calibrate the storage subsystem across nodes.

Calibrate IO Test Run on One Database At-a-Time

In this test scenario, we ran the calibrate IO on each OLTP and DSS database one at a time to get the result. We observed the following number of IOPS, throughput and latency.

- For OLTPDB1 Database
  - max_iops = 450827
  - latency = 1
  - max_mbps = 8648
- For SOEDB1 Database
  - max_iops = 445286
  - latency = 2
  - max_mbps = 8662
- For DSSDB1 Database
  - max_iops = 384166
  - latency = 3
  - max_mbps = 8184

As these tests were ran one at a time, each database was able to get the maximum IOPS and bandwidth at lower latency.

Calibrate IO Test Run on All Databases at Once

In this test scenario, we ran the calibrate IO on all three database (2 OLTP and one DSS) at the same time to get the result. We observed the following number of IOPS, throughput and latency:

- For OLTPDB1 Database
  - max_iops = 249580
  - latency = 13
  - max_mbps = 7786
- For SOEDB1 Database
  - max_iops = 150149
  - latency = 35
  - max_mbps = 4329
- For DSSDB1 Database
  - max_iops = 135476
  - latency = 33
  - max_mbps = 8524

As expected, when calibrate IO was simultaneously across all three databases, all of them ended in getting a portion of the IOPS that was possible. Similarly, all three received the maximum possible bandwidth but with higher latency due to higher queues of IO requests pushing the array to its limit. In comparison to the standalone calibrate IO tests, the bandwidth results from the above test stands out as couple of databases were able to achieve 7.7 and 8.5 GBps simultaneously. The latency is the end-to-end latency seen by Oracle and not necessarily reflect just the storage latency.

Scalability Performance on FlashArray //X70

The first step after the databases creation is calibration; about the number of concurrent users, nodes, OS and database optimization. For Pure Storage FlashArray //X70, we will scale the system from a 1 to 8 Oracle RAC Nodes. Also, for this FlashStack solution, we tested system performance with different databases running at a time and capture the results as explained in the following sections.

One (OLTP) Database Performance

For OLTP database workload featuring Order Entry schema, we used one SOEDB1 database. For the SOEDB1 database (5 TB), we used 64GB size of System Global Area (SGA). We also ensured that HugePages were in use. The OLTP Database scalability test was run for at least 12 hours and made sure that results are consistent for the duration of the full run.

We ran the SwingBench scripts on each node to start SOEDB1 database and generate AWR reports for each scenario as shown below:

- User Scalability

Total number of Reads and Write IOPS, TPM for various users are as shown below with system utilization under 25% all the time.

Users	Read IOPS	Write IOPS	Total IOPS	TPM	System Utilization (%)
100 (8 Nodes)	39,084	23,886	62,969	455,855	5.2
200 (8 Nodes)	69,046	39,662	108,708	763,248	10.3
300 (8 Nodes)	93,765	56,734	150,499	1,064,874	14.0
400 (8 Nodes)	130,635	73,493	204,128	1,552,266	19.0
600 (8 Nodes)	160,107	86,035	246,142	1,967,604	23.4
800 (8 Nodes)	161,166	96,800	257,966	2,094,720	25.4

The graph below illustrates the TPM for SOEDB1 database user scale on 8 node:

Figure 25 User Scalability – TPM & System Utilization (%)

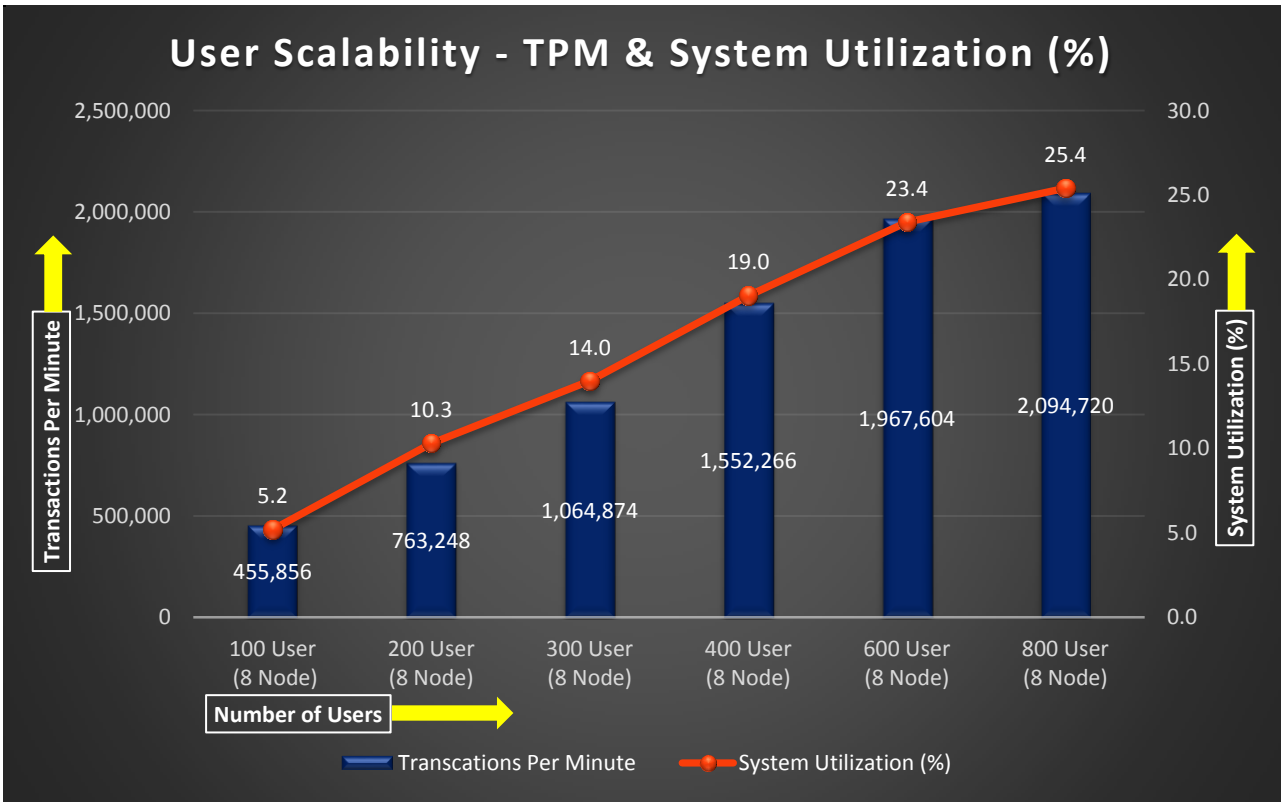

The graph illustrates steady scalability till 600 users. Beyond 600 users, the additional users yield higher TPM but not at the same TPM/user rate. The below graph illustrates the total number of IOPS for SOEDB1 database user scale on all 8 nodes.

Figure 26 User Scalability – IOPS

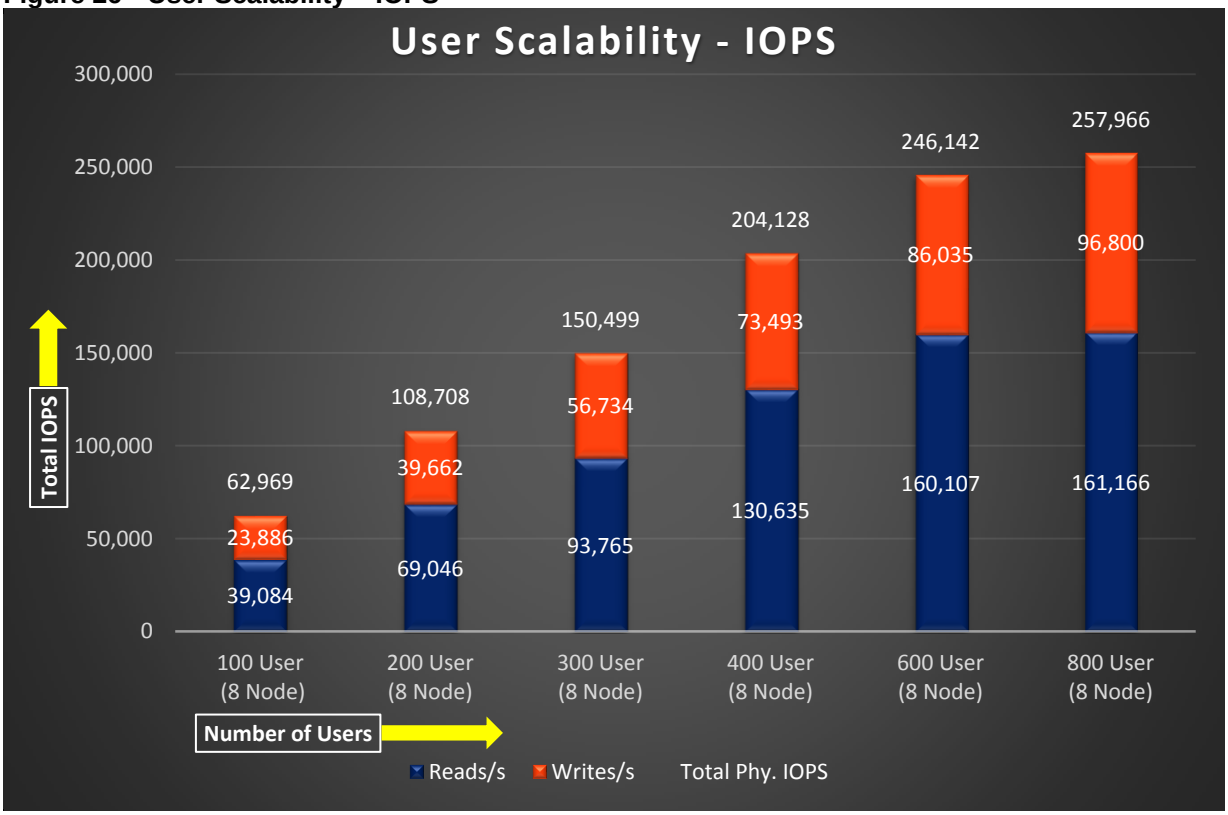

The screenshot below was captured from Pure Storage GUI for the 800 Users Scale Test scenario while running Swingbench workload on one database.

Figure 27 User Scalability – System Statistics Per Second

The screenshot shown below was captured from the 800 Users Scale Test scenario while running Swingbench workload on one database. The snapshot shows a section from 3-hour window of AWR Global report from the run that highlights Physical Reads/Sec and Physical Writes/Sec for each instance. Notice that IO load is distributed across all the cluster nodes performing workload operations. Even though the FlashArray //X is capable of

achieving higher IOPS, the application was benefiting from the SGA and global cache and not all requests for data was sent to the storage array.

Figure 28 User Scalability – System Statistics Per Second

System Statistics - Per Second DB/Inst: SOEDB1/soedb11 Snaps: 27-31

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	442,198.97	19,141.8	11,323.2	14,318.4	90,139.8	12,339.4	42,715.3	4,888.5	0.12	4,104.6
2	460,705.17	20,033.0	11,883.2	15,085.5	94,919.9	12,984.9	44,946.9	5,141.3	0.12	4,319.4
3	473,471.17	20,365.5	12,223.5	15,316.6	96,443.1	13,204.8	45,707.9	5,235.4	0.12	4,392.5
4	457,944.66	19,843.7	11,853.4	14,986.3	94,274.5	12,902.6	44,669.0	5,115.3	0.12	4,292.1
5	455,656.07	19,738.2	11,913.0	14,923.6	93,906.2	12,852.1	44,483.1	5,093.2	0.12	4,275.2
6	460,117.57	19,897.2	12,017.6	15,057.4	94,785.3	12,974.7	44,918.0	5,128.3	0.12	4,315.9
7	456,932.35	19,808.4	11,816.9	14,943.7	94,098.7	12,886.8	44,610.8	5,108.3	0.12	4,286.6
8	524,797.71	22,337.9	13,769.2	17,111.4	108,079.1	14,806.1	51,257.3	5,871.4	0.12	4,925.7
Sum	3,731,823.67	161,165.7	96,800.1	121,743.0	766,646.6	104,951.4	363,308.4	41,581.7	0.97	34,912.0
Avg	466,477.96	20,145.7	12,100.0	15,217.9	95,830.8	13,118.9	45,413.5	5,197.7	0.12	4,364.0
Std	25,055.36	949.2	720.6	816.2	5,262.2	724.3	2,508.4	289.0	0.00	241.1

The AWR screenshot shown below shows latency for the same 800 Users Scale Test while swingbench test was running.

Figure 29 User Scalability – Top Timed Events

Top Timed Events DB/Inst: SOEDB1/soedb11 Snaps: 27-31
-> Instance '*' - cluster wide summary
-> '*' Waits, %Timeouts, Wait Time Total(s) : Cluster-wide total for the wait event
-> '*' 'Wait Time Avg' : Cluster-wide average computed as (Wait Time Total / Event Waits)
-> '*' Summary 'Avg Wait Time ' : Per-instance 'Wait Time Avg ' used to compute the following statistics
-> '*' [Avg/Min/Max/Std Dev] : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
-> '*' Cnt : count of instances with wait times for the event

	Wait		Event		Wait Time			Summary Avg Wait Time				
I#	Class	Event	Waits	%Timeouts	Total(s)	Avg Wait	%DB time	Avg	Min	Max	Std Dev	Cnt
*	User I/O	db file sequential read	1.679115E+09	0.0	2,018,414.54	1.20ms	34.76	1.20ms	1.18ms	1.21ms	10.77us	8
		DB CPU	N/A	N/A	1,263,753.80		21.76					8
	Commit	log file sync	373,925,161	0.0	791,233.31	2.12ms	13.62	2.12ms	2.03ms	2.16ms	41.98us	8
	Cluster	gc cr block 3-way	506,608,241	0.0	274,258.24	541.36us	4.72	540.10us	329.62us	606.27us	86.39us	8
	Cluster	gc current block busy	639,480,825	0.0	265,645.86	415.41us	4.57	417.46us	308.85us	443.99us	44.99us	8
	Cluster	gc current block 3-way	479,055,766	0.0	247,249.70	516.12us	4.26	514.93us	324.09us	568.54us	77.89us	8
	Cluster	gc current grant 2-way	855,533,922	0.0	214,135.72	250.29us	3.69	250.26us	242.92us	264.62us	7.06us	8
	Cluster	gc current block congested	53,410,383	0.0	177,337.22	3.32ms	3.05	3.27ms	2.92ms	3.33ms	142.48us	8
	Cluster	gc cr grant 2-way	706,498,677	0.0	147,518.65	208.80us	2.54	209.02us	197.59us	215.39us	6.41us	8
	Cluster	gc current grant busy	112,218,694	0.0	130,489.24	1.16ms	2.25	1.16ms	1.15ms	1.19ms	12.51us	8

The AWR screenshot shown below shows the Interconnect device statistics for the same 800 Users Scale Test while swingbench test running. This confirms the caching phenomenon that reduced the amount of IO to the storage.

Figure 30 Interconnect Device Statistics

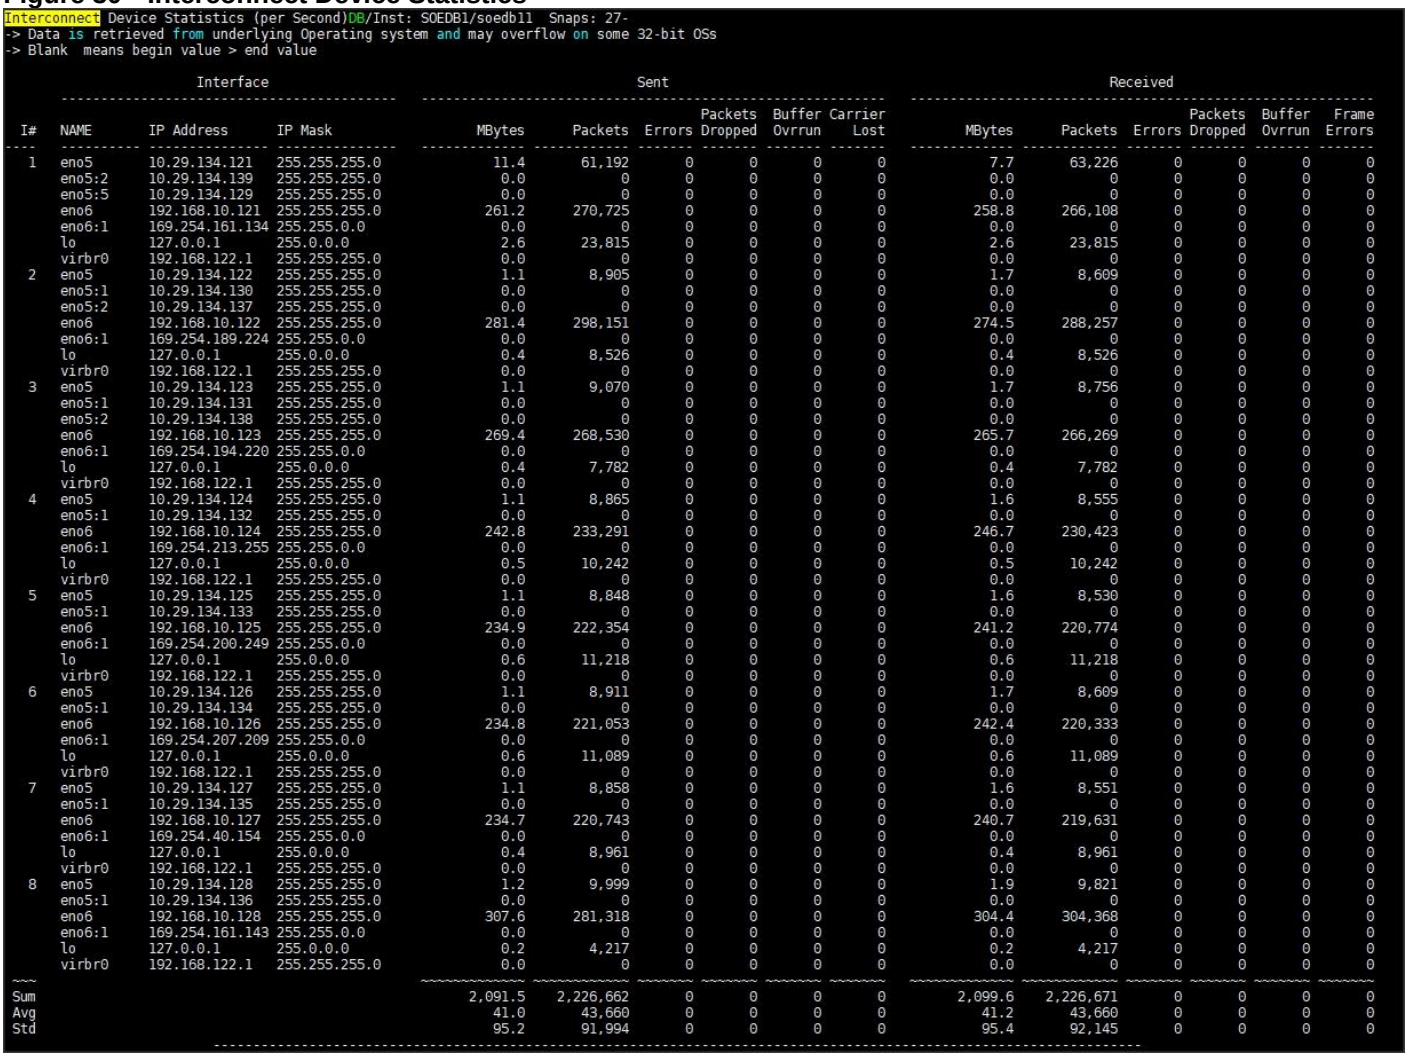

Interconnect Device Statistics (per Second)DB/Inst: SOEDB1/soedb11 Snaps: 27-
-> Data is retrieved from underlying Operating system and may overflow on some 32-bit OSs
-> Blank means begin value > end value

	Interface			Sent						Received					
I#	NAME	IP Address	IP Mask	MBytes	Packets	Errors	Packets Dropped	Buffer Ovrrun	Carrier Lost	MBytes	Packets	Errors	Packets Dropped	Buffer Ovrrun	Frame Errors
1	eno5	10.29.134.121	255.255.255.0	11.4	61,192	0	0	0	0	7.7	63,226	0	0	0	0
	eno5:2	10.29.134.139	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno5:5	10.29.134.129	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.121	255.255.255.0	261.2	270,725	0	0	0	0	258.8	266,108	0	0	0	0
	eno6:1	169.254.161.134	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	2.6	23,815	0	0	0	0	2.6	23,815	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
2	eno5	10.29.134.122	255.255.255.0	1.1	8,905	0	0	0	0	1.7	8,609	0	0	0	0
	eno5:1	10.29.134.130	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno5:2	10.29.134.137	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.122	255.255.255.0	281.4	298,151	0	0	0	0	274.5	288,257	0	0	0	0
	eno6:1	169.254.189.224	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.4	8,526	0	0	0	0	0.4	8,526	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
3	eno5	10.29.134.123	255.255.255.0	1.1	9,070	0	0	0	0	1.7	8,756	0	0	0	0
	eno5:1	10.29.134.131	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno5:2	10.29.134.138	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.123	255.255.255.0	269.4	268,530	0	0	0	0	265.7	266,269	0	0	0	0
	eno6:1	169.254.194.220	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.4	7,782	0	0	0	0	0.4	7,782	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
4	eno5	10.29.134.124	255.255.255.0	1.1	8,865	0	0	0	0	1.6	8,555	0	0	0	0
	eno5:1	10.29.134.132	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.124	255.255.255.0	242.8	233,291	0	0	0	0	246.7	230,423	0	0	0	0
	eno6:1	169.254.213.255	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.5	10,242	0	0	0	0	0.5	10,242	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
5	eno5	10.29.134.125	255.255.255.0	1.1	8,848	0	0	0	0	1.6	8,530	0	0	0	0
	eno5:1	10.29.134.133	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.125	255.255.255.0	234.9	222,354	0	0	0	0	241.2	220,774	0	0	0	0
	eno6:1	169.254.200.249	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.6	11,218	0	0	0	0	0.6	11,218	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
6	eno5	10.29.134.126	255.255.255.0	1.1	8,911	0	0	0	0	1.7	8,609	0	0	0	0
	eno5:1	10.29.134.134	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.126	255.255.255.0	234.8	221,053	0	0	0	0	242.4	220,333	0	0	0	0
	eno6:1	169.254.207.209	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.6	11,089	0	0	0	0	0.6	11,089	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
7	eno5	10.29.134.127	255.255.255.0	1.1	8,858	0	0	0	0	1.6	8,551	0	0	0	0
	eno5:1	10.29.134.135	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.127	255.255.255.0	234.7	220,743	0	0	0	0	240.7	219,631	0	0	0	0
	eno6:1	169.254.40.154	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.4	8,961	0	0	0	0	0.4	8,961	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
8	eno5	10.29.134.128	255.255.255.0	1.2	9,999	0	0	0	0	1.9	9,821	0	0	0	0
	eno5:1	10.29.134.136	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.128	255.255.255.0	307.6	281,318	0	0	0	0	304.4	304,368	0	0	0	0
	eno6:1	169.254.161.143	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.2	4,217	0	0	0	0	0.2	4,217	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
Sum				2,091.5	2,226,662	0	0	0	0	2,099.6	2,226,671	0	0	0	0
Avg				41.0	43,660	0	0	0	0	41.2	43,660	0	0	0	0
Std				95.2	91,994	0	0	0	0	95.4	92,145	0	0	0	0

Two (OLTP) Database Performance

For two OLTP database workload featuring Order Entry schema, we used SOEDB1 and OLTPDB1databases. For both the databases, we used 64GB of System Global Area (SGA). We also made sure that HugePages were in use all the time while databases were running. The SOEDB1 + OLTPDB1 Database scalability test were run for at least 12 hours and ensured that results are consistent for the duration of the full run.

We ran the SwingBench scripts on each node to start SOEDB1 and OLTPDB1 database and generate AWR reports for each scenario as shown below:

- User Scalability

The table below illustrates the TPM for SOEDB1 + OLTPDB1 database user scale on all 8 node:

Users	TPM for SOEDB1	TPM for OLTPDB1	Total TPM	System Utilization (%)
100	231,174	231,840	463,014	5.7
200	455,244	453,300	908,544	11.5
300	654,330	659,574	1,313,904	17.9
400	737,880	850,128	1,588,008	21.7
600	931,932	977,994	1,909,926	29.7
800	1,095,000	1,108,542	2,203,542	31.6

The graph below illustrates the Total Transactions Per Minute (TPM) for SOEDB1 + OLTPDB1 database user scale on 8 nodes.

Figure 31 User Scalability – TPM & System Utilization

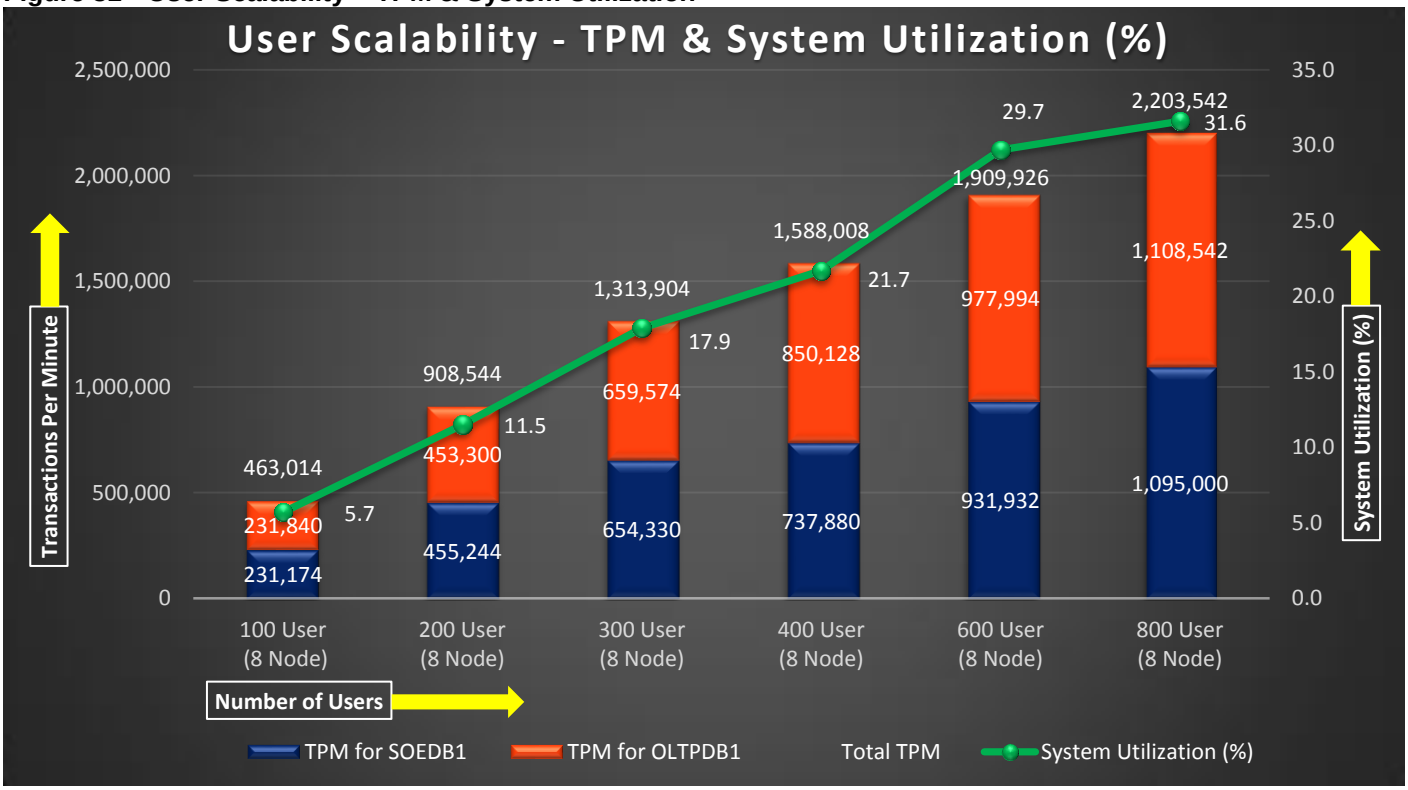

The table below illustrates the TPM for SOEDB1 + OLTPDB1 database user scale on all 8 node.

Users	IOPS for SOEDB1	IOPS for OLTPDB1	Total IOPS
100	31,732	32,701	64,433
200	58,774	60,936	119,710
300	81,514	82,965	164,479
400	92,883	100,591	193,474
600	119,714	120,655	240,369
800	133,672	125,647	259,319

The results were in line with prior assessments where the user scalability was almost linear till 600 users and beyond 600 users the rate of IOPS increase slowed down.

The graph below illustrates the total IOPS for SOEDB1 + OLTPDB1 database user scale on all the 8 Oracle RAC nodes.

Figure 32 User Scalability – IOPS

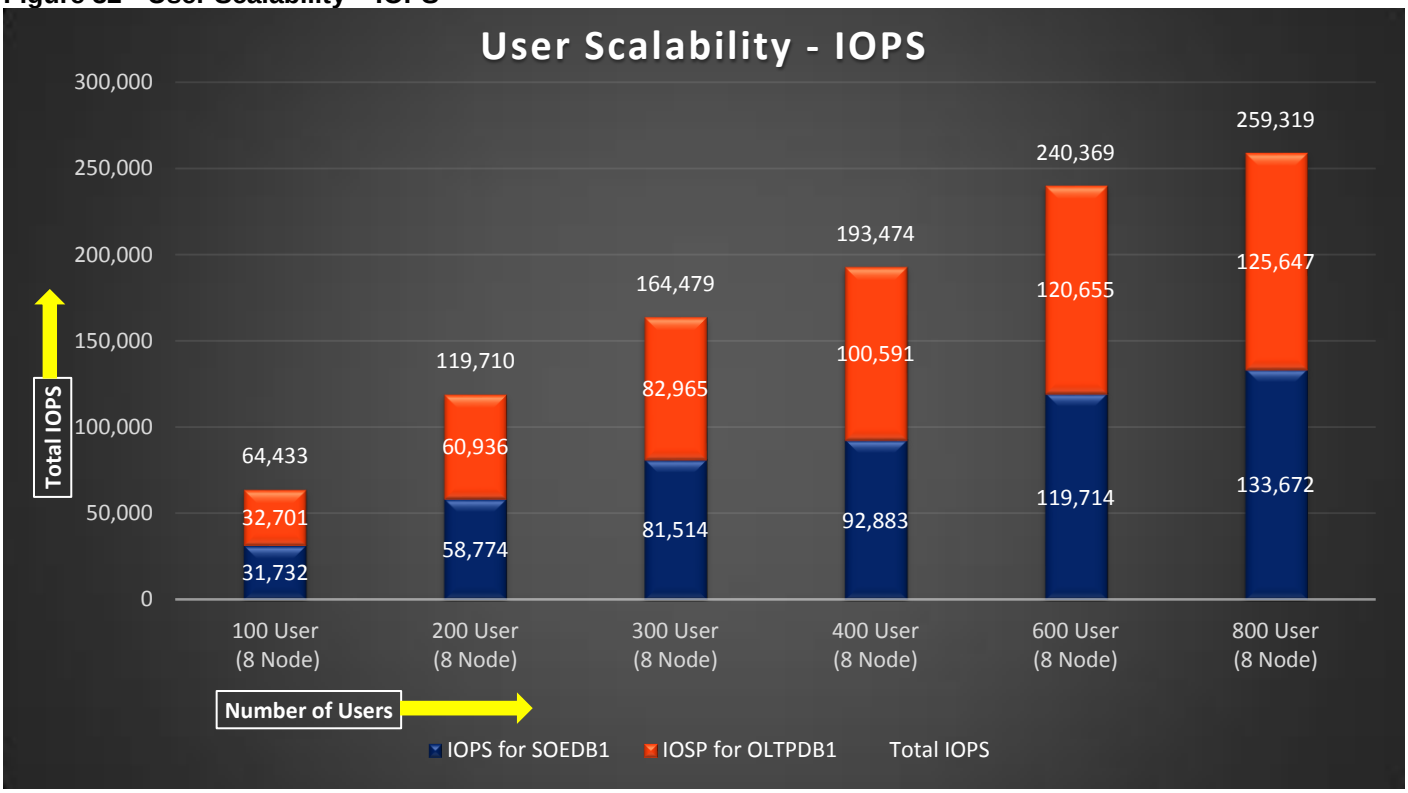

The screenshot shown below was captured from the 800 Users Scale Test scenario while running Swingbench workload on two database at the same time. The snapshot shows a section from 3-hour window of AWR Global report from the run that highlights Physical Reads/Sec, Physical Writes/Sec and Transactions per Seconds for each instance. Notice that IO load is distributed across all the cluster nodes performing workload operations.

Figure 33 User Scalability – System Statistics Per Second for SOEDB1 Database

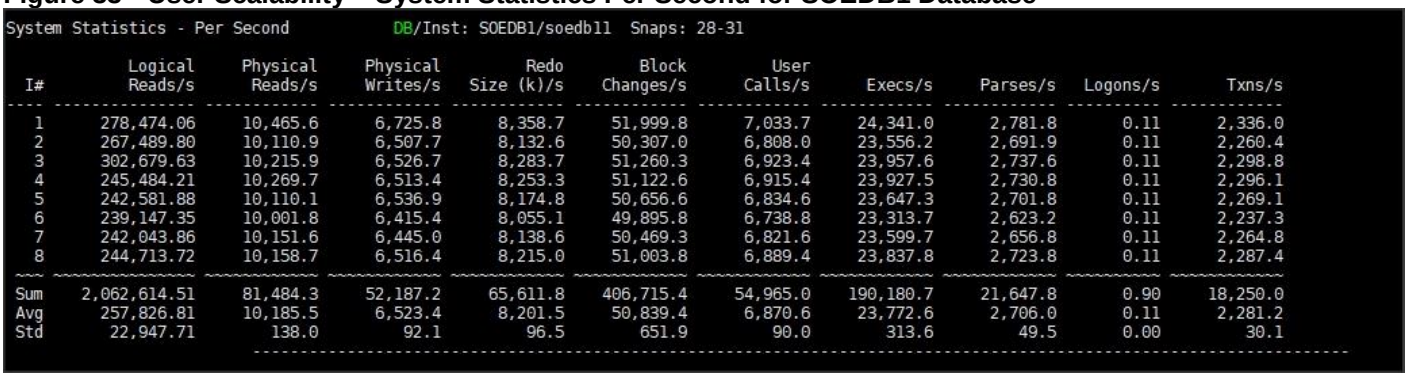

System Statistics - Per Second DB/Inst: SOEDB1/soedb11 Snaps: 28-31

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	278,474.06	10,465.6	6,725.8	8,358.7	51,999.8	7,033.7	24,341.0	2,781.8	0.11	2,336.0
2	267,489.80	10,110.9	6,507.7	8,132.6	50,307.0	6,808.0	23,556.2	2,691.9	0.11	2,260.4
3	302,679.63	10,215.9	6,526.7	8,283.7	51,260.3	6,923.4	23,957.6	2,737.6	0.11	2,298.8
4	245,484.21	10,269.7	6,513.4	8,253.3	51,122.6	6,915.4	23,927.5	2,730.8	0.11	2,296.1
5	242,581.88	10,110.1	6,536.9	8,174.8	50,656.6	6,834.6	23,647.3	2,701.8	0.11	2,269.1
6	239,147.35	10,001.8	6,415.4	8,055.1	49,895.8	6,738.8	23,313.7	2,623.2	0.11	2,237.3
7	242,043.86	10,151.6	6,445.0	8,138.6	50,469.3	6,821.6	23,599.7	2,656.8	0.11	2,264.8
8	244,713.72	10,158.7	6,516.4	8,215.0	51,003.8	6,889.4	23,837.8	2,723.8	0.11	2,287.4
Sum	2,062,614.51	81,484.3	52,187.2	65,611.8	406,715.4	54,965.0	190,180.7	21,647.8	0.90	18,250.0
Avg	257,826.81	10,185.5	6,523.4	8,201.5	50,839.4	6,870.6	23,772.6	2,706.0	0.11	2,281.2
Std	22,947.71	138.0	92.1	96.5	651.9	90.0	313.6	49.5	0.00	30.1

Figure 34 User Scalability – System Statistics Per Second for OLTPDB1 Database

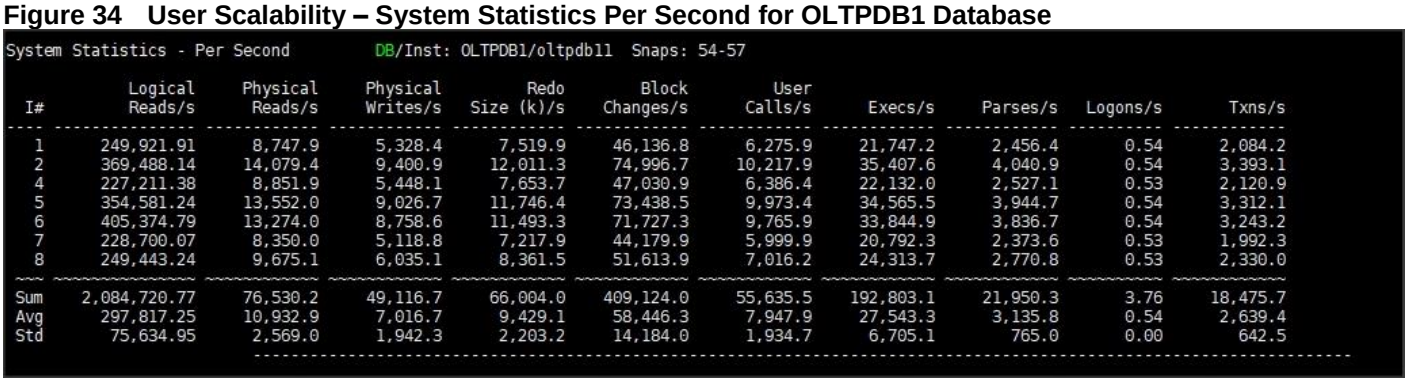

System Statistics - Per Second DB/Inst: OLTPDB1/oltpdb11 Snaps: 54-57

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	249,921.91	8,747.9	5,328.4	7,519.9	46,136.8	6,275.9	21,747.2	2,456.4	0.54	2,084.2
2	369,488.14	14,079.4	9,400.9	12,011.3	74,996.7	10,217.9	35,407.6	4,040.9	0.54	3,393.1
4	227,211.38	8,851.9	5,448.1	7,653.7	47,030.9	6,386.4	22,132.0	2,527.1	0.53	2,120.9
5	354,581.24	13,552.0	9,026.7	11,746.4	73,438.5	9,973.4	34,565.5	3,944.7	0.54	3,312.1
6	405,374.79	13,274.0	8,758.6	11,493.3	71,727.3	9,765.9	33,844.9	3,836.7	0.54	3,243.2
7	228,700.07	8,350.0	5,118.8	7,217.9	44,179.9	5,999.9	20,792.3	2,373.6	0.53	1,992.3
8	249,443.24	9,675.1	6,035.1	8,361.5	51,613.9	7,016.2	24,313.7	2,770.8	0.53	2,330.0
Sum	2,084,720.77	76,530.2	49,116.7	66,004.0	409,124.0	55,635.5	192,803.1	21,950.3	3.76	18,475.7
Avg	297,817.25	10,932.9	7,016.7	9,429.1	58,446.3	7,947.9	27,543.3	3,135.8	0.54	2,639.4
Std	75,634.95	2,569.0	1,942.3	2,203.2	14,184.0	1,934.7	6,705.1	765.0	0.00	642.5

The screenshot shown below shows latency for the same 800 Users Scale Test while swingbench test running for SOEDB1 Database.

Figure 35 User Scalability – Top Timed Events

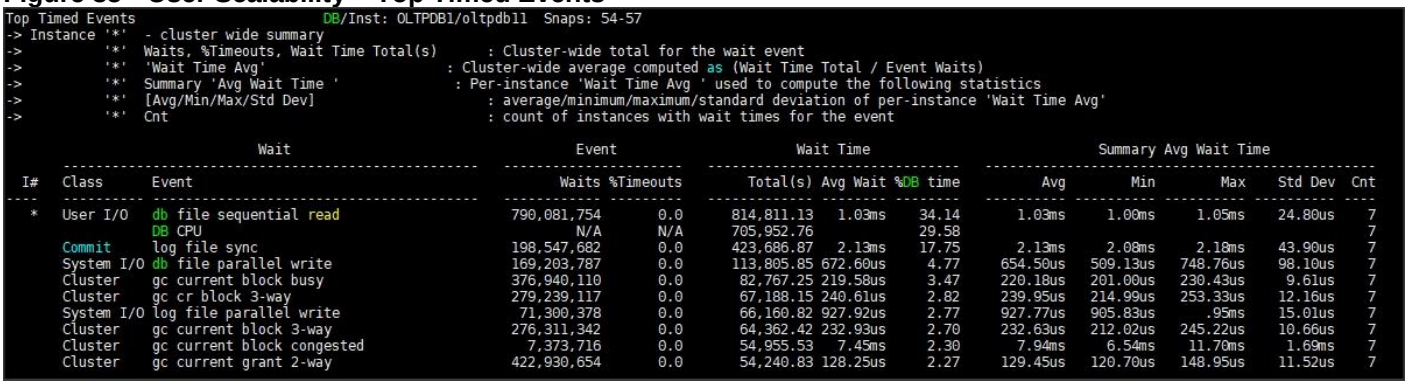

Top Timed Events DB/Inst: OLTPDB1/oltpdb11 Snaps: 54-57

-> Instance '*' - cluster wide summary
-> '*' Waits, %Timeouts, Wait Time Total(s) : Cluster-wide total for the wait event
-> '*' 'Wait Time Avg' : Cluster-wide average computed as (Wait Time Total / Event Waits)
-> '*' Summary 'Avg Wait Time ' : Per-instance 'Wait Time Avg ' used to compute the following statistics
-> '*' [Avg/Min/Max/Std Dev] : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
-> '*' Cnt : count of instances with wait times for the event

			Wait Event		Wait Time			Summary Avg Wait Time				
I#	Class	Event	Waits	%Timeouts	Total(s)	Avg Wait	%DB time	Avg	Min	Max	Std Dev	Cnt
*	User I/O	db file sequential read	790,081,754	0.0	814,811.13	1.03ms	34.14	1.03ms	1.00ms	1.05ms	24.80us	7
		DB CPU	N/A	N/A	705,952.76		29.58					7
	Commit	log file sync	198,547,682	0.0	423,686.87	2.13ms	17.75	2.13ms	2.08ms	2.18ms	43.90us	7
	System I/O	db file parallel write	169,203,787	0.0	113,805.85	672.60us	4.77	654.50us	509.13us	748.76us	98.10us	7
	Cluster	gc current block busy	376,940,110	0.0	82,767.25	219.58us	3.47	220.18us	201.00us	230.43us	9.61us	7
	Cluster	gc cr block 3-way	279,239,117	0.0	67,188.15	240.61us	2.82	239.95us	214.99us	253.33us	12.16us	7
	System I/O	log file parallel write	71,300,378	0.0	66,160.82	927.92us	2.77	927.77us	905.83us	.95ms	15.01us	7
	Cluster	gc current block 3-way	276,311,342	0.0	64,362.42	232.93us	2.70	232.63us	212.02us	245.22us	10.66us	7
	Cluster	gc current block congested	7,373,716	0.0	54,955.53	7.45ms	2.30	7.94ms	6.54ms	11.70ms	1.69ms	7
	Cluster	gc current grant 2-way	422,930,654	0.0	54,240.83	128.25us	2.27	129.45us	120.70us	148.95us	11.52us	7

The screenshot shown below shows the Interconnect device statistics for the same 800 Users Scale Test while swingbench test running on both OLTP databases.

Figure 36 Interconnect Device Statistics

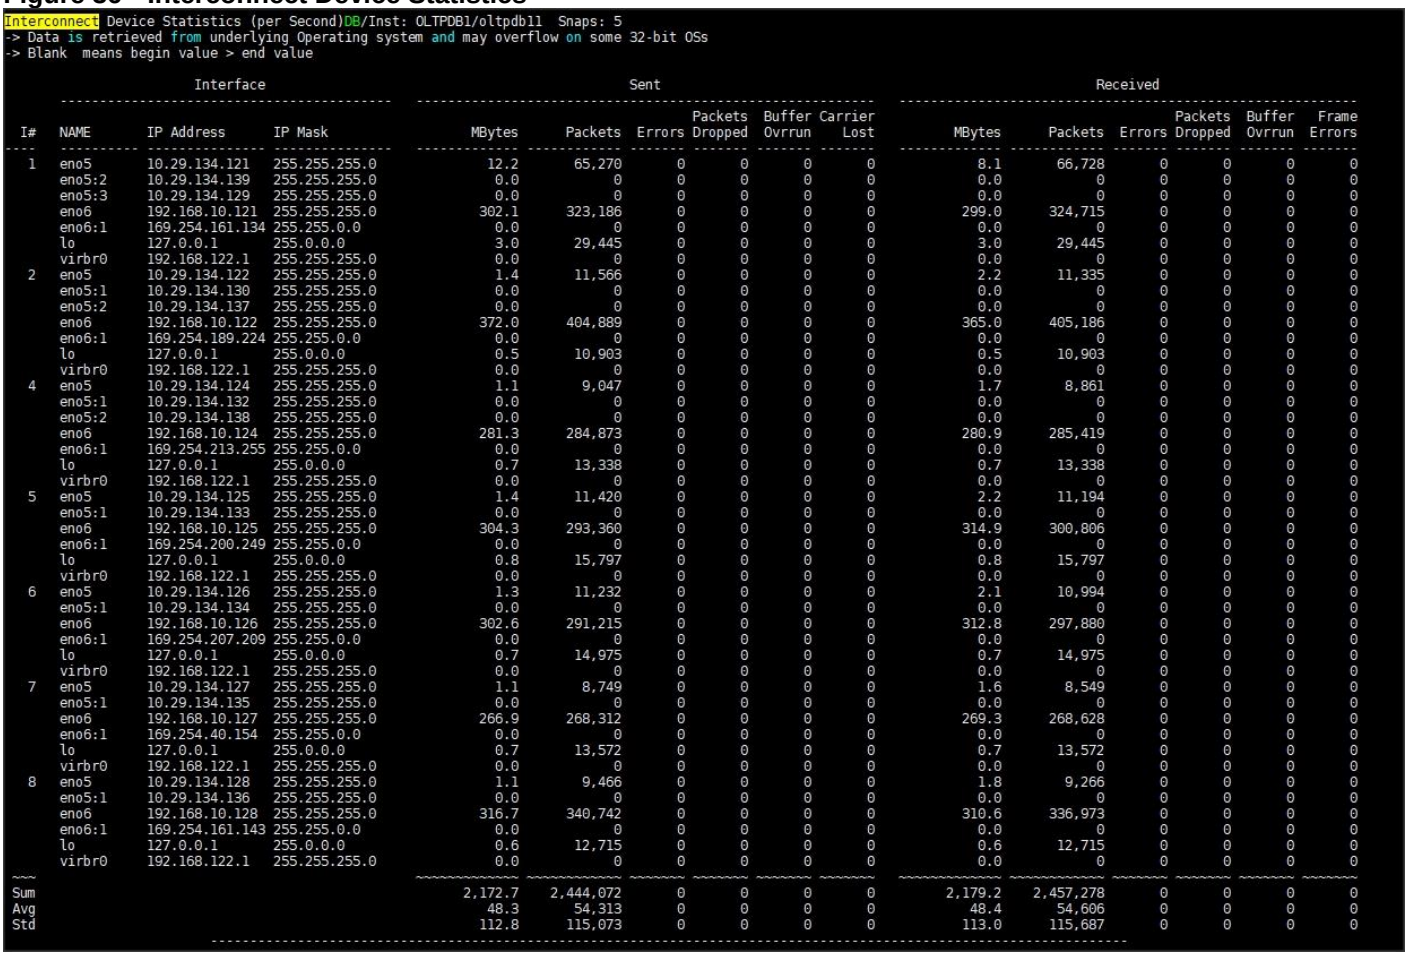

Interconnect Device Statistics (per Second)DB/Inst: OLTPDB1/oltpdb11 Snaps: 5
-> Data is retrieved from underlying Operating system and may overflow on some 32-bit OSs
-> Blank means begin value > end value

	Interface			Sent						Received					
I#	NAME	IP Address	IP Mask	MBytes	Packets	Errors	Packets Dropped	Buffer Ovrrun	Carrier Lost	MBytes	Packets	Errors	Packets Dropped	Buffer Ovrrun	Frame Errors
1	eno5	10.29.134.121	255.255.255.0	12.2	65,270	0	0	0	0	8.1	66,728	0	0	0	0
	eno5:2	10.29.134.139	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno5:3	10.29.134.129	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.121	255.255.255.0	302.1	323,186	0	0	0	0	299.0	324,715	0	0	0	0
	eno6:1	169.254.161.134	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	3.0	29,445	0	0	0	0	3.0	29,445	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
2	eno5	10.29.134.122	255.255.255.0	1.4	11,566	0	0	0	0	2.2	11,335	0	0	0	0
	eno5:1	10.29.134.130	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno5:2	10.29.134.137	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.122	255.255.255.0	372.0	404,889	0	0	0	0	365.0	405,186	0	0	0	0
	eno6:1	169.254.189.224	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.5	10,903	0	0	0	0	0.5	10,903	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
4	eno5	10.29.134.124	255.255.255.0	1.1	9,047	0	0	0	0	1.7	8,861	0	0	0	0
	eno5:1	10.29.134.132	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno5:2	10.29.134.138	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.124	255.255.255.0	281.3	284,873	0	0	0	0	280.9	285,419	0	0	0	0
	eno6:1	169.254.213.255	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.7	13,338	0	0	0	0	0.7	13,338	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
5	eno5	10.29.134.125	255.255.255.0	1.4	11,420	0	0	0	0	2.2	11,194	0	0	0	0
	eno5:1	10.29.134.133	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.125	255.255.255.0	304.3	293,360	0	0	0	0	314.9	300,806	0	0	0	0
	eno6:1	169.254.200.249	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.8	15,797	0	0	0	0	0.8	15,797	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
6	eno5	10.29.134.126	255.255.255.0	1.3	11,232	0	0	0	0	2.1	10,994	0	0	0	0
	eno5:1	10.29.134.134	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.126	255.255.255.0	302.6	291,215	0	0	0	0	312.8	297,880	0	0	0	0
	eno6:1	169.254.207.209	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.7	14,975	0	0	0	0	0.7	14,975	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
7	eno5	10.29.134.127	255.255.255.0	1.1	8,749	0	0	0	0	1.6	8,549	0	0	0	0
	eno5:1	10.29.134.135	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.127	255.255.255.0	266.9	268,312	0	0	0	0	269.3	268,628	0	0	0	0
	eno6:1	169.254.40.154	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.7	13,572	0	0	0	0	0.7	13,572	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
8	eno5	10.29.134.128	255.255.255.0	1.1	9,466	0	0	0	0	1.8	9,266	0	0	0	0
	eno5:1	10.29.134.136	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	eno6	192.168.10.128	255.255.255.0	316.7	340,742	0	0	0	0	310.6	336,973	0	0	0	0
	eno6:1	169.254.161.143	255.255.0.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
	lo	127.0.0.1	255.0.0.0	0.6	12,715	0	0	0	0	0.6	12,715	0	0	0	0
	virbr0	192.168.122.1	255.255.255.0	0.0	0	0	0	0	0	0.0	0	0	0	0	0
Sum				2,172.7	2,444,072	0	0	0	0	2,179.2	2,457,278	0	0	0	0
Avg				48.3	54,313	0	0	0	0	48.4	54,606	0	0	0	0
Std				112.8	115,073	0	0	0	0	113.0	115,687	0	0	0	0

DSSDB1 Database Performance

DSS database workloads are generally sequential in nature, read intensive and exercise large IO size. DSS database workload runs a small number of users that typically exercise extremely complex queries that run for hours. We configured 9 TB of DSS database by loading Swingbench sh schema into Datafile Tablespace. DSS Database activity is captured for four Oracle RAC Instances using Oracle Enterprise Manager for 24 hours workload test.

Figure 37 DSS Performance – Bandwidth

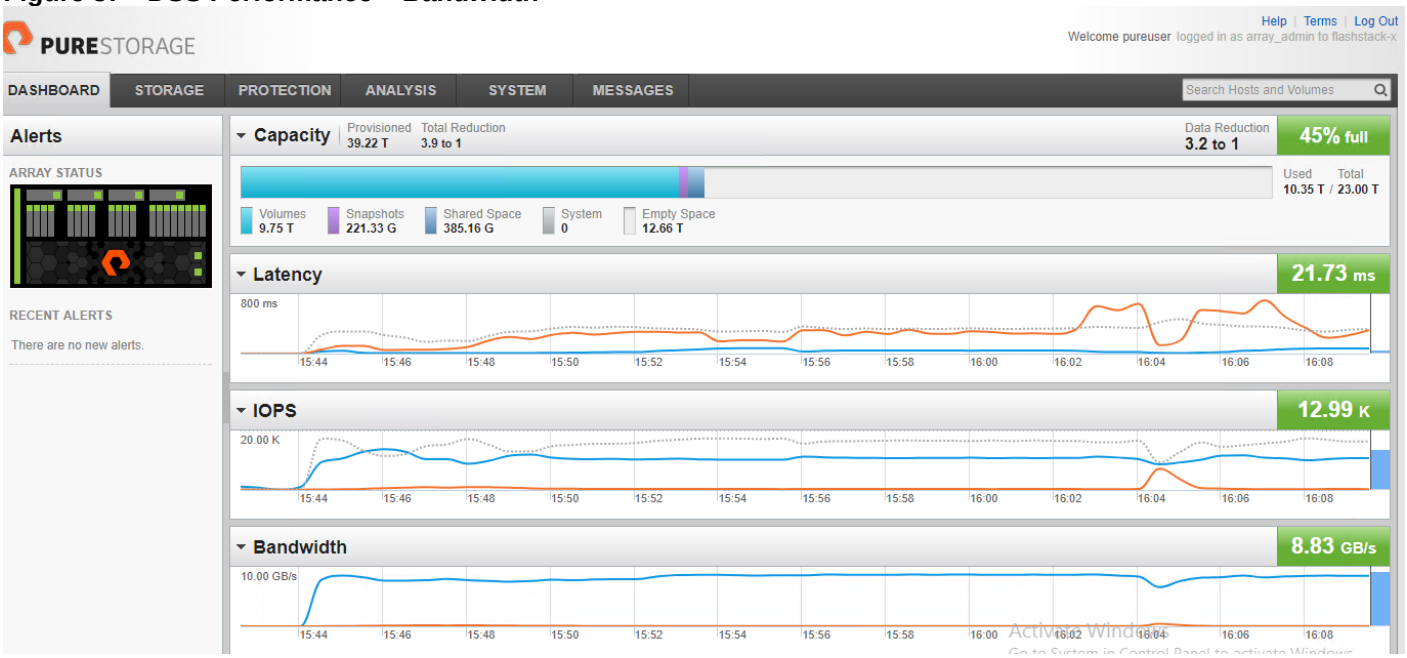

For 24 hours DSS workload test, we observed the total sustained IO bandwidth was up to 8.7 GB/sec after the initial ramp up workload. As indicated on the charts, the IO was consistent throughout the run and we did not observe any significant dips in performance for complete period of time.

The screenshot shown below shows latency while swingbench test running for DSSDB1 Database.

Figure 38 DSS Performance – Latency

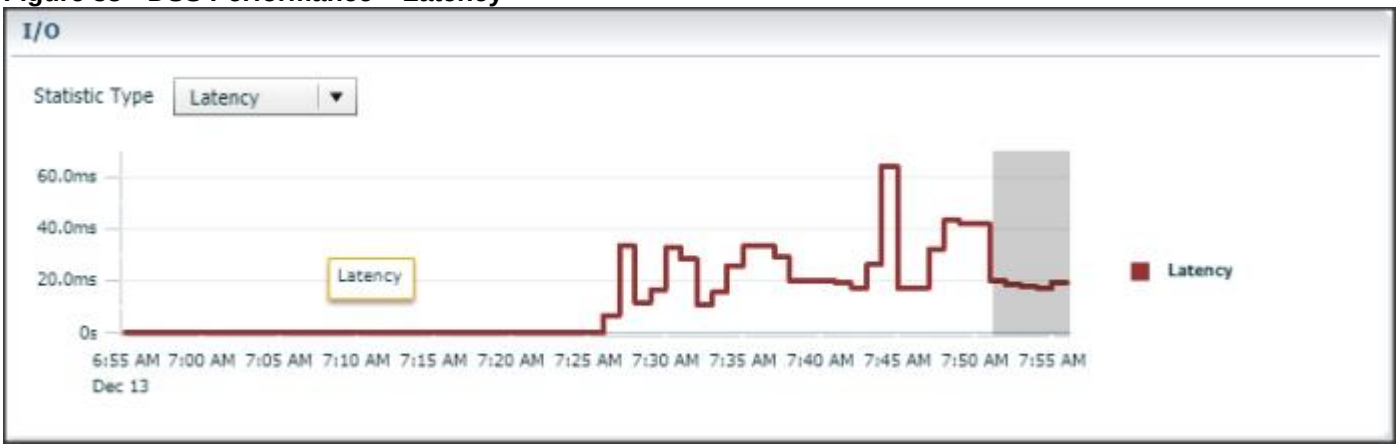

The screenshot shown below shows system utilization of each instance while swingbench test running for DSSDB1 Database.

Figure 39 DSS Performance – All Nodes System Utilization

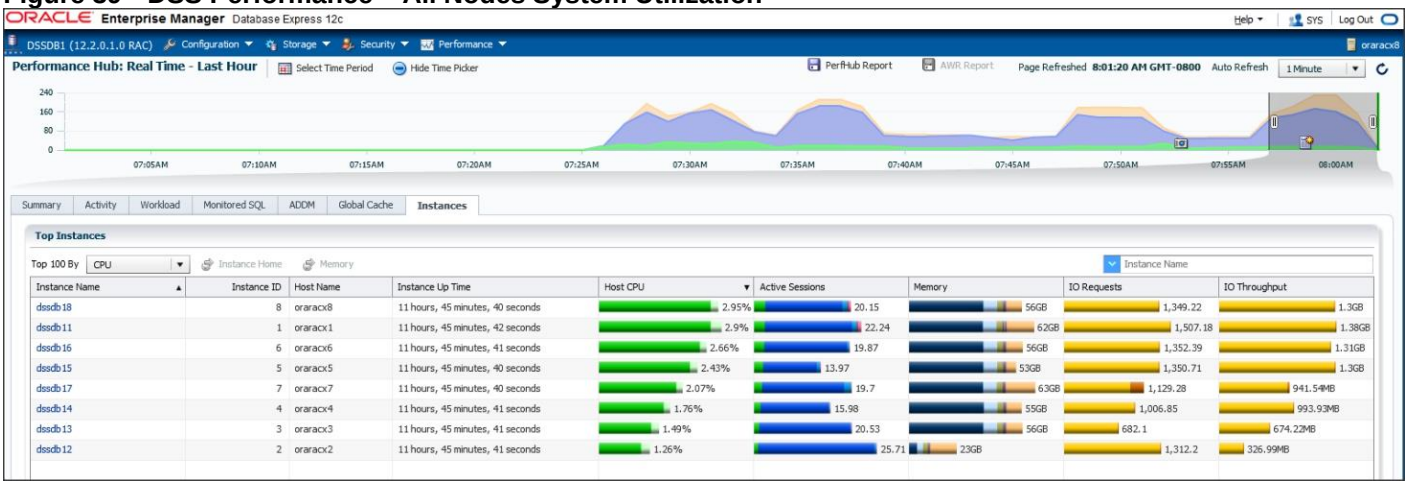

Mixed Workload

The next test is to run both OLTP and DSS database workloads simultaneously. This test will make sure that configuration in this test is able to sustain small random queries presented via OLTP database along with large and sequential transactions submitted via DSS database workload.

The screenshot shown below shows IOPS, Latency and Throughput of the FlashArray //X70 system while all three databases (2 OLTP and 1 DSS) running swingbench workload.

Figure 40 All three Database Performance

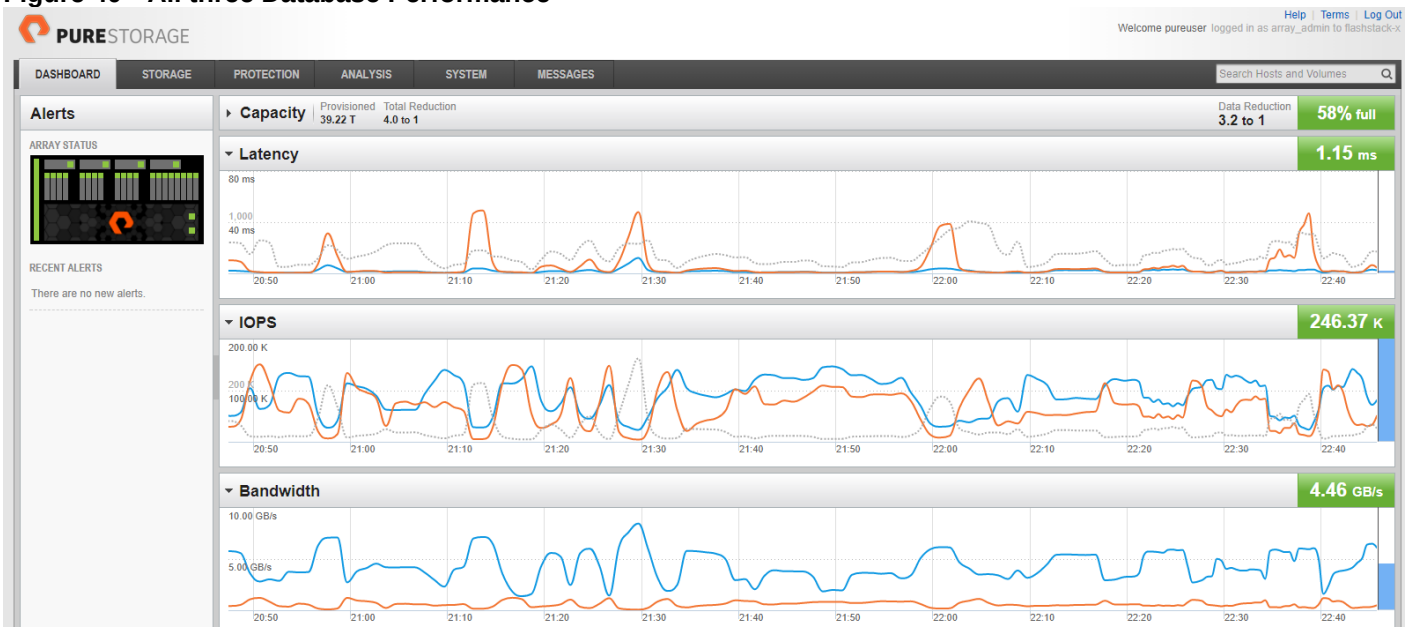

DSSDB1 Database activity was captured for all eight node Oracle RAC Instances using Oracle Enterprise Manager for 24 hours mixed workload test. The screenshot below shows Throughput GB/s for DSSDB1 database.

Figure 41 DSS Performance – Bandwidth

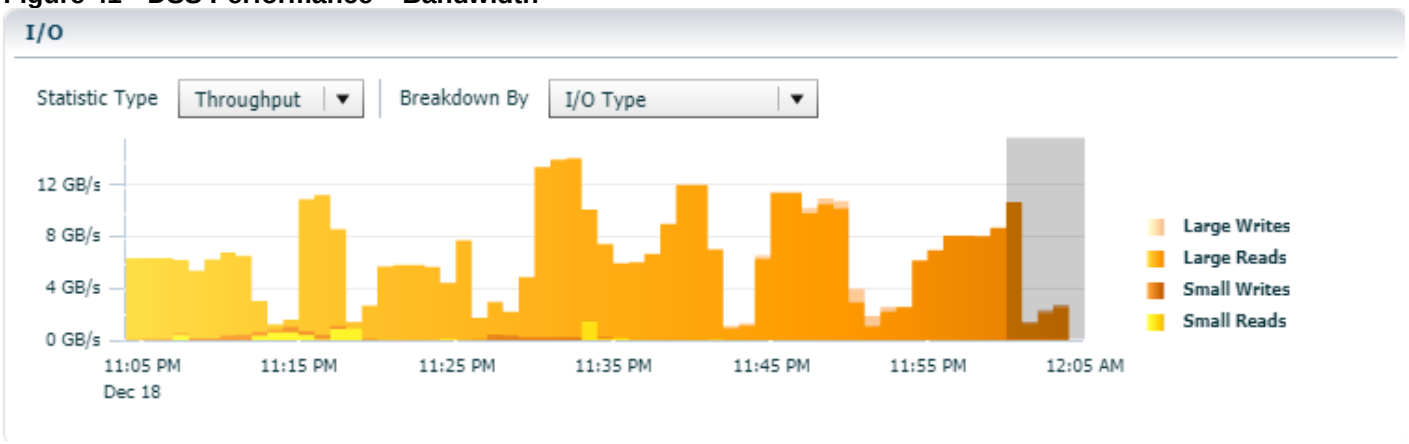

SOEDB1 Database activity was captured for all eight node Oracle RAC Instances using Oracle Enterprise Manager for 24 hours mixed workload test. The screenshot below shows Physical I/Os per second for SOEDB1 database.

Figure 42 SOEDB1 Performance – IOPS

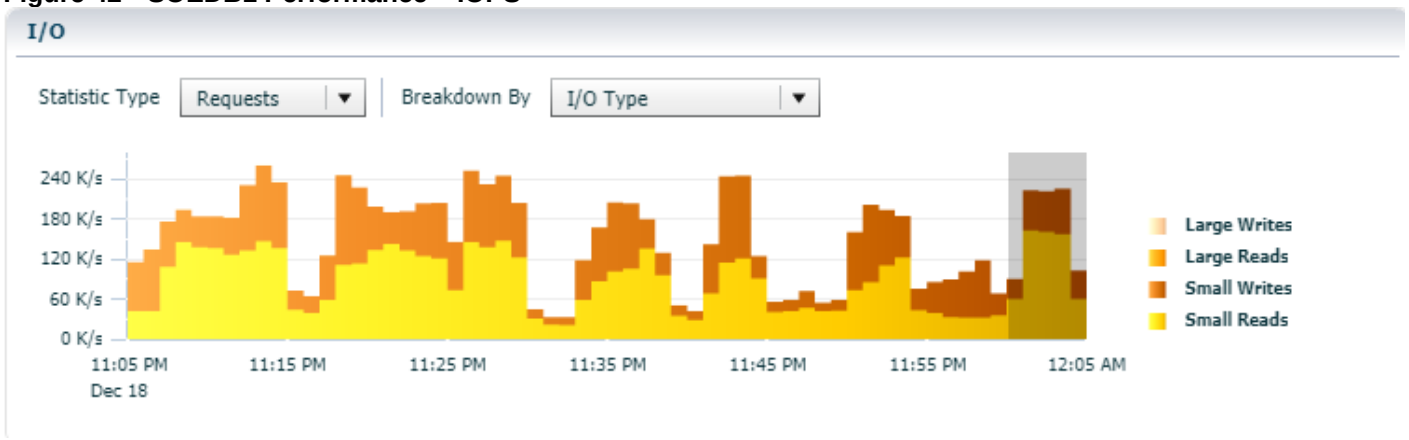

OLTPDB1 Database activity was captured for all eight node Oracle RAC Instances using Oracle Enterprise Manager for 24 hours mixed workload test. The screenshot below shows Physical I/Os per second for OLTPDB1 database.

Figure 43 OLTPDB1 Performance – IOPS

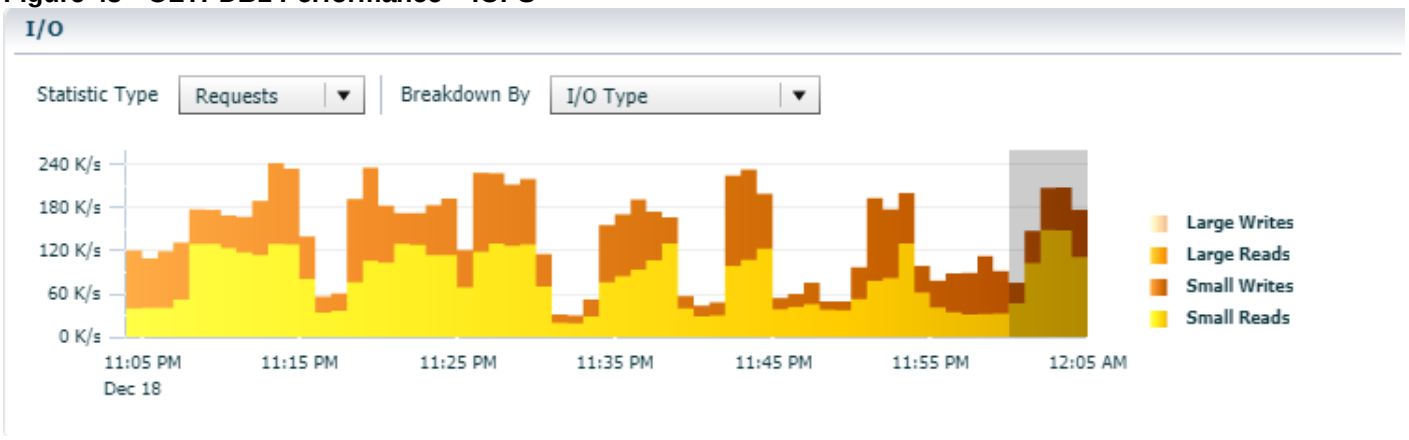

The Mixed workload results were in line with the simultaneous calibrate IO tests that were performed earlier and clearly showcases the level of performance that can be achieved with this FlashStack solution.

Resiliency and Failure Tests

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

Table 12 Hardware Failover Tests

Scenario	Test	Status
Test 1 – UCS FI – A Failure	Run the system on Full Database work Load. Power Off Fabric Interconnect – A and check network traffic on Fabric Interconnect – B.	Fabric Interconnect Failover did not cause any disruption to Private, Public and Storage Traffic
Test 2 – UCS FI – B Failure	Run the system on Full Database work Load. Power Off Fabric Interconnect – B and check network traffic on Fabric Interconnect – A	Fabric Interconnect Failover did not cause any disruption to Private, Public and Storage Traffic
Test 3 – UCS Nexus Switch – A Failure	Run the system on Full Database work Load. Power Off Nexus Switch – A and check network traffic on Nexus Switch – B.	Nexus Switch Failover did not cause any disruption to Private and Public network Traffic
Test 4 – UCS Nexus Switch – B Failure	Run the system on Full Database work Load. Power Off Nexus Switch – B and check network traffic on Nexus Switch – A.	Nexus Switch Failover did not cause any disruption to Private and Public network Traffic
Test 5 – UCS MDS Switch – A Failure	Run the system on Full Database work Load. Power Off MDS Switch – A and check storage traffic on MDS Switch – B	MDS Switch Failover did not cause any disruption to Storage network Traffic
Test 6 – UCS MDS Switch – B Failure	Run the system on Full Database work Load. Power Off MDS Switch – B and check storage traffic on MDS Switch – A	MDS Switch Failover did not cause any disruption to Storage network Traffic
Test 7 – UCS Chassis 1 and Chassis 2 IOM Links Failure	Run the system on full Database work Load. Disconnect two links from each Chassis 1 IOM and Chassis 2 IOM by pulling it out and reconnect it after 5 minutes.	No disruption in network traffic.

Figure 44 FlashStack Infrastucture

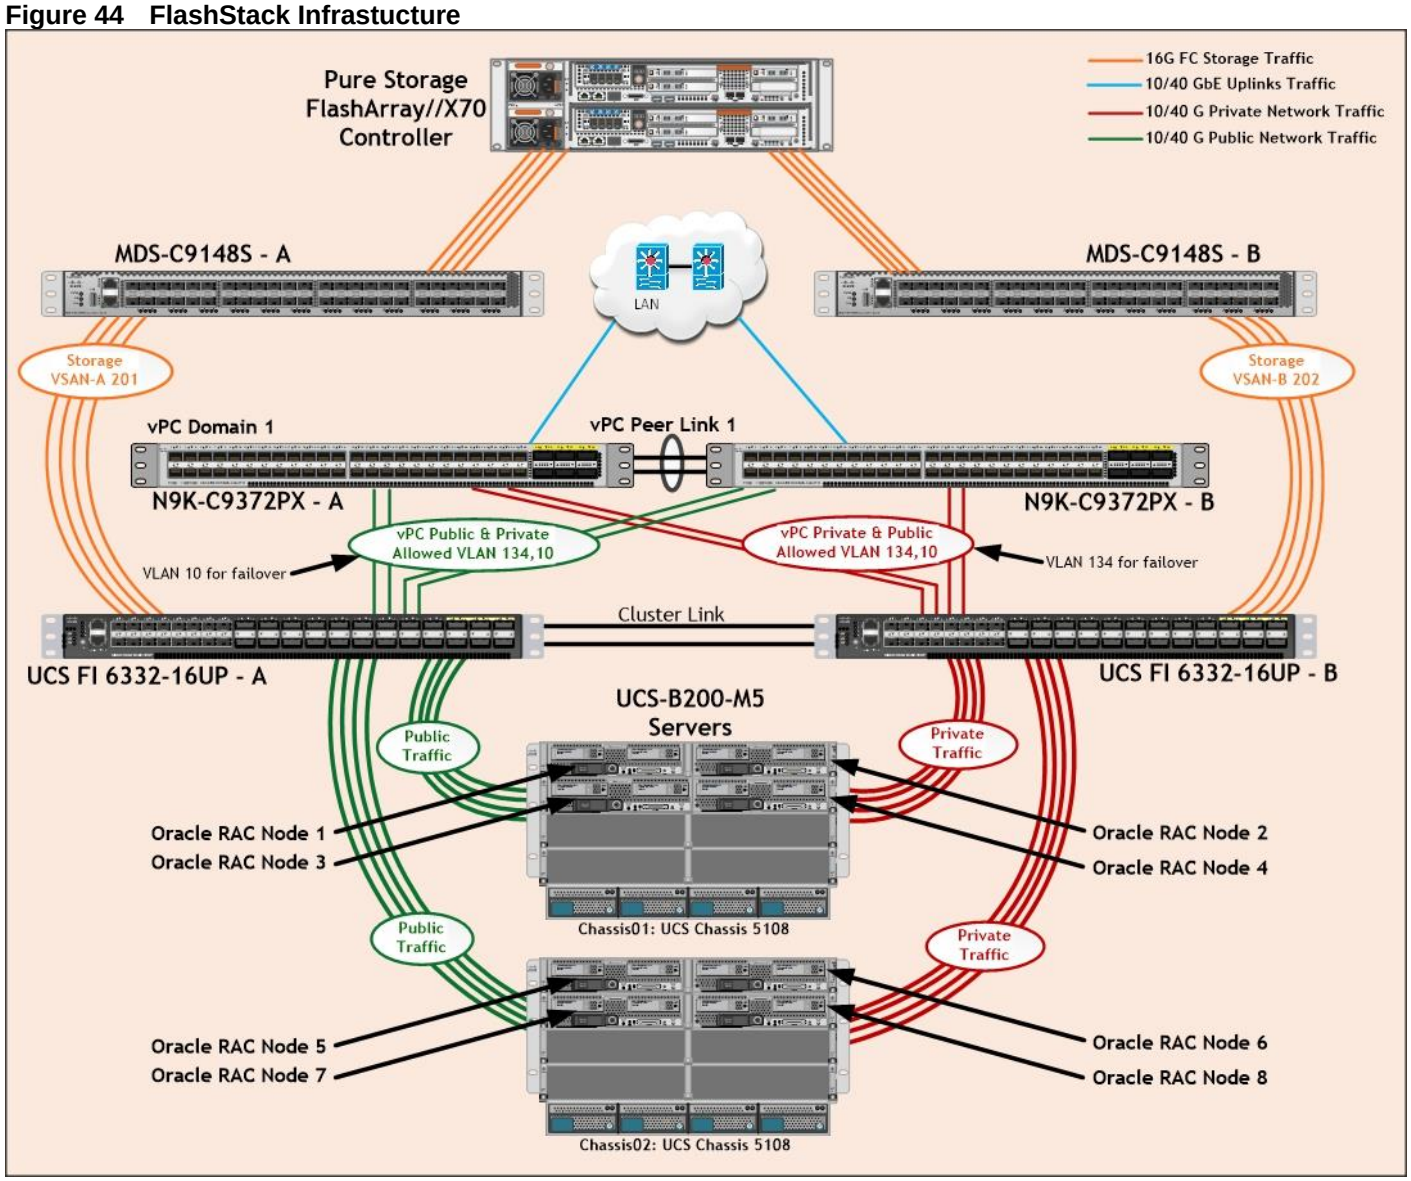

0 illustrates the FlashStack solution infrastructure diagram under normal operating conditions. Cisco UCS 6332-16UP Fabric Interconnects carries both storage and network traffic from the blades with the help of Cisco Nexus 9372PX-E and Cisco MDS 9148S switches. Two virtual Port-Channels (vPCs) are configured to provide public network and private network paths for the blades to northbound switches. Eight (four per chassis) links go to Fabric Interconnect – A. Similarly, eight links go to Fabric Interconnect – B. Fabric Interconnect – A links are used for Oracle Public network traffic shown as green lines. Fabric Interconnect – B links are used for Oracle Private Interconnect traffic shown as red lines. FC Storage access from Fabric Interconnect – A and Fabric Interconnect – B shown as an orange line.

The figure below shows a complete infrastructure details of MAC address, VLAN information and Server connections for Cisco UCS Fabric Interconnect – A switch before failover test.

Log into Cisco Fabric Interconnect – A and “connect nxos a” then type “show mac address-table” to see all VLAN connection on Fabric Interconnect – A as shown below:

Figure 45 Fabric Interconnect – A Network Traffic

```
ORARAC-X-FI-A(nxos)#
ORARAC-X-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 134      0025.b5cc.0000    static    0          F    F  Veth1029
* 134      0025.b5cc.0001    static    0          F    F  Veth1037
* 134      0025.b5cc.0002    static    0          F    F  Veth1133
* 134      0025.b5cc.0003    static    0          F    F  Veth1053
* 134      0025.b5cc.0004    static    0          F    F  Veth1061
* 134      0025.b5cc.0005    static    0          F    F  Veth1069
* 134      0025.b5cc.0006    static    0          F    F  Veth1077
* 134      0025.b5cc.0007    static    0          F    F  Veth1085
* 4044     025d.7358.c8b3    dynamic   0          F    F  Eth2/1/33
```

As shown in the above screenshot, Fabric Interconnect – A carry Oracle Public Network traffic on VLAN 134 under normal operating conditions before failover test.

Log in to Cisco Fabric Interconnect – B and “connect nxos b” then type “show mac address-table” to see all VLAN connection on Fabric – B as shown in the screenshot below:

Figure 46 Fabric Interconnect – B Network Traffic

```
ORARAC-X-FI-B(nxos)#
ORARAC-X-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b5dd.0000    static    0          F    F  Veth1031
* 10       0025.b5dd.0001    static    0          F    F  Veth1039
* 10       0025.b5dd.0002    static    0          F    F  Veth1135
* 10       0025.b5dd.0003    static    0          F    F  Veth1055
* 10       0025.b5dd.0004    static    0          F    F  Veth1063
* 10       0025.b5dd.0005    static    0          F    F  Veth1071
* 10       0025.b5dd.0006    static    0          F    F  Veth1079
* 10       0025.b5dd.0007    static    0          F    F  Veth1087
* 4044     025d.7358.c8b2    dynamic   0          F    F  Eth2/1/33
```

As shown in the above screenshot, Fabric Interconnect – B carry Oracle Private Network traffic on VLAN 10 under normal operating conditions before failover test.

Test 1 – Cisco UCS 6332-16UP Fabric Interconnect – A Failure Test

We conducted a hardware failure test on Fabric Interconnect – A by disconnecting power cable to the switch as explained below.

The figure below illustrates how during Fabric Interconnect – A switch failure, the respective blades (ORARACX1, ORARACX2, ORARACX3 and ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 and ORARACX8) on chassis 2 will fail over the public network interface MAC addresses and its VLAN network traffic to fabric interconnect – B.

Figure 47 Fabric Interconnect – A Failure

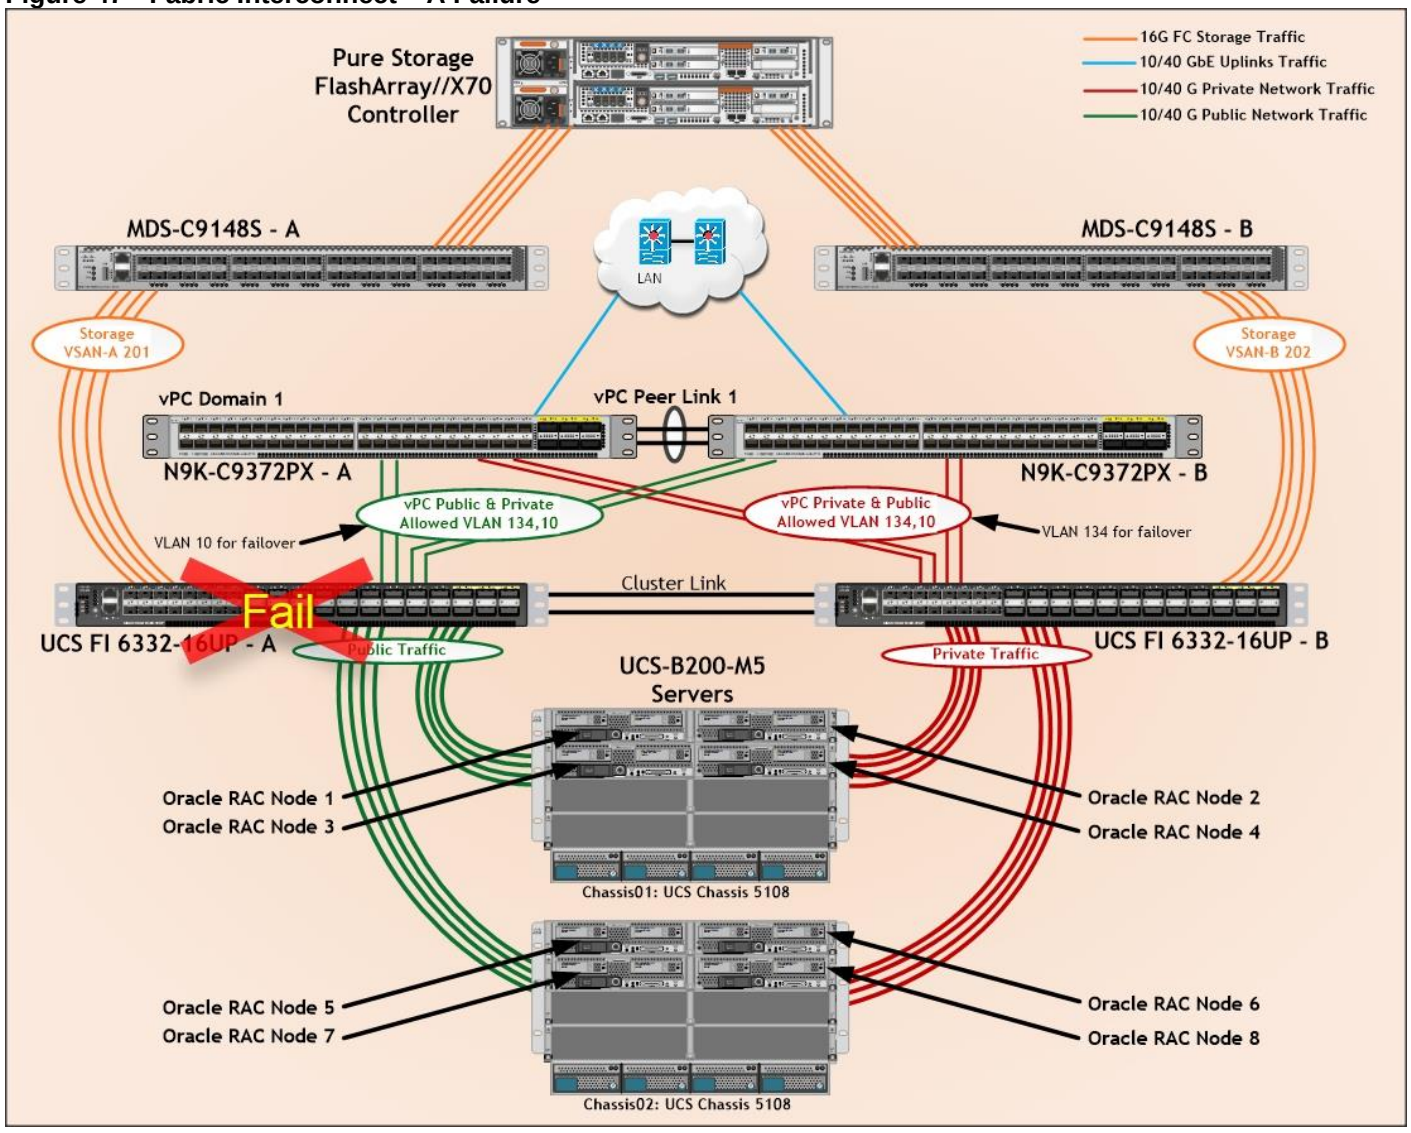

Unplug power cable from Fabric Interconnect – A and check the MAC address and VLAN information on Cisco UCS Fabric Interconnect – B.

Figure 48 Fabric Interconnect – B Network Traffic

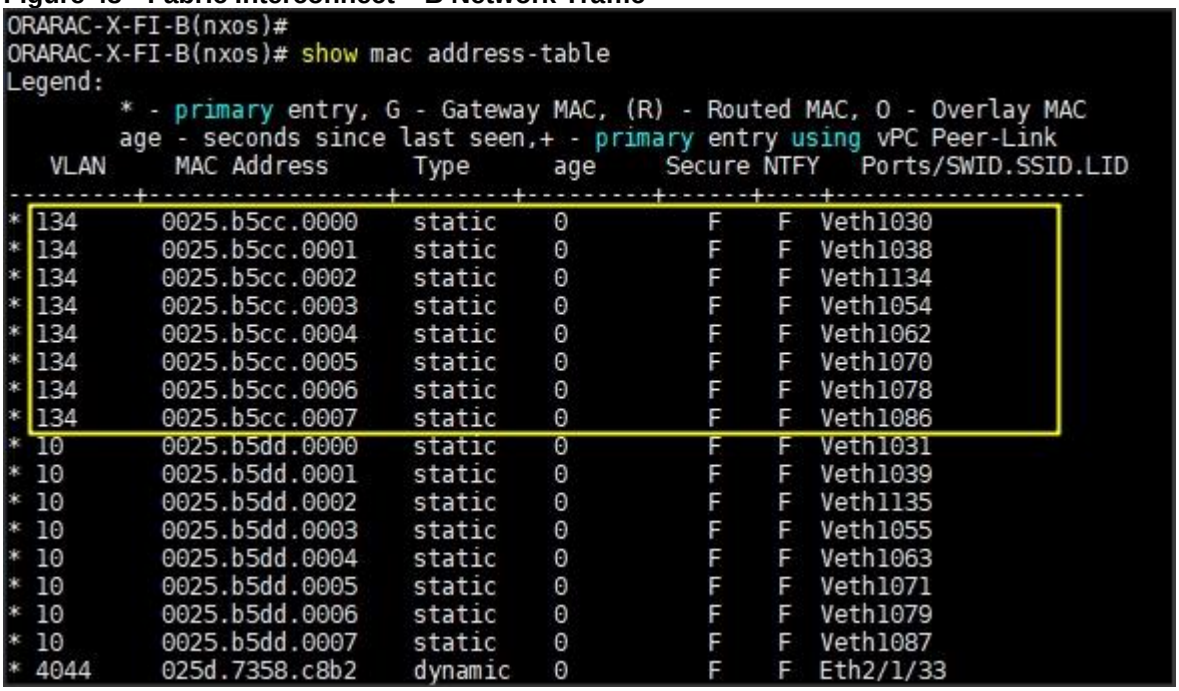

```
ORARAC-X-FI-B(nxos)#
ORARAC-X-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 134      0025.b5cc.0000    static    0          F    F  Veth1030
* 134      0025.b5cc.0001    static    0          F    F  Veth1038
* 134      0025.b5cc.0002    static    0          F    F  Veth1134
* 134      0025.b5cc.0003    static    0          F    F  Veth1054
* 134      0025.b5cc.0004    static    0          F    F  Veth1062
* 134      0025.b5cc.0005    static    0          F    F  Veth1070
* 134      0025.b5cc.0006    static    0          F    F  Veth1078
* 134      0025.b5cc.0007    static    0          F    F  Veth1086
* 10       0025.b5dd.0000    static    0          F    F  Veth1031
* 10       0025.b5dd.0001    static    0          F    F  Veth1039
* 10       0025.b5dd.0002    static    0          F    F  Veth1135
* 10       0025.b5dd.0003    static    0          F    F  Veth1055
* 10       0025.b5dd.0004    static    0          F    F  Veth1063
* 10       0025.b5dd.0005    static    0          F    F  Veth1071
* 10       0025.b5dd.0006    static    0          F    F  Veth1079
* 10       0025.b5dd.0007    static    0          F    F  Veth1087
* 4044     025d.7358.c8b2    dynamic   0          F    F  Eth2/1/33
```

We noticed in the figure above, when the Fabric Interconnect – A failed, it would route all the Public Network traffic of VLAN 134 to Fabric Interconnect – B. So Fabric Interconnect – A Failover did not cause any disruption to Private, Public and Storage Network Traffic.

After plug back power cable to Fabric Interconnect – A Switch, the respective blades (ORARACX1, ORARACX2, ORARACX3 & ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 & ORARACX8) on chassis 2 will route back the MAC addresses and its VLAN traffic to Fabric Interconnect – A.

The figure below shows details of MAC address, VLAN information and Server connections for Cisco UCS Fabric Interconnect – A switch under normal operating condition.

Figure 49 Fabric Interconnect – A Network Traffic

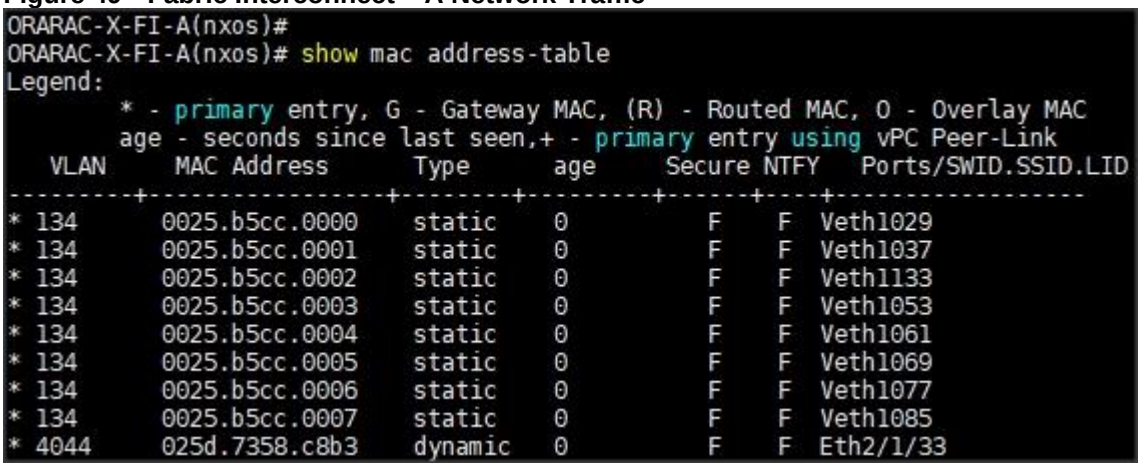

```
ORARAC-X-FI-A(nxos)#
ORARAC-X-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 134      0025.b5cc.0000    static    0          F    F  Veth1029
* 134      0025.b5cc.0001    static    0          F    F  Veth1037
* 134      0025.b5cc.0002    static    0          F    F  Veth1133
* 134      0025.b5cc.0003    static    0          F    F  Veth1053
* 134      0025.b5cc.0004    static    0          F    F  Veth1061
* 134      0025.b5cc.0005    static    0          F    F  Veth1069
* 134      0025.b5cc.0006    static    0          F    F  Veth1077
* 134      0025.b5cc.0007    static    0          F    F  Veth1085
* 4044     025d.7358.c8b3    dynamic   0          F    F  Eth2/1/33
```

The figure below shows details of MAC address, VLAN information and Server connections for Cisco UCS Fabric Interconnect – B switch under normal operating condition.

Figure 50 Fabric Interconnect – B Network Traffic

```
ORARAC-X-FI-B(nxos)#
ORARAC-X-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b5dd.0000    static    0          F    F  Veth1031
* 10       0025.b5dd.0001    static    0          F    F  Veth1039
* 10       0025.b5dd.0002    static    0          F    F  Veth1135
* 10       0025.b5dd.0003    static    0          F    F  Veth1055
* 10       0025.b5dd.0004    static    0          F    F  Veth1063
* 10       0025.b5dd.0005    static    0          F    F  Veth1071
* 10       0025.b5dd.0006    static    0          F    F  Veth1079
* 10       0025.b5dd.0007    static    0          F    F  Veth1087
* 4044     025d.7358.c8b2    dynamic   0          F    F  Eth2/1/33
```

Test 2 – Cisco UCS 6332-16UP Fabric Interconnect – B Failure Test

The below figure illustrates how during Fabric Interconnect – B switch failure, the respective blades (ORARACX1, ORARACX2, ORARACX3 and ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 and ORARACX8) on chassis 2 will fail over the public network interface MAC addresses and its VLAN network traffic to fabric interconnect – A.

Figure 51 Fabric Interconnect – B Failure

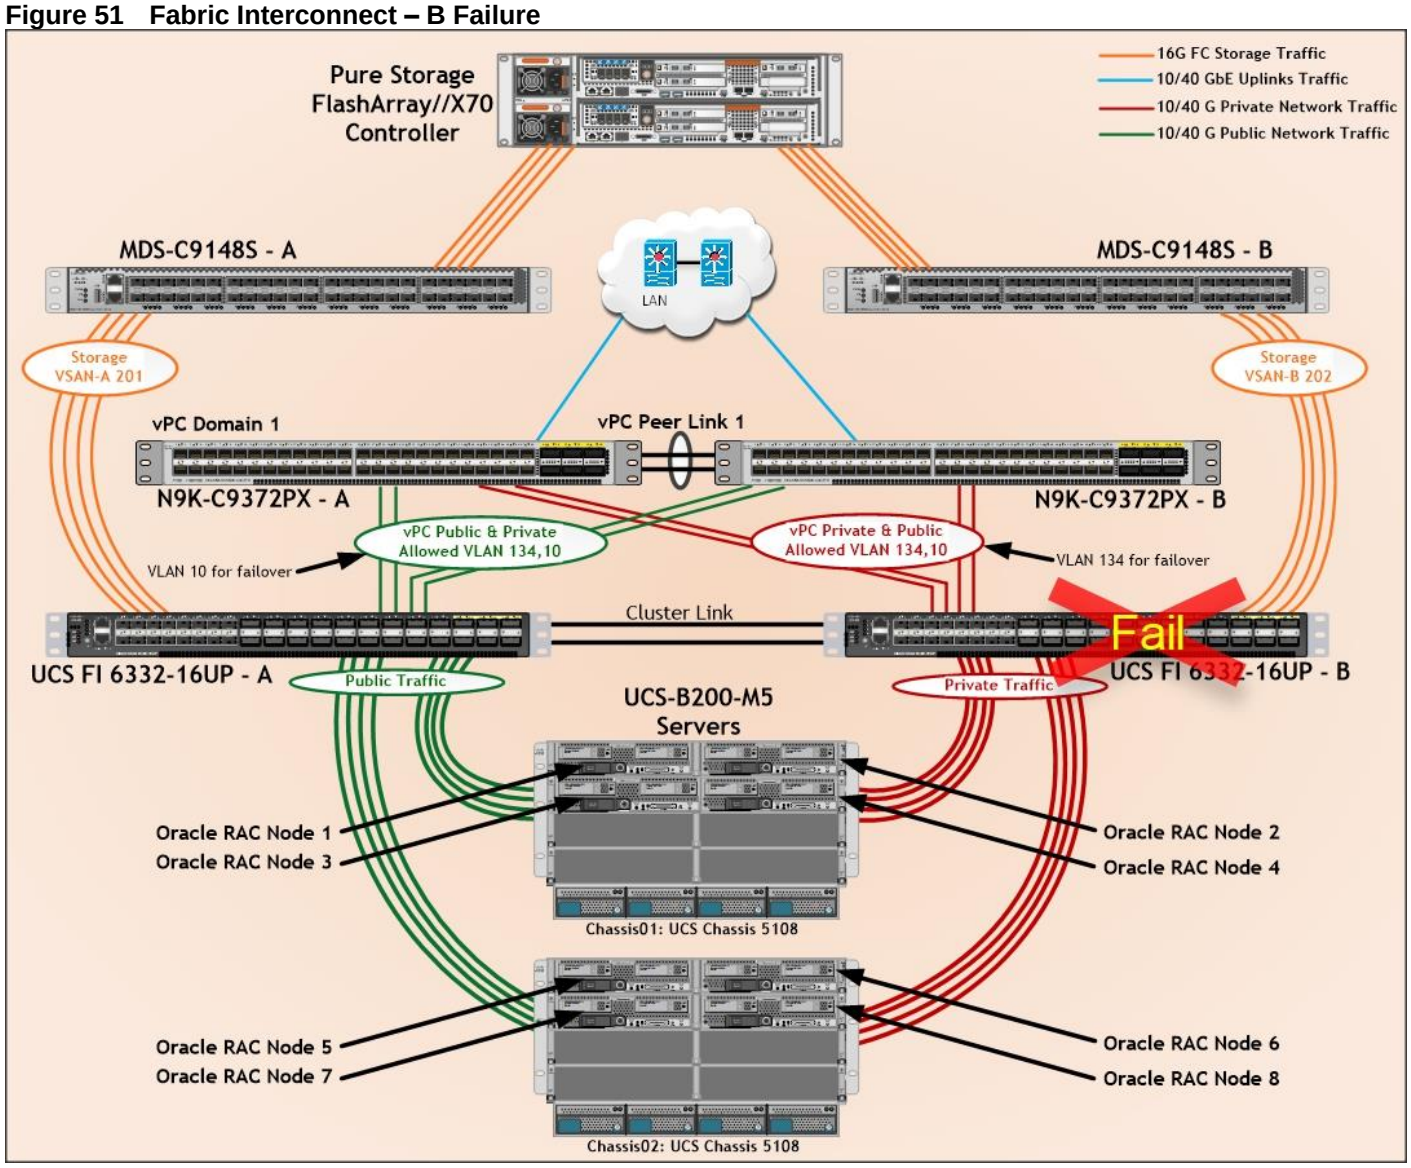

Unplug power cable from Fabric Interconnect – B and check the MAC address and VLAN information on Cisco UCS Fabric Interconnect – A.

Figure 52 Fabric Interconnect – A Network Traffic

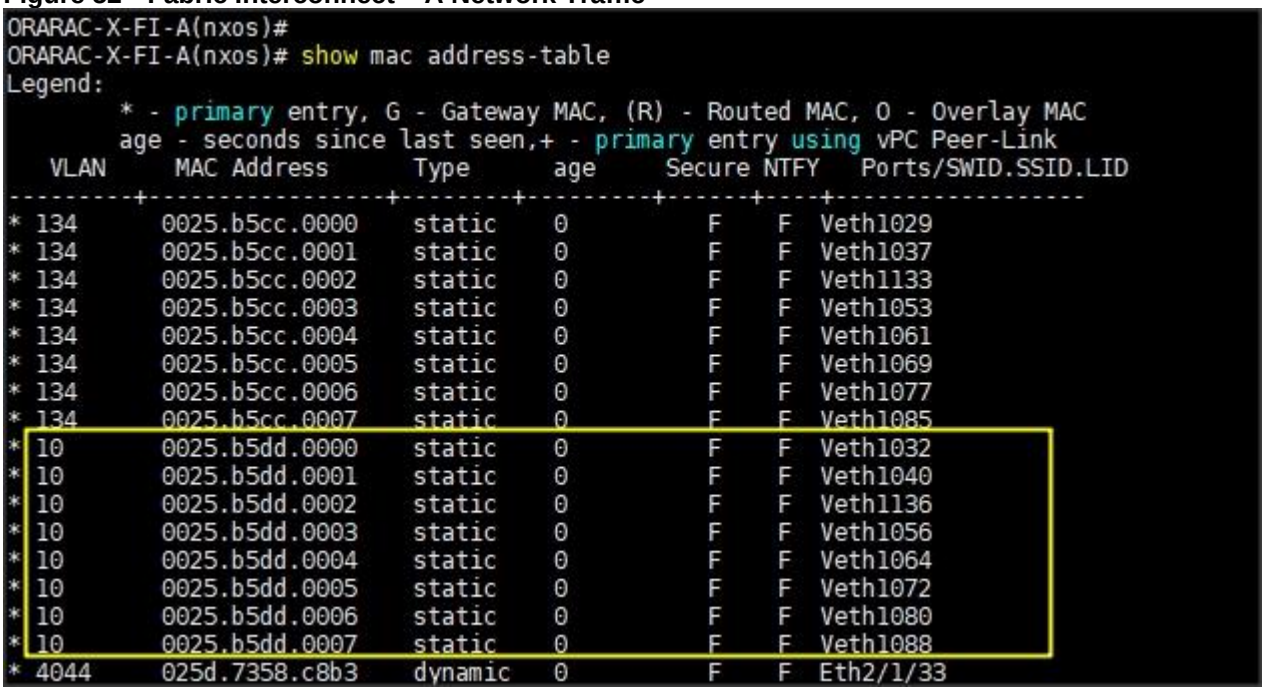

```
ORARAC-X-FI-A(nxos)#
ORARAC-X-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 134      0025.b5cc.0000    static    0          F    F  Veth1029
* 134      0025.b5cc.0001    static    0          F    F  Veth1037
* 134      0025.b5cc.0002    static    0          F    F  Veth1133
* 134      0025.b5cc.0003    static    0          F    F  Veth1053
* 134      0025.b5cc.0004    static    0          F    F  Veth1061
* 134      0025.b5cc.0005    static    0          F    F  Veth1069
* 134      0025.b5cc.0006    static    0          F    F  Veth1077
* 134      0025.b5cc.0007    static    0          F    F  Veth1085
* 10       0025.b5dd.0000    static    0          F    F  Veth1032
* 10       0025.b5dd.0001    static    0          F    F  Veth1040
* 10       0025.b5dd.0002    static    0          F    F  Veth1136
* 10       0025.b5dd.0003    static    0          F    F  Veth1056
* 10       0025.b5dd.0004    static    0          F    F  Veth1064
* 10       0025.b5dd.0005    static    0          F    F  Veth1072
* 10       0025.b5dd.0006    static    0          F    F  Veth1080
* 10       0025.b5dd.0007    static    0          F    F  Veth1088
* 4044     025d.7358.c8b3    dynamic   0          F    F  Eth2/1/33
```

As seen in the screenshot above, When the Fabric Interconnect – B failed, it will route all the Private Network traffic of VLAN 10 to Fabric Interconnect – A. So Fabric Interconnect – B Failover did not cause any disruption to Private, Public and Storage Network Traffic.

After plug back power cable to Fabric Interconnect – B Switch, the respective blades (ORARACX1, ORARACX2, ORARACX3 & ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 & ORARACX8) on chassis 2 will route back the MAC addresses and its VLAN traffic to Fabric Interconnect – B.

The figure below shows details of MAC address, VLAN information and Server connections for Cisco UCS Fabric Interconnect – A switch under normal operating condition.

Figure 53 Fabric Interconnect – A Network Traffic

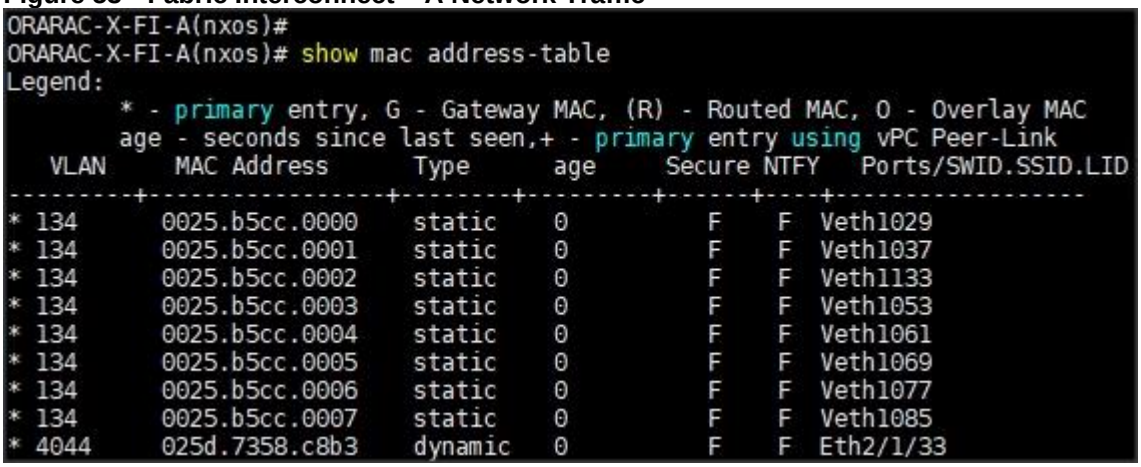

```
ORARAC-X-FI-A(nxos)#
ORARAC-X-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 134      0025.b5cc.0000    static    0          F    F  Veth1029
* 134      0025.b5cc.0001    static    0          F    F  Veth1037
* 134      0025.b5cc.0002    static    0          F    F  Veth1133
* 134      0025.b5cc.0003    static    0          F    F  Veth1053
* 134      0025.b5cc.0004    static    0          F    F  Veth1061
* 134      0025.b5cc.0005    static    0          F    F  Veth1069
* 134      0025.b5cc.0006    static    0          F    F  Veth1077
* 134      0025.b5cc.0007    static    0          F    F  Veth1085
* 4044     025d.7358.c8b3    dynamic   0          F    F  Eth2/1/33
```

The figure below shows details of MAC address, VLAN information and Server connections for Cisco UCS Fabric Interconnect – B switch.

Figure 54 Fabric Interconnect – B Network Traffic

```
ORARAC-X-FI-B(nxos)#
ORARAC-X-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b5dd.0000    static    0          F    F  Veth1031
* 10       0025.b5dd.0001    static    0          F    F  Veth1039
* 10       0025.b5dd.0002    static    0          F    F  Veth1135
* 10       0025.b5dd.0003    static    0          F    F  Veth1055
* 10       0025.b5dd.0004    static    0          F    F  Veth1063
* 10       0025.b5dd.0005    static    0          F    F  Veth1071
* 10       0025.b5dd.0006    static    0          F    F  Veth1079
* 10       0025.b5dd.0007    static    0          F    F  Veth1087
* 4044     025d.7358.c8b2    dynamic   0          F    F  Eth2/1/33
```

Test 3 – Cisco Nexus 9372PX-E Switch – A Failure Test

We conducted a hardware failure test on Nexus Switch – A by disconnecting power cable to the switch as explained below.

The figure below illustrates how during Nexus Switch – A failure, the respective blades (ORARACX1, ORARACX2, ORARACX3 and ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 and ORARACX8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to Nexus Switch – B.

Figure 55 Nexus Switch – A Failure

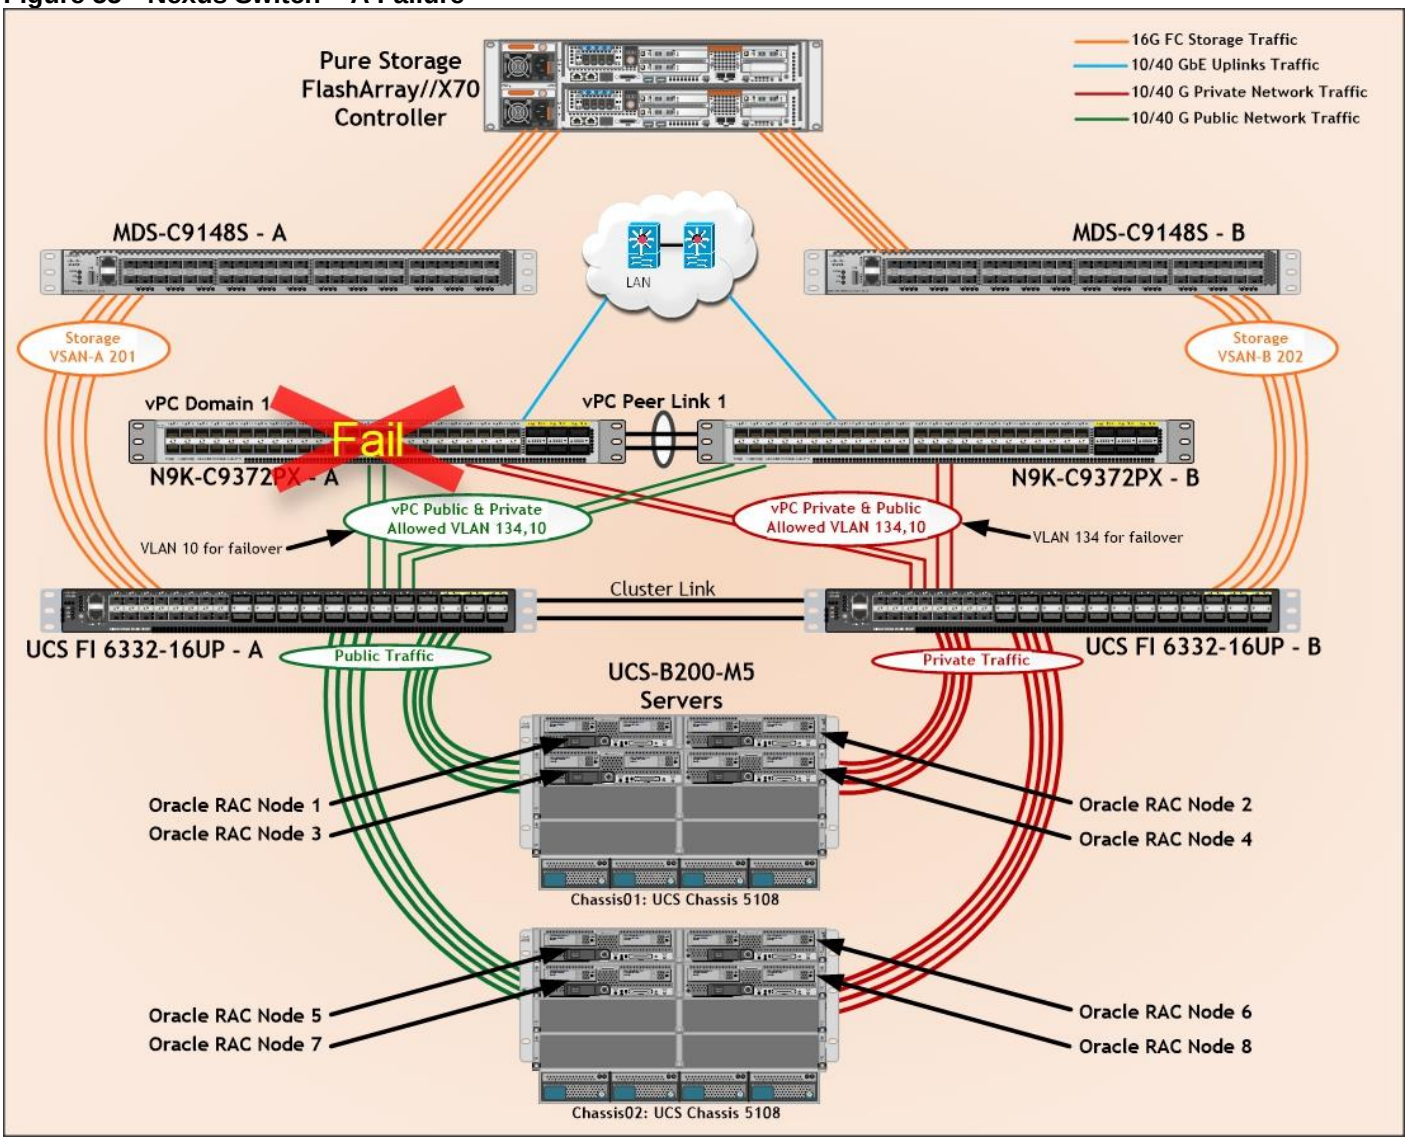

Unplug the power cable from Nexus Switch – A, and check the MAC address and VLAN information on Cisco UCS Nexus Switch – B. We noticed when the Nexus Switch – A failed, it would route all the Private Network and Public Network Traffic of VLAN 10 and VLAN 134 to Nexus Switch – B. So, Nexus Switch – A Failover did not cause any disruption to Private, Public and Storage Network Traffic.

After plug back power cable to Nexus Switch – A Switch, the respective blades on chassis 1 and chassis 2 will route back the MAC addresses and its VLAN traffic to Nexus Switch – A.

Test 4 – Cisco Nexus 9372PX-E Switch – B Failure Test

We conducted a hardware failure test on Nexus Switch – B by disconnecting power cable to the switch as explained below.

The figure below illustrates how during Nexus Switch – B failure, the respective blades (ORARACX1, ORARACX2, ORARACX3 and ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 and ORARACX8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to Nexus Switch – A.

Figure 56 Cisco Nexus Switch – B Failure

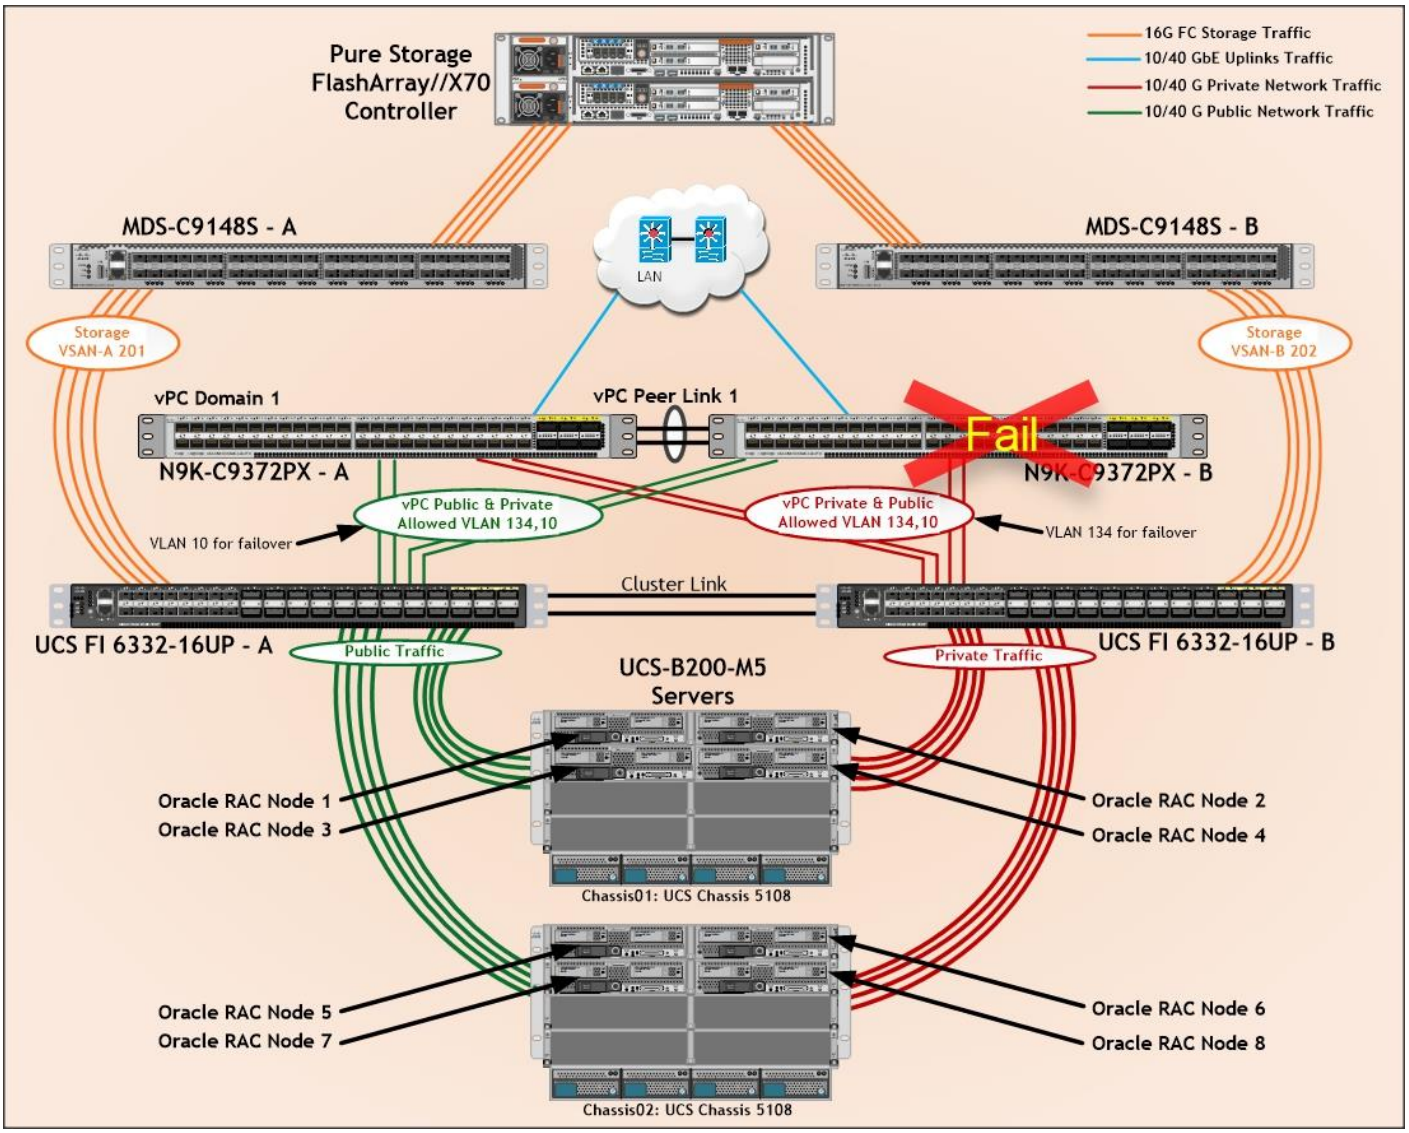

Unplug the power cable from Nexus Switch – B, and check the MAC address and VLAN information on Cisco UCS Nexus Switch – A. We noticed when the Nexus Switch – B failed, it will route all the Private Network and Public Network Traffic of VLAN 10 and VLAN 134 to Nexus Switch – A. So Nexus Switch – B Failover did not cause any disruption to Private, Public and Storage Network Traffic.

After plug back power cable to Nexus Switch – B Switch, the respective blades on chassis 1 and chassis 2 will route back the MAC addresses and its VLAN traffic to Nexus Switch – B.

Test 5 – MDS 9148S Switch – A Failure Test

We conducted hardware failure test on MDS Switch – A by disconnecting power cable to the Switch as explained below.

The figure below illustrates how during MDS Switch – A failure, the respective blades (ORARACX1, ORARACX2, ORARACX3, & ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 & ORARACX8) on chassis 2 will failover the MAC addresses and its storage traffic to MDS Switch B same way as Fabric Switch failure.

Figure 57 MDS Switch A Failure

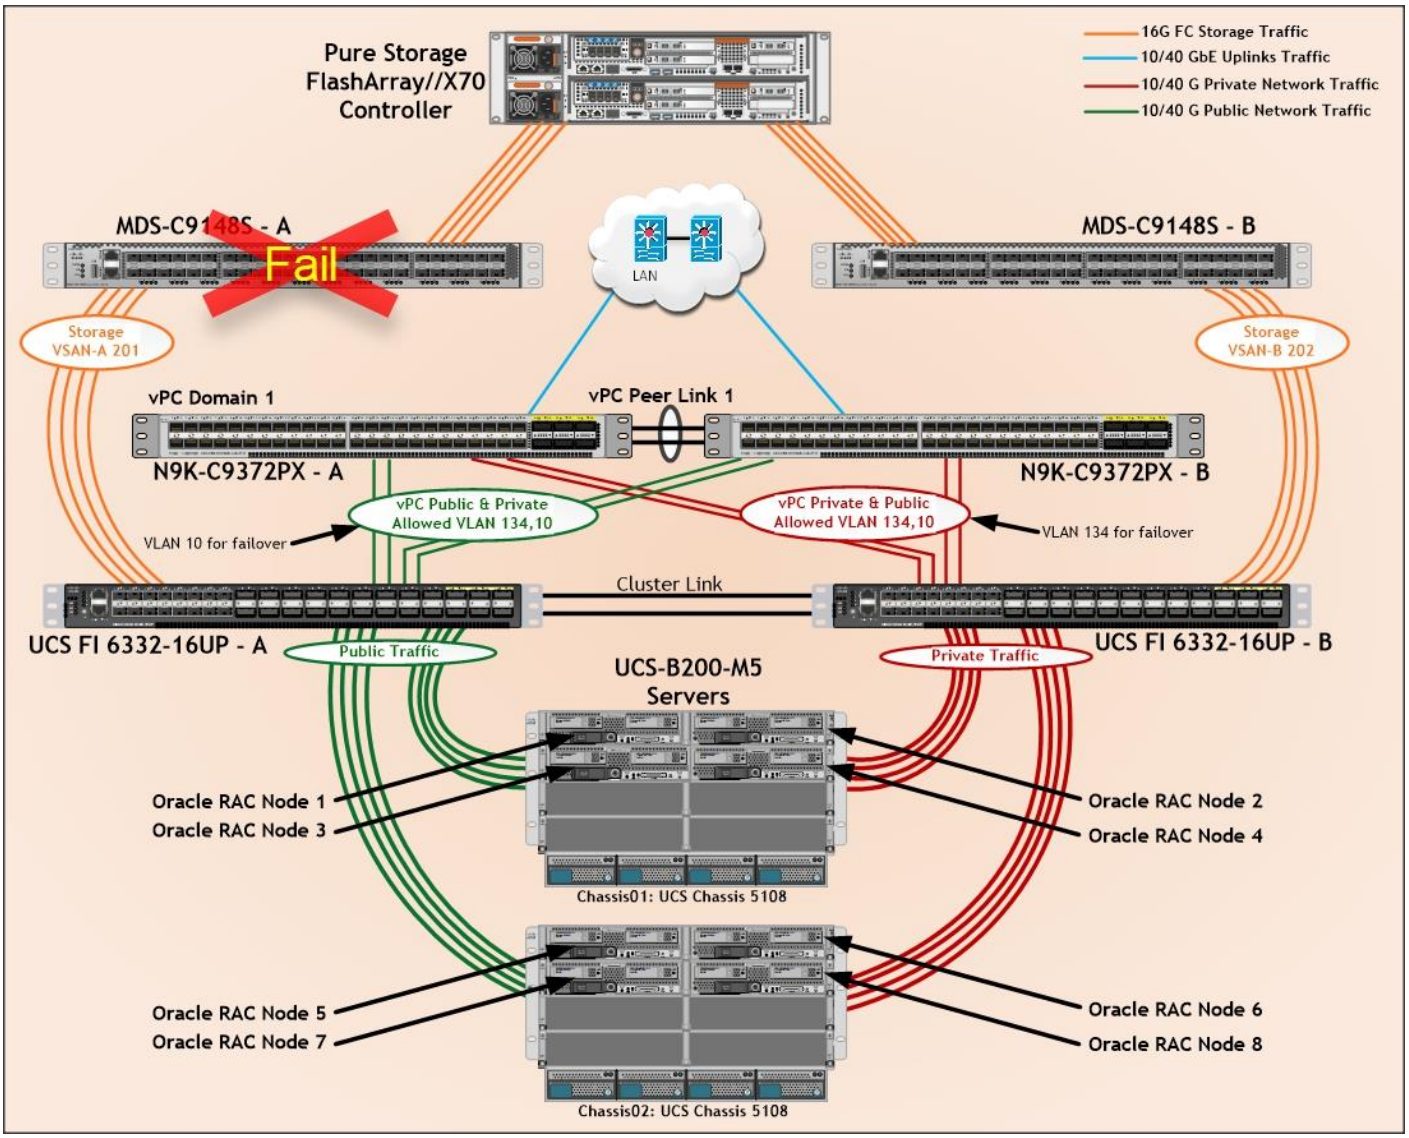

Test 6 – MDS 9148S Switch – B Failure Test

We conducted a hardware failure test on MDS Switch – B by disconnecting power cable to the Switch as explained below.

The figure below illustrates how during MDS Switch – B failure, the respective blades (ORARACX1, ORARACX2, ORARACX3, & ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 & ORARACX8) on chassis 2 will failover the MAC addresses and its storage traffic to MDS Switch B same way as Fabric Switch failure.

Figure 58 1 MDS Switch B Failure

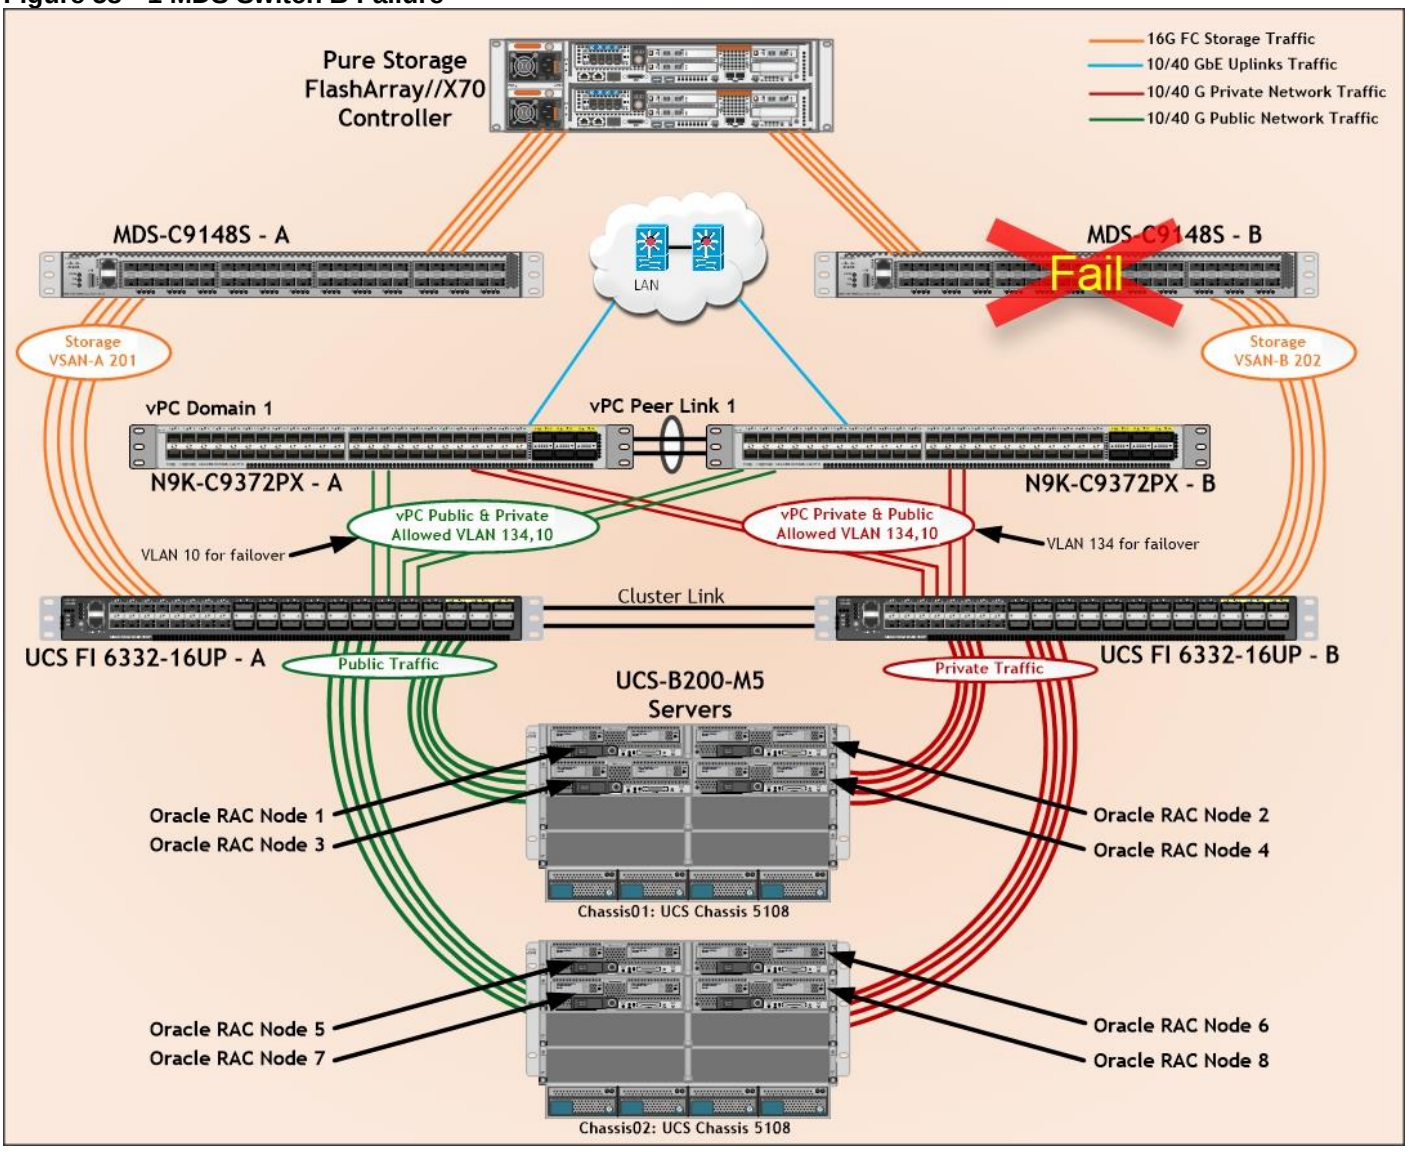

Test 7 – Cisco UCS Chassis 1 and 2 IOM Links Failure

We conducted a Cisco UCS Chassis 1 and Chassis 2 IOM Link Failure test by disconnecting two of the server port link cables from the Chassis as explained below.

The figure below illustrates how during UCS Chassis 1 and Chassis 2 IOM Links failure, the respective blades (ORARACX1, ORARACX2, ORARACX3 and ORARACX4) on chassis 1 and (ORARACX5, ORARACX6, ORARACX7 and ORARACX8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to fabric interconnect – B.

Figure 59 Chassis 1 and 2 IOM Link Failure

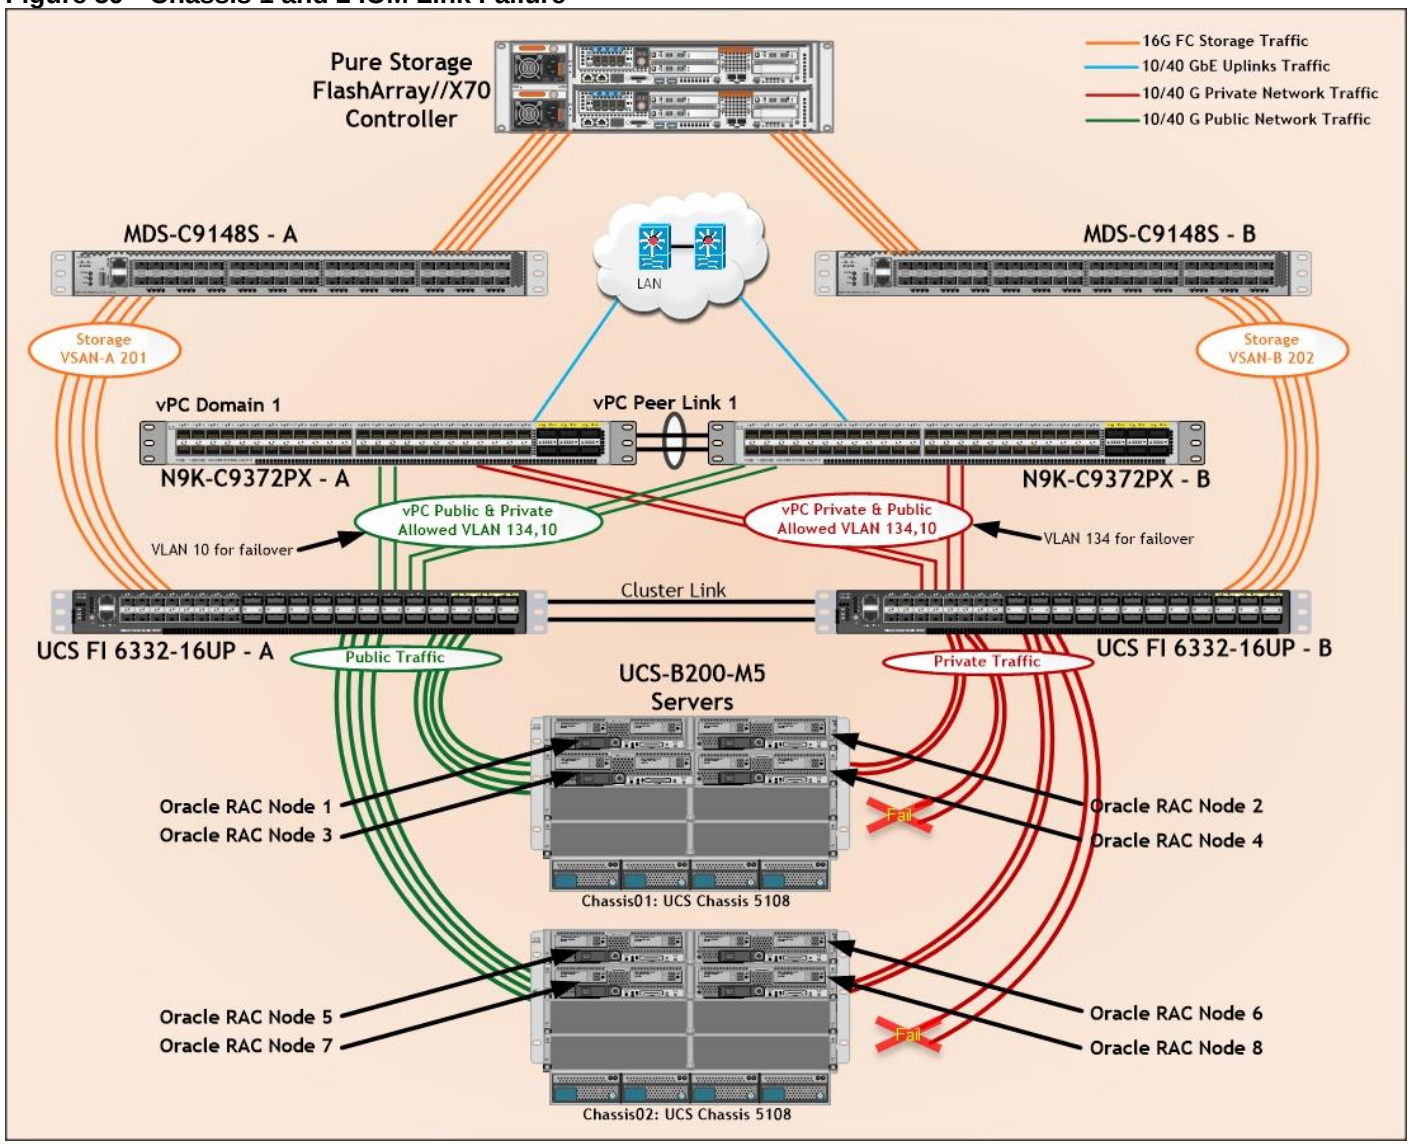

Unplug two server port cables from Chassis 1 and Chassis 2 and check the MAC address and VLAN traffic information on both UCS Fabric Interconnects. The screenshot below shows network traffic on Fabric Interconnect A when two links from Chassis 1 and two links from Chassis 2 IOM Failed.

Figure 60 Fabric Interconnect – A Network Traffic

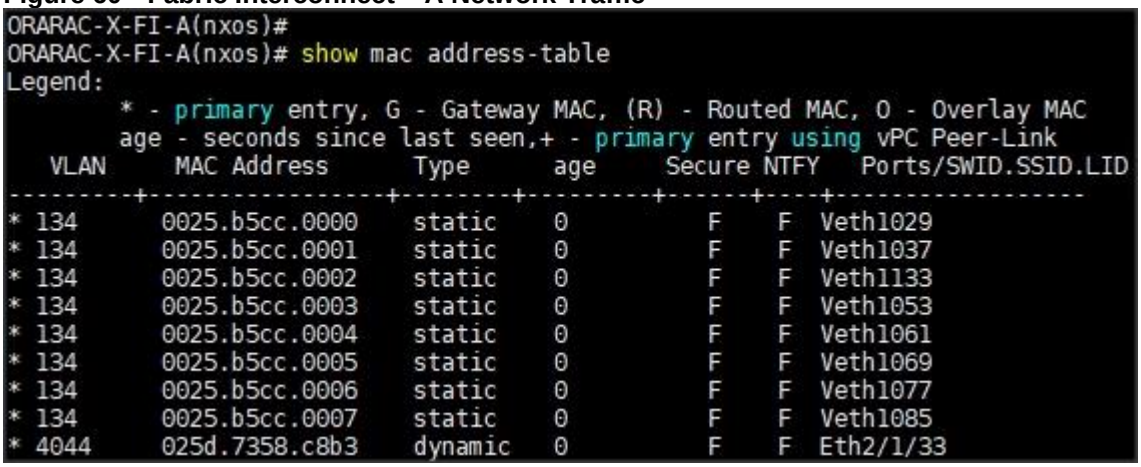

```
ORARAC-X-FI-A(nxos)#
ORARAC-X-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 134      0025.b5cc.0000    static    0          F    F  Veth1029
* 134      0025.b5cc.0001    static    0          F    F  Veth1037
* 134      0025.b5cc.0002    static    0          F    F  Veth1133
* 134      0025.b5cc.0003    static    0          F    F  Veth1053
* 134      0025.b5cc.0004    static    0          F    F  Veth1061
* 134      0025.b5cc.0005    static    0          F    F  Veth1069
* 134      0025.b5cc.0006    static    0          F    F  Veth1077
* 134      0025.b5cc.0007    static    0          F    F  Veth1085
* 4044     025d.7358.c8b3    dynamic   0          F    F  Eth2/1/33
```

The screenshot below shows network traffic on Fabric Interconnect B when two links from Chassis 1 and two links from Chassis 2 IOM Failed.

Figure 61 Fabric Interconnect – B Network Traffic

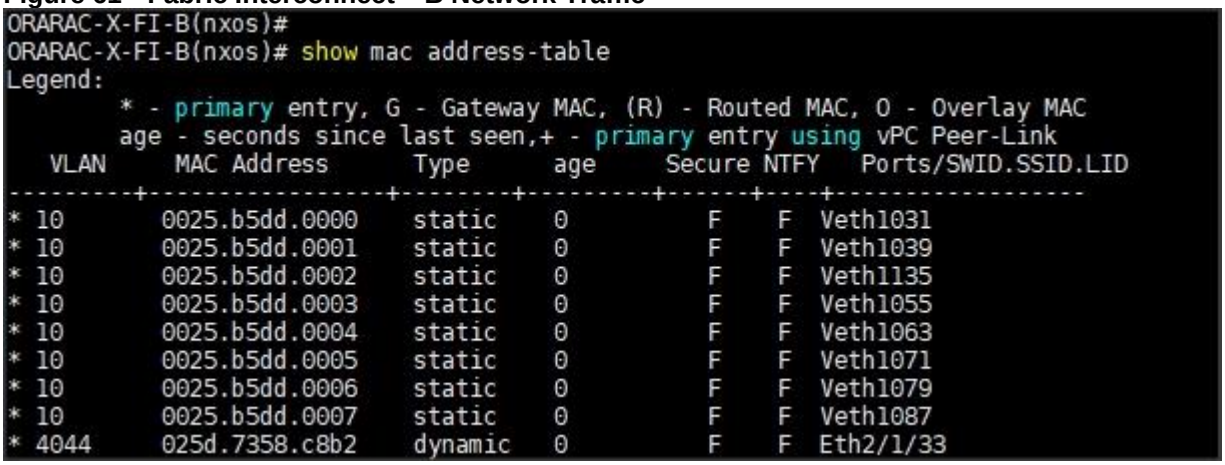

```
ORARAC-X-FI-B(nxos)#
ORARAC-X-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b5dd.0000    static    0          F    F  Veth1031
* 10       0025.b5dd.0001    static    0          F    F  Veth1039
* 10       0025.b5dd.0002    static    0          F    F  Veth1135
* 10       0025.b5dd.0003    static    0          F    F  Veth1055
* 10       0025.b5dd.0004    static    0          F    F  Veth1063
* 10       0025.b5dd.0005    static    0          F    F  Veth1071
* 10       0025.b5dd.0006    static    0          F    F  Veth1079
* 10       0025.b5dd.0007    static    0          F    F  Veth1087
* 4044     025d.7358.c8b2    dynamic   0          F    F  Eth2/1/33
```

We noticed no disruption in public and private network traffic even after two failed traffic links from both the Chassis because of the port-channel feature.

We completed additional failure scenario and validated that there are no single point of failure in this reference design.

Summary

Cisco and Pure Storage have partnered to deliver the FlashStack solution, that uses best-in-class storage, server, and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed. FlashStack Datacenter is predesigned to provide agility to large enterprise data centers with high availability and storage scalability. With a FlashStack solution, customers can leverage a secure, integrated, and optimized stack that includes compute, network, and storage resources that are sized, configured, and deployed as a fully tested unit running industry standard applications such as Oracle RAC Database 12c R2

The following factors make the combination of Cisco UCS with Pure Storage so powerful for Oracle environments:

- Cisco UCS stateless computing architecture provided by the Service Profile capability of Cisco UCS allows fast, non-disruptive workload changes to be executed simply and seamlessly across the integrated UCS infrastructure and Cisco x86 servers.
- Cisco UCS, combined with Pure Storage's highly scalable FlashArray storage system provides the ideal combination for Oracle's unique, scalable, and highly available FAS technology.
- Hardware level redundancy for all major components using Cisco UCS and Pure Storage availability features.

FlashStack is a flexible infrastructure platform composed of pre-sized storage, networking, and server components. It is designed to ease your IT transformation and operational challenges with maximum efficiency and minimal risk.

FlashStack differs from other solutions by providing:

- Integrated, validated technologies from industry leaders and top-tier software partners.
- A single platform built from unified compute, fabric, and storage technologies, allowing you to scale to large-scale data centers without architectural changes.
- Centralized, simplified management of infrastructure resources, including end-to-end automation.
- A flexible Cooperative Support Model that resolves issues rapidly and spans across new and legacy products.

Appendix

Cisco Nexus 9372PX-E Running Configuration

```
PURESTG-NEXUS-A# show running-config

!Command: show running-config
!Time: Wed Jan 10 19:59:21 2018

version 6.1(2)I2(2a)
hostname PURESTG-NEXUS-A
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
vdc PURESTG-NEXUS-A id 1
  allocate interface Ethernet1/1-48
  allocate interface Ethernet2/1-12
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

cfs eth distribute
feature lacp
feature vpc

system qos
  service-policy type network-qos jumbo

vlan 1,10,134
vlan 10
  name Oracle_Private_Traffic
vlan 134
  name Oracle_Public_Traffic

spanning-tree port type edge bpduguard default
spanning-tree port type network default
vrf context management
  ip route 0.0.0.0/0 10.29.134.1
port-channel load-balance src-dst l4port
vpc domain 1
  role priority 10
  peer-keepalive destination 10.29.134.154 source 10.29.134.153
  auto-recovery

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

interface port-channel21
```

```
  description connect to Fabric Interconnect A
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  vpc 21

interface port-channel22
  description connect to Fabric Interconnect B
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  vpc 22

interface Ethernet1/1
  description Nexus5k-B-Cluster-Interconnect
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  channel-group 1 mode active

interface Ethernet1/2
  description Nexus5k-B-Cluster-Interconnect
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  channel-group 1 mode active

interface Ethernet1/3
..

interface Ethernet1/11
  description Fabric-Interconnect-A:11
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  channel-group 21 mode active

interface Ethernet1/12
  description Fabric-Interconnect-A:12
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  channel-group 21 mode active

interface Ethernet1/13
  description Fabric-Interconnect-B:11
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  channel-group 22 mode active

interface Ethernet1/14
  description Fabric-Interconnect-B:12
  switchport mode trunk
  switchport trunk allowed vlan 1,10,134
  spanning-tree port type edge trunk
  channel-group 22 mode active

interface Ethernet1/15
  description connect to uplink switch
```

```
  switchport access vlan 134
  speed 1000

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43
```

```
interface Ethernet1/44

interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet2/1

interface Ethernet2/2

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6

interface Ethernet2/7

interface Ethernet2/8

interface Ethernet2/9

interface Ethernet2/10

interface Ethernet2/11

interface Ethernet2/12

interface mgmt0
  vrf member management
  ip address 10.29.134.153/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I2.2a.bin
```

MDS 9148S FC Switch Running Configuration

```
PURESTG-MDS-A# show running-config

!Command: show running-config
!Time: Mon Jan  8 22:36:38 2018

version 6.2(9)
power redundancy-mode redundant
feature npiv
feature telnet
no feature http-server
ip domain-lookup
ip host PURESTG-MDS-A  10.29.134.155
vsan database
```

```
  vsan 101
  vsan 201
device-alias database
  device-alias name oraracx1-hba0 pwwn 20:00:00:25:b5:6a:00:00
  device-alias name oraracx1-hba2 pwwn 20:00:00:25:b5:6a:00:01
  device-alias name oraracx2-hba0 pwwn 20:00:00:25:b5:6a:00:02
  device-alias name oraracx2-hba2 pwwn 20:00:00:25:b5:6a:00:03
  device-alias name oraracx3-hba0 pwwn 20:00:00:25:b5:6a:00:04
  device-alias name oraracx3-hba2 pwwn 20:00:00:25:b5:6a:00:05
  device-alias name oraracx4-hba0 pwwn 20:00:00:25:b5:6a:00:06
  device-alias name oraracx4-hba2 pwwn 20:00:00:25:b5:6a:00:07
  device-alias name oraracx5-hba0 pwwn 20:00:00:25:b5:6a:00:08
  device-alias name oraracx5-hba2 pwwn 20:00:00:25:b5:6a:00:09
  device-alias name oraracx6-hba0 pwwn 20:00:00:25:b5:6a:00:0a
  device-alias name oraracx6-hba2 pwwn 20:00:00:25:b5:6a:00:0b
  device-alias name oraracx7-hba0 pwwn 20:00:00:25:b5:6a:00:0c
  device-alias name oraracx7-hba2 pwwn 20:00:00:25:b5:6a:00:0d
  device-alias name oraracx8-hba0 pwwn 20:00:00:25:b5:6a:00:0e
  device-alias name oraracx8-hba2 pwwn 20:00:00:25:b5:6a:00:0f
  device-alias name FLASHSTACK-X-CT0-FC0 pwwn 52:4a:93:7b:25:8b:4d:00
  device-alias name FLASHSTACK-X-CT0-FC6 pwwn 52:4a:93:7b:25:8b:4d:06
  device-alias name FLASHSTACK-X-CT1-FC0 pwwn 52:4a:93:7b:25:8b:4d:10
  device-alias name FLASHSTACK-X-CT1-FC6 pwwn 52:4a:93:7b:25:8b:4d:16

device-alias commit

fcdomain fcid database
  vsan 1 wwn 52:4a:93:7a:b3:18:ce:02 fcid 0x3a0000 dynamic
!            [Pure-STG-CT0-FC2]
  vsan 1 wwn 52:4a:93:7a:b3:18:ce:12 fcid 0x3a0100 dynamic
!            [Pure-STG-CT1-FC2]
  vsan 1 wwn 20:01:8c:60:4f:bd:31:80 fcid 0x3a0200 dynamic
  vsan 1 wwn 20:02:8c:60:4f:bd:31:80 fcid 0x3a0300 dynamic
  vsan 1 wwn 20:01:8c:60:4f:bd:64:80 fcid 0x3a0400 dynamic
  vsan 1 wwn 52:4a:93:7b:25:8b:4d:00 fcid 0x3a0500 dynamic
!            [FLASHSTACK-X-CT0-FC0]
  vsan 201 wwn 20:04:8c:60:4f:bd:64:80 fcid 0x570000 dynamic
  vsan 201 wwn 20:01:8c:60:4f:bd:64:80 fcid 0x570100 dynamic
  vsan 201 wwn 20:03:8c:60:4f:bd:64:80 fcid 0x570200 dynamic
  vsan 201 wwn 52:4a:93:7b:25:8b:4d:16 fcid 0x570300 dynamic
!              [FLASHSTACK-X-CT1-FC6]
  vsan 201 wwn 52:4a:93:7b:25:8b:4d:06 fcid 0x570400 dynamic
!              [FLASHSTACK-X-CT0-FC6]
  vsan 201 wwn 52:4a:93:7b:25:8b:4d:10 fcid 0x570500 dynamic
!              [FLASHSTACK-X-CT1-FC0]
  vsan 201 wwn 20:02:8c:60:4f:bd:64:80 fcid 0x570600 dynamic
  vsan 201 wwn 52:4a:93:7b:25:8b:4d:00 fcid 0x570700 dynamic
!              [FLASHSTACK-X-CT0-FC0]
  vsan 201 wwn 20:00:00:25:b5:aa:00:00 fcid 0x570102 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:06 fcid 0x570605 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:02 fcid 0x570206 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:04 fcid 0x570003 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:08 fcid 0x570107 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:0a fcid 0x570608 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:0c fcid 0x570202 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:0e fcid 0x570002 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:01 fcid 0x570601 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:10 fcid 0x570208 dynamic
```

```
  vsan 201 wwn 20:00:00:25:b5:aa:00:03 fcid 0x570103 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:05 fcid 0x570203 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:07 fcid 0x570006 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:09 fcid 0x570205 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:0b fcid 0x570005 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:0d fcid 0x570602 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:0f fcid 0x570606 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:12 fcid 0x570603 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:13 fcid 0x57010a dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:11 fcid 0x570109 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:14 fcid 0x570207 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:15 fcid 0x570105 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:16 fcid 0x570108 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:17 fcid 0x570001 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:18 fcid 0x570101 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:19 fcid 0x570607 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:1a fcid 0x570604 dynamic
  vsan 201 wwn 20:00:00:25:b5:aa:00:1b fcid 0x570201 dynamic
  vsan 201 wwn 20:00:00:25:b5:6a:00:00 fcid 0x570110 dynamic
!              [oraracx1-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:02 fcid 0x570609 dynamic
!              [oraracx2-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:04 fcid 0x570204 dynamic
!              [oraracx3-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:06 fcid 0x570007 dynamic
!              [oraracx4-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:08 fcid 0x570009 dynamic
!              [oraracx5-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:0a fcid 0x570106 dynamic
!              [oraracx6-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:0c fcid 0x57020c dynamic
!              [oraracx7-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:0e fcid 0x57060d dynamic
!              [oraracx8-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:10 fcid 0x57060a dynamic
!              [test-ora1-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:12 fcid 0x570104 dynamic
!              [test-ora2-hba0]
  vsan 201 wwn 20:00:00:25:b5:6a:00:01 fcid 0x57020d dynamic
!              [oraracx1-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:03 fcid 0x57010f dynamic
!              [oraracx2-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:05 fcid 0x57010d dynamic
!              [oraracx3-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:0f fcid 0x57060f dynamic
!              [oraracx8-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:11 fcid 0x57010c dynamic
!              [test-ora1-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:13 fcid 0x570209 dynamic
!              [test-ora2-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:07 fcid 0x57000a dynamic
!              [oraracx4-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:0b fcid 0x57060e dynamic
!              [oraracx6-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:09 fcid 0x570004 dynamic
!              [oraracx5-hba2]
  vsan 201 wwn 20:00:00:25:b5:6a:00:0d fcid 0x57020a dynamic
!              [oraracx7-hba2]
```

```
vsan database
  vsan 201 interface fc1/25
  vsan 201 interface fc1/26
  vsan 201 interface fc1/27
  vsan 201 interface fc1/28
  vsan 201 interface fc1/33
  vsan 201 interface fc1/34
  vsan 201 interface fc1/35
  vsan 201 interface fc1/36
switchname PURESTG-MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9.bin
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
```

```
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48

!Active Zone Database Section for vsan 201
zone name oraracx1 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:00
!               [oraracx1-hba0]
    member pwwn 20:00:00:25:b5:6a:00:01
!               [oraracx1-hba2]

zone name oraracx2 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:02
!               [oraracx2-hba0]
    member pwwn 20:00:00:25:b5:6a:00:03
!               [oraracx2-hba2]

zone name oraracx3 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:04
!               [oraracx3-hba0]
    member pwwn 20:00:00:25:b5:6a:00:05
!               [oraracx3-hba2]

zone name oraracx4 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
```

```
    member pwwn 20:00:00:25:b5:6a:00:06
!               [oraracx4-hba0]
    member pwwn 20:00:00:25:b5:6a:00:07
!               [oraracx4-hba2]

zone name oraracx5 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:08
!               [oraracx5-hba0]
    member pwwn 20:00:00:25:b5:6a:00:09
!               [oraracx5-hba2]

zone name oraracx6 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:0a
!               [oraracx6-hba0]
    member pwwn 20:00:00:25:b5:6a:00:0b
!               [oraracx6-hba2]

zone name oraracx7 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:0c
!               [oraracx7-hba0]
    member pwwn 20:00:00:25:b5:6a:00:0d
!               [oraracx7-hba2]

zone name oraracx8 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:0e
!               [oraracx8-hba0]
```

```
    member pwwn 20:00:00:25:b5:6a:00:0f
!               [oraracx8-hba2]

zoneset name oraracx vsan 201
    member oraracx1
    member oraracx2
    member oraracx3
    member oraracx4
    member oraracx5
    member oraracx6
    member oraracx7
    member oraracx8

zoneset activate name oraracx vsan 201
do clear zone database vsan 201
!Full Zone Database Section for vsan 201
zone name oraracx1 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:00
!               [oraracx1-hba0]
    member pwwn 20:00:00:25:b5:6a:00:01
!               [oraracx1-hba2]

zone name oraracx2 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:02
!               [oraracx2-hba0]
    member pwwn 20:00:00:25:b5:6a:00:03
!               [oraracx2-hba2]

zone name oraracx3 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:04
!               [oraracx3-hba0]
    member pwwn 20:00:00:25:b5:6a:00:05
!               [oraracx3-hba2]
```

```
zone name oraracx4 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:06
!               [oraracx4-hba0]
    member pwwn 20:00:00:25:b5:6a:00:07
!               [oraracx4-hba2]

zone name oraracx5 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:08
!               [oraracx5-hba0]
    member pwwn 20:00:00:25:b5:6a:00:09
!               [oraracx5-hba2]

zone name oraracx6 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:0a
!               [oraracx6-hba0]
    member pwwn 20:00:00:25:b5:6a:00:0b
!               [oraracx6-hba2]

zone name oraracx7 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:0c
!               [oraracx7-hba0]
    member pwwn 20:00:00:25:b5:6a:00:0d
!               [oraracx7-hba2]

zone name oraracx8 vsan 201
    member pwwn 52:4a:93:7b:25:8b:4d:00
```

```
!               [FLASHSTACK-X-CT0-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:06
!               [FLASHSTACK-X-CT0-FC6]
    member pwwn 52:4a:93:7b:25:8b:4d:10
!               [FLASHSTACK-X-CT1-FC0]
    member pwwn 52:4a:93:7b:25:8b:4d:16
!               [FLASHSTACK-X-CT1-FC6]
    member pwwn 20:00:00:25:b5:6a:00:0e
!               [oraracx8-hba0]
    member pwwn 20:00:00:25:b5:6a:00:0f
!               [oraracx8-hba2]

zoneset name oraracx vsan 201
    member oraracx1
    member oraracx2
    member oraracx3
    member oraracx4
    member oraracx5
    member oraracx6
    member oraracx7
    member oraracx8

interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/26
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/27
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/28
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/29
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/30
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/31
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/32
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/33
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/34
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/35
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/36
  switchport trunk allowed vsan 201
  switchport trunk mode off
  port-license acquire
```

```
  no shutdown

interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface mgmt0
  ip address 10.29.134.155 255.255.255.0
no system default switchport shutdown
ip default-gateway 10.29.134.1
```

Multipath Configuration "/etc/multipath.conf"

```
[root@oraracx1 ~]# cat /etc/multipath.conf
blacklist {
       devnode "^(ram|zram|raw|loop|fd|md|sr|scd|st)[0-9]*"
}
defaults {
        find_multipaths yes
        polling_interval 1
}
devices {
        device {
                vendor                  "PURE"
                path_grouping_policy    multibus
                path_checker            tur
                path_selector           "queue-length 0"
                fast_io_fail_tmo        10
                dev_loss_tmo            30
                no_path_retry           0
        }
}
multipaths {
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011066
                alias                   orarax1_os
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011084
                alias                   dg_orarac_crs
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d80001107f
                alias                   dg_oradata_oltp1
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011080
                alias                   dg_oraredo_oltp1
        }
        multipath {
```

```
                wwid                    3624a93701c0d5dfa58fa45d800011091
                alias                   dg_oradata_dss1
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011093
                alias                   dg_oraredo_dss1
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011096
                alias                   dg_oradata_soe1
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011097
                alias                   dg_oraredo_soe1
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011125
                alias                   dg_oradata_slob01
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011126
                alias                   dg_oradata_slob02
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011127
                alias                   dg_oradata_slob03
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011128
                alias                   dg_oradata_slob04
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011129
                alias                   dg_oraredo_slob
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011130
                alias                   dg_oradata_slob05
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011131
                alias                   dg_oradata_slob06
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011132
                alias                   dg_oradata_slob07
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011133
                alias                   dg_oradata_slob08
        }
        multipath {
                wwid                    3624a93701c0d5dfa58fa45d800011134
                alias                   dg_oradata_slob09
        }

}
```

Configuration of "/etc/sysctl.conf"

```
### File located "/etc/sysctl.conf" directory
[root@oraracx1 ~]# cat /etc/sysctl.conf
# oracle-database-server-12cR2-preinstall setting for fs.file-max is 6815744
fs.file-max = 6815744

# oracle-database-server-12cR2-preinstall setting for kernel.sem is '250 32000 100 128'
kernel.sem = 250 32000 100 128

# oracle-database-server-12cR2-preinstall setting for kernel.shmmni is 4096
kernel.shmmni = 4096

# oracle-database-server-12cR2-preinstall setting for kernel.shmall is 1073741824 on
x86_64
kernel.shmall = 1073741824

# oracle-database-server-12cR2-preinstall setting for kernel.shmmax is 4398046511104 on
x86_64
kernel.shmmax = 4398046511104

# oracle-database-server-12cR2-preinstall setting for kernel.panic_on_oops is 1 per
Orabug 19212317
kernel.panic_on_oops = 1

# oracle-database-server-12cR2-preinstall setting for net.core.rmem_default is 262144
net.core.rmem_default = 262144

# oracle-database-server-12cR2-preinstall setting for net.core.rmem_max is 4194304
net.core.rmem_max = 4194304

# oracle-database-server-12cR2-preinstall setting for net.core.wmem_default is 262144
net.core.wmem_default = 262144

# oracle-database-server-12cR2-preinstall setting for net.core.wmem_max is 1048576
net.core.wmem_max = 1048576

# oracle-database-server-12cR2-preinstall setting for net.ipv4.conf.all.rp_filter is 2
net.ipv4.conf.all.rp_filter = 2

# oracle-database-server-12cR2-preinstall setting for net.ipv4.conf.default.rp_filter
is 2
net.ipv4.conf.default.rp_filter = 2

# oracle-database-server-12cR2-preinstall setting for fs.aio-max-nr is 1048576
fs.aio-max-nr = 1048576

# oracle-database-server-12cR2-preinstall setting for net.ipv4.ip_local_port_range is
9000 65500
net.ipv4.ip_local_port_range = 9000 65500

# Huge Page Setting for Oracle
vm.nr_hugepages=125000
```

Configuration of "oracle-database-server-12cR2-preinstall.conf"

```
### File located "/etc/security/limits.d/oracle-database-server-12cR2-preinstall.conf"
directory
```

```
[root@oraracx1 ~]# cat /etc/security/limits.d/oracle-database-server-12cR2-
preinstall.conf

# oracle-database-server-12cR2-preinstall setting for nofile soft limit is 1024
oracle   soft   nofile    1024

# oracle-database-server-12cR2-preinstall setting for nofile hard limit is 65536
oracle   hard   nofile    65536

# oracle-database-server-12cR2-preinstall setting for nproc soft limit is 16384
# refer orabug15971421 for more info.
oracle   soft   nproc    16384

# oracle-database-server-12cR2-preinstall setting for nproc hard limit is 16384
oracle   hard   nproc    16384

# oracle-database-server-12cR2-preinstall setting for stack soft limit is 10240KB
oracle   soft   stack    10240

# oracle-database-server-12cR2-preinstall setting for stack hard limit is 32768KB
oracle   hard   stack    32768

# oracle-database-server-12cR2-preinstall setting for memlock hard limit is maximum of
128GB on x86_64 or 3GB on x86 OR 90 % of RAM
oracle   hard   memlock    237114345

# oracle-database-server-12cR2-preinstall setting for memlock soft limit is maximum of
128GB on x86_64 or 3GB on x86 OR 90% of RAM
oracle   soft   memlock    237114345

grid   soft   nofile    1024
grid   hard   nofile    65536
grid   soft   nproc    16384
grid   hard   nproc    16384
grid   soft   stack    10240
grid   hard   stack    32768
grid   soft   memlock  237114345
grid   hard   memlock  237114345
```

Configuration of “/etc/udev/rules.d/99-oracle-asmdevices.rules”

```
### File located “/etc/udev/rules.d/” directory

[root@oraracx1 ~]# cat /etc/udev/rules.d/99-oracle-asmdevices.rules
#All volumes which starts with dg_orarac_* #
ENV{DM_NAME}=="dg_orarac_crs", OWNER:="grid", GROUP:="oinstall", MODE:="660"

#All volumes which starts with dg_oradata_* #
ENV{DM_NAME}=="dg_oradata_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"

#All volumes which starts with dg_oraredo_* #
ENV{DM_NAME}=="dg_oraredo_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"

#All volumes which starts with dg_oraarchive_* #
ENV{DM_NAME}=="dg_oraarchive_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"
```

Configuration of “/etc/udev/rules.d/99-pure-storage.rules”

```
### File located “/etc/udev/rules.d/” directory
```

```
[root@oraracx1 ~]# cat /etc/udev/rules.d/99-pure-storage.rules
# Recommended settings for Pure Storage FlashArray.

# Use noop scheduler for high-performance solid-state storage
ACTION=="add|change", KERNEL=="sd*[!0-9]", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/scheduler}="noop"

# Reduce CPU overhead due to entropy collection
ACTION=="add|change", KERNEL=="sd*[!0-9]", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/add_random}="0"

# Spread CPU load by redirecting completions to originating CPU
ACTION=="add|change", KERNEL=="sd*[!0-9]", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/rq_affinity}="2"

# Set the HBA timeout to 60 seconds
ACTION=="add", SUBSYSTEMS=="scsi", ATTRS{model}=="FlashArray      ", RUN+="/bin/sh -c
'echo 60 > /sys/$DEVPATH/device/timeout'"
```

About the Authors

Tushar Patel, Principal Engineer, CSPG UCS Product Management and Data Center Solutions Engineering Group, Cisco Systems, Inc.

Tushar Patel is a Principal Engineer in Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group and a specialist in Flash Storage technologies and Oracle RAC RDBMS. Tushar has over 23 years of experience in Flash Storage architecture, Database architecture, design and performance. Tushar also has strong background in Intel X86 architecture, hyper converged systems, Storage technologies and Virtualization. He has worked with large number of enterprise customers, evaluate and deploy mission critical database solutions. Tushar has presented to both internal and external audiences at various conferences and customer events.

Hardikkumar Vyas, Solution Engineer, CSPG UCS Product Management and Data Center Solutions Engineering Group, Cisco Systems, Inc.

Hardikkumar Vyas is a Solution Engineer in Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group for developing and validating infrastructure best practices for Oracle RAC and Standalone databases on Cisco UCS Servers, Cisco Nexus Products and Storage Technologies. Hardikkumar Vyas holds a Master's degree in Electrical Engineering and has over 5 years of experience in Oracle Database and applications. Hardikkumar Vyas's focus is developing Oracle Database solutions on Cisco UCS Platform.

Somu Rajarathinam, Oracle Solutions Architect, Pure Storage

Somu Rajarathinam is the Oracle Solutions Architect at Pure Storage responsible for defining database solution based on the company's products, performing benchmarks, preparing reference architecture and technical papers for Oracle databases on Pure. Somu has over 20 years of Oracle database experience, including as a member of Oracle Corporation's Systems Performance and Oracle Applications Performance Groups. His career also included assignments with Logitech, Inspirage, and Autodesk, ranging from providing database and performance solutions to managing infrastructure, to delivering database and application support, both in-house and in the cloud.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Radhakrishna Manga, Sr. Director, Pure Storage