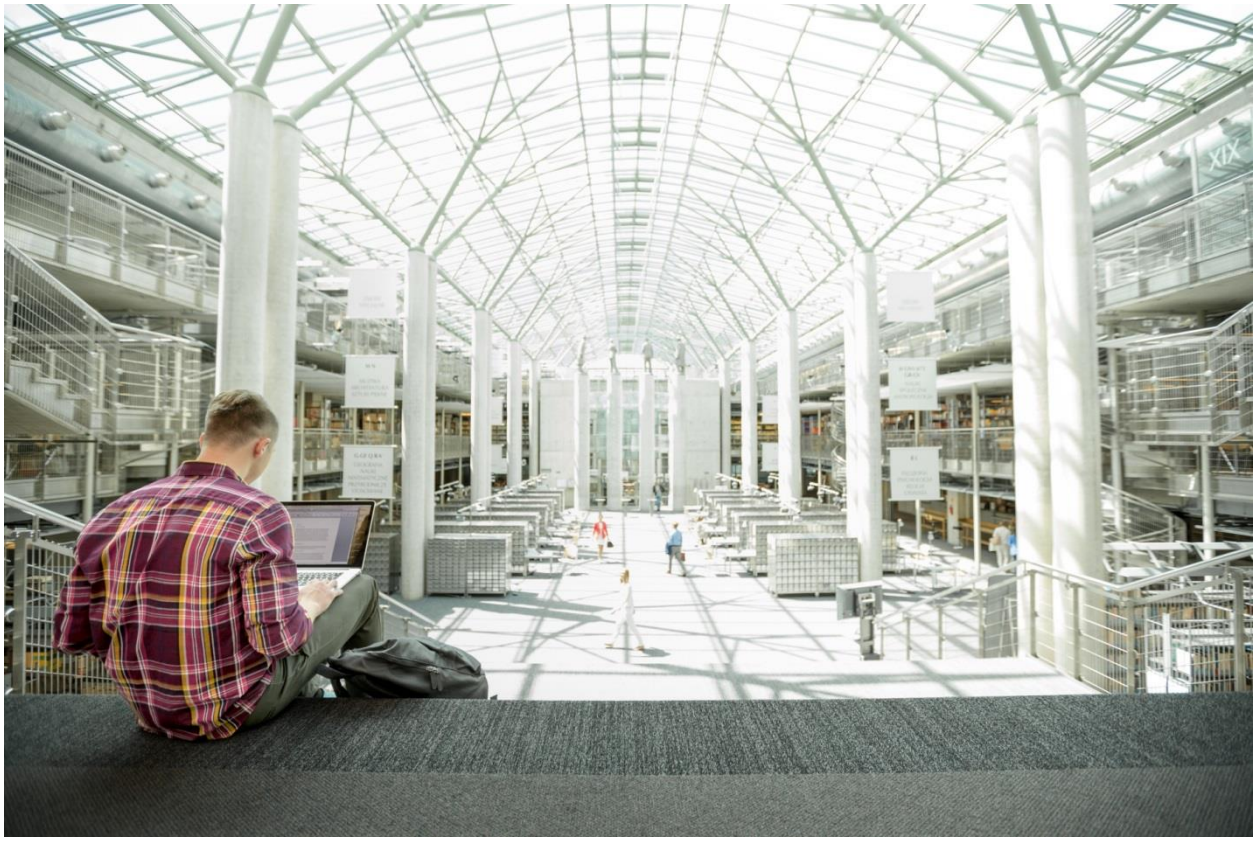

# FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series

## Deployment Guide for Oracle RAC 12cR2 Databases on Cisco Unified Computing System and NetApp AFF A-700s Storage using Fibre Channel

**Last Updated:** January 8, 2019

CISCO VALIDATED DESIGN

# 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.

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series. Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2019 Cisco Systems, Inc. All rights reserved.

# Table of Contents


Executive Summary ..... 6
Solution Overview ..... 7
Introduction ..... 7
Audience ..... 7
Purpose of this Document ..... 7
FlexPod System Overview ..... 8
What's New in this Release? ..... 9
Solution Components ..... 11
Cisco UCS 5108 Blade Server Chassis ..... 11
Cisco UCS 2304 Fabric Extender ..... 11
Cisco UCS B200 M5 Blade Servers ..... 12
Cisco UCS Virtual Interface Card (VIC) 1340 ..... 12
Cisco UCS 6332-16UP Fabric Interconnect ..... 12
Cisco UCS MDS 9148S Fabric Switch ..... 13
Cisco Nexus 93180LC-EX Switch ..... 13
NetApp AFF All Flash Array A700s ..... 13
Solution Design ..... 16
Physical Infrastructure ..... 16
Physical Topology ..... 16
Design Topology ..... 18
Solution Configuration ..... 21
Configure Cisco Nexus 93180LC-EX Switch ..... 21
Set Up Initial Configuration ..... 21
Configure Global Settings for Cisco Nexus A and Cisco Nexus B ..... 23
Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches ..... 24
Virtual Port Channel (vPC) Summary for Data and Storage Network ..... 24
Create vPC Peer-Link Between the Two Nexus Switches ..... 25
Create vPC Configuration Between Nexus 9372PX-E and Fabric Interconnects ..... 26
Verify All vPC Status Is Up on Both the Nexus Switches ..... 30
Cisco UCS Configuration Overview ..... 33
Perform Initial Setup of Cisco UCS 6332-16UP Fabric Interconnects for a Cluster Setup ..... 33
Log in to Cisco UCS Manager ..... 35
Upgrade Cisco UCS Manager Software to Version 3.2 ..... 36
Anonymous Reporting ..... 36
Configure Cisco UCS Call Home ..... 37
High Level Steps to Configure Base Cisco UCS ..... 37

Set Fabric Interconnects to Fibre Channel End Host Mode ..... 38
Synchronize Cisco UCS to NTP ..... 39
Configure Fabric Interconnect for Chassis and Blade Discovery ..... 39
Configure LAN and SAN on Cisco UCS Manager ..... 42
Configure IP, UUID, Server, MAC, WWNN and WWPN Pools ..... 49
Set Jumbo Frames in both the Cisco Fabric Interconnect ..... 54
Configure Server BIOS Policy ..... 55
Create Adapter Policy ..... 56
Configure Update Default Maintenance Policy ..... 58
Configure vNIC and vHBA Template ..... 58
Create Server Boot Policy for SAN Boot ..... 62
Configure and Create a Service Profile Template ..... 69
Create Service Profile Template ..... 69
Create Service Profiles from Template and Associate to Servers ..... 80
Create Service Profiles from Template ..... 80
Associate Service Profiles to the Servers ..... 81
Configure Cisco MDS 9148S Switch ..... 83
Configure feature for MDS Switch A and MDS Switch B ..... 84
Configure VSANs for MDS Switch A and MDS Switch B ..... 85
Create and Configure Fibre Channel Zoning ..... 86
Create Zoning ..... 88
Verify FC Ports on MDS Switch ..... 91
Configure NetApp AFF A700s Storage ..... 93
NetApp Storage Connectivity ..... 93
Operating System and Database Deployment ..... 96
Operating System Configuration ..... 96
Operating System Prerequisites for Oracle Software Installation ..... 98
Configure BIOS ..... 98
Prerequisites Automatic Installation ..... 99
Additional Prerequisites Setup ..... 100
Configure Multipath Setup ..... 100
Oracle Database 12c GRID Infrastructure Deployment ..... 103
Install and Configure Oracle Database Grid Infrastructure Software ..... 105
Install Oracle Database Software ..... 113
Oracle Flex ASM ..... 117
Scalability Test and Results ..... 120
Hardware Calibration Test using FIO ..... 120
4k and 8k Random Read/Write IOPs Tests ..... 120

Bandwidth Tests ..... 121
Database Creation with DBCA ..... 123
SLOB Test ..... 124
Swingbench Test ..... 128
Resiliency and Failure Tests ..... 140
Summary ..... 151
Appendix ..... 152
References ..... 179
Cisco Unified Computing System ..... 179
Cisco UCS Data Center Design Guides ..... 179
FlexPod Converged Infrastructure ..... 179
Cisco UCS B200 M5 Servers ..... 179
Oracle Database 12c Release 2 ..... 179
NetApp AFF A-Series All Flash Storage ..... 179
NetApp Support ..... 179
About the Authors ..... 180
Acknowledgements ..... 180

# 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 such as Oracle RAC. The combination of Cisco UCS, NetApp and Oracle Real Application Cluster Database 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. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

The FlexPod Datacenter with NetApp All Flash AFF system is a converged infrastructure platform that combines best-of breed technologies from Cisco and NetApp into a powerful converged platform for enterprise applications. Cisco and NetApp works closely with Oracle to support the most demanding transactional and response-time-sensitive databases required by today's businesses.

This Cisco Validated Design (CVD) describes the reference FlexPod Datacenter architecture using Cisco UCS and NetApp® All Flash AFF Storage for deploying a highly available Oracle RAC Database environment. This document shows the hardware and software configuration of the components involved, results of various tests and offers implementation and best practices guidance using Cisco UCS Compute Servers, Cisco Fabric Interconnect Switches, Cisco MDS Switches, Cisco Nexus Switches, NetApp AFF Storage and Oracle RAC Database.

# Solution Overview

## Introduction

Data powers essentially every operation in a modern enterprise, from keeping the supply chain operating efficiently to managing relationships with customers. Modern data centers face increasing demands for agile, high-performance service delivery. Digital transformations are driving an increased number of new applications, with more sources of data. Organizations of all kinds rely on their relational databases for both transaction processing (OLTP) and analytics (OLAP), but many still have challenges in meeting their goals of high availability, security, and performance. Applications must be able to move quickly from development to a reliable, scalable platform. An ideal solution integrates best-in-class components that can scale compute and storage independently to meet the needs of dynamic business requirements.

Like all the FlexPod Systems, the FlexPod Datacenter with NetApp All Flash AFF is comprised of compute (database, application and management servers from Cisco), network (three-layer network and SAN technologies from Cisco), and storage (NetApp All Flash AFF storage systems).

This CVD describes how the Cisco Unified Computing System™ (Cisco UCS®) can be used in conjunction with NetApp® AFF storage systems to implement a mission-critical applications such as an Oracle Real Application Clusters (RAC) 12cR2 database solution. This CVD documents validation of the real world performance, ease of management, and agility of the FlexPod Datacenter with Cisco UCS and All Flash AFF in high-performance Oracle RAC Databases environments.

## Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, database administrators, IT managers, oracle database architects, and customers who want to deploy Oracle RAC 12cR2 database solution on FlexPod Converged Infrastructure with NetApp clustered Data ONTAP® and the Cisco UCS platform. 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. This FlexPod solution for Oracle databases delivers industry-leading storage, unprecedented scalability, continuous data access, and automated data management for immediate responses to business opportunities.

The goal of this document is to determine the Oracle database server read latency, peak sustained throughput and IOPS of this FlexPod reference architecture system while running the Oracle OLTP and OLAP workloads.

This document provides a step-by-step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS Compute Servers, Cisco Fabric Interconnect Switches, Cisco MDS Switches, Cisco Nexus Switches and NetApp AFF Storage to deploy an Oracle RAC Database solution. Here are the objectives we would like to accomplish in this reference document

1. Provide reference FlexPod architecture design guidelines for the Oracle RAC Databases solution.

2. Demonstrate simplicity and agility with the software-driven architecture and high performance of Cisco UCS compute Servers.

3. Build, validate and predict performance of Servers, Network and Storage platform on a per workload basis.

## FlexPod System Overview

Built on innovative technology from NetApp and Cisco, the FlexPod converged infrastructure platform meets and exceeds the challenges of simplifying deployments for best-in-class data center infrastructure. FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. Composed of pre-validated storage, networking, and server technologies, FlexPod is designed to increase IT responsiveness to organizational needs and reduce the cost of computing with maximum uptime and minimal risk. Simplifying the delivery of data center platforms gives enterprises an advantage in delivering new services and applications.

FlexPod provides the following differentiators:

1. Flexible design with a broad range of reference architectures and validated designs.

4. Elimination of costly, disruptive downtime through Cisco UCS and NetApp® ONTAP®.

5. Leverage a pre-validated platform to minimize business disruption and improve IT agility and reduce deployment time from months to weeks.

6. Cisco Validated Designs (CVDs) and NetApp Validated Architectures (NVAs) covering a variety of use cases.

Cisco and NetApp have carefully validated and verified the FlexPod 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.

Figure 1 FlexPod System Overview

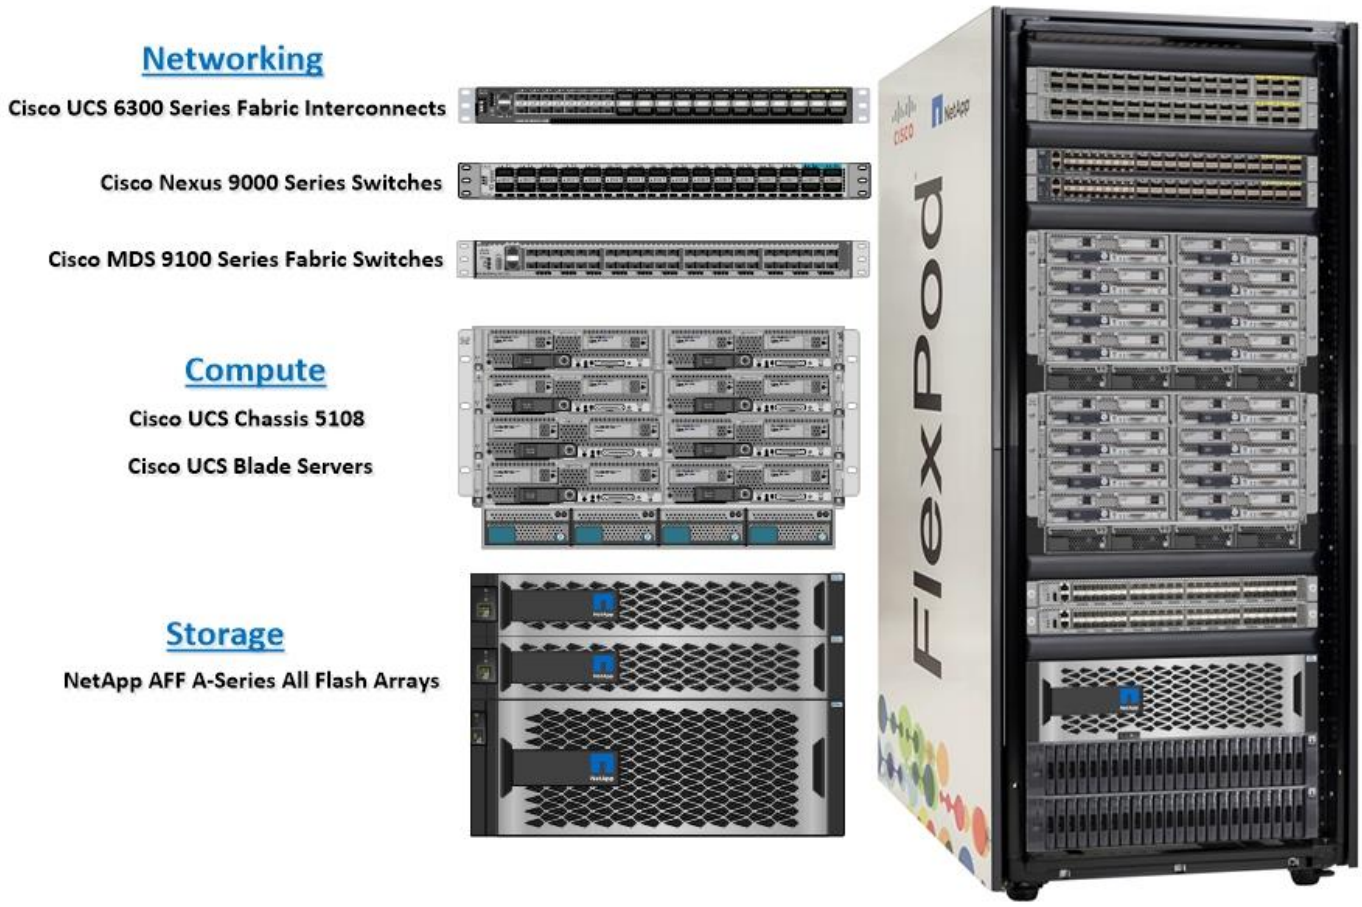


FlexPod datacenter architecture includes three components:

- Cisco Unified Computing System (Cisco UCS)
- Cisco MDS and Nexus Switches
- NetApp AFF Storage Systems

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). This document highlights the resiliency, cost benefit, and ease of deployment of a Fibre Channel storage solution to deploy Oracle RAC Database environments on FlexPod Infrastructure.

## What's New in this Release?

This version of FlexPod CVD introduces new hardware with NetApp A-Series All Flash Storage AFF A700s along with Cisco UCS B200 M5 Blade Server featuring Intel Xeon Scalable Family of CPUs.

It incorporates the following features:

- Support for the Cisco UCS 3.2 unified software release and Cisco UCS B200 M5 Blade Servers
- Support for the latest release of NetApp ONTAP® 9.3

- Fibre channel storage design
- Validation of Oracle RAC Database 12c Release 2

# Solution Components

## 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, 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 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 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 Databases.

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 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.

## 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.

## 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 93180LC-EX Switch

The Cisco Nexus 93180LC-EX Switch is the industry's first 50-Gbps-capable 1RU switch that supports 3.6 Tbps of bandwidth and over 2.6 bpps across up to 32 fixed 40- and 50-Gbps QSFP+ ports or up to 18 fixed 100-Gbps ports (Figure 3). The switch can support up to 72 10-Gbps ports using breakout cables. A variety of flexible port configurations are supported using templates.

## NetApp AFF All Flash Array A700s

This solution includes the NetApp All flash fabric-attached storage AF 700s unified scale-out storage system for Oracle Database 12cR2.

Built on ONTAP software, AFF speeds up the operations required to meet your business requirements, without compromising efficiency or reliability, while providing great flexibility and scalability. As true enterprise-class, all-flash arrays, these systems accelerate, manage, and protect business-critical data, now and in the future.

NetApp AFF provides industry-leading performance while continuing to provide a full suite of enterprise-grade data management and data protection features. Powered by NetApp clustered Data ONTAP, the AF 700s series unifies the SAN and NAS storage infrastructure. Systems architects can choose from a range of models representing a spectrum of cost-versus-performance points. Every model, however, provides the following core benefits:

- HA and fault tolerance. Non-disruptive operations are achieved through clustering, HA pairing of controllers, hot-swappable components, NetApp RAID-TEC® and disk protection (allowing two independent disk failures without data loss), and network interface redundancy.
- Data Protection. Create thousands of instantaneous backups with NetApp Snapshot® copies, replication of both active data and snapshot backups with NetApp SnapMirror® software, and application backup integration with NetApp SnapCenter® storage management software.
- Storage efficiency. Users can store more data with less physical media. This is achieved with thin provisioning (unused space is shared among volumes), NetApp Snapshot® copies (zero-storage, read-only replicas of data over time), NetApp FlexClone® volumes, and LUNs (read/write copies of data in which only changes are stored), deduplication (dynamic detection and removal of redundant data), and data compression.
- Unified Data Management. Every model runs the same software (clustered Data ONTAP); supports all popular storage protocols (CIFS, NFS, iSCSI, FCP, and FCoE). This allows freedom of choice in upgrades and expansions, without the need to redesign the solution or retraining operations personnel.
- Advanced clustering. Storage controllers are grouped into clusters for both availability and performance pooling. Workloads can be moved between controllers, permitting dynamic load balancing and zero-downtime maintenance, upgrades, and even complete hardware refreshes. Physical media and storage controllers can be added as needed to support growing demand without downtime.

This solution utilizes the NetApp AFF A700s as seen in the above figure. This controller provides the high-performance benefits of 40GbE and all flash SSDs.

Combined with the 3 disk shelfs and 72 1TB Solid State disks (SSD), this solution can provide over ample horsepower and over 60TB of raw capacity, all while taking up minimum valuable rack space. This makes it an ideal controller for a shared workload converged infrastructure. For situations where more performance is needed, the A700s would be an ideal fit. ONTAP 9.3 is used with the AFF A700s storage platform in our design.

ONTAP data management software offers unified storage for applications that read and write data over block or file-access protocols in storage configurations that range from high-speed flash to lower-priced spinning media or cloud-based object storage. The ONTAP platform used on the AFF systems also delivers interconnection with a larger ONTAP-based storage environment. In addition to AFF hardware systems, ONTAP is also available on commodity hardware (ONTAP Select), and in private, public, or hybrid clouds (NetApp Private Storage and Cloud Volumes ONTAP).

Together these implementations form the basic framework of the NetApp Data Fabric, with a common software-defined approach to data management and fast, efficient replication across platforms. FlexPod and ONTAP can serve as the foundation for both hybrid cloud and private cloud designs

Starting with ONTAP 9, NetApp guarantees that the use of NetApp storage efficiency technologies on AFF systems reduce the total logical capacity used to store customer data by 75 percent, a data reduction ratio of 4:1. This space reduction is enabled by a combination of several different technologies, including deduplication, compression, and compaction.

Beginning with ONTAP 9.1, NetApp has extended the encryption capabilities further with NetApp Volume Encryption (NVE), a software-based mechanism for encrypting data. It allows a user to encrypt data at the per-volume level instead of requiring encryption of all data in the cluster, thereby providing more flexibility and granularity to the ONTAP administrators. This encryption extends to Snapshot copies and NetApp FlexClone® volumes that are created in the cluster. One benefit of NVE is that it executes after the implementation of the storage efficiency features, and, therefore, it does not interfere with the ability of ONTAP to create space savings.

# Solution Design

## Physical Infrastructure

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. FlexPod components are connected and configured according to best practices of both Cisco and NetApp to provide the ideal platform for running a variety of enterprise workloads with confidence. This solution provides an end-to-end architecture with Cisco Unified Computing System, Oracle, and NetApp technologies and demonstrates the FlexPod configuration benefits for running Oracle RAC Databases 12cR2 workloads with high availability and server redundancy.

The reference FlexPod architecture covered in this document is built on the NetApp All Flash AFF A700s for Storage, Cisco B200 M5 Blade Servers for Compute, Cisco Nexus 93180LC-EX Switches, Cisco MDS 9148S 16G Multilayer Fabric Switches and Cisco Fabric Interconnects 6332-16UP Fabric Interconnects for System Management in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design.

The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

### Physical Topology

This section describes the design considerations for the Oracle RAC Database 12c Release 2 on FlexPod deployment. In this solution design, we have used two Cisco UCS Blade Server Chassis with 4 identical Intel Xeon CPU based Cisco UCS B-Series B200 M5 Blade Servers for hosting the 4-Node Oracle RAC Databases. The Cisco UCS B200 M5 Server has Virtual Interface Card (VIC) 1340 with port expander and they were connected two 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 NetApp AFF Storage A700s.

Figure 2 shows the architecture diagram of the FlexPod components and the network connections to deploy a four-node Oracle RAC 12cR2 Databases solution. This reference design is a typical network configuration that can be deployed in a customer's environments.

Figure 2 FlexPod Architecture Design

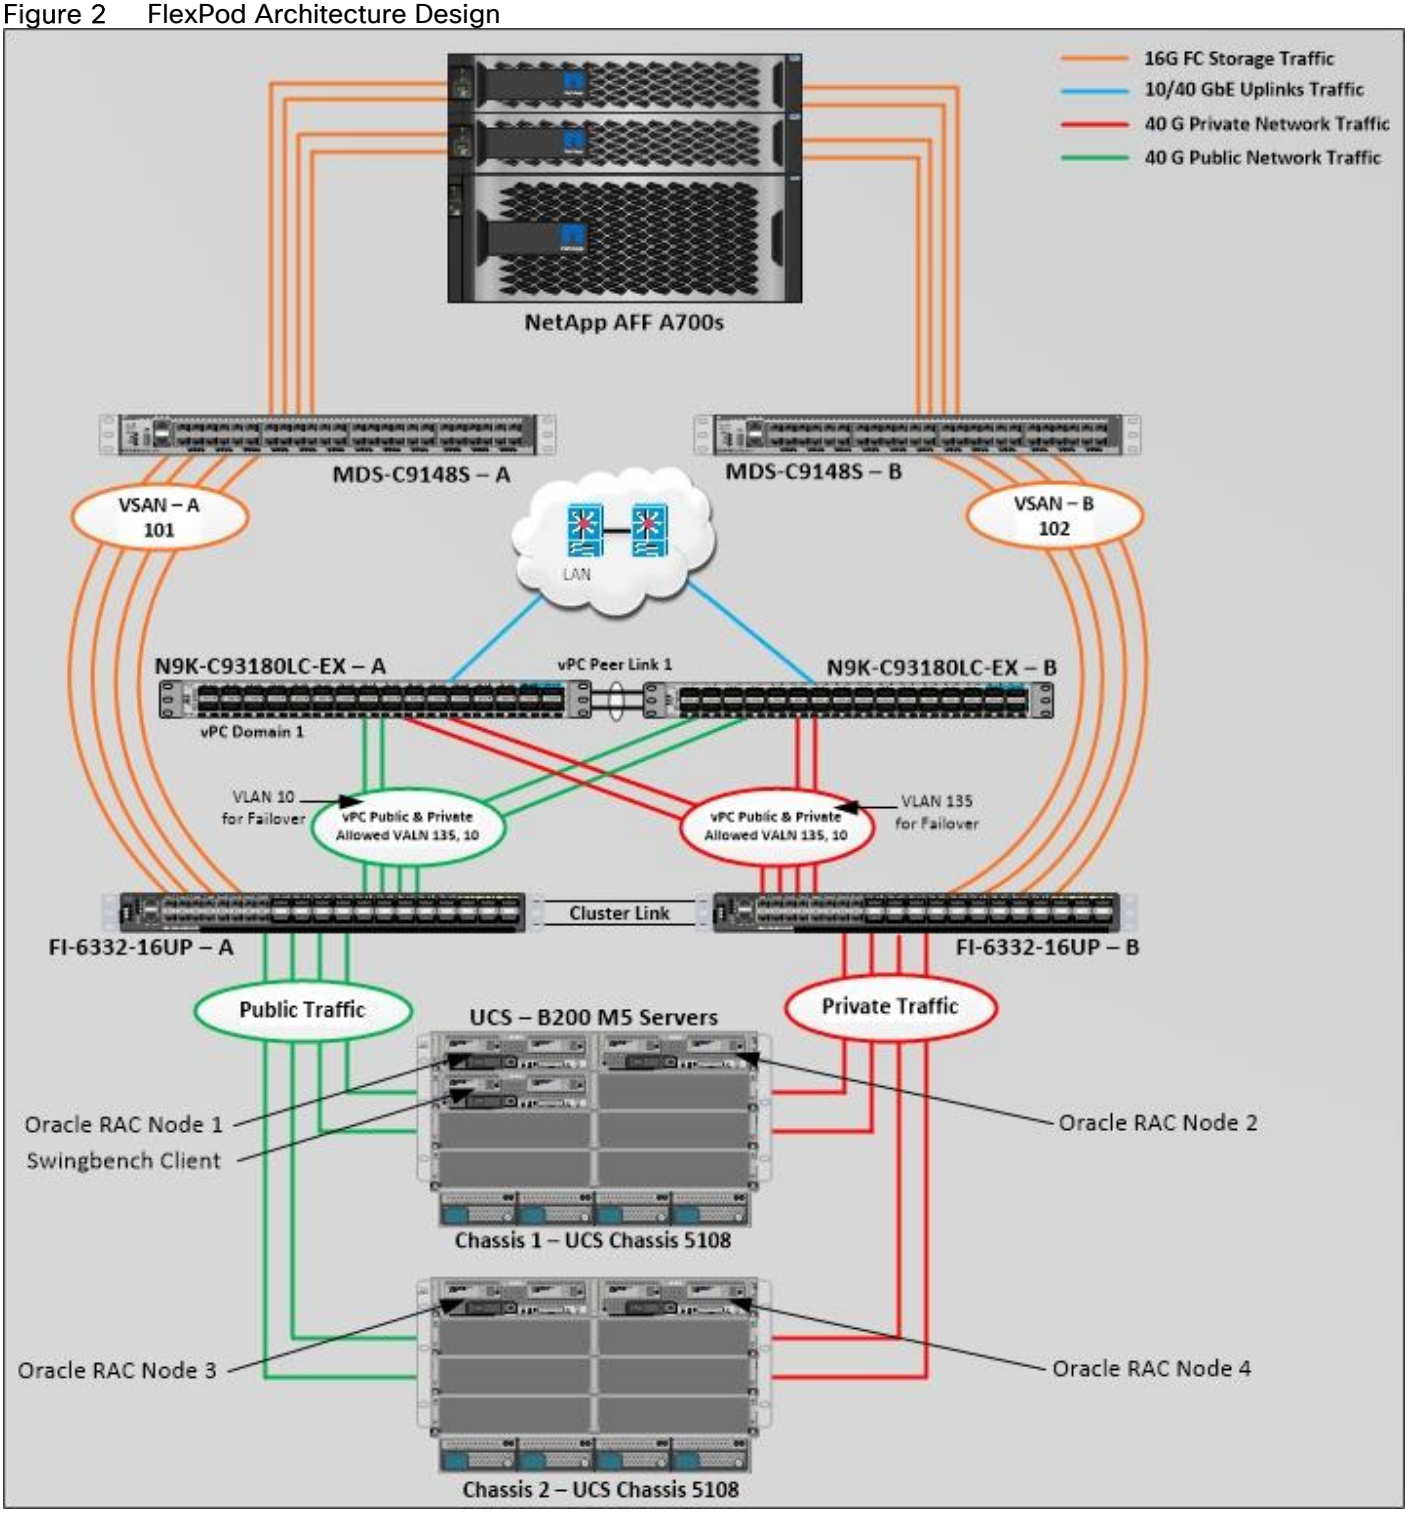


Figure 2 details the cable connections used in the validation lab for the 40Gb end-to-end with Fibre Channel topology based on the Cisco UCS 6332-16UP fabric interconnect. As shown in above architecture, a pair of Cisco UCS 6332-16UP Fabric Interconnects (FI) carries both storage and network traffic from the Cisco UCS B200 M5 server blades with the help of Cisco Nexus 93180LC-EX and Cisco MDS 9148S switches. Both the Fabric Interconnects and the Cisco Nexus switches 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 the above architecture, four (2 x 40G link per chassis) links go to Fabric Interconnect – A. Similarly, four (2 x 40G link per chassis) 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 orange lines.

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

Four 16Gb uplinks are connected from Cisco UCS Fabric Interconnect A to MDS switch A. Similarly, four 16Gb uplinks are connected from Cisco UCS Fabric Interconnect B to MDS switch B. The NetApp Storage AFF A700s have eight active FC connection goes to the Cisco MDS switches. Four FC ports are connected to Cisco MDS-A, and other four FC ports are connected to Cisco MDS-B Switches. The SAN Ports FC-Port-0-Slot-2 and FC-Port-0-Slot-3 of NetApp AFF A700s Controller – 1 are connected to Cisco MDS Switch A and FC-Port-1-Slot-2 and FC-Port-1-Slot-3 are connected to Cisco MDS Switch B. Similarly, the SAN Ports FC-Port-0-Slot-2 and FC-Port-0-Slot-3 of NetApp AFF A700s Controller – 2 are connected to Cisco MDS Switch A and FC-Port-1-Slot-2 and FC-Port-1-Slot-3 are connected to Cisco MDS Switch B.

Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has two connections to the out-of-band network switch.

The reference architecture includes the following hardware:

- Two Cisco UCS 5108 Blade Server Chassis
- Four Cisco UCS B200 M5 Blade Servers with Cisco Virtual Interface Cards (VIC)
- Two Cisco UCS 6332-16UP Fabric Interconnects
- Two Cisco MDS 9148S Multilayer Fabric Switch
- Two Cisco Nexus 93180LC-EX Switch
- One NetApp AFF A700s (HA pair) running ONTAP with Disk shelves and Solid State Drives (SSD)

Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps for deploying the base architecture, as shown in Figure 2. These procedures cover everything from physical cabling to network, compute and storage device configurations.

## Design Topology

This section describes the physical and logical high-level design considerations for this Oracle RAC 12cR2 Databases on FlexPod deployment.

The inventory of the components used in this Solution stack is listed in Table 1 .

Table 1 Hardware Components

| Name | Model | Description | Quantity |
|---|---|---|---|
| Cisco Nexus 93180LC-EX Switch | N9K-C93180LC-EX | Cisco Nexus 9300 Series Switches | 2 |
| Cisco MDS 9148S 16G Fabric Switch | DS-C9148S-12PK9 | Cisco MDS 9100 Series Multilayer Fabric Switches | 2 |
| Cisco UCS 6332-16UP Fabric Interconnect | UCS-FI-6332-16UP | Cisco 6300 Series Fabric Interconnects | 2 |
| Cisco UCS 5108 Blade Server Chassis | UCSB-5108-AC2 | Cisco UCS 5100 Series Blade Server AC2 Chassis | 2 |
| Cisco UCS Fabric Extender | UCS-IOM-2304 | Cisco UCS 2304XP I/O Module (4 External, 8 Internal 40Gb Ports) | 4 |
| Cisco UCS B200 M5 Blade Server | UCSB-B200-M5 | Cisco UCS B-Series Blade Servers | 4 |
| Cisco UCS VIC 1340 | UCSB-MLOM-40G-03 | Cisco UCS Virtual Interface Card 1340 | 4 |
| Cisco UCS Port Expander Card | UCSB-MLOM-PT-01 | Port Expander Card for Cisco UCS MLOM | 4 |
| NetApp Storage | AFF A700s | AFF A-Series All Flash Arrays | 1 |
| NetApp Disk | DS-224C | Disk Shelves and Storage Media for NetApp AFF and FAS systems | 2 |

Table 2 lists the Cisco UCS B200 M5 Blade Server Configuration.

Table 2 Cisco UCS B200 M5 Blade Server Configuration

| Server Configuration | | |
|---|---|---|
| Processor | 2 x Intel(R) Xeon(R) Gold 6148 2.40 GHz 150W 20C 27.50MB Cache DDR4 2666MHz 768GB | UCS-CPU-6148 |
| Memory | 16 x 32GB DDR4-2666-MHz RDIMM/dual rank/x4/1.2v | UCS-MR-X32G2RS-H |
| Cisco UCS VIC 1340 | Cisco UCS VIC 1340 Blade MLOM | UCSB-MLOM-40G-03 |
| Cisco UCS Port Expander Card | Port Expander Card for Cisco UCS MLOM | UCSB-MLOM-PT-01 |

For this FlexPod solution design, we configured two VLAN and two VSAN as listed below:

Table 3 VLAN and VSAN Configurations

| VLAN & VSAN Configuration | | |
|---|---|---|
| VLAN | | |
| Name | ID | Description |
| • Default VLAN | 1 | Native VLAN |
| • Public VLAN | 135 | VLAN for Public Network Traffic |
| • Private VLAN | 10 | VLAN for Private Network Traffic |
| VSAN | | |
| Name | ID | Description |
| • VSAN – A | 101 | SAN Communication through for Fabric Interconnect A |
| • VSAN – B | 102 | SAN Communication through for Fabric Interconnect B |

This FlexPod solution consists of NetApp All-Flash AFF-Series Storage as listed in Table 4 .

Table 4 NetApp AFF A700s Storage Configuration

| Storage Components | Description |
|---|---|
| Flash Arrays | NetApp All Flash AFF A700s Storage Array (24 x 960GB SSD Drives ) |
| Disk Shelves | 2 x NetApp DS224C (24 x 960GB SSD Drives per Shelves) |
| Capacity | 62.63 TB |
| Connectivity | 8 x 16 Gb/s redundant Fibre Channel<br>1 Gb/s redundant Ethernet (Management port ) |
| Physical | 8U Rack Units |

For this FlexPod solution, we used the following versions of the software and firmware releases.

Table 5 Software and Firmware Revisions

| Software and Firmware | Version |
|---|---|
| Cisco UCS Manager System | 3.2 (3c) |
| Cisco UCS Adapter VIC 1340 | 4.2 (3b) |
| Cisco eNIC (modinfo enic) (Cisco VIC Ethernet NIC Driver) | 2.3.0.53 |
| Cisco fNIC (modinfo fnic) (Cisco FCoE HBA Driver) | 1.6.0.34 |
| Cisco Nexus 93180LC-EX NXOS Version | 7.0(3)I6(1) |
| Cisco MDS 9148S Switch System Version | 7.3(0)D1(1) |
| Oracle Linux Server Release 7.5 (64 bit) operating System | Linux 4.1.12-124.16.4.el7uek.x86_64 |
| Oracle Database 12c Release 2 Grid Infrastructure for Linux x86-64 | 12.2.0.1.0 |
| Oracle Database 12c Release 2 for Linux x86-64 | 12.2.0.1.0 |
| NetApp Storage AFF A700s System Version | ONTAP 9.3 |
| Oracle Swingbench | 2.5.971 |
| SLOB | 2.4.2 |

# Solution Configuration

## Configure Cisco Nexus 93180LC-EX Switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 93180LC-EX 7.0(3)I6(1) switches deployed with the 40Gb end-to-end topology.

### Set Up Initial Configuration

#### Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>, complete the following steps:

1. Configure the switch.

> **On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.**

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes

Do you want to enforce secure password standard (yes/no) [y]: Enter

Enter the password for "admin": <password>

Confirm the password for "admin": <password>

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name: <nexus-A-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>

Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway: <nexus-A-mgmt0-gw>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter
```

```
Configure the ntp server? (yes/no) [n]: y

NTP server IPv4 address: <global-ntp-server-ip>

Configure default interface layer (L3/L2) [L3]: L2

Configure default switchport interface state (shut/noshut) [noshut]: Enter

Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter

Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

## Cisco Nexus B

To set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>, complete the following steps:

1. Configure the switch.

> **On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.**

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes

Do you want to enforce secure password standard (yes/no) [y]: Enter

Enter the password for "admin": <password>

Confirm the password for "admin": <password>

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name: <nexus-B-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>

Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway: <nexus-B-mgmt0-gw>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
```

```
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [noshut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

## Configure Global Settings for Cisco Nexus A and Cisco Nexus B

To set global configuration, follow these steps on both the nexus switches:

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

```
configure terminal
feature interface-vlan
feature hsrp
feature lacp
feature vpc
spanning-tree port type network default
spanning-tree port type edge bpduguard default
port-channel load-balance src-dst l4port
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
system qos
  service-policy type network-qos jumbo
vrf context management
  ip route 0.0.0.0/0 10.29.135.1
copy run start
```

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

**Make sure to run copy run start to save the configuration on each switch after the configuration is completed.**

## Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

To create the necessary virtual local area networks (VLANs), follow these steps on both Nexus switches:

1. Log in as admin user into the Nexus Switch A.
2. Create VLAN 135 for Public Network Traffic and VLAN 10 for Private Network Traffic.

```
configure terminal
vlan 135
name Oracle_RAC_Public_Traffic
no shutdown
vlan 10
name Oracle_RAC_Private_Traffic
no shutdown
copy run start
```

3. Log in as admin user into the Nexus Switch B and create VLAN 135 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 93180LC-EX switch topology, a single vPC feature is enabled to provide HA, faster convergence in the event of a failure, and greater throughput. Cisco Nexus vPC configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers is shown in Table 6

**Table 6 vPC Summary**

| vPC Domain | vPC Name | vPC ID |
|---|---|---|
| 1 | Peer-Link | 1 |
| 1 | vPC Prublic | 51 |
| 1 | vPC Private | 52 |

As listed in Table 6 , a single vPC domain with Domain ID 1 is created across two Cisco Nexus 93180LC-EX 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.

To define vPC IDs 51 and 52 for public and private network traffic from Cisco UCS fabric interconnects, complete the steps in the following sections.

## Create vPC Peer-Link Between the Two Nexus Switches

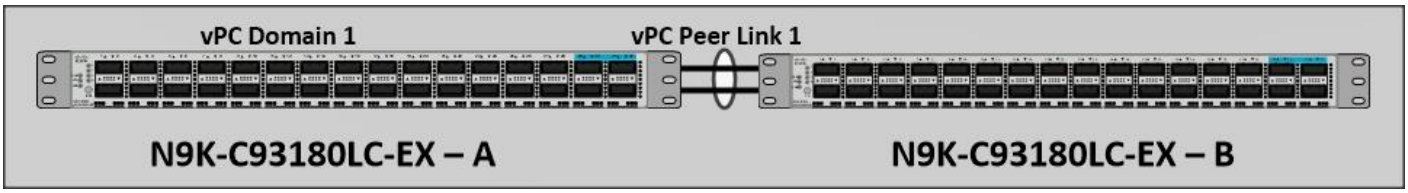


1. Login as admin user into the Nexus Switch A.

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

2. Create the necessary port channels between devices on both Nexus Switches:

```
configure terminal
vpc domain 1
peer-keepalive destination 10.29.135.104 source 10.29.135.103
auto-recovery

interface port-channel1
description vPC peer-link
switchport mode trunk
switchport trunk allowed vlan 1,10,135
spanning-tree port type network
vpc peer-link
exit

interface Ethernet1/1
description Peer link connected to N9K-B-Eth1/1
switchport mode trunk
switchport trunk allowed vlan 1,10,135
channel-group 1 mode active
no shutdown
interface Ethernet1/2
description Peer link connected to N9K-B-Eth1/2
switchport mode trunk
switchport trunk allowed vlan 1,10,135
```

```
channel-group 1 mode active
no shutdown

interface Ethernet1/29
description connect to uplink switch
switchport access vlan 135
speed 1000
exit
copy run start
```

3. Log in 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

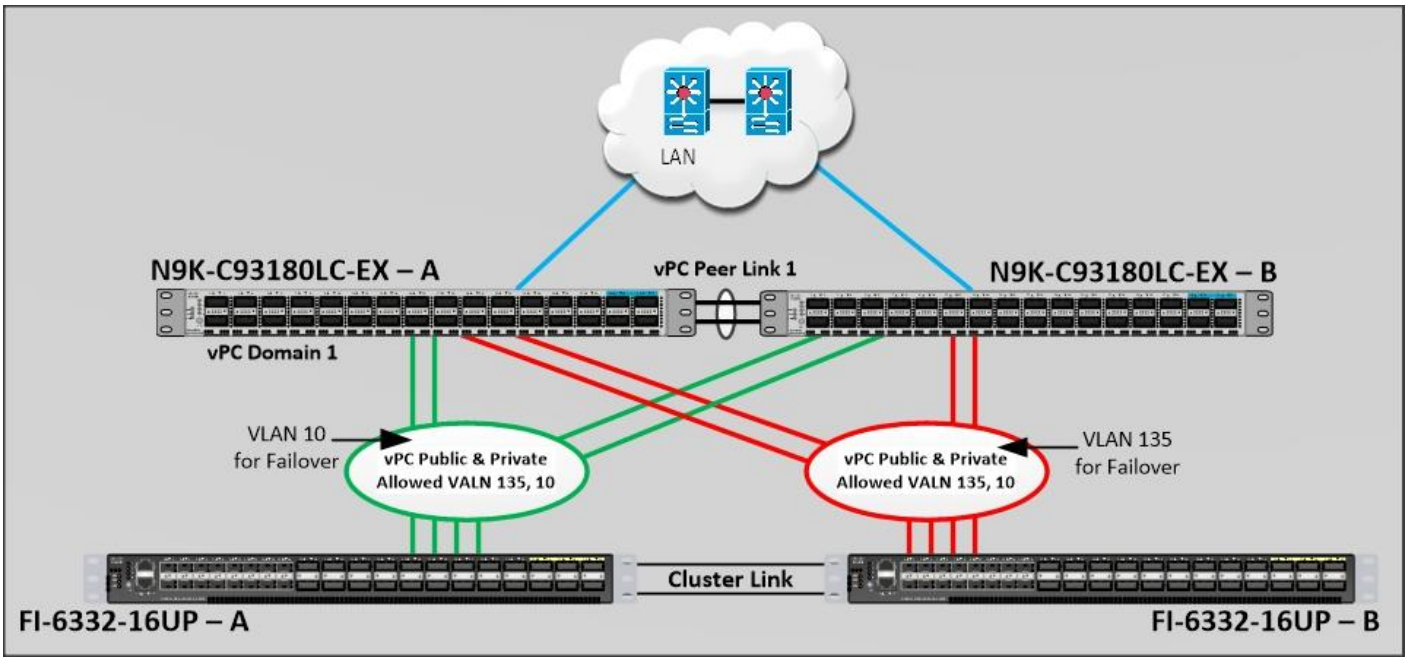


Next, you will create and configure vPC 51 and 52 for Data network between Nexus switches and Fabric Interconnects.

Table 7 lists the vPC IDs, allowed VLAN IDs, and Ethernet uplink ports.

**Table 7 vPC IDs and VLAN IDs**

| vPC Description | vPC ID | Fabric Interconnects Ports | Nexus Ports | Allowed VLANs |
|---|---|---|---|---|
| Port Channel FI-A | 51 | FI-A Port 1/31 | N9K-A Port 1/21 | 135, 10 |
| | | FI-A Port 1/32 | N9K-A Port 1/22 | |

| vPC Description | vPC ID | Fabric Interconnects Ports | Nexus Ports | Allowed VLANs |
|---|---|---|---|---|
| | | FI-A Port 1/33 | N9K-B Port 1/21 | Note: VLAN 10 Needed for Failover |
| | | FI-A Port 1/34 | N9K-B Port 1/22 | |
| Port Channel FI-B | 52 | FI-B Port 1/31 | N9K-A Port 1/23 | 10, 135 Note: VLAN 135 Needed for Failover |
| | | FI-B Port 1/32 | N9K-A Port 1/24 | |
| | | FI-B Port 1/33 | N9K-B Port 1/23 | |
| | | FI-B Port 1/34 | N9K-B Port 1/24 | |

To create the necessary port channels between devices, follow these steps on both the Nexus Switches:

1. Log in as admin user into the Nexus Switch A and follow these steps:

```
configure terminal
interface port-channel51
description Port-Channel FI-A
switchport mode trunk
switchport trunk allowed vlan 1,10,135
spanning-tree port type edge trunk
mtu 9216
vpc 51
no shutdown

interface port-channel52
description Port-Channel FI-B
switchport mode trunk
switchport trunk allowed vlan 1,10,135
spanning-tree port type edge trunk
mtu 9216
vpc 52
no shutdown

interface Ethernet1/21
  description Fabric-Interconnect-A-31
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active
```

```
  no shutdown

interface Ethernet1/22
  description Fabric-Interconnect-A-32
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active
  no shutdown

interface Ethernet1/23
  description Fabric-Interconnect-B-31
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 52 mode active
  no shutdown

interface Ethernet1/24
  description Fabric-Interconnect-B-32
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 52 mode active
  no shutdown

copy run start
```

2. Log in as admin user into the Nexus Switch B and run the following commands to complete second switch configuration:

```
configure terminal
interface port-channel51
description Port-Channel FI-A
switchport mode trunk
switchport trunk allowed vlan 1,10,135
spanning-tree port type edge trunk
mtu 9216
vpc 51
no shutdown

interface port-channel52
description Port-Channel FI-B
switchport mode trunk
switchport trunk allowed vlan 1,10,135
spanning-tree port type edge trunk
mtu 9216
vpc 52
no shutdown

interface Ethernet1/21
  description Fabric-Interconnect-A-33
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active
  no shutdown

interface Ethernet1/22
  description Fabric-Interconnect-A-34
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
```

```
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active
  no shutdown

interface Ethernet1/23
  description Fabric-Interconnect-B-33
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 52 mode active
  no shutdown

interface Ethernet1/24
  description Fabric-Interconnect-B-34
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 52 mode active
  no shutdown

copy run start
```

**Note: Make sure to change peer-keepalive destination and source IP address appropriately for Nexus Switch B**

## Verify All vPC Status Is Up on Both the Nexus Switches

1. Verify the port-channel summary of Nexus Switch A as shown below:

```
FLEXPOD-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
        b - BFD Session Wait
        S - Switched    R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        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)
51    Po51(SU)    Eth      LACP      Eth1/21(P)   Eth1/22(P)
52    Po52(SU)    Eth      LACP      Eth1/23(P)   Eth1/24(P)
FLEXPOD-NEXUS-A#
```

2. Verify the port-channel summary of Nexus Switch B as shown below:

```
FLEXPOD-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
        b - BFD Session Wait
        S - Switched    R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        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)
51    Po51(SU)    Eth      LACP      Eth1/21(P)   Eth1/22(P)
52    Po52(SU)    Eth      LACP      Eth1/23(P)   Eth1/24(P)
FLEXPOD-NEXUS-B#
```

3. Verify the vPC summary of Nexus Switch A as shown below:

```
FLEXPOD-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 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id     Port   Status Active vlans
--     ----   ------ -------------------------------------------------
1      Po1    up     1,10,135

vPC status
----------------------------------------------------------------------------
Id     Port          Status Consistency Reason                Active vlans
--     ------------  ------ ----------- ------                ---------------
51     Po51          up     success     success               1,10,135

52     Po52          up     success     success               1,10,135
```

4. Verify the vPC summary of Nexus Switch B as shown below:

```
FLEXPOD-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 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id     Port   Status Active vlans
--     ----   ------ -------------------------------------------------
1      Po1    up     1,10,135

vPC status
----------------------------------------------------------------------------
Id     Port          Status Consistency Reason                Active vlans
--     ------------  ------ ----------- ------                ---------------
51     Po51          up     success     success               1,10,135

52     Po52          up     success     success               1,10,135
```

# Cisco UCS Configuration Overview

This section details the Cisco UCS configuration that was completed as part of the infrastructure build out. The racking, power, and installation of the chassis are described in the installation 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 the 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

We will list all the tasks to configure Cisco UCS system but only include some of the screenshots in this document.

## Perform Initial Setup of Cisco UCS 6332-16UP Fabric Interconnects for a Cluster Setup

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlexPod environment. The steps are necessary to provision the Cisco UCS B-Series and C-Series servers and should be followed precisely to avoid improper configuration.

### Cisco UCS Fabric Interconnect A and Fabric Interconnect B

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.

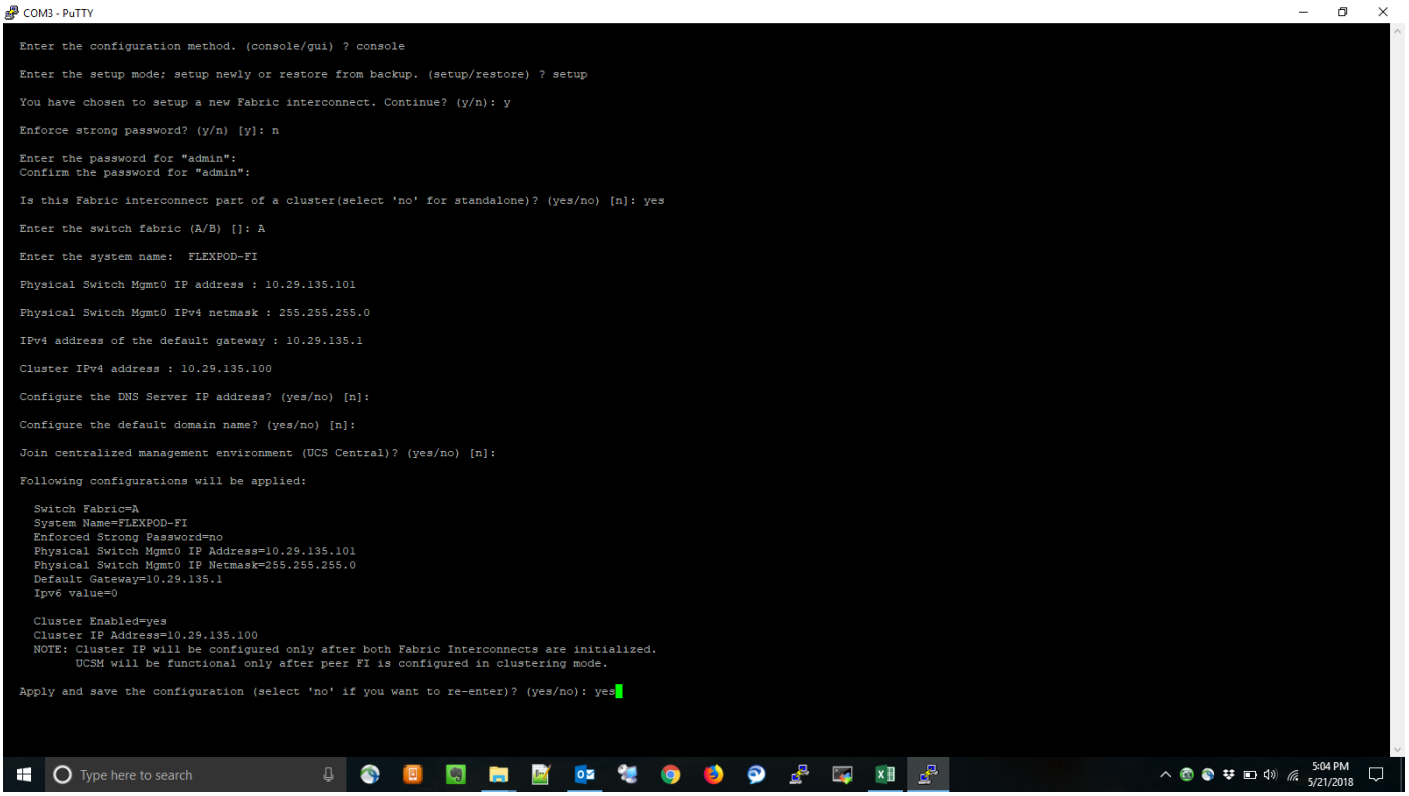


3. Review the settings printed to the console. Answer yes to apply and save the configuration.

4. Wait for the login prompt to make the configuration has been saved to Fabric Interconnect A.

5. Connect the console port on the second Fabric Interconnect and do the following:

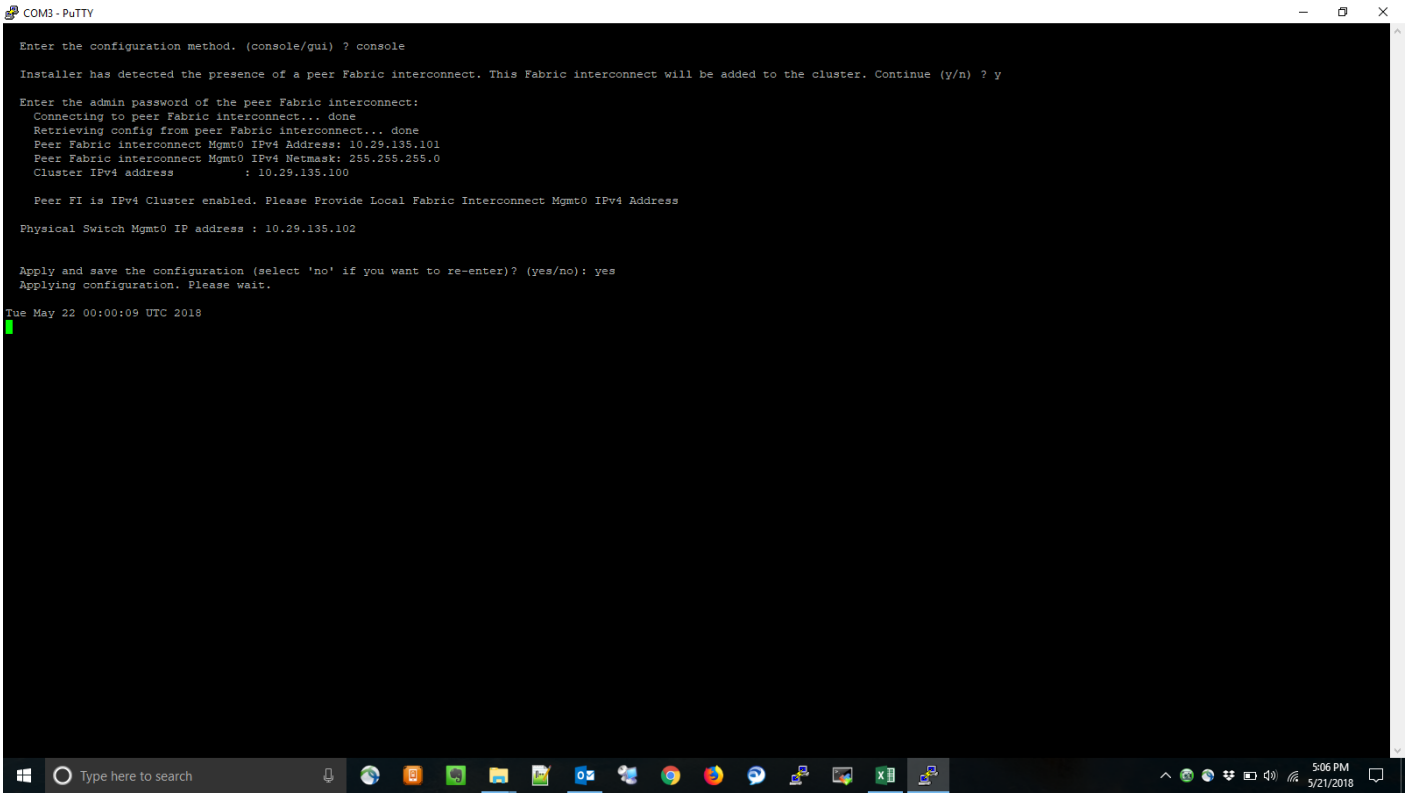


6. Review the settings printed to the console. Answer yes to apply and save the configuration.

7. Wait for the login prompt to make the configuration has been saved to Fabric Interconnect B.

## Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS fabric interconnect cluster address.

> You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to come up.

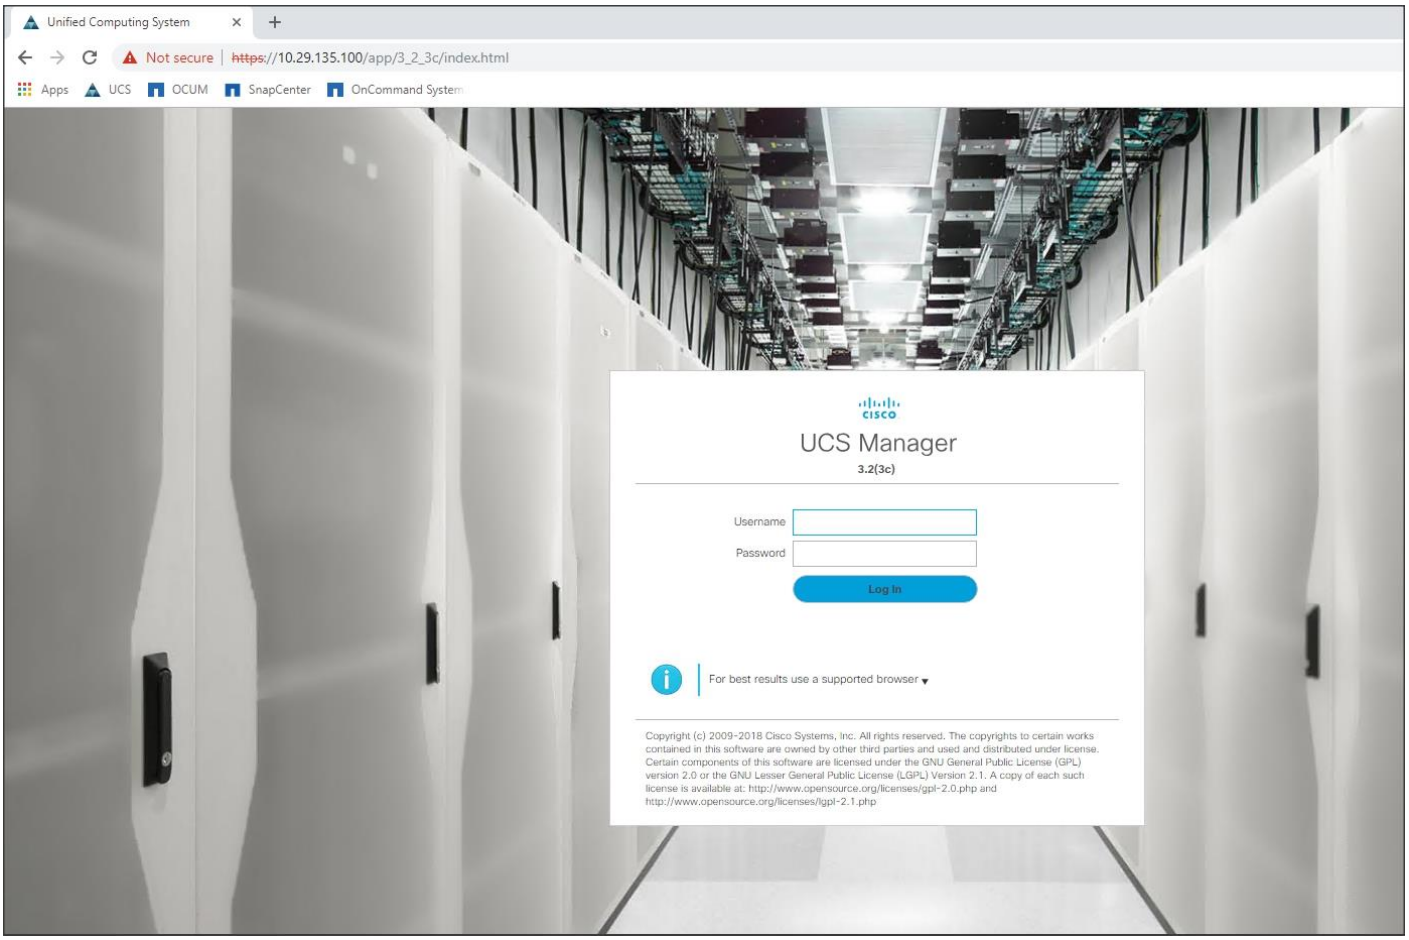


2. Click the Launch UCS Manager link under HTML to launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log into Cisco UCS Manager.

## Upgrade Cisco UCS Manager Software to Version 3.2

This document assumes the use of Cisco UCS 3.2(3c). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.2(3c), refer to Cisco UCS Manager Install and Upgrade Guides.

## Anonymous Reporting

To create anonymous reporting, complete the following step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products. If you select Yes, enter the IP address of your SMTP Server. Click OK.

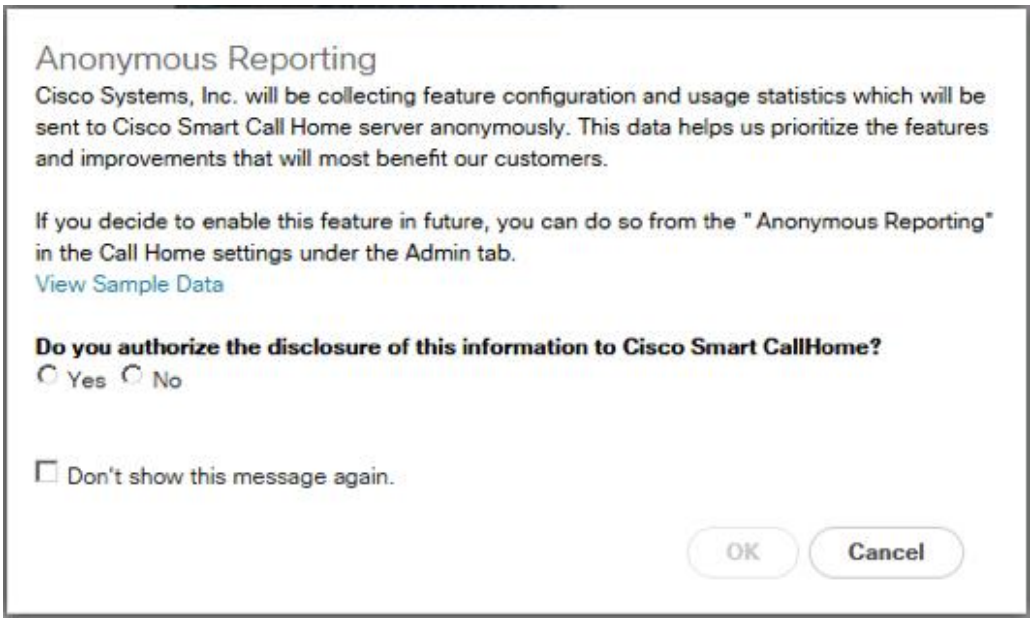


## Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in Cisco UCS Manager. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, follow these steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

# High Level Steps to Configure Base Cisco UCS

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.

The following are the high-level steps for configuring the Cisco UCS for this FlexPod solution.

1. Set Fabric Interconnect to Fibre Channel End Host Mode.
2. Synchronize Cisco UCS to NTP.
3. Configure Fabric Interconnect for Chassis and Blade Discovery.

   a. Configure Global Policies
   b. Configure Server Ports
4. 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
5. 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
6. Set Jumbo Frames in both the Cisco Fabric Interconnect
7. Configure Server BIOS Policy
8. Create Adapter Policy
9. Configure Update Default Maintenance Policy
10. Configure vNIC and vHBA Template
    a. Create Public and Private vNIC Template
    b. Create Storage vHBA Template
11. Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

## 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 Interconnect 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 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 insures that Cisco UCS Manager uses the maximum number of IOM uplinks available.

To configure Global Policies, follow these steps:

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

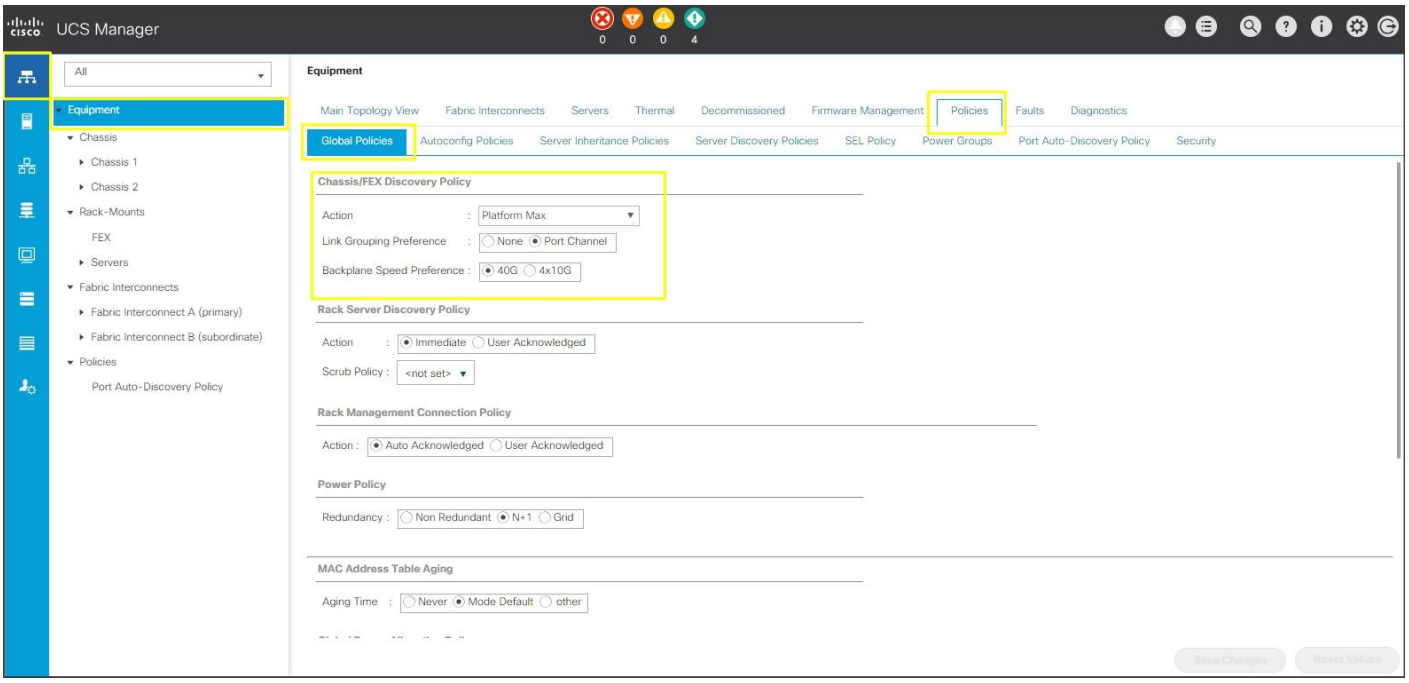


The difference between Discrete mode vs Port Channel mode is shown below:

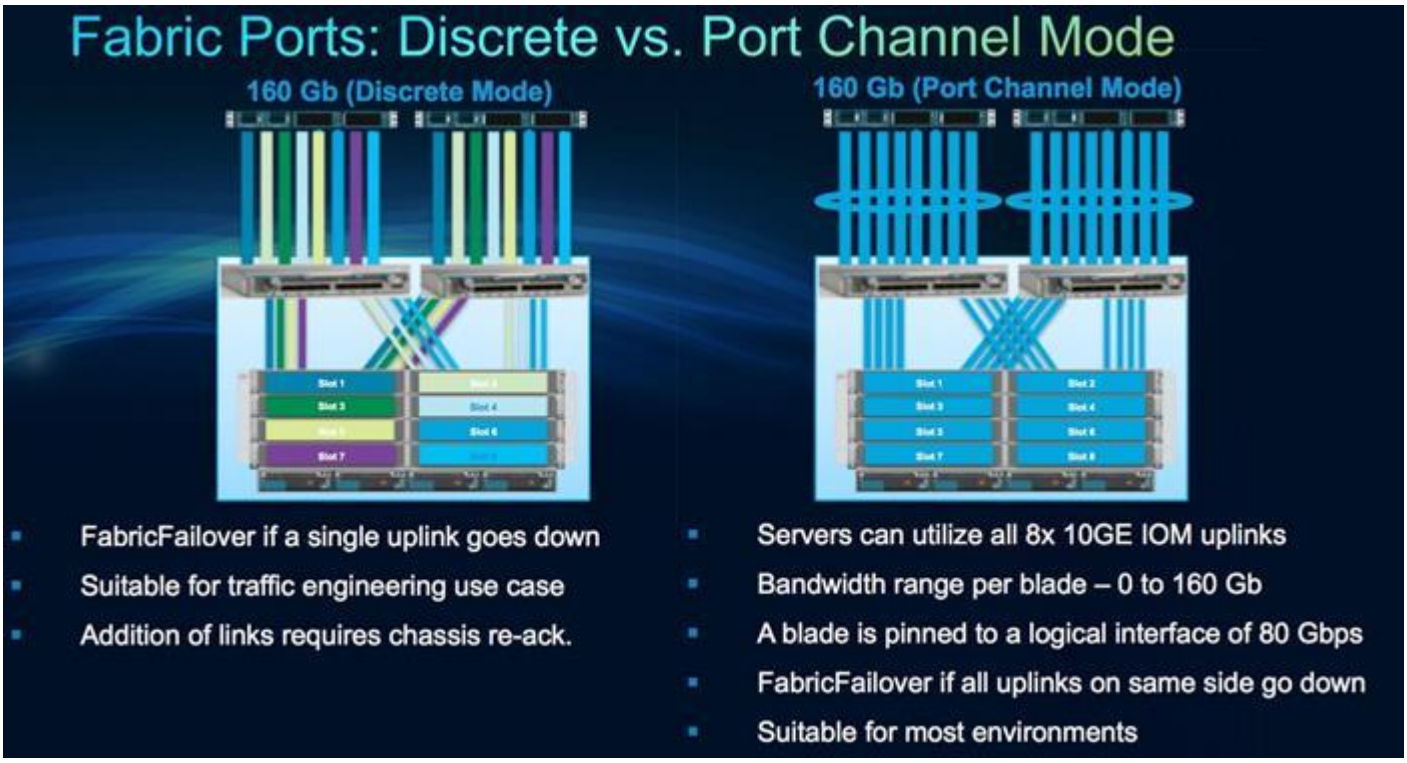


## Configure Server Ports

Configure Server Ports to initiate Chassis and Blade discovery. To configure server ports, follow these steps:

1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > Ethernet Ports.

2. Select the ports (for this solution ports are 17-20) 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.

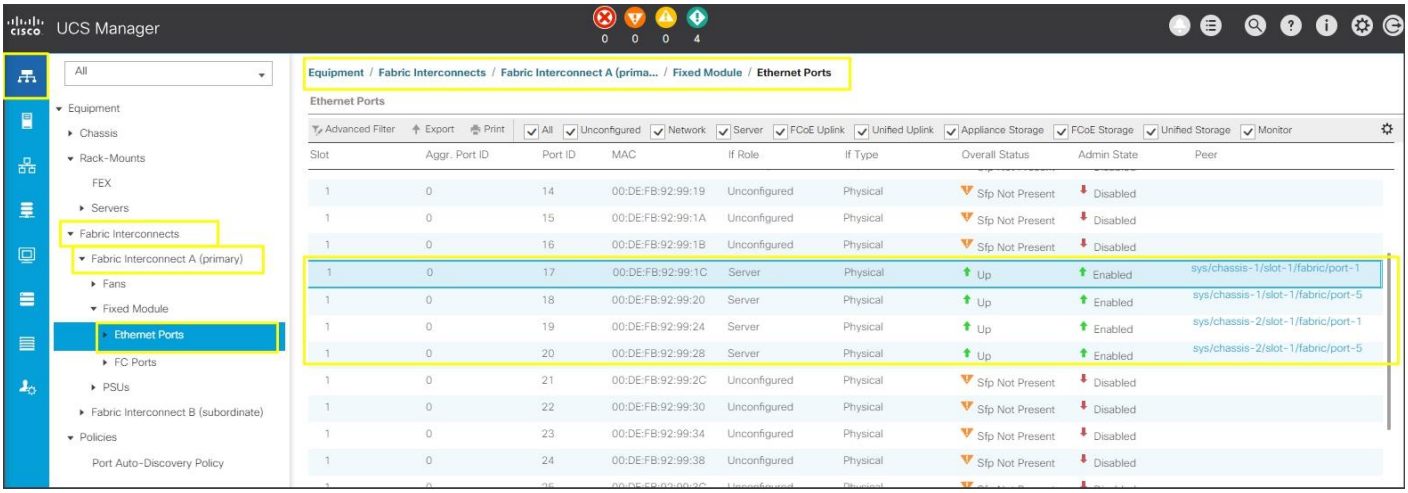


5. Repeat the same task for Fabric Interconnect B.

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. Similarly, repeat the process to Re-acknowledge all the eight Servers.

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

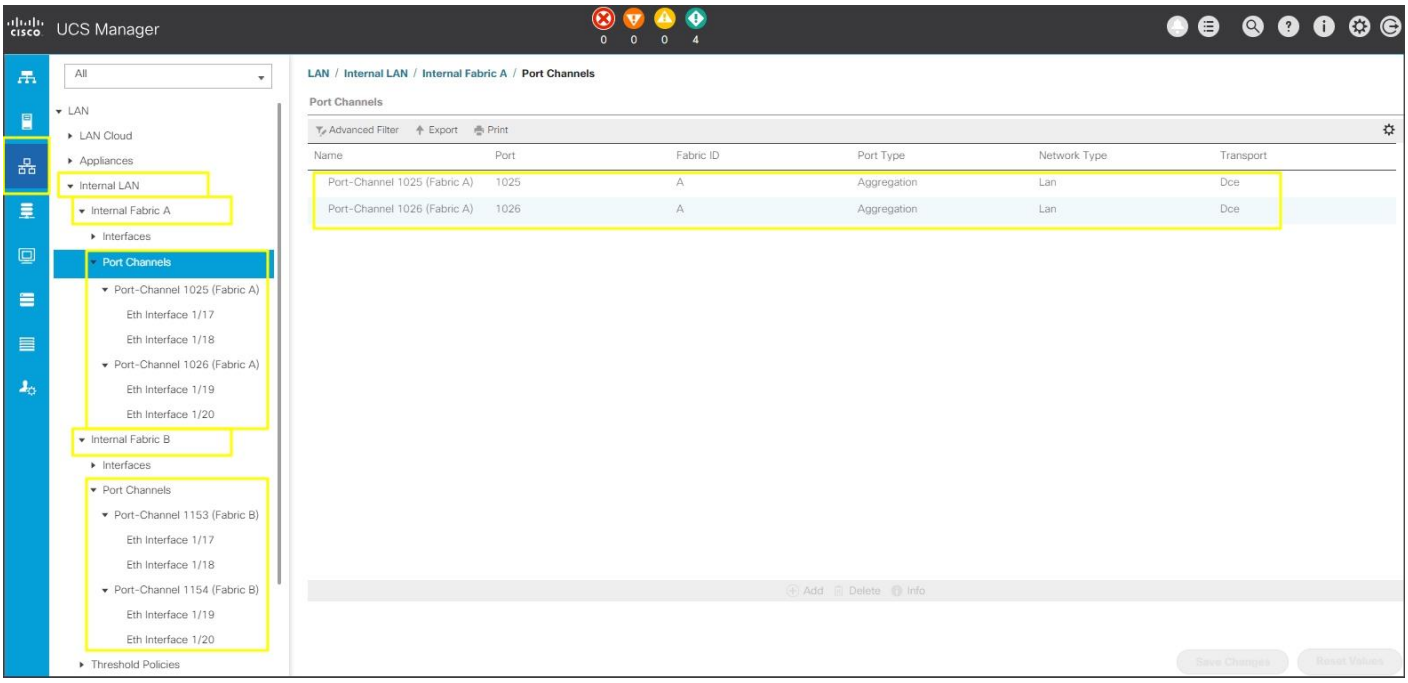


9. Verify the same for Internal Fabric B.

**The last 6 ports of the UCS 6332 and UCS 6332-16UP FIs will only work with optical based QSFP transceivers and AOC cables, so they can be better utilized as uplinks to upstream resources that might be optical only.**

## Configure LAN and SAN on Cisco UCS Manager

Configure Ethernet Uplink Ports and Fibre Channel (FC) Storage ports on Cisco UCS as explained below.

### Configure Ethernet LAN Uplink Ports

To configure network ports used to uplink the Fabric Interconnects to the Cisco Nexus switches, follow these 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 31-34) 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 shows the network uplink ports for Fabric A.

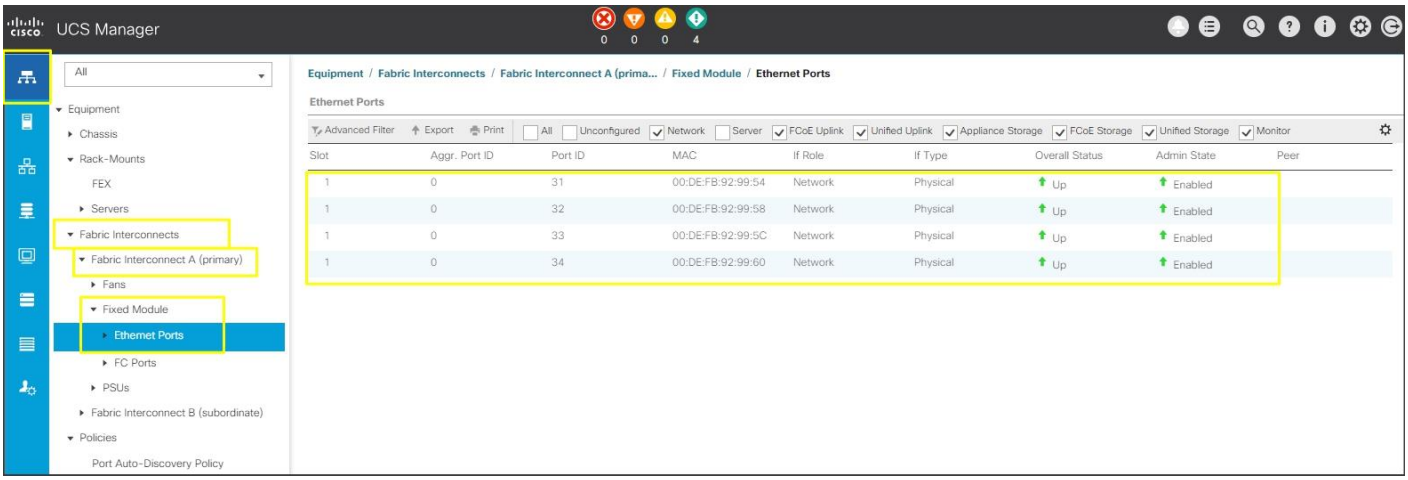


We created four uplink ports on each Fabric Interconnect as shown above. These ports will be used to create Virtual Port Channel in the next section.

**The last 6 ports of the UCS 6332 and UCS 6332-16UP FIs only work with optical based QSFP transceivers and AOC cables, so they can be better utilized as uplinks to upstream resources that might be optical only.**

## Create Uplink Port Channels to Cisco Nexus Switches

In this procedure, two port channels were created: one from Fabric A to both Cisco Nexus switch and one from Fabric B to both Cisco Nexus switch. To configure the necessary port channels in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, expand node Fabric A tree:
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 51 as the unique ID of the port channel.

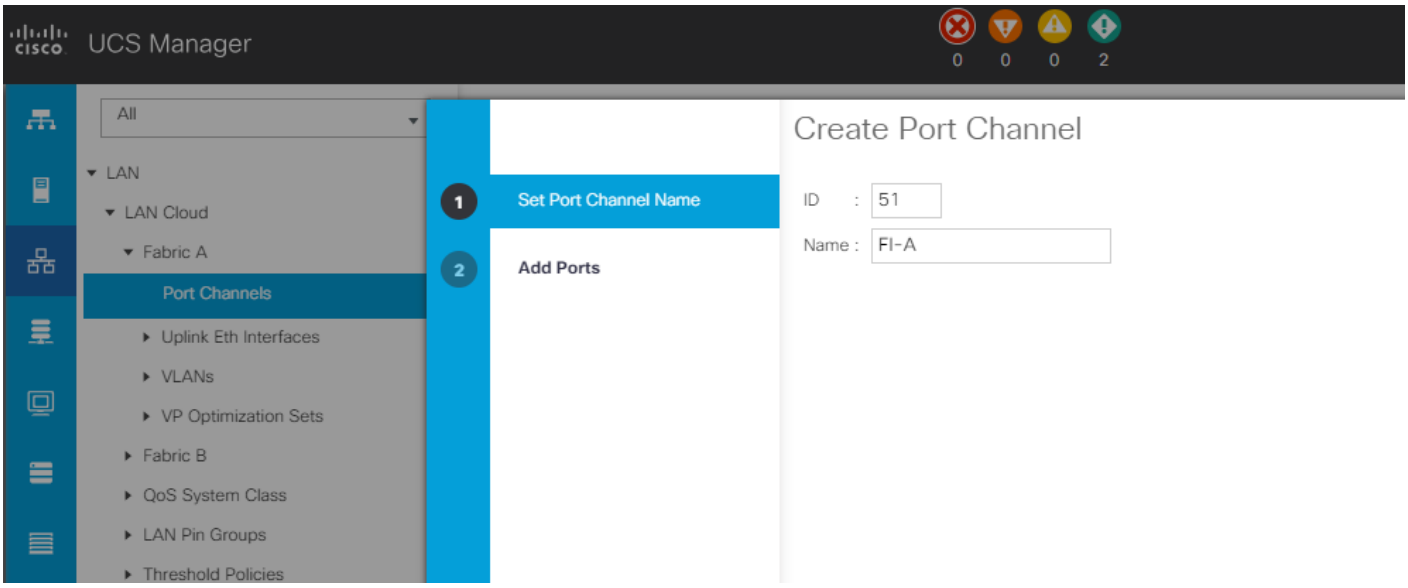


6. Enter FI-A as the name of the port channel.

7. Click Next.

8. Select Ethernet ports 31-34 for the port channel.

9. Click >> to add the ports to the port channel.

10. Click Finish to create the port channel and then click OK.

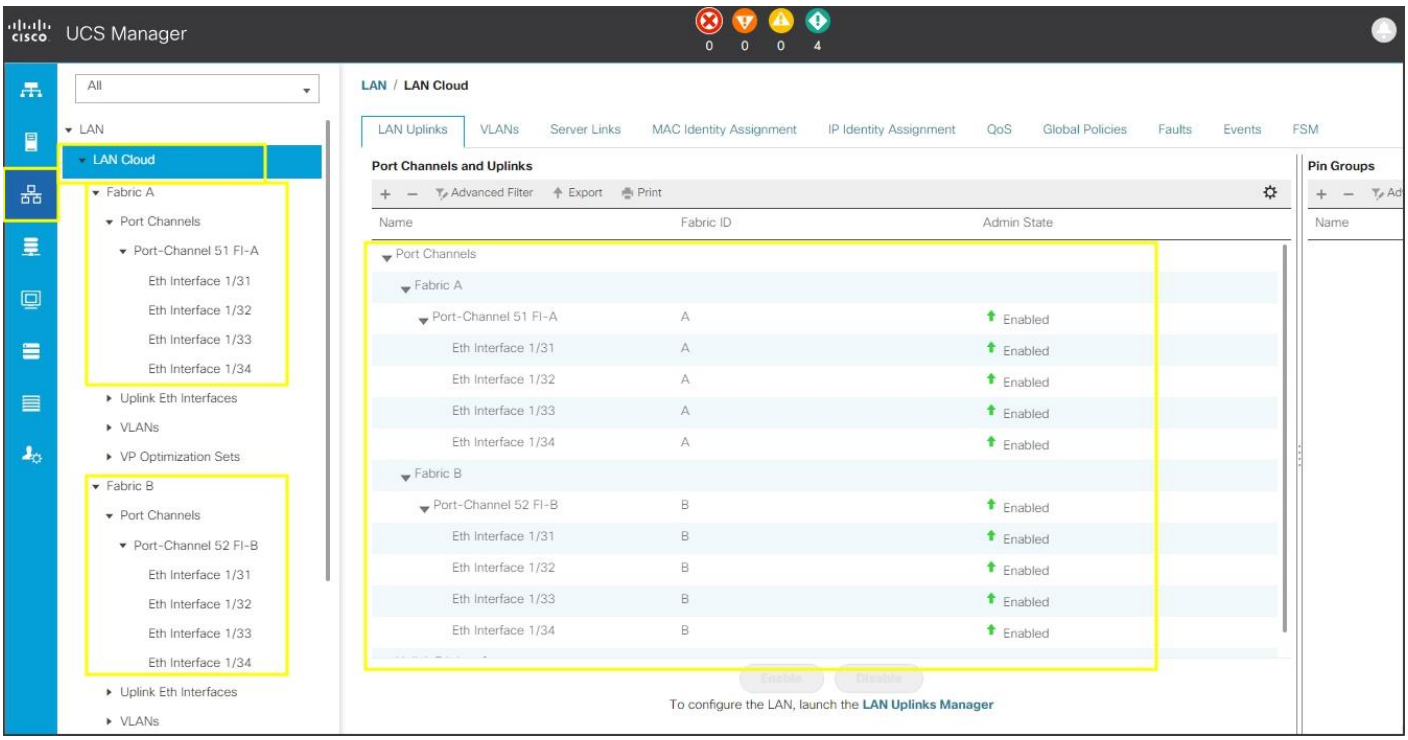


11. Repeat steps 1-10 for Fabric Interconnect B, substituting 52 for the port channel number and FI-B for the name. The resulting configuration should look like the screenshot shown above.

## Configure FC SAN Uplink Ports

The fibre channel port selection options for the 6332-16UP are from the first 16 ports starting from the first port on the left, and configured in increments of the first 6, 12, or all 16 of the unified ports.

To enable the Fibre channel ports, follow these steps for the 6332-16UP:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. Select Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
5. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 6, 12, or 16 ports to be set as FC Uplinks.

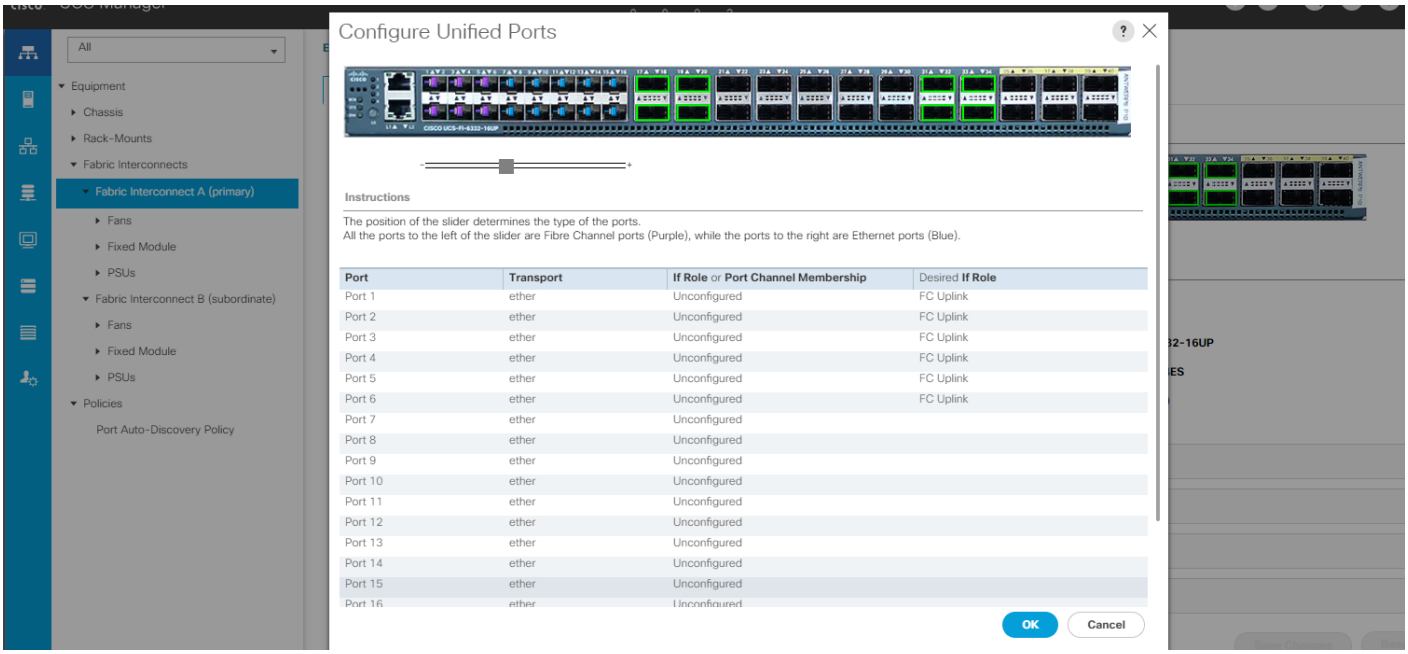


6. For this solution, we configured the first six ports on the FI as FC Uplink ports. Click OK, then click Yes, then click OK to continue

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

7. Select Equipment > Fabric Interconnects > Fabric Interconnect B (primary).
8. Select Configure Unified Ports.
9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.

10. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 6, 12, or 16 ports to be set as FC Uplinks.

11. Click OK, then Yes, then OK to continue.

12. Wait for both Fabric Interconnects to reboot.

13. Log back into Cisco UCS Manager.

## Configure VLAN

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, follow these 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 135) 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

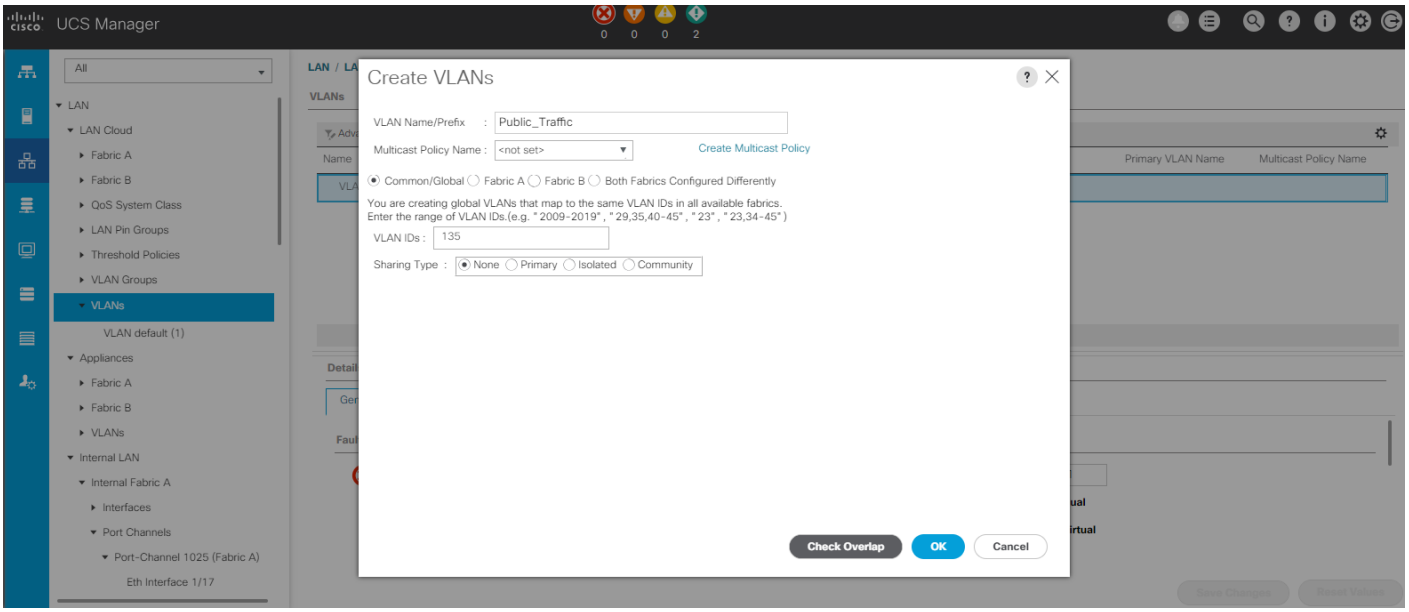


5. Enter Public_Traffic as the name of the VLAN to be used for Public Network Traffic.

6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 135 as the ID of the VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.

**We have also created the second VLAN: for private network (VLAN 10) traffic.**

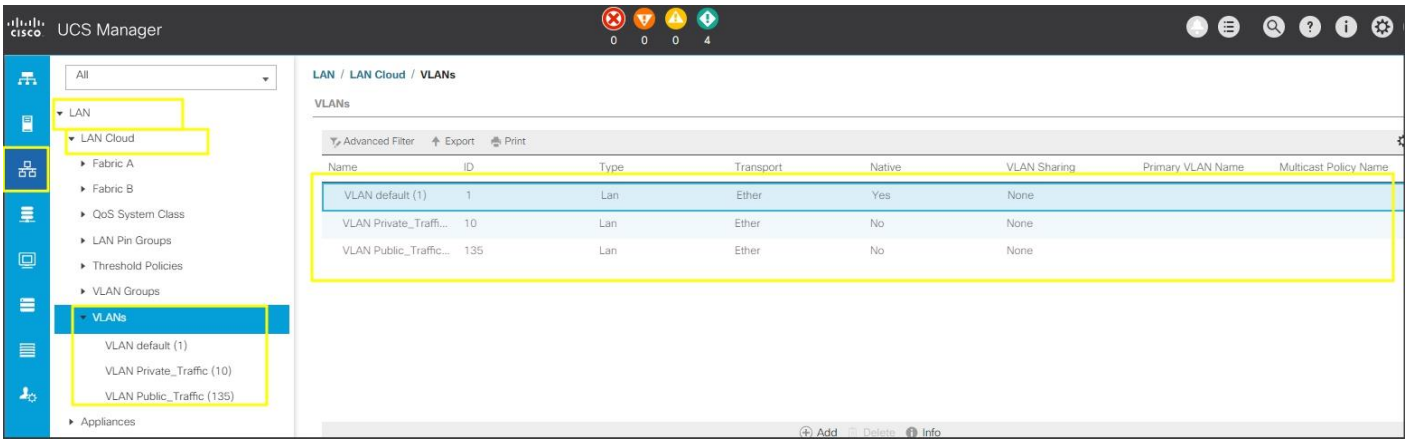


**These two VLANs will be used in the vNIC templates that are discussed later.**

## Configure VSAN

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

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

**In this solution, we have created two VSANs. ORA-VSAN-A 101 and ORA-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 ORA-VSAN-A as the name of the VSAN.
6. Leave FC Zoning set at Disabled.

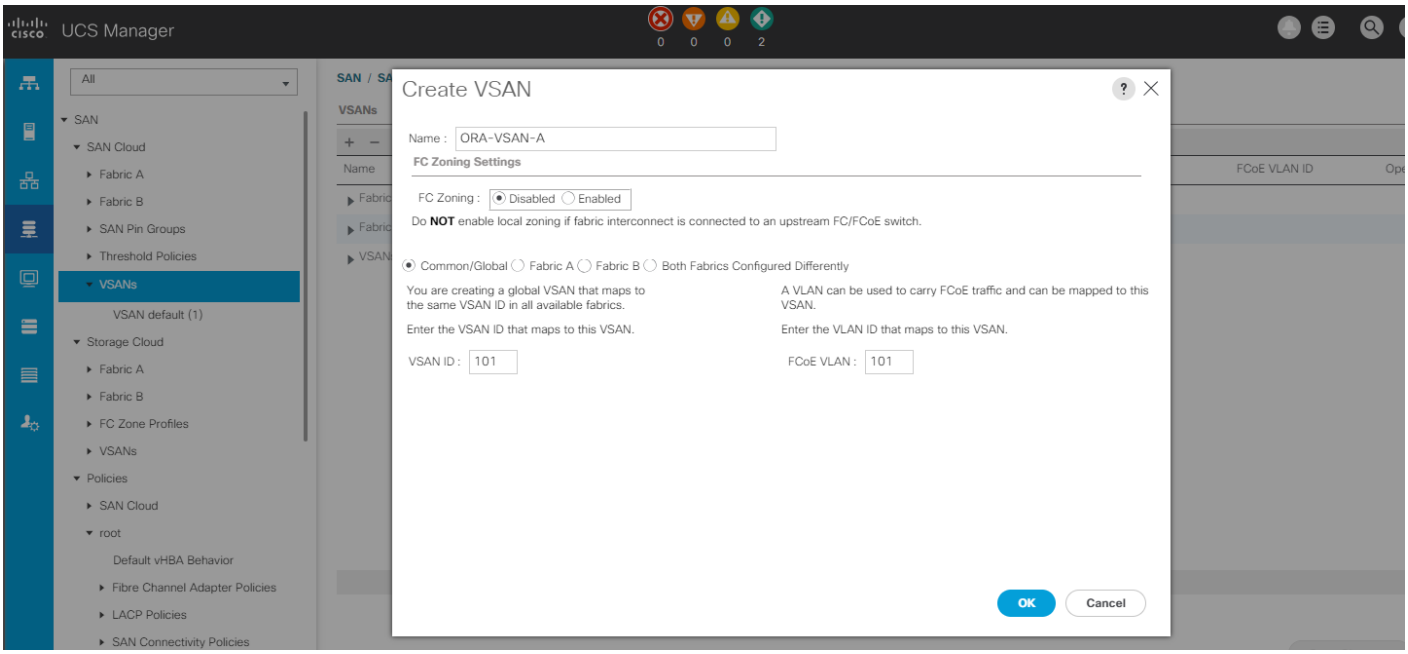


7. Select Fabric A for the scope of the VSAN.

8. Enter 101 as the ID of the VSAN.

> Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

9. Click OK and then click OK again.

10. Repeat these steps to create the VSANs necessary for this solution. VSAN 101 and 102 are configured as shown below:

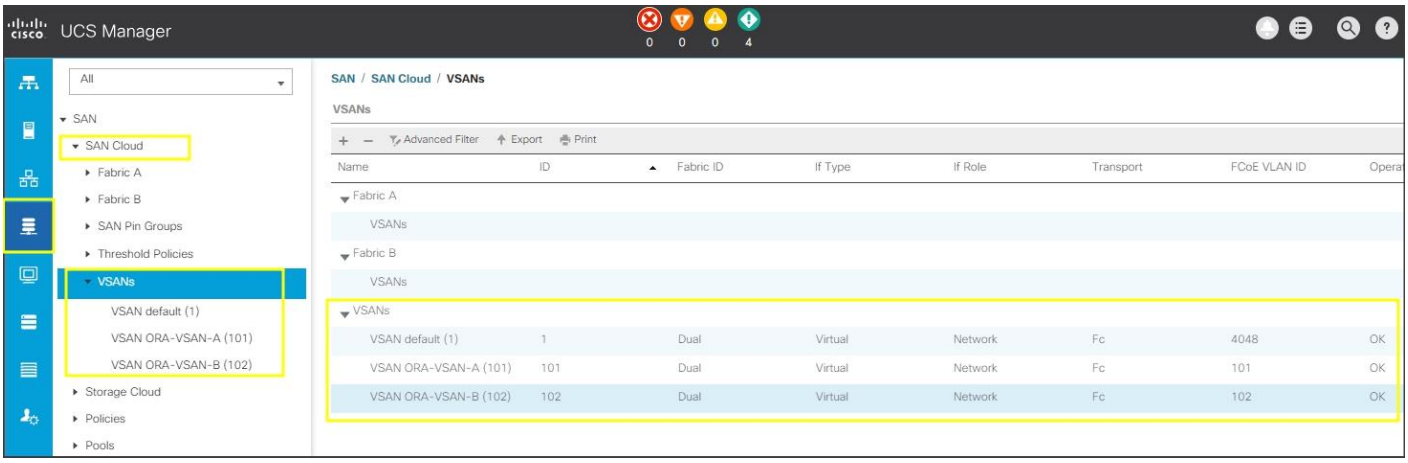

## Configure IP, UUID, Server, MAC, WWNN and WWPN Pools

### 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, follow these steps:

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Pools > root > IP Pools >Click Create IP Pool.
3. We have named IP Pool as ORA-IP-Pool for this solution.
4. Select option Sequential to assign IP in sequential order then click next.
5. Click Add IPv4 Block.
6. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information as shown below.

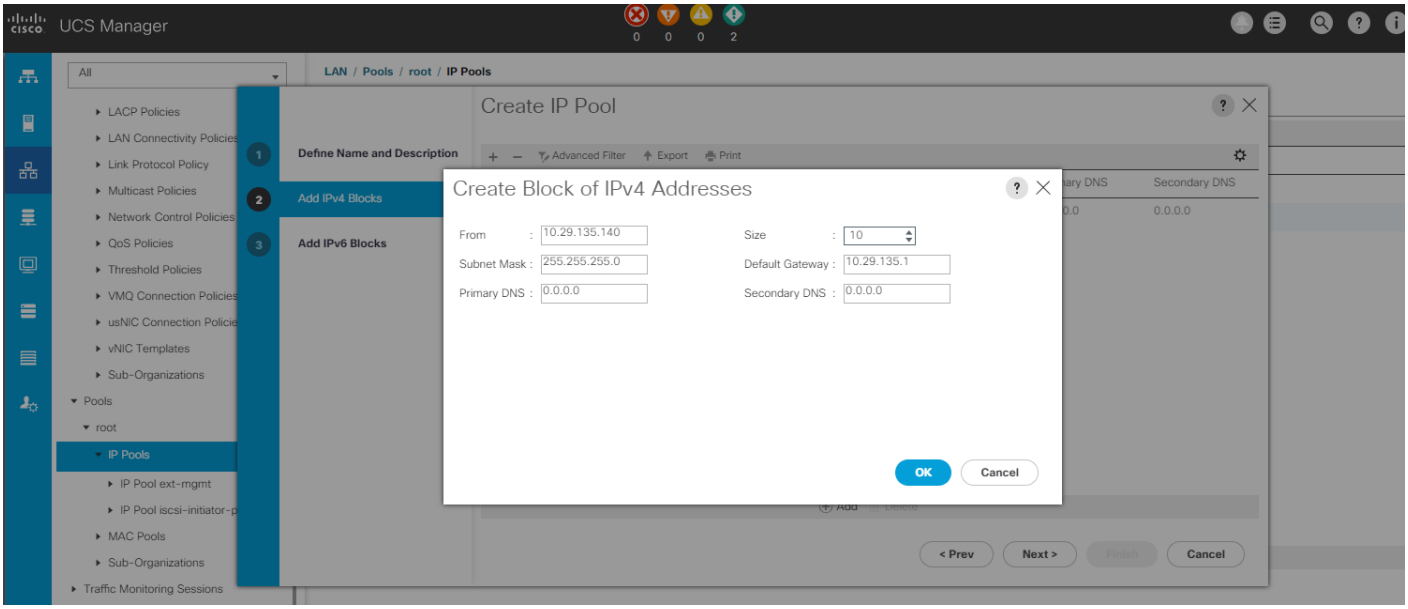


7. Click Next and then click Finish to create the IP block.

### UUID Suffix Pool Creation

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, follow these 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.

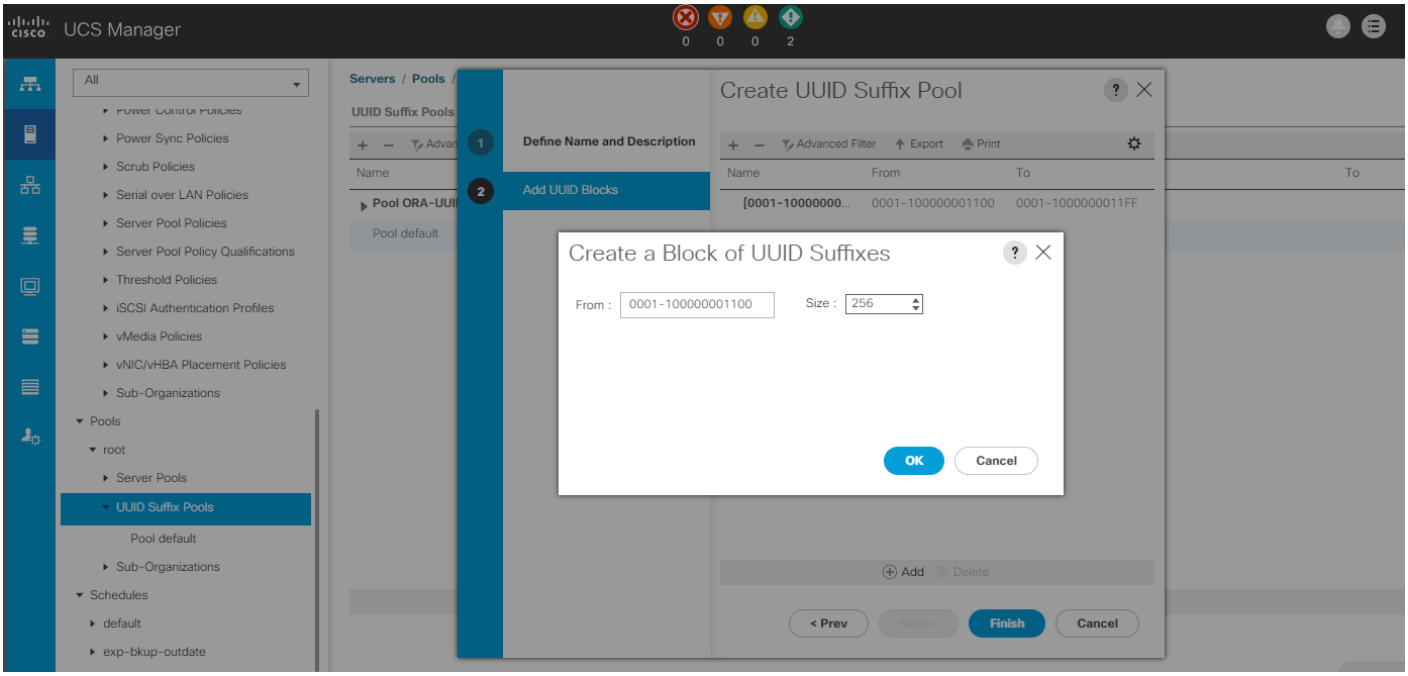


9. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

## Server Pool Creation

To configure the necessary server pool for the Cisco UCS environment, follow these 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, follow these 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.

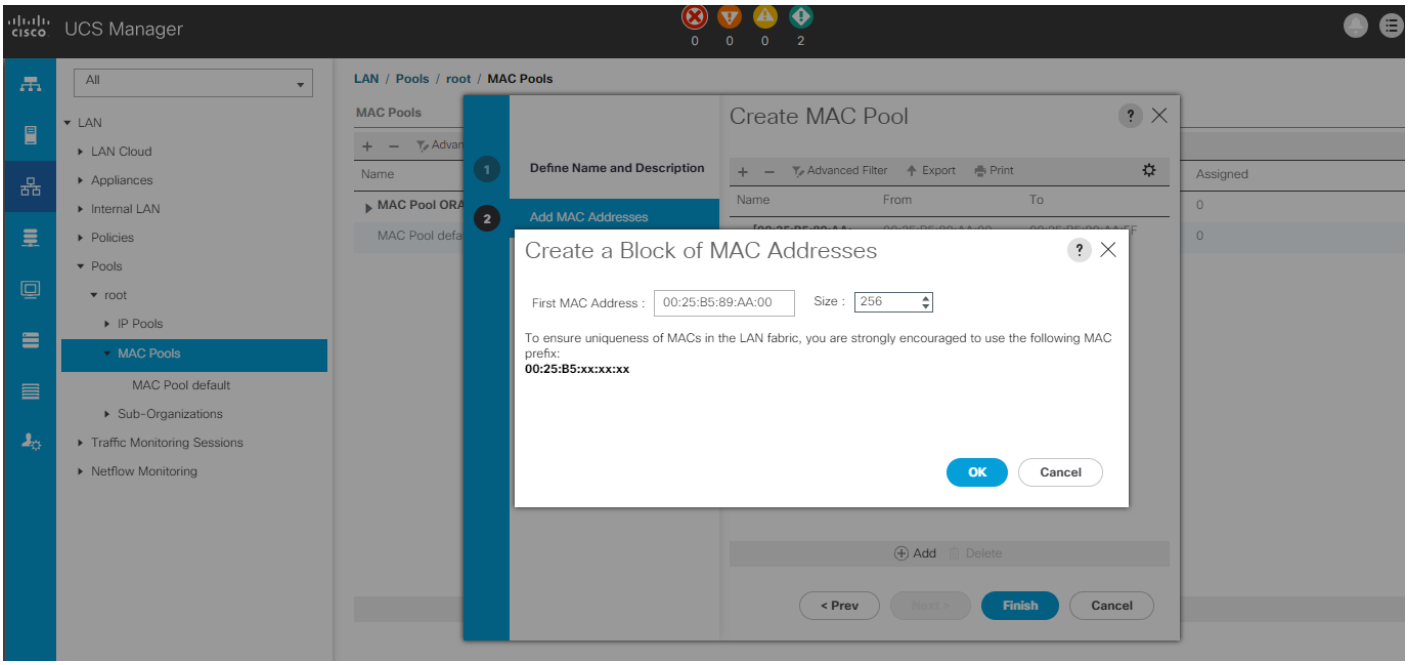


6. Click OK and then click Finish.
7. In the confirmation message, click OK.
8. Create MAC Pool B and assign unique MAC Addresses as shown below.

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

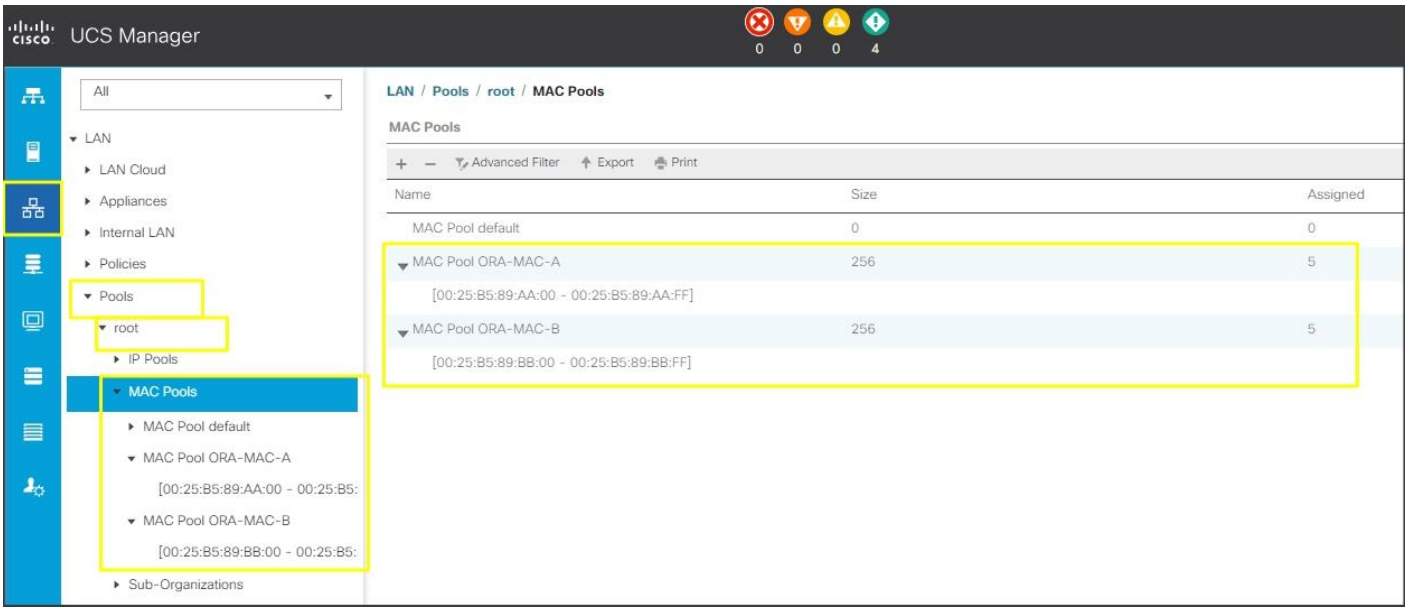


## WWNN and WWPN Pool Creation

To configure the necessary WWNN pools for the Cisco UCS environment, follow these 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.

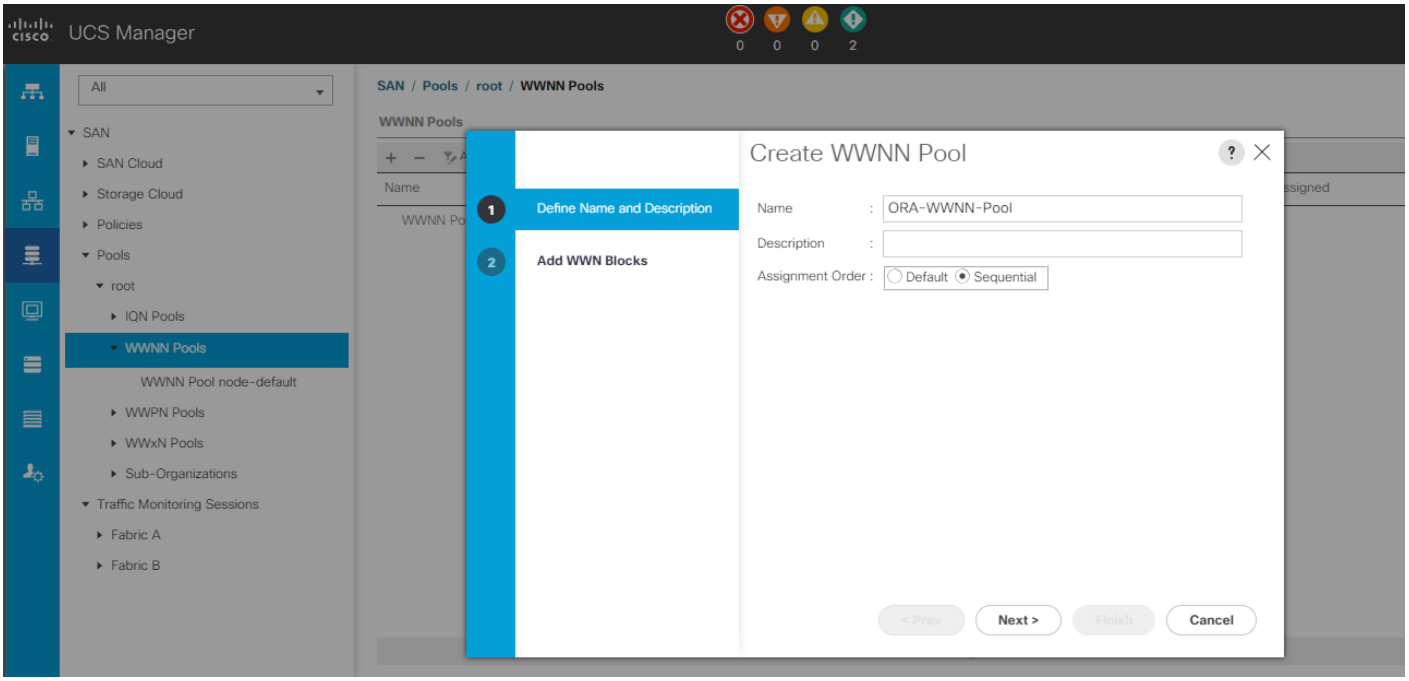


4. Click Next and then click Add to add block of Ports.
5. Enter Block for WWN and size of WWNN Pool as shown below.

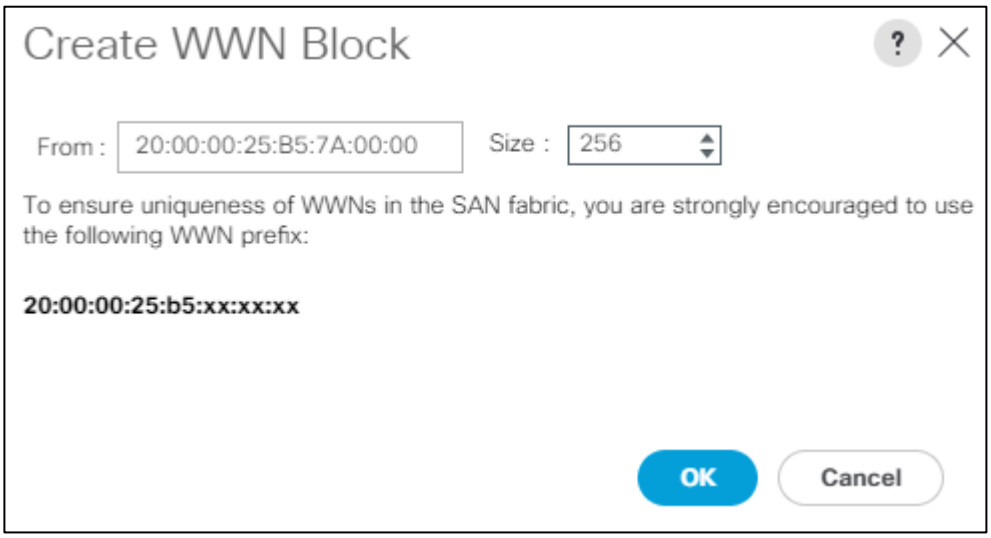


6. Click OK and then click Finish.

To configure the necessary WWPN pools for the Cisco UCS environment, follow these 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.

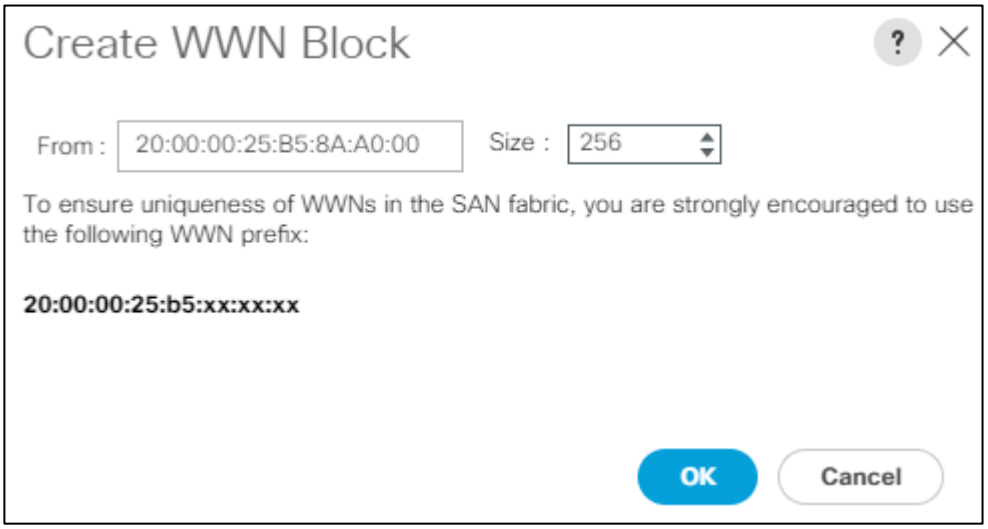


4. Click Next and then click Add to add block of ports.
5. Enter Block for WWN and size.
6. Click OK and then Finish.

7. Configure the ORA-WWPN-B Pool and assign the unique block IDs as shown below.

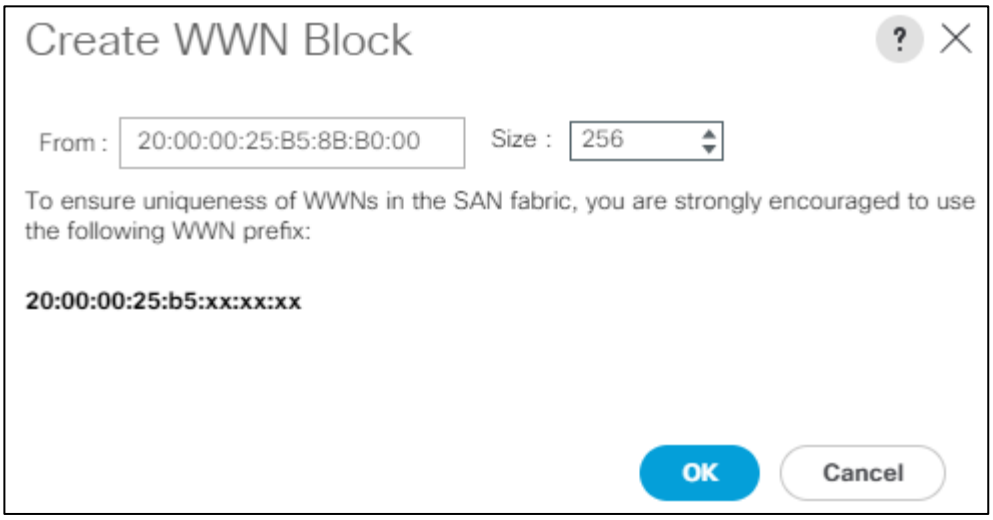


When there are multiple UCS domains sitting in adjacency, it is important that these blocks, the WWNN, WWPN, and MAC, hold differing values between each set.

## Set Jumbo Frames in both the Cisco Fabric Interconnect

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, follow these 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.

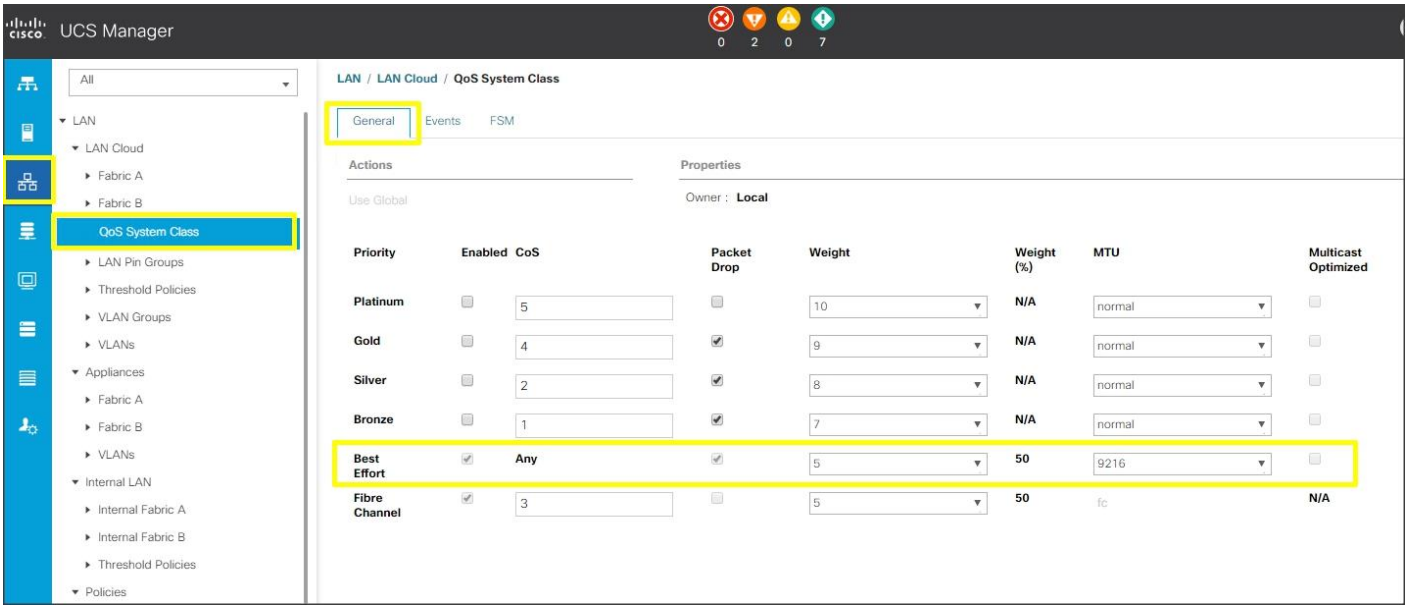

5. Click Save Changes in the bottom of the window.
6. Click OK.

Only the Fibre Channel and Best Effort QoS System Classes are enabled in this FlexPod implementation. The UCS and Nexus switches are intentionally configured this way so that all IP traffic within the FlexPod will be treated as Best Effort. Enabling the other QoS System Classes without having a comprehensive, end-to-end QoS setup in place can cause difficult to troubleshoot issues. For example, NetApp storage controllers by default mark IP-based storage protocol packets with a CoS value of 4. With the default configuration on the Nexus switches in this implementation, storage packets pass through the switches and into the UCS Fabric Interconnects with CoS 4 set in the packet header. If the Gold QoS System Class in Cisco UCS is enabled, these storage packets will be treated according to that class; if Jumbo Frames are being used for the storage protocols but the MTU of the Gold QoS System Class is not set to Jumbo, packet drops will occur.

## Configure Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter OLTP_BIOS as the BIOS policy name
6. Select and click the newly created BIOS Policy.
7. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab.
8. Set the following within the Processor tab:

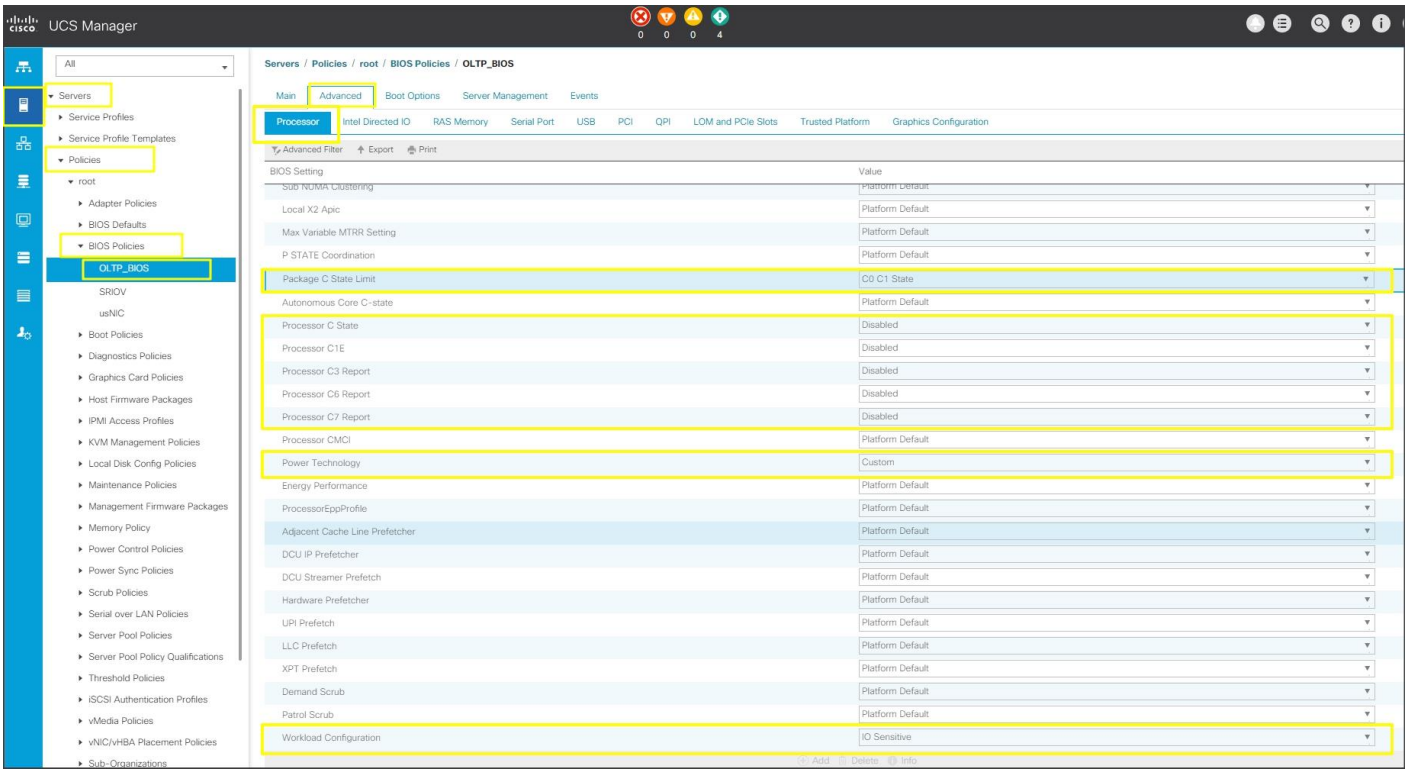


9. Click Save Changes and then click OK

All of them may have to be required on your setup. Please follow the steps appropriately according to your environment and requirement. The following changes were made on the test bed where Oracle RAC installed. Please validate and change as needed.

For further details on BIOS settings, please refer to this document https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/whitepaper_c11-740098.pdf

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.

## Create Adapter Policy

To create an Adapter Policy for the Cisco UCS environment, follow these 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:

   a. Resources

      i. Transmit Queues: 8
      ii. Ring Size: 4096
      iii. Receive Queues: 8
      iv. Ring Size: 4096
      v. Completion Queues: 16
      vi. Interrupts: 32

   b. Options

      i. Receive Side Scaling (RSS): Enabled

5. Configure adapter policy as shown below.

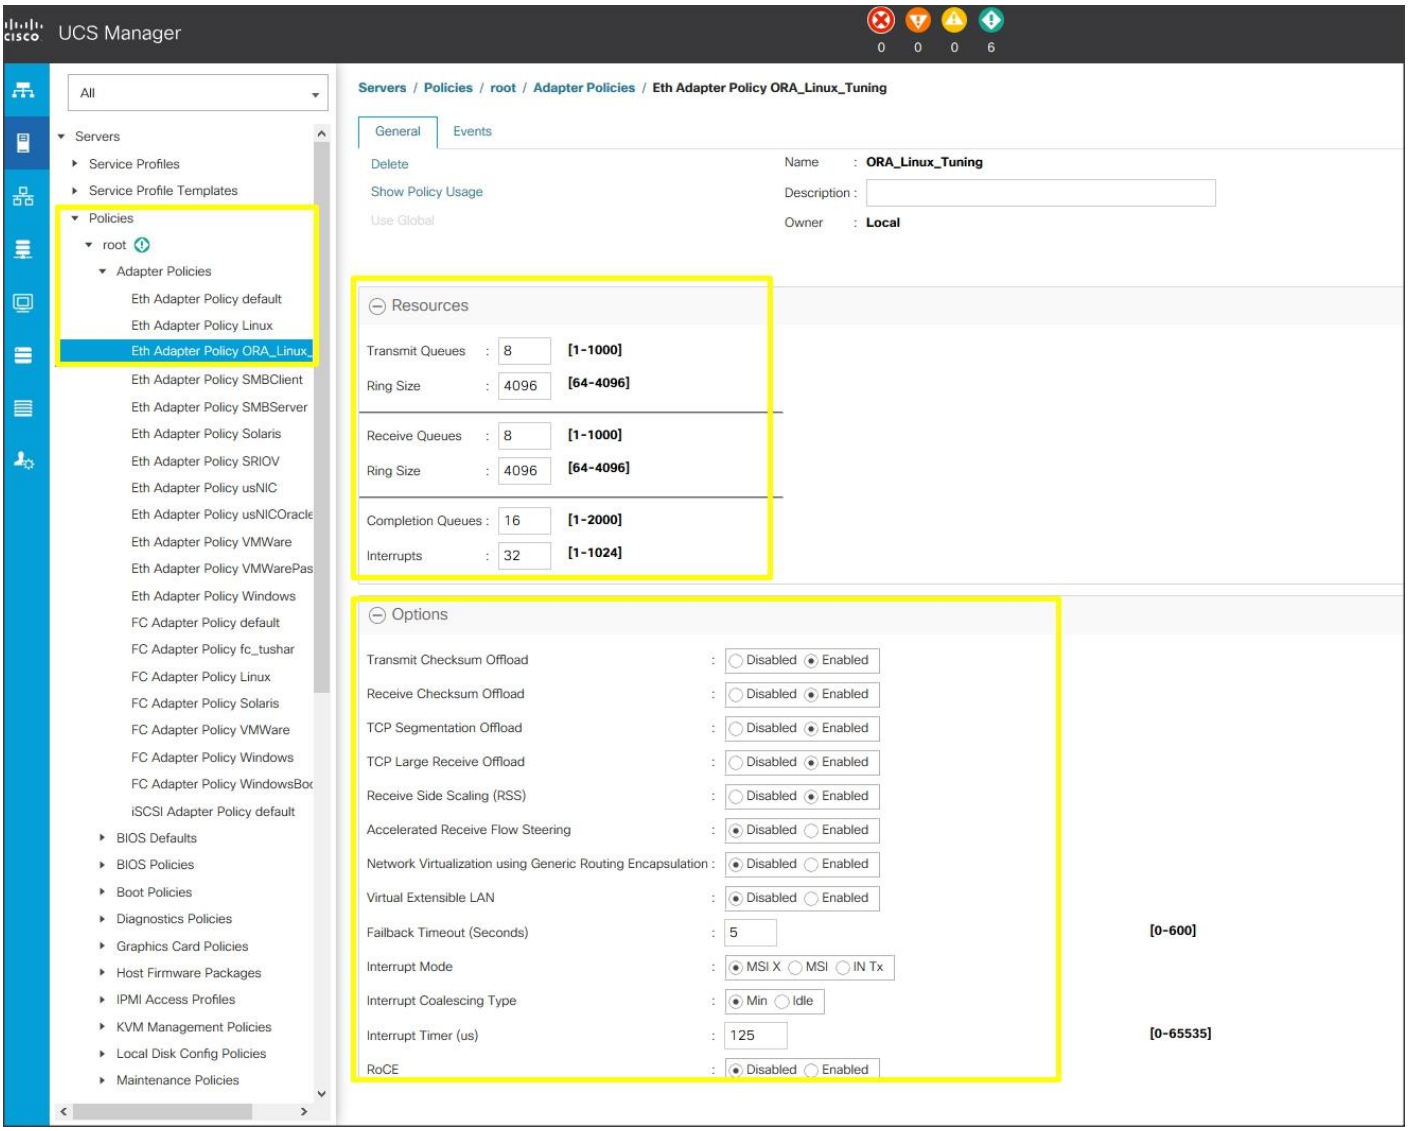

> **RSS distributes network receive processing across multiple CPUs in multiprocessor systems. This can be one of the following:**
> **Disabled—Network receive processing is always handled by a single processor even if additional processors are available.**
>
> **Enabled—Network receive processing is shared across processors whenever possible.**

6. Click OK to finish.

## Configure Update Default Maintenance Policy

To update the default Maintenance Policy, follow these 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

For this solution, we created two vNIC template for Public Network and Private Network Traffic.

### Create Public and Private vNIC Template

To create vNIC (virtual network interface card) template for the Cisco UCS environment, follow these 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.

> **Selecting Failover is a critical step to improve link failover time by handling it at the hardware level, and to guard against NIC any potential for NIC failure not being detected by the virtual switch.**

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.

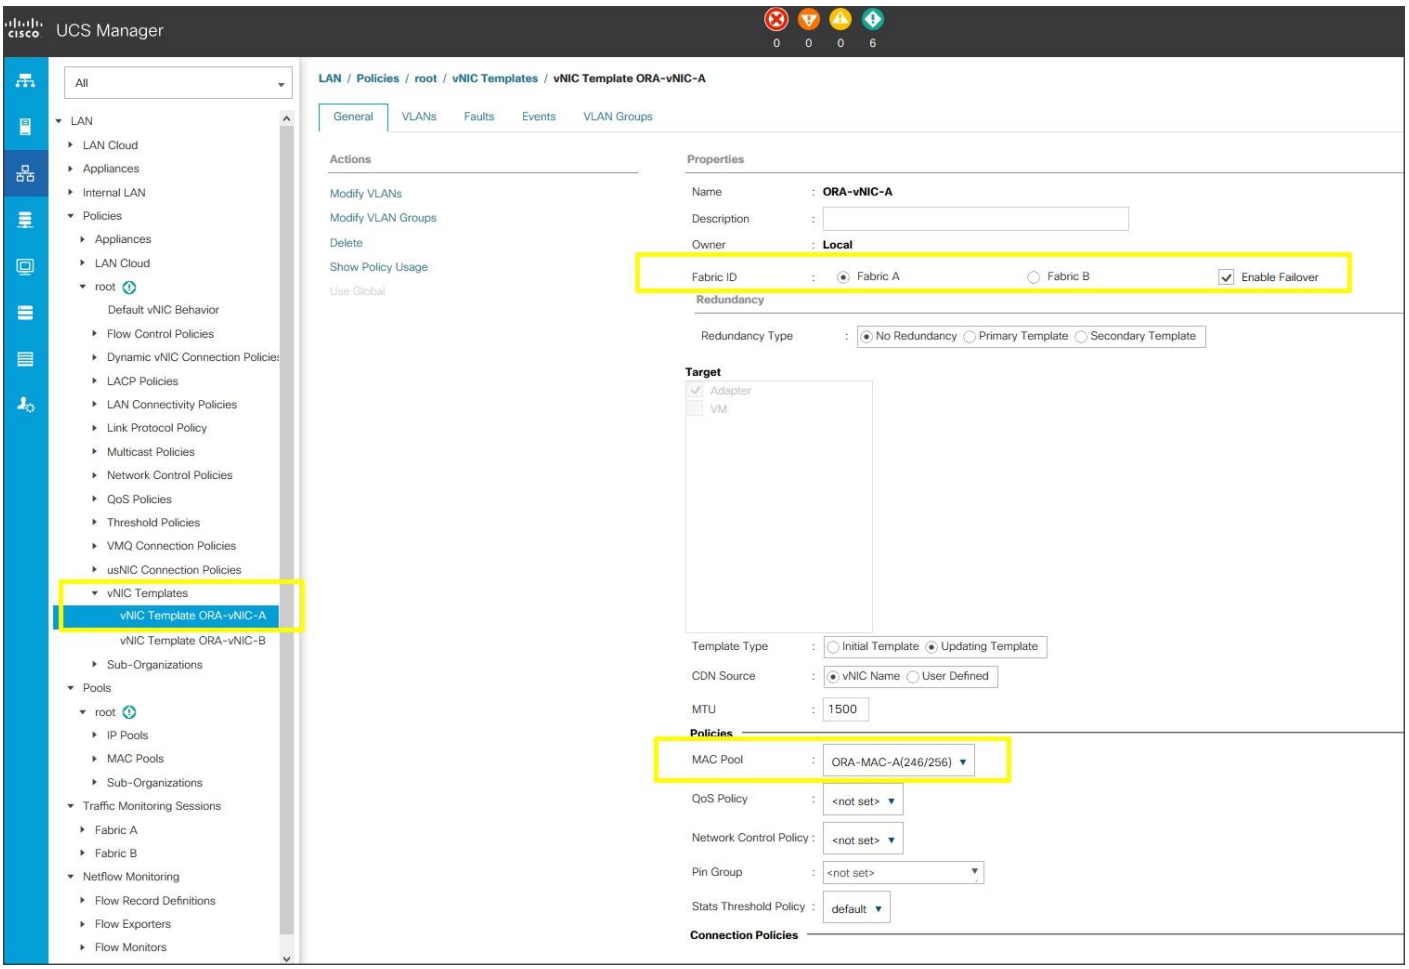


10. Click OK to finish.

11. Create another vNIC template for Private Network Traffic with few changes:

    a. Enter ORA-vNIC-B as the vNIC template name for Private Network Traffic.

    b. Select the Fabric B and Enable Failover for Fabric ID options.

    c. Select Template Type as Updating Template.

    d. Under VLANs, select the checkboxes default and Private_Traffic and set Native-VLAN as the Private_Traffic.

    e. Set MTU value to 9000 and MAC Pool as ORA-MAC-B.

12. Click OK to create the vNIC template as shown below.

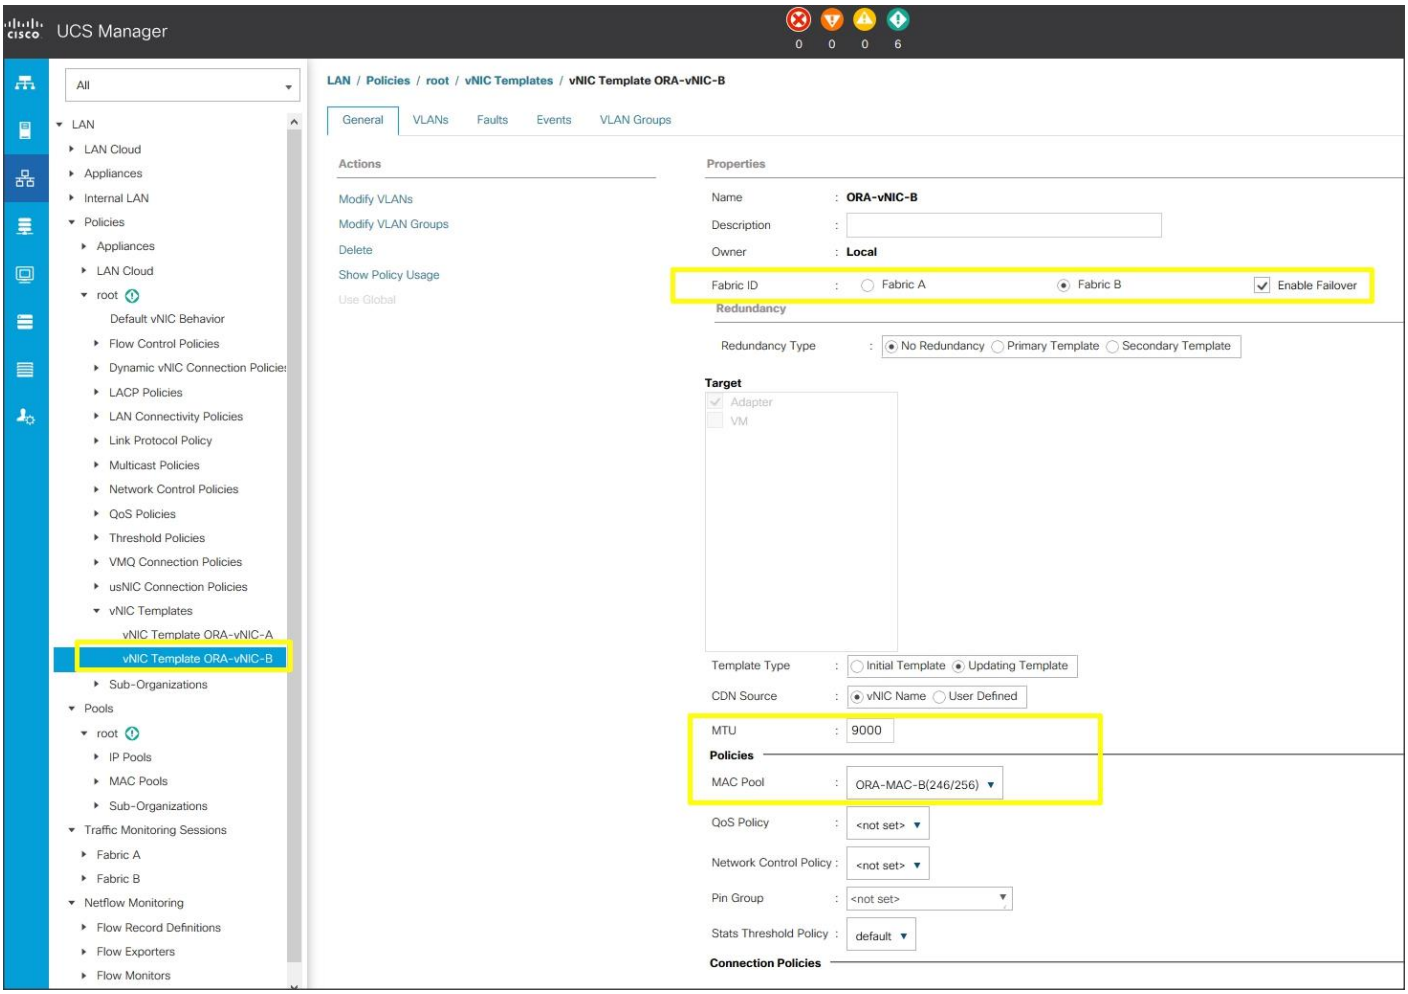


## Create Storage vHBA Template

To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, follow these 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-vHBA-A and keep Fabric A selected.
4. Select VSAN as ORA-VSAN-A and template type to Updating Template.
5. Select WWPN Pool as ORA-WWPN-A from the drop-down list as shown below.

Create vHBA Template

Name : ORA-vHBA-A

Description :

Fabric ID : A B

**Redundancy**

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : ORA-VSAN-A Create VSAN

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : ORA-WWPN-A(256/256)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

OK Cancel

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

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

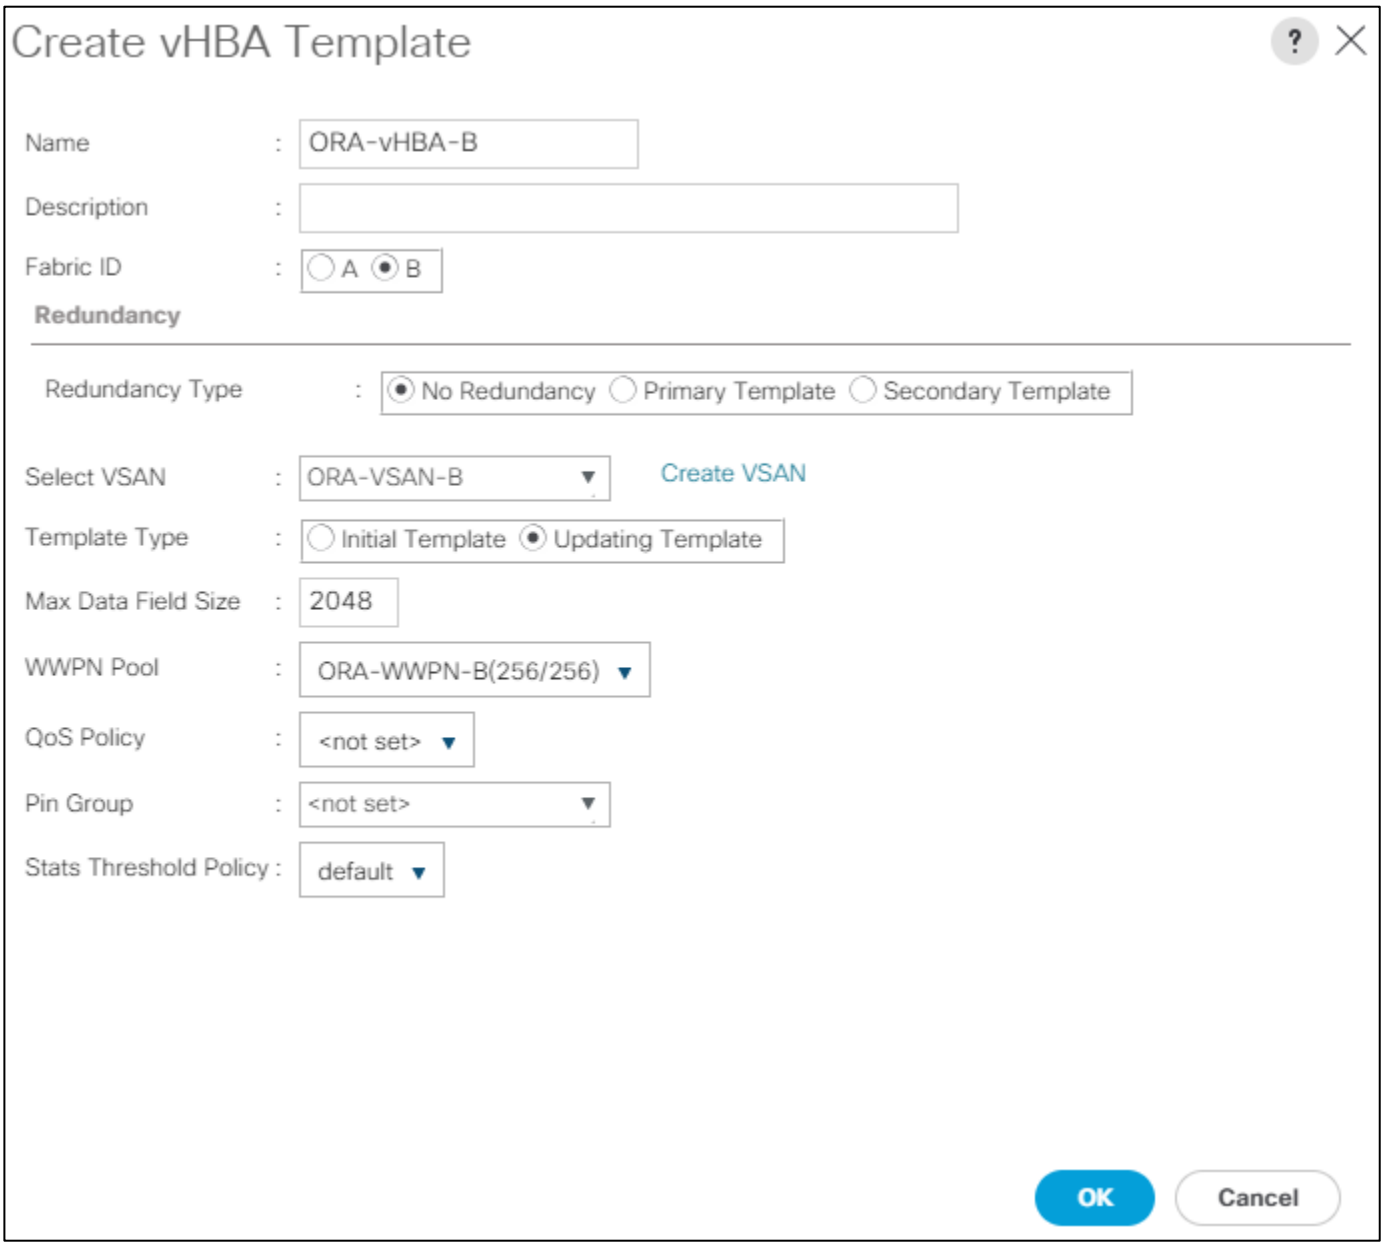


## Create Server Boot Policy for SAN Boot

All Oracle nodes were set to boot from SAN for the 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.

**We strongly recommend to use "Boot from SAN" to realize 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 as detailed in the following sections.

## Create Local Disk Configuration Policy

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, follow these 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 below.

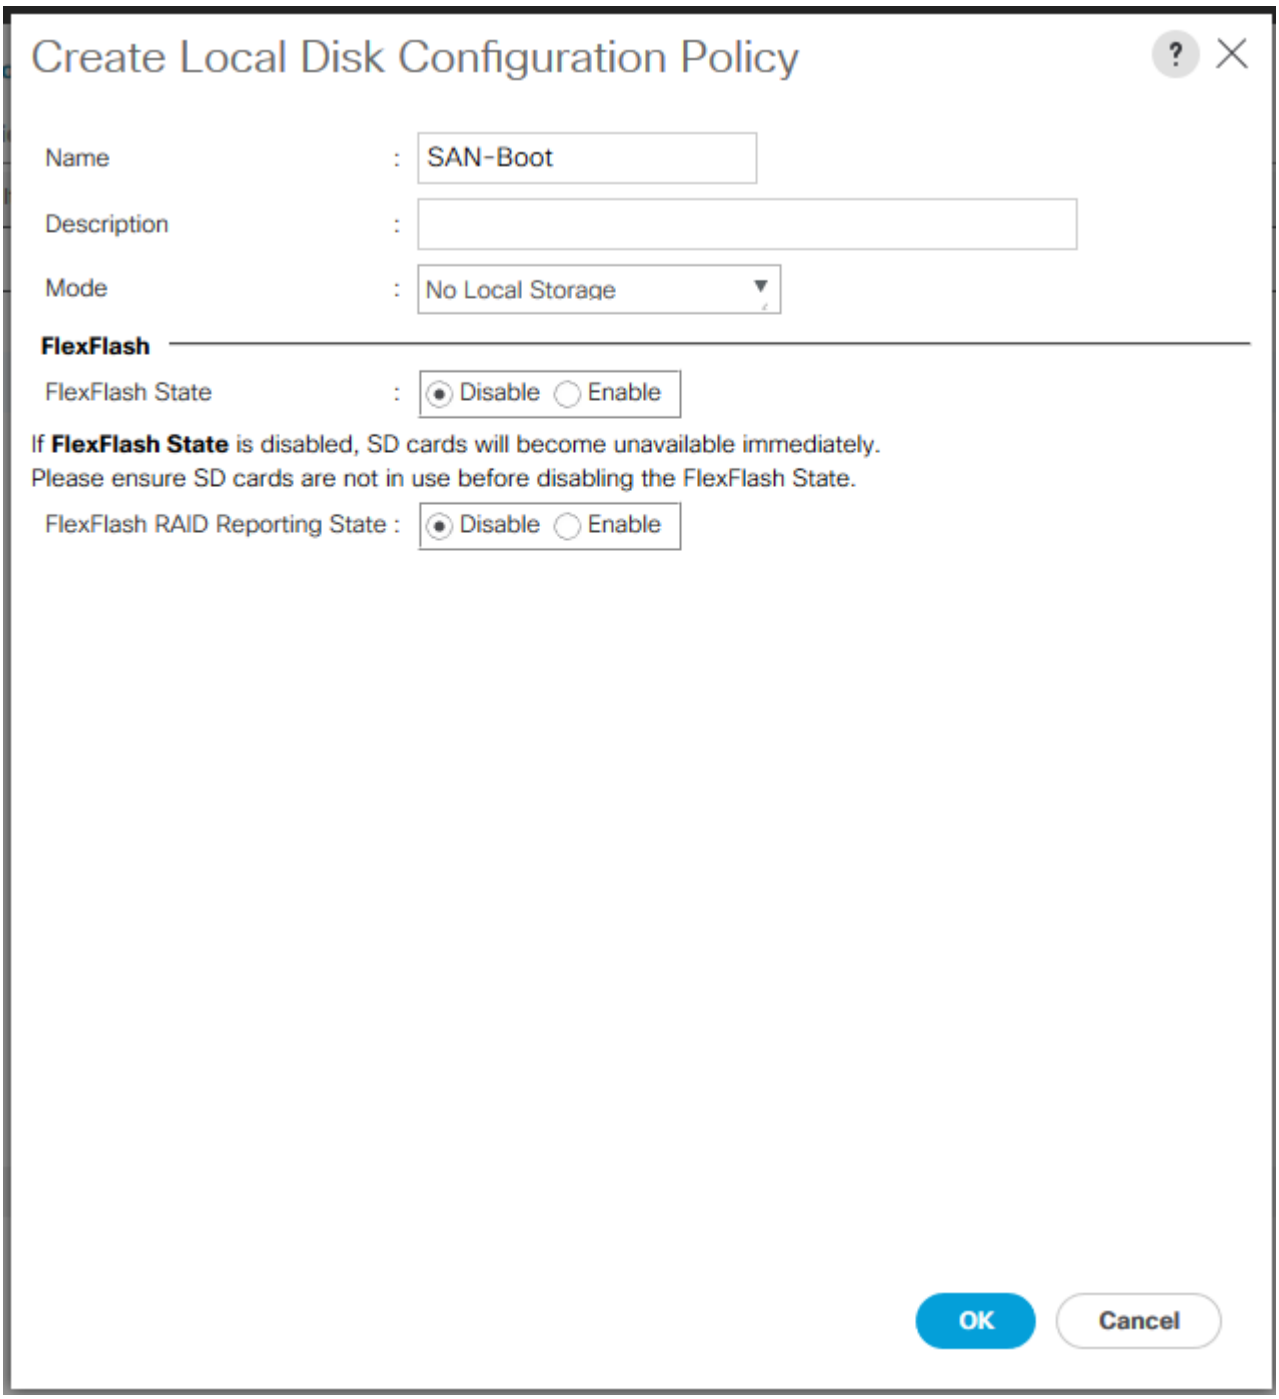

## Create SAN Boot Policy

The screenshot below shows the network interface, WWPN and ports connectivity configured for NetApp AFF A700s controller. Four Fibre Channel logical interfaces (LIFs) are created on storage controller cluster node 1 (node1_lif02a, node1_lif02b, node1_lif03a and node1_lif03b) and four Fibre Channel LIFs are created on storage controller cluster node 2 (node2_lif02a, node2_lif02b, node2_lif03a and node2_lif03b).

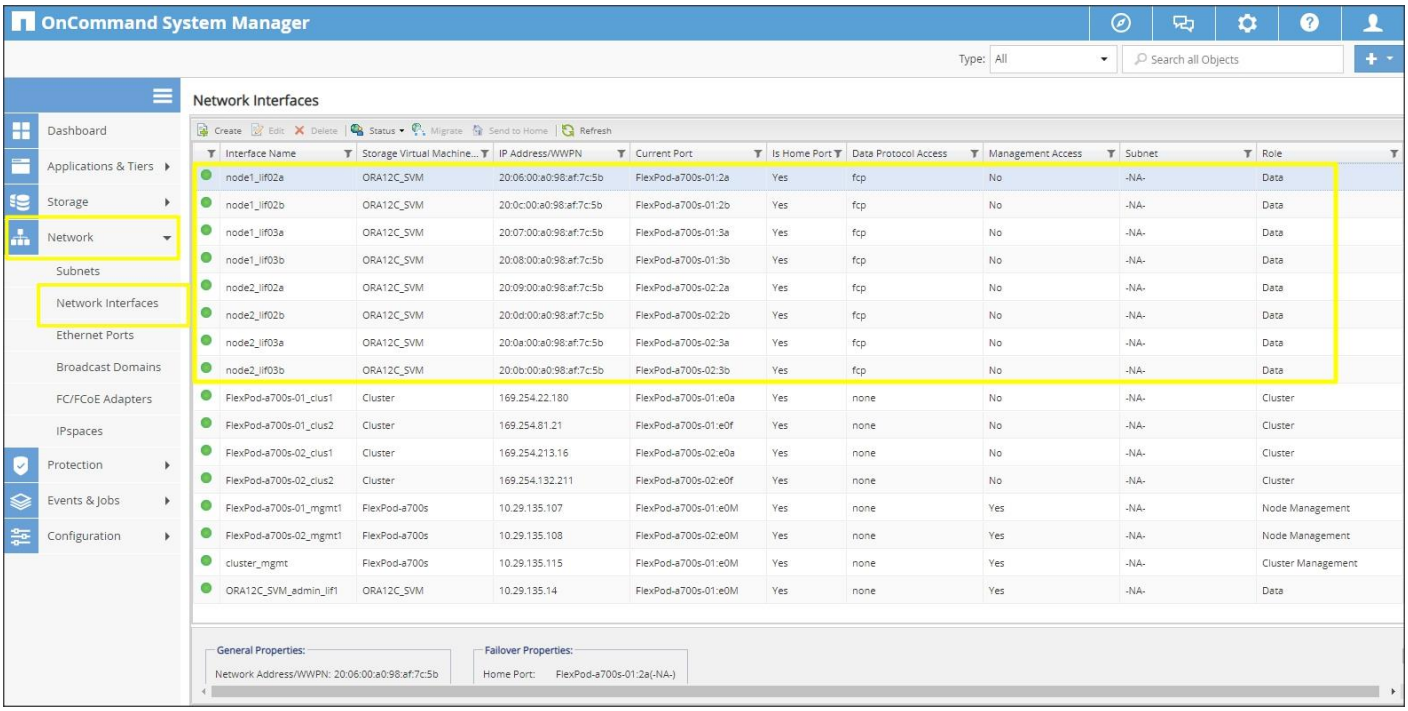


**You can also obtain this information by login to the storage cluster and run the network interface show command**

The SAN boot policy configures the SAN Primary's primary-target to be network interface node1_lif02a on NetApp storage cluster and SAN Primary's secondary-target to be network interface node1_lif03a on NetApp storage cluster. Similarly, the SAN Secondary's primary-target to be network interface node2_lif02b on NetApp storage cluster and SAN Secondary's secondary-target to be network interface node2_lif03b on NetApp storage cluster. Login into NetApp storage controller and verify all the port information is correct. This information can be found in the NetApp Storage GUI under Network > Network Interfaces.

**You have to create SAN Primary (hba0) and SAN Secondary (hba1) in SAN Boot Policy by entering WWPN of NetApp Storage LIFs as explained below.**

To create Boot Policies for the Cisco UCS environments follow these 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_Boot as the name of the boot policy as shown below.

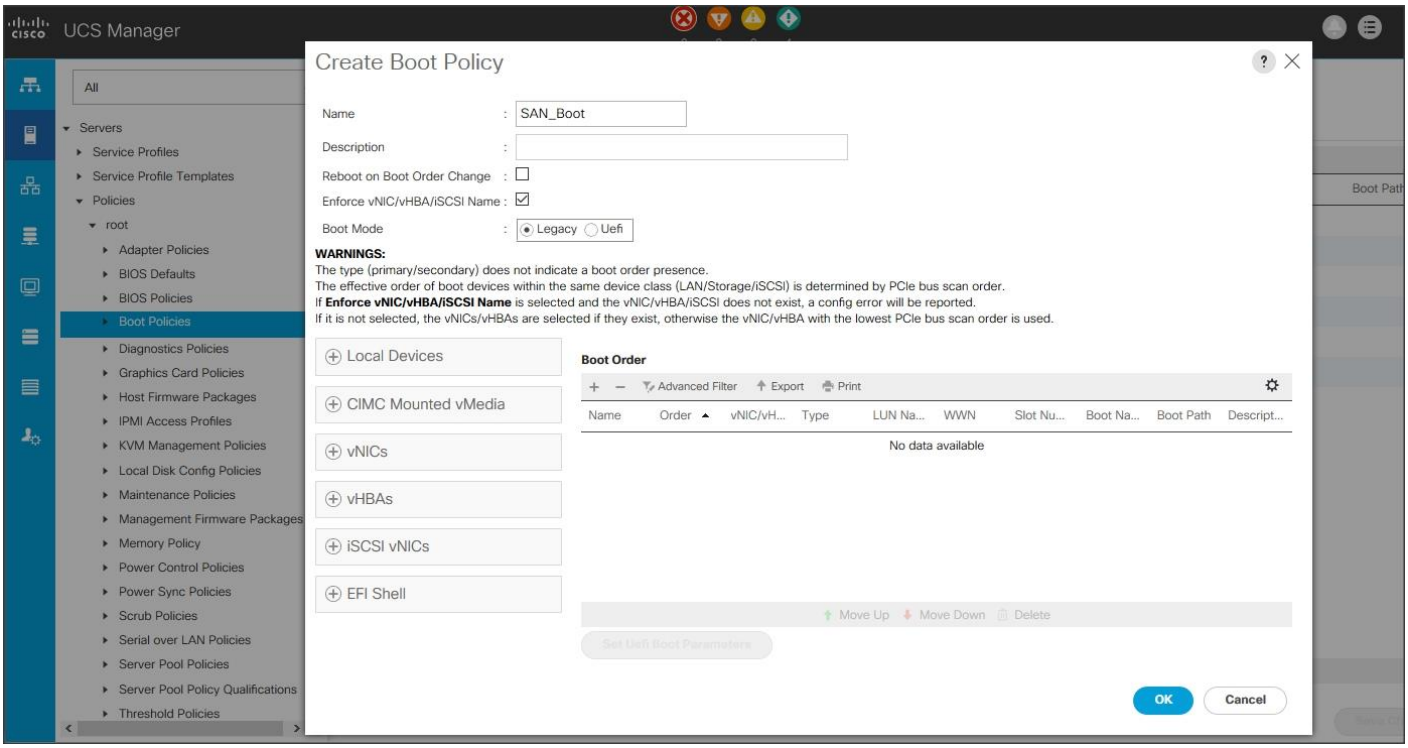


7. Expand the Local Devices drop-down menu and Choose Add CD/DVD. Expand the vHBAs drop-down menu and select Add SAN Boot.

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

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

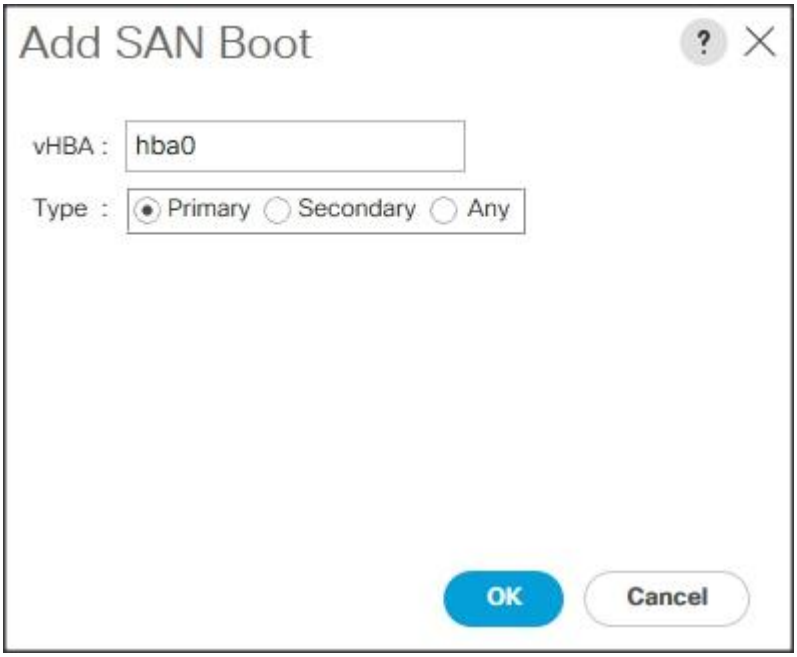

9. Select add SAN Boot Target to enter WWPN address of storage LIF. Keep 0 as the value for Boot Target LUN. Enter the WWPN of NetApp Storage cluster interface node1_lif02a and add SAN Boot Primary Target.

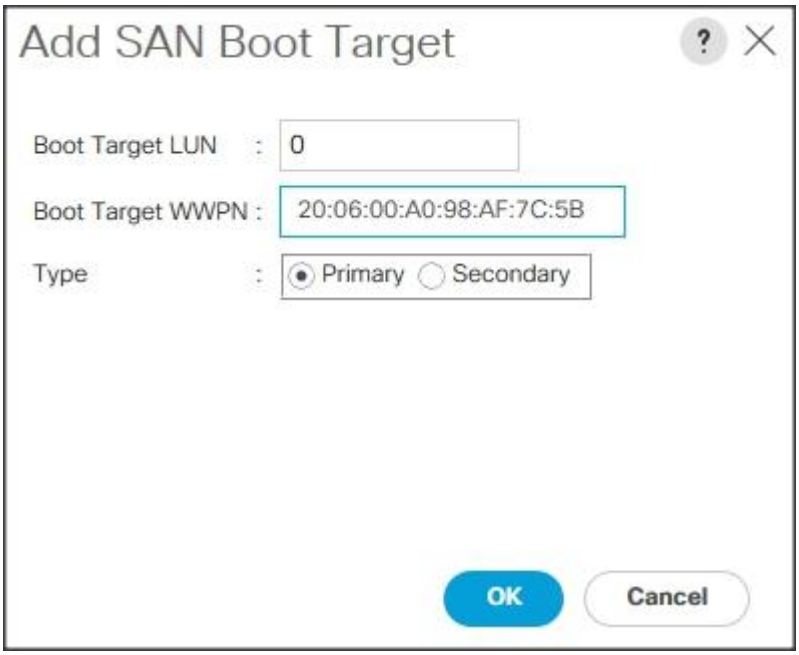


10. Add secondary SAN Boot target into same hba0, enter boot target LUN as 0 and WWPN of NetApp Storage cluster interface node1_lif03a and add SAN Boot Secondary Target.

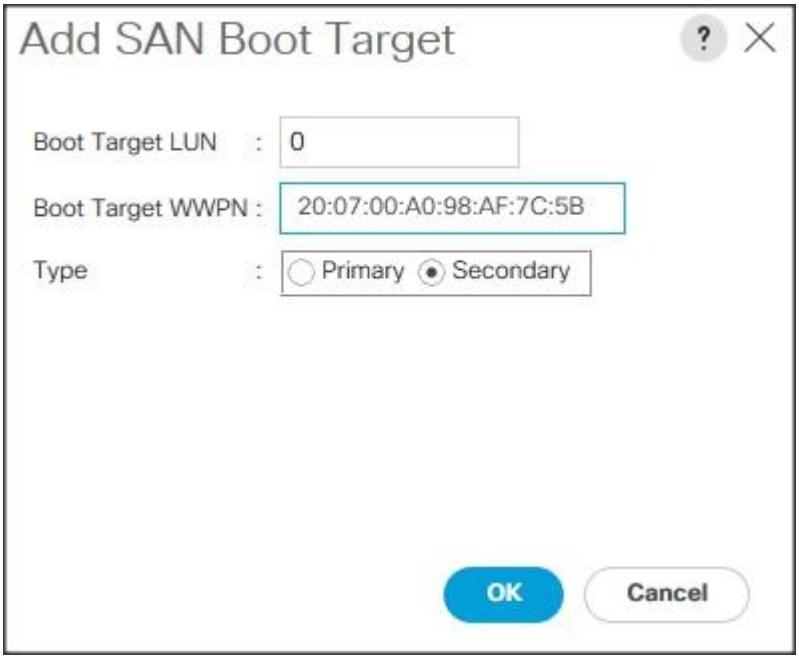


11. From the vHBA drop-down list and Choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba1" in the vHBA field.

12. Click OK to SAN Boot, then choose add SAN Boot Target.

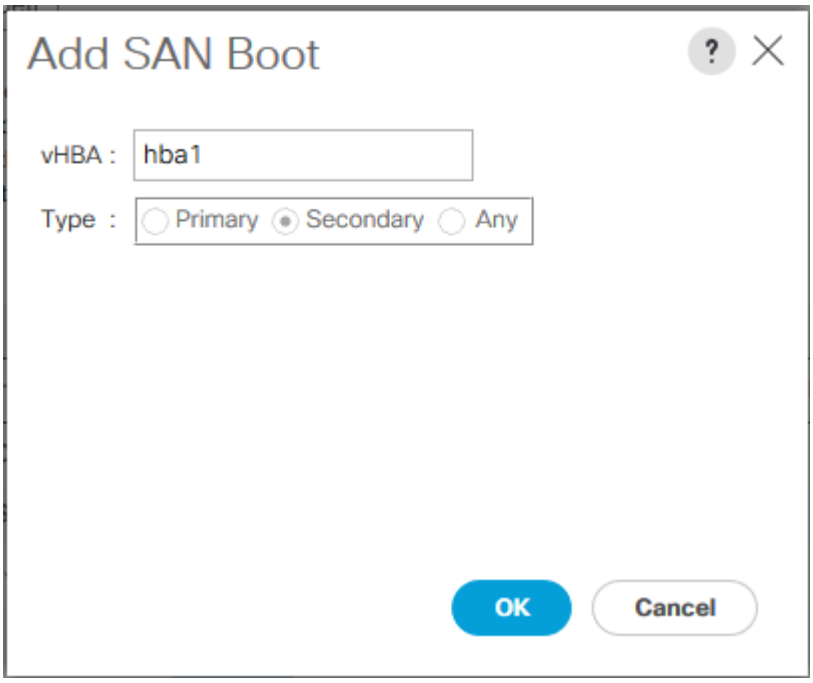


13. Enter 0 as the value for Boot Target LUN. Enter the WWPN of NetApp Storage cluster interface node2_lif02b and add SAN Boot Primary Target.

14. Add secondary SAN Boot target into same hba1, and enter boot target LUN as 0 and WWPN of NetApp Storage cluster interface node2_lif03b and add SAN Boot Secondary Target.

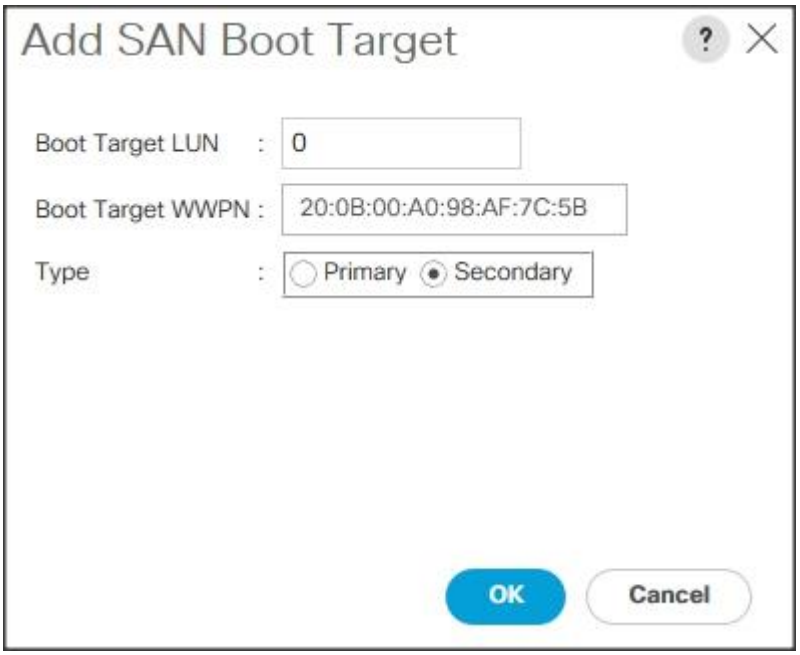


15. After creating the FC boot policies, you can view the boot order in the UCS Manager GUI. To view the boot order, navigate to Servers > Policies > Boot Policies. Click Boot Policy SAN-Boot-A to view the boot order in the right pane of the Cisco UCS Manager as shown below.

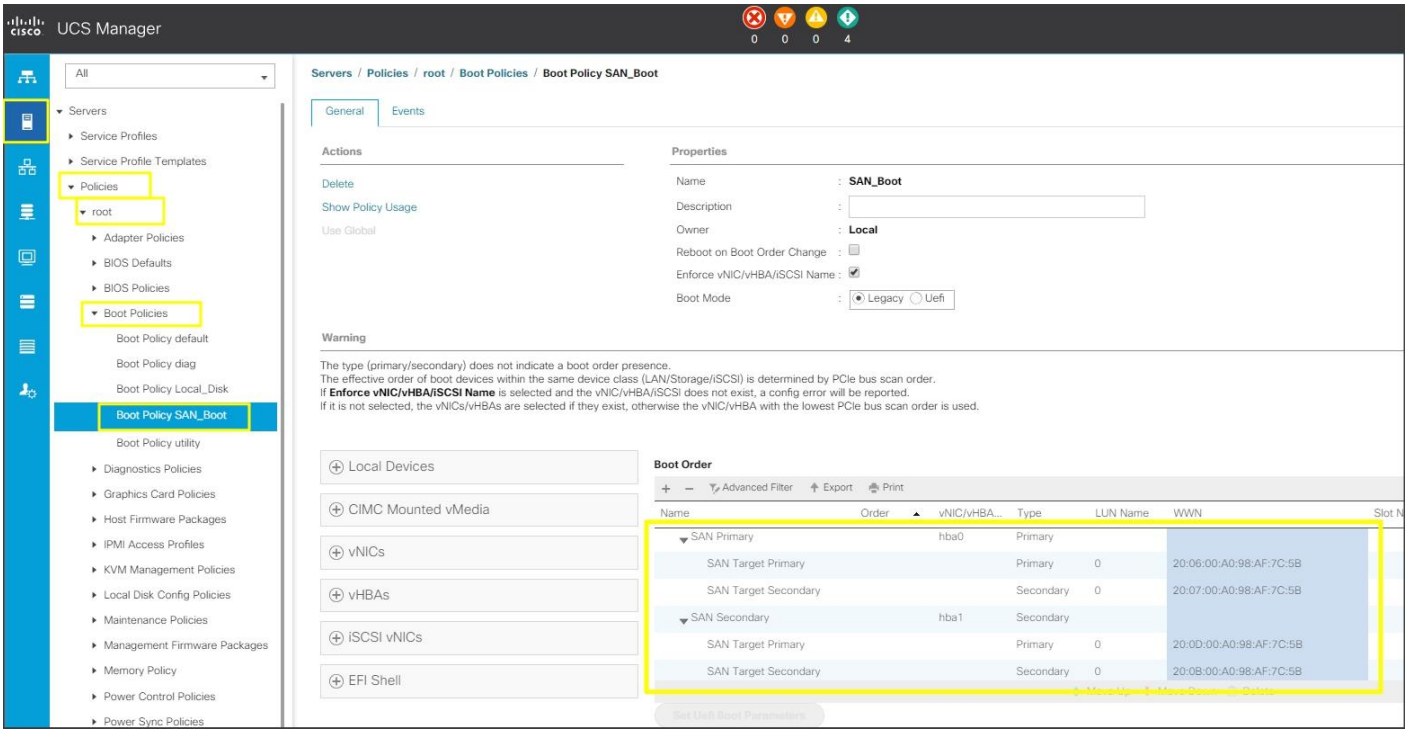


For this solution, we created one Boot Policy as “SAN_Boot”. For all 4 Oracle Database RAC Nodes(flex1, flex2, flex3 and flex4), we will assign this boot policy to the Service Profiles as explained in the following section.

# 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.

The Cisco UCS service profiles 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.
- 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.

You will create one Service Profile Template "ORA_FLEXPOD" using boot policy created earlier to utilize four LIF ports from NetApp Storage for high-availability in case of any FC links go down.

The following sections detail how to create ORA_FLEXPOD.

## Create Service Profile Template

To create a service profile template, follow these 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.

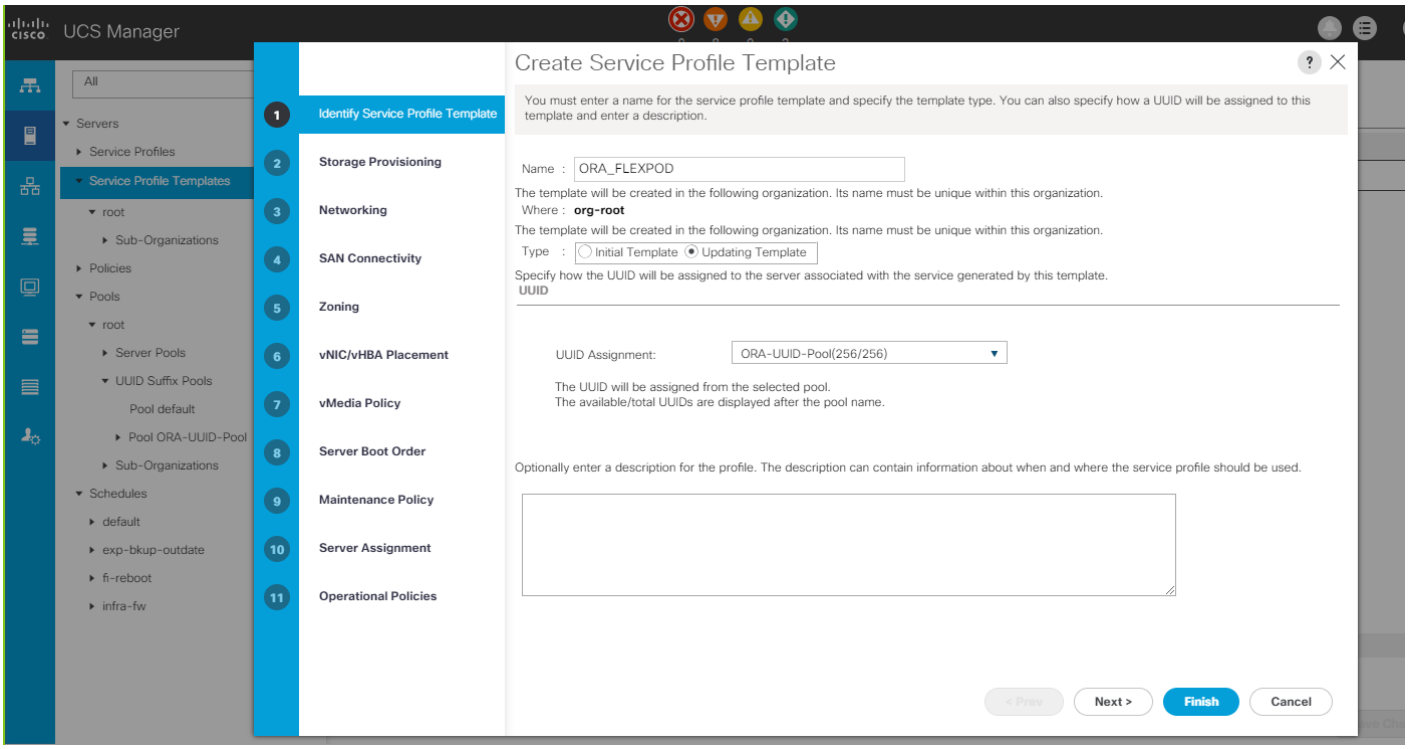


3. Select Local Disk Configuration Policy to SAN-Boot as no Local Storage.

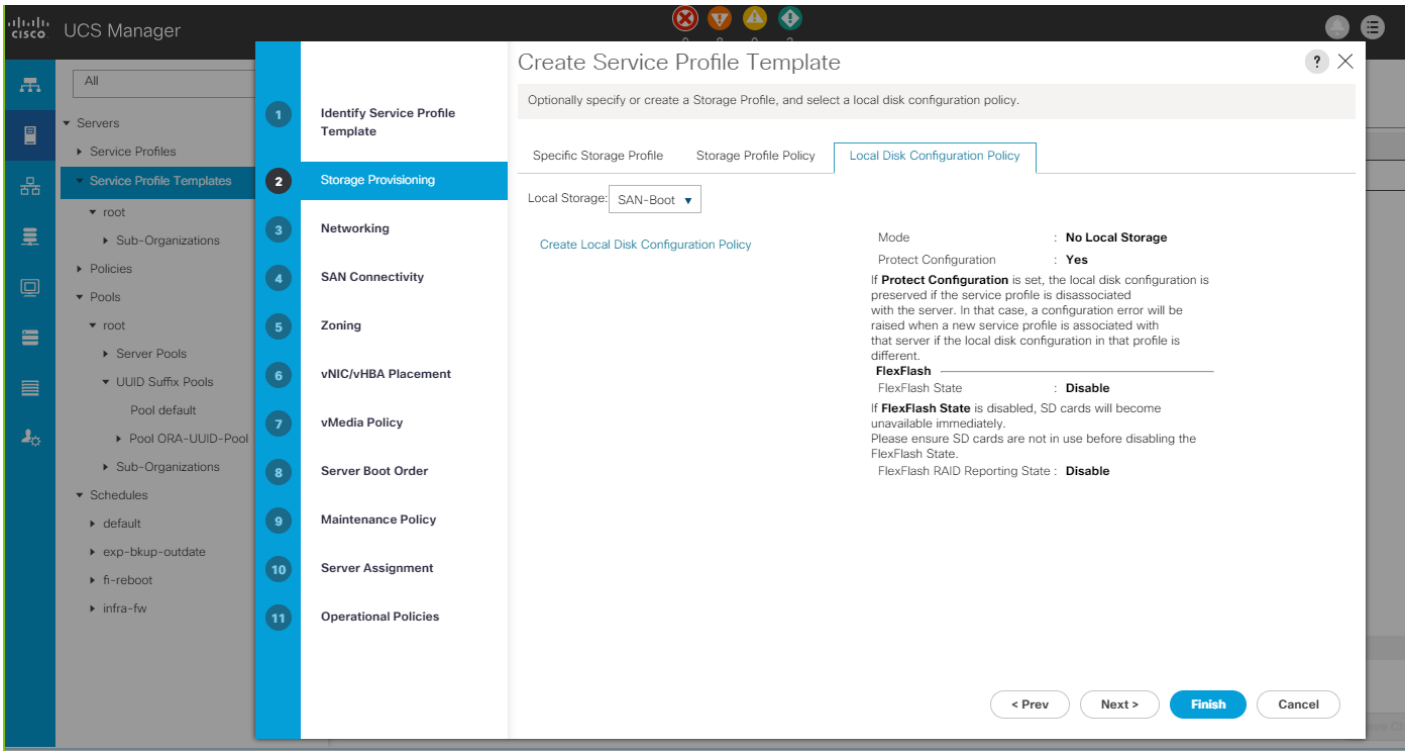


4. In the networking window, select “Expert” and click “Add” to create vNICs. Add one or more vNICs that the server should use to connect to the LAN.

> We created two vNIC in the create vNIC menu. We have given name to first vNIC as “eth0” and second vNIC as “eth1”.

5. Select vNIC Template as ORA-vNIC-A and Adapter Policy as ORA_Linux_Tuning which was created earlier for vNIC “eth0”.

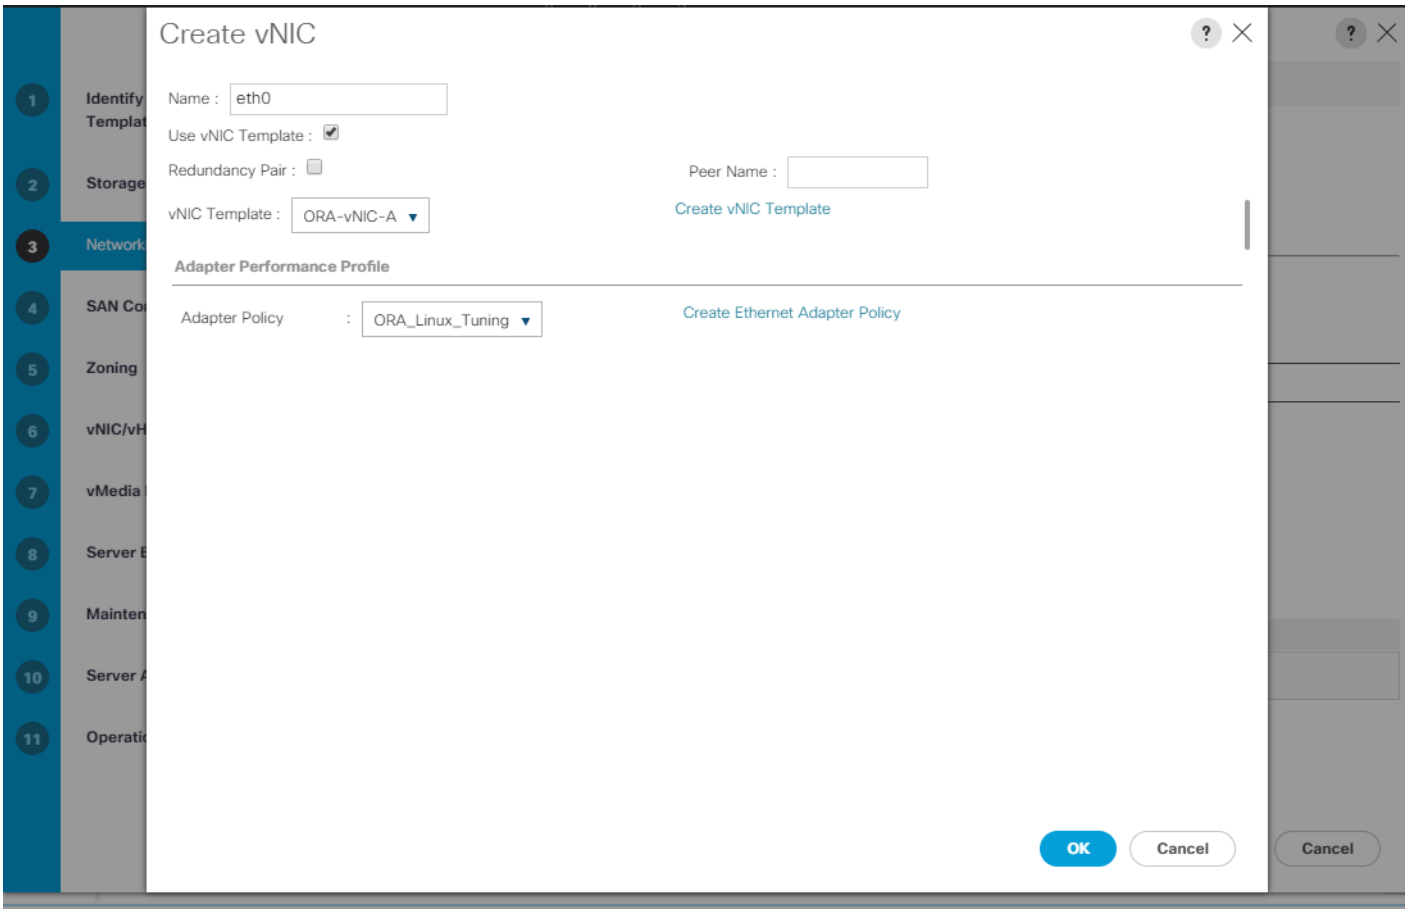


6. Select vNIC Template as ORA-vNIC-B and Adapter Policy as ORA_Linux_Tuning for vNIC “eth1”. eth0 and eth1 vNICs are created so that Servers should use to connect to the LAN.

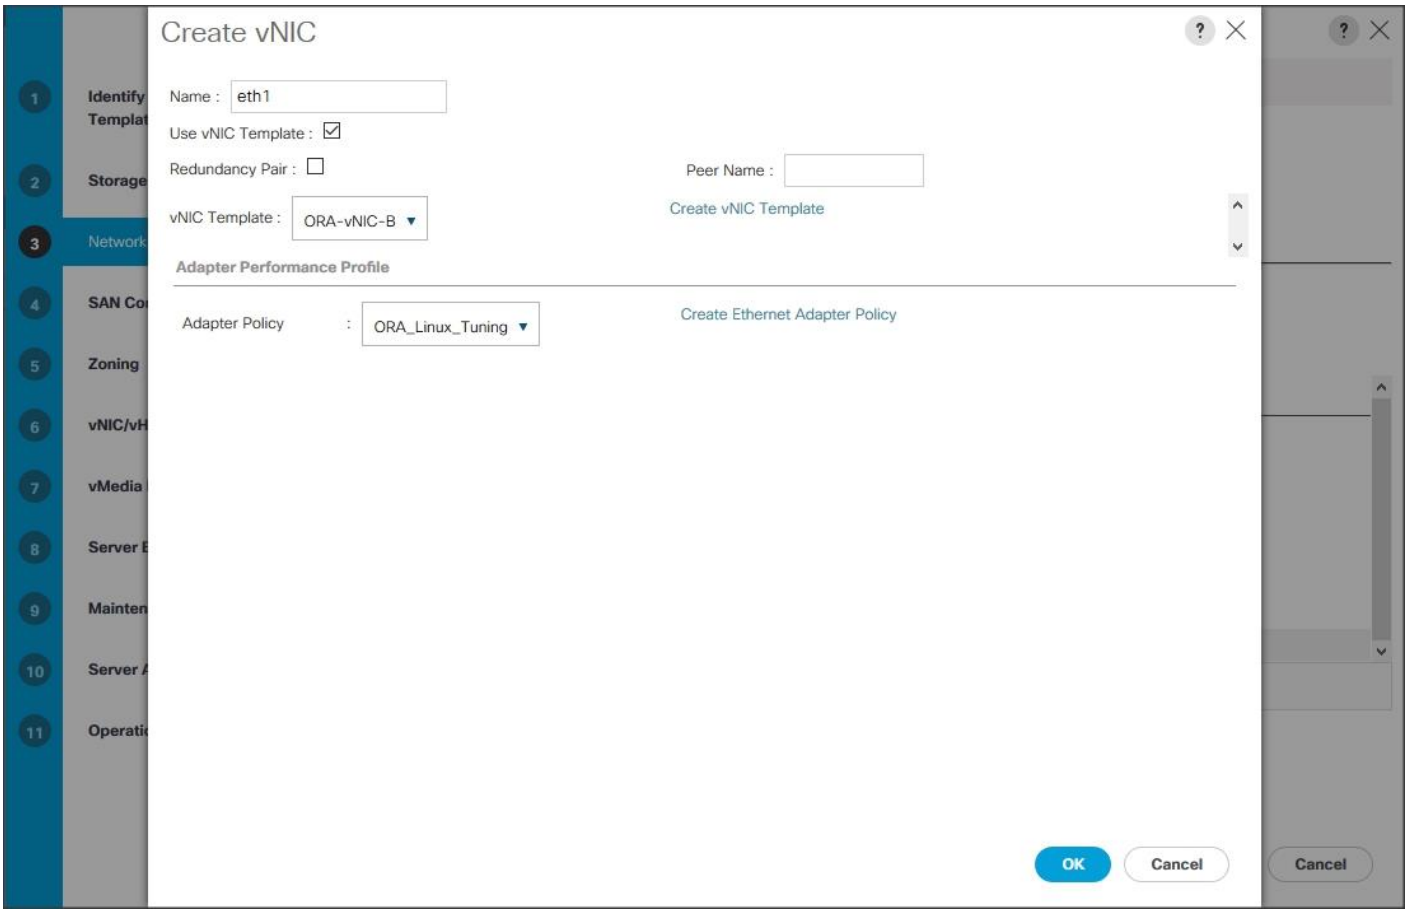
Create vNIC
Name : eth1
Use vNIC Template :
Redundancy Pair :
Peer Name :
vNIC Template : ORA-vNIC-B
Create vNIC Template
Adapter Performance Profile
Adapter Policy : ORA_Linux_Tuning
Create Ethernet Adapter Policy
OK
Cancel
Cancel

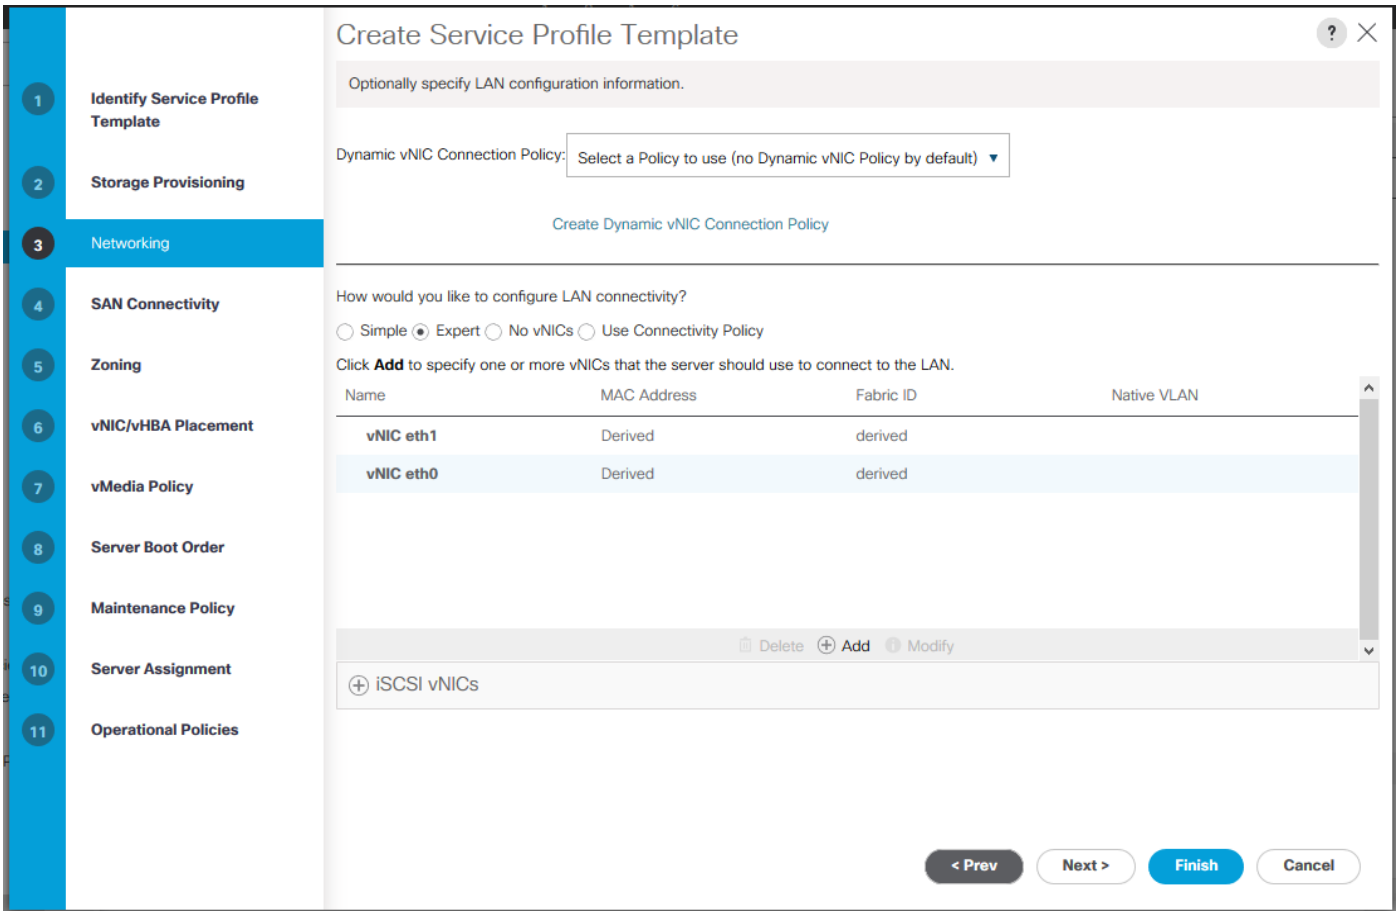


7. When the vNICs are created, you need to create vHBAs. Click Next.

8. In the SAN Connectivity menu, select “Expert” to configure as SAN connectivity. Select WWNN (World Wide Node Name) pool, created previously. Click “Add” to add vHBAs as shown below. The following four HBA have been created:

   a. hba0 using vHBA Template Oracle-HBA-A
   b. hba1 using vHBA Template Oracle-HBA-B
   c. hba2 using vHBA Template Oracle-HBA-A
   d. hba3 using vHBA Template Oracle-HBA-B

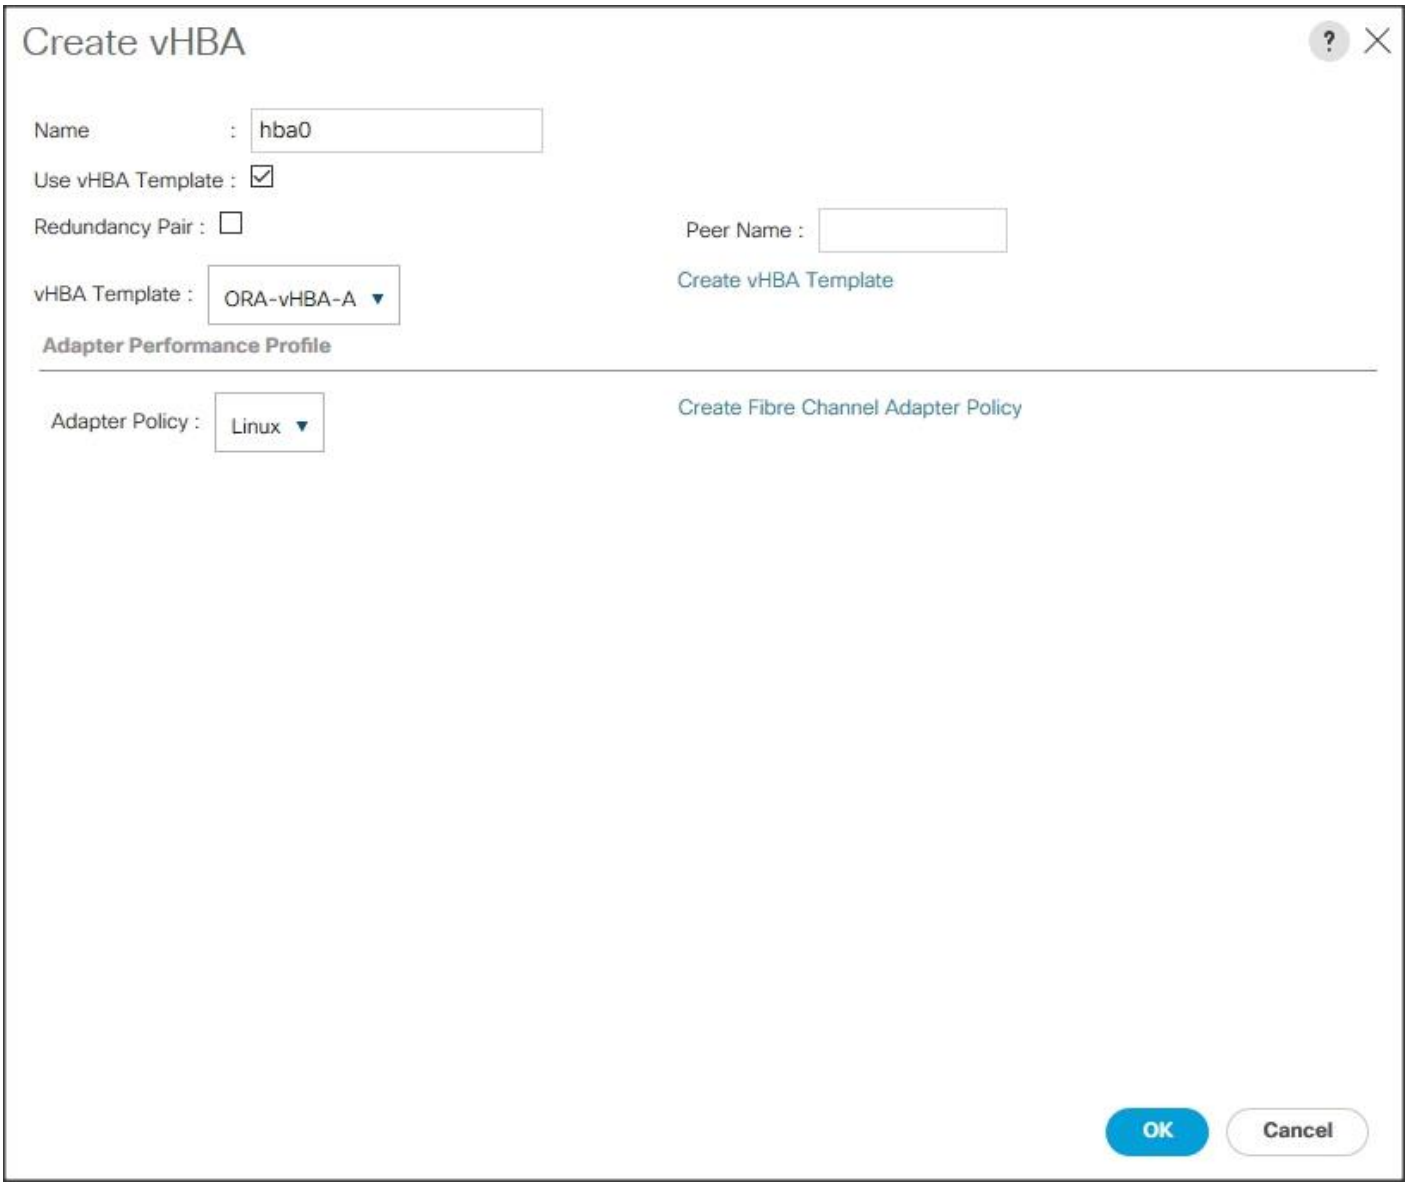
Create vHBA
Name : hba0
Use vHBA Template :
Redundancy Pair :
Peer Name :
vHBA Template : ORA-vHBA-A
Create vHBA Template
Adapter Performance Profile
Adapter Policy : Linux
Create Fibre Channel Adapter Policy
OK
Cancel

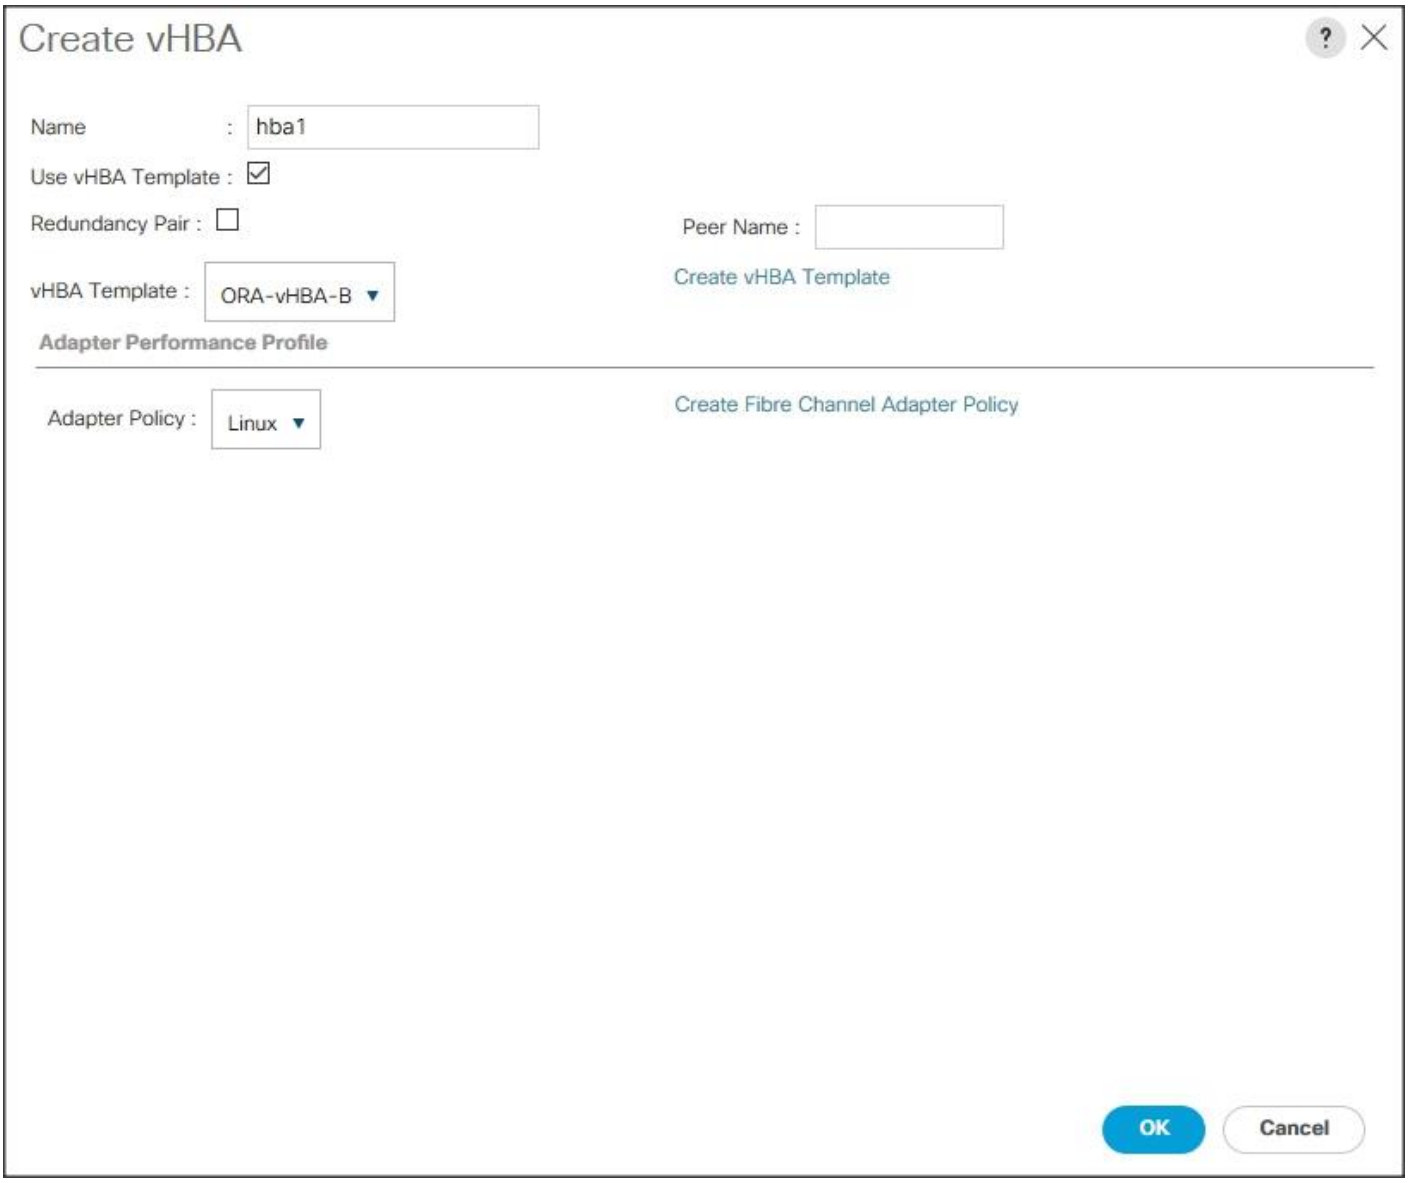
Create vHBA
Name : hba1
Use vHBA Template :
Redundancy Pair :
Peer Name :
vHBA Template : ORA-vHBA-B
Create vHBA Template
Adapter Performance Profile
Adapter Policy : Linux
Create Fibre Channel Adapter Policy
OK
Cancel

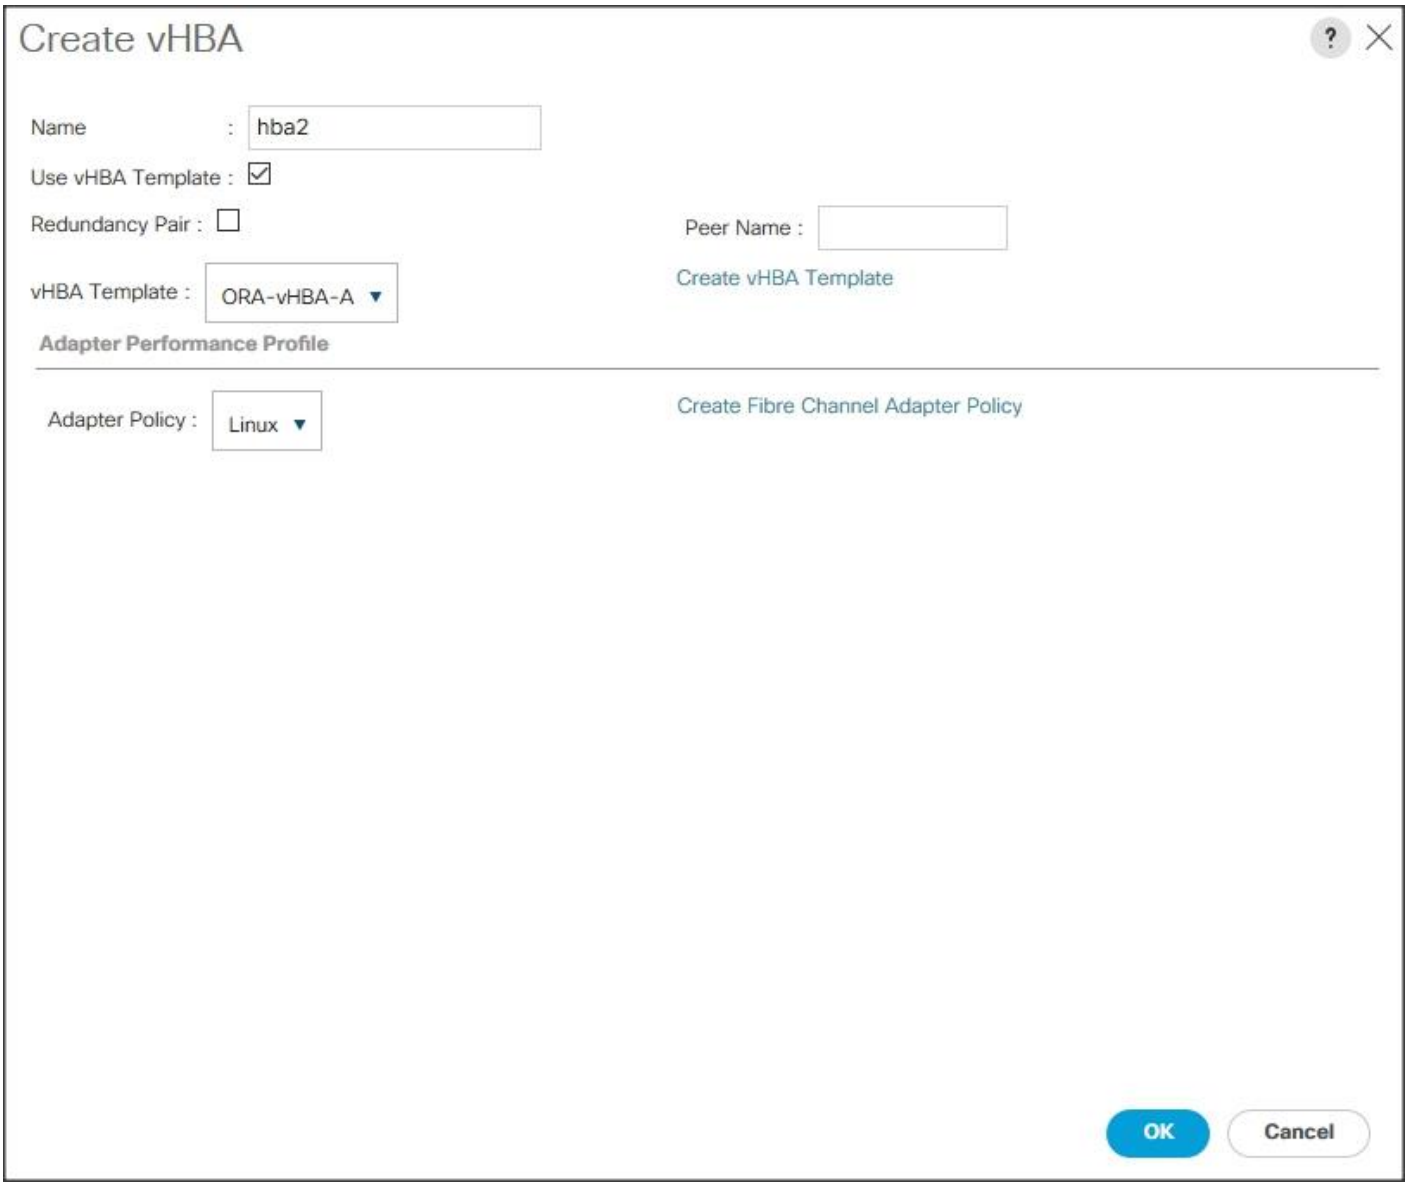
Create vHBA
Name : hba2
Use vHBA Template :
Redundancy Pair :
Peer Name :
vHBA Template : ORA-vHBA-A
Create vHBA Template
Adapter Performance Profile
Adapter Policy : Linux
Create Fibre Channel Adapter Policy
OK
Cancel

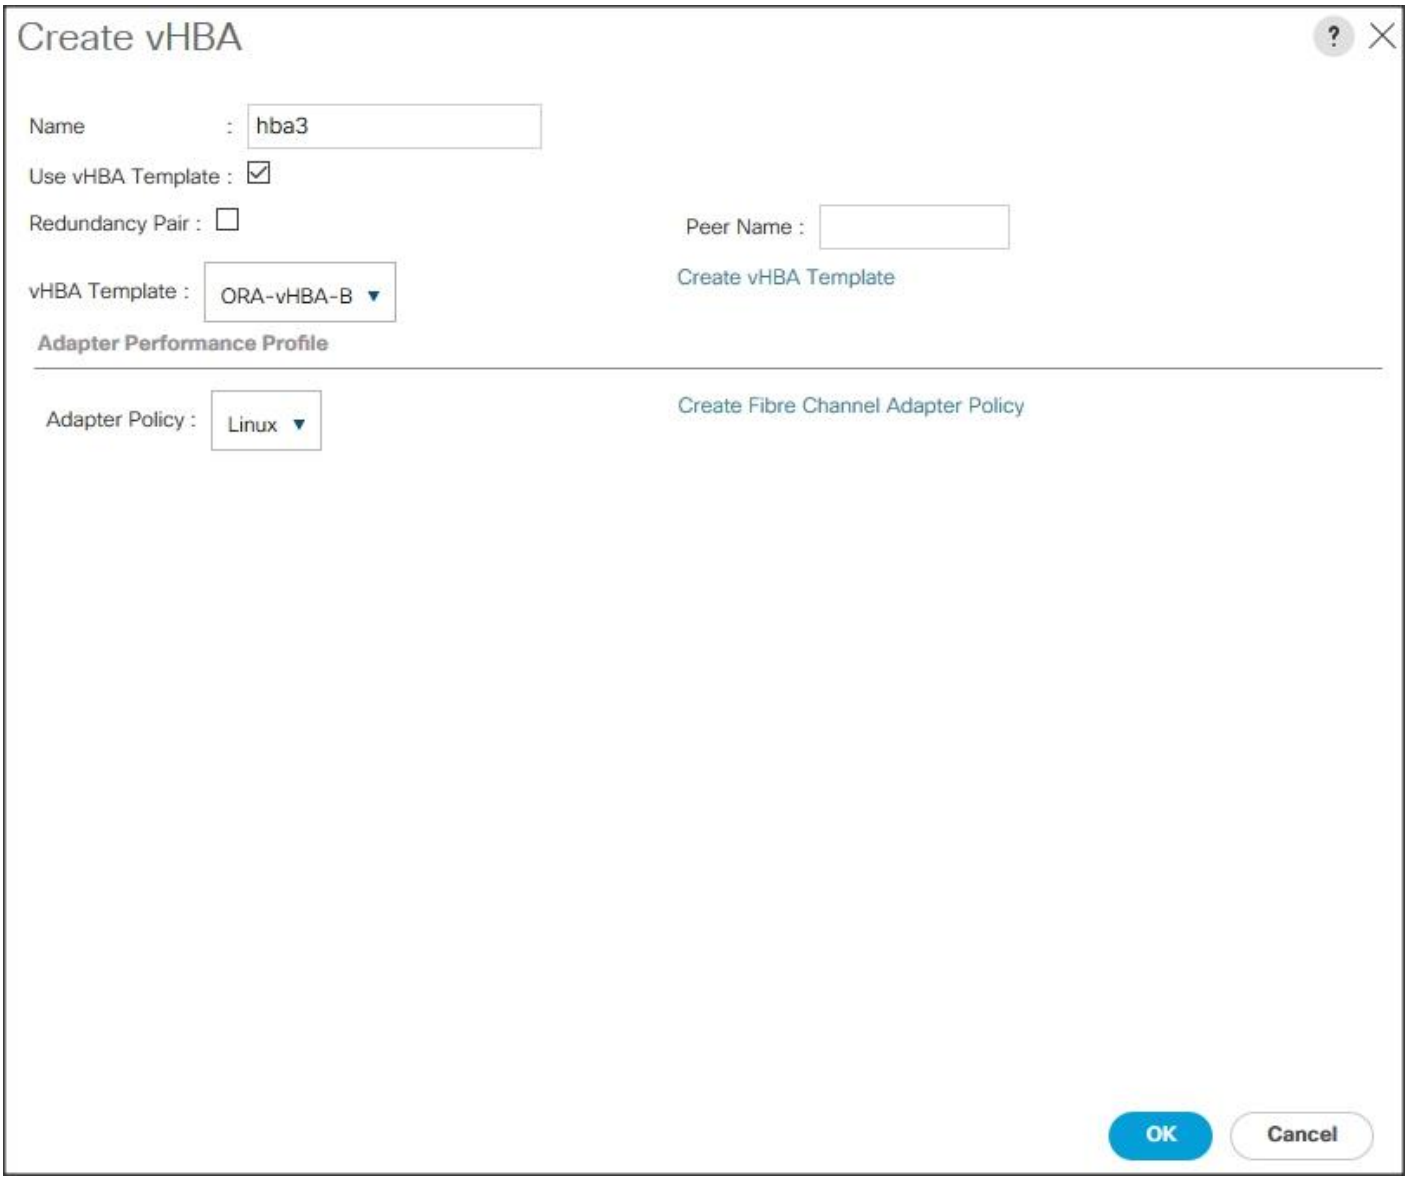
Create vHBA
Name : hba3
Use vHBA Template :
Redundancy Pair :
Peer Name :
vHBA Template : ORA-vHBA-B
Create vHBA Template
Adapter Performance Profile
Adapter Policy : Linux
Create Fibre Channel Adapter Policy
OK
Cancel

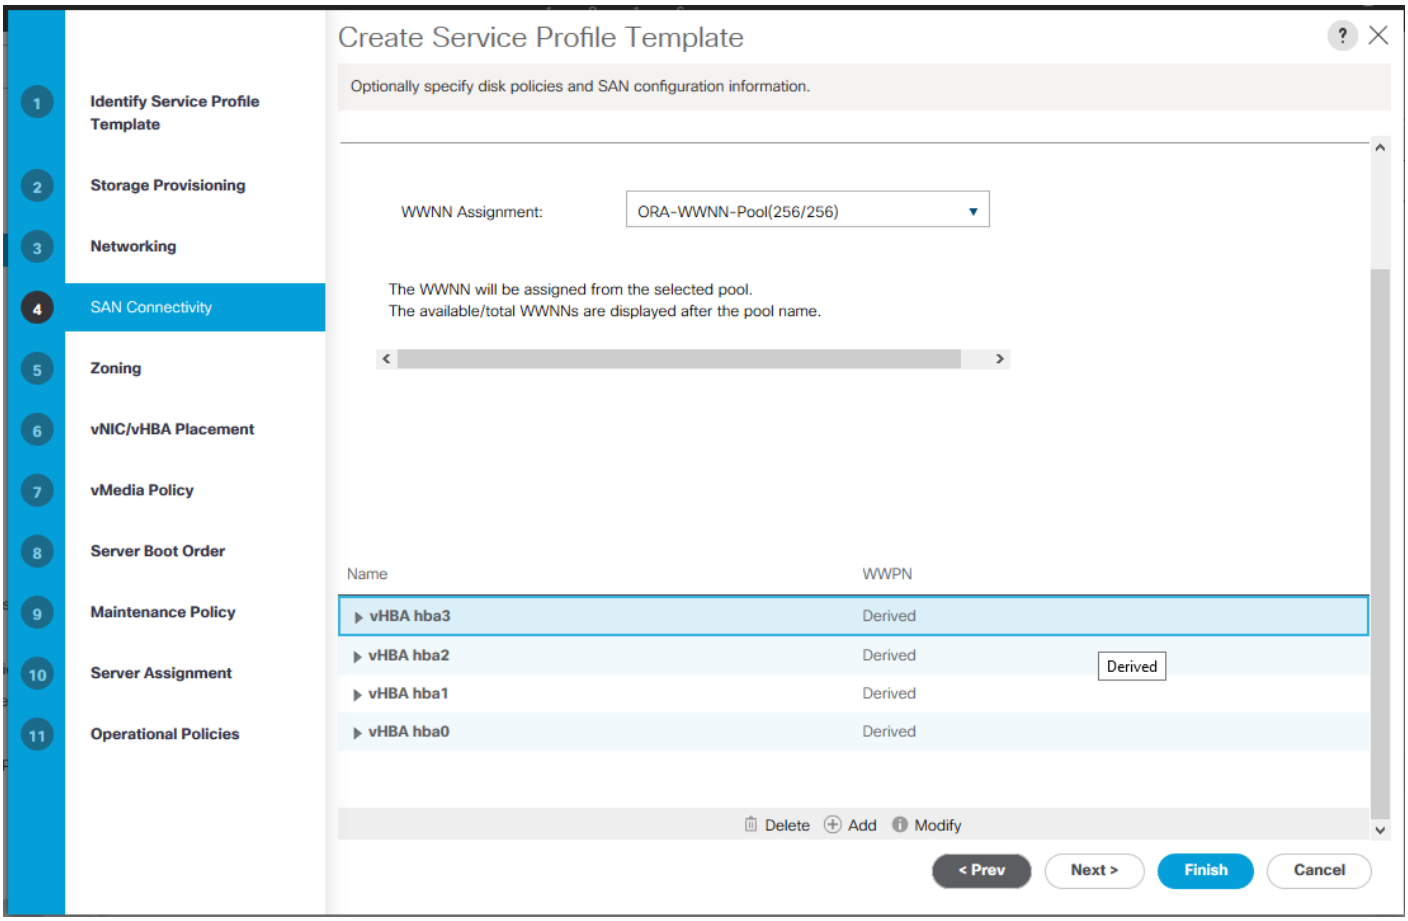


For this Oracle RAC Configuration, the Cisco MDS 9148S is used for zoning. Skip zoning and go to the next step.

9. In vNIC/vHBA Placement menu, keep option as Let System Perform Placement.

10. For this solution, do not configure any vMedia Policy. Click Next.

11. In the Server Boot Order, select “SAN_Boot” as Boot Policy created previously.

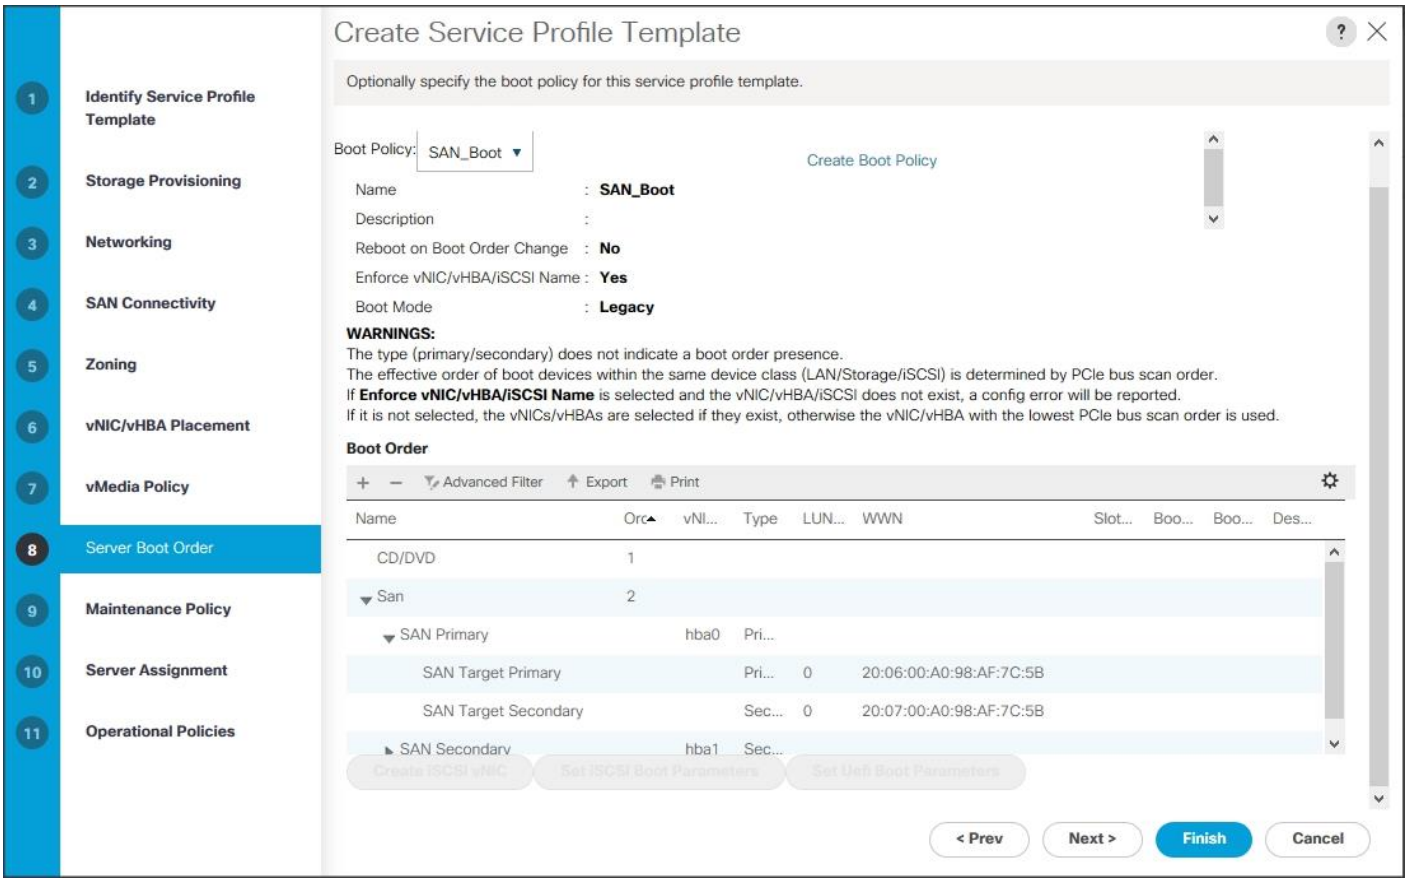


12. The maintenance policy and server assignment options were left as default in the configuration.

13. In operational policies menu, select OLTP_BIOS as BIOS Policy, created previously.

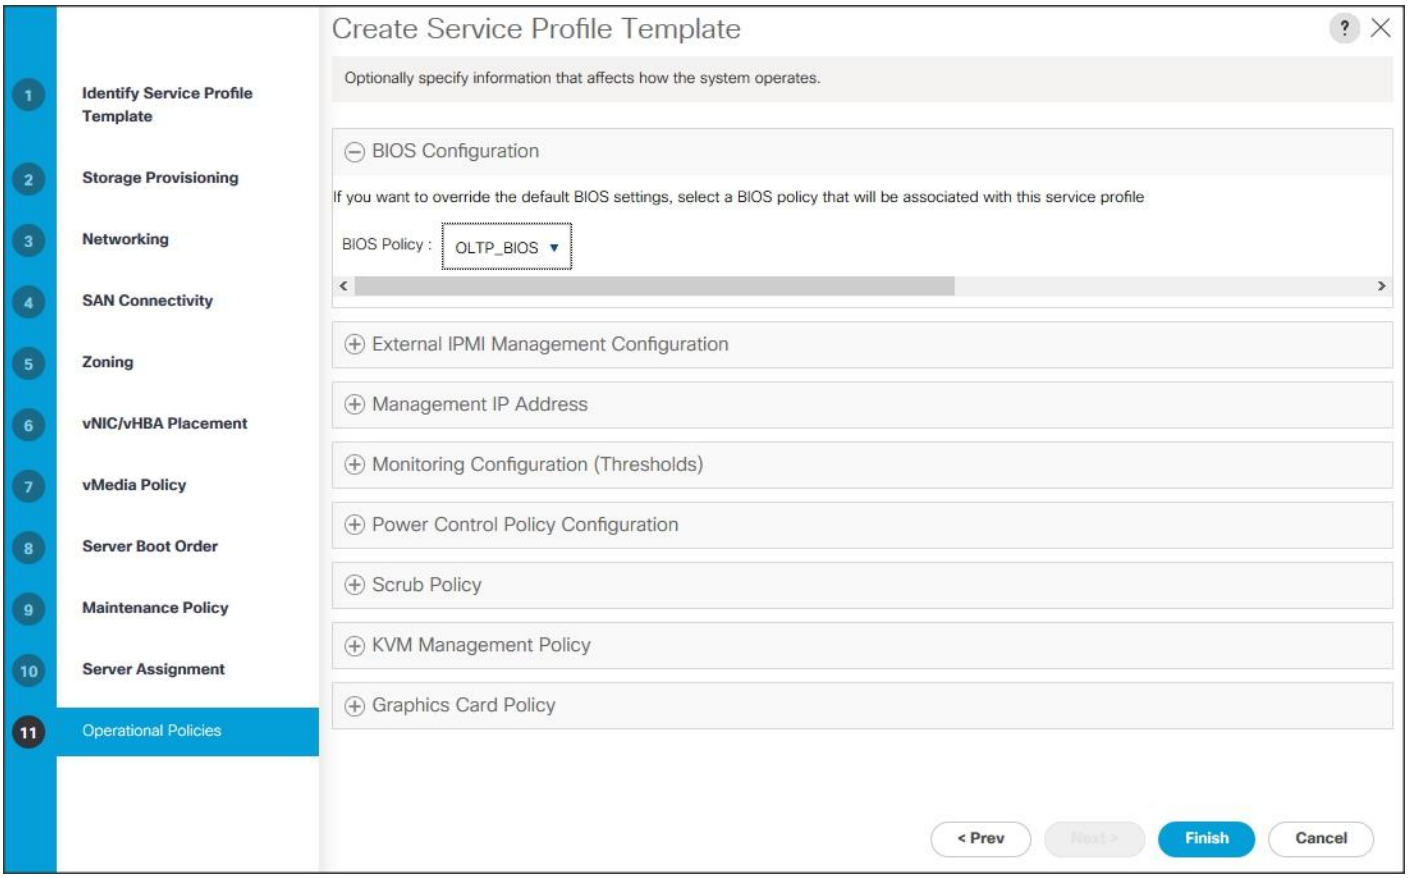


14. The rest of the configuration is left as default in the configuration. However, it may vary from site-to-site depending on workloads, best practices, and policies.

15. Click Finish to create service profile template as “ORA_FLEXPOD”. This service profile template will be used to create four service profiles for four oracle RAC nodes(flex1, flex2, flex3 and flex4).

Now you have one Service profile template as “ORA_FLEXPOD” having four vHBAs and two vNICs.

# Create Service Profiles from Template and Associate to Servers

## Create Service Profiles from Template

For all four Oracle RAC Nodes (flex1, flex2, flex3 and flex4), you will create four Service Profiles as FLEX1, FLEX2, FLEX3 and FLEX4 from Template “ORA_FLEXPOD”.

To create Service Profiles from Template, follow these steps:

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

2. Select the Service profile template as “ORA_FLEXPOD” which we created earlier and name the service profile as “FLEX”.

3. To create four service profiles, enter “Number of Instances” as 4 as shown below. This process creates service profiles as “FLEX1”, “FLEX2”, “FLEX3” and “FLEX4”.

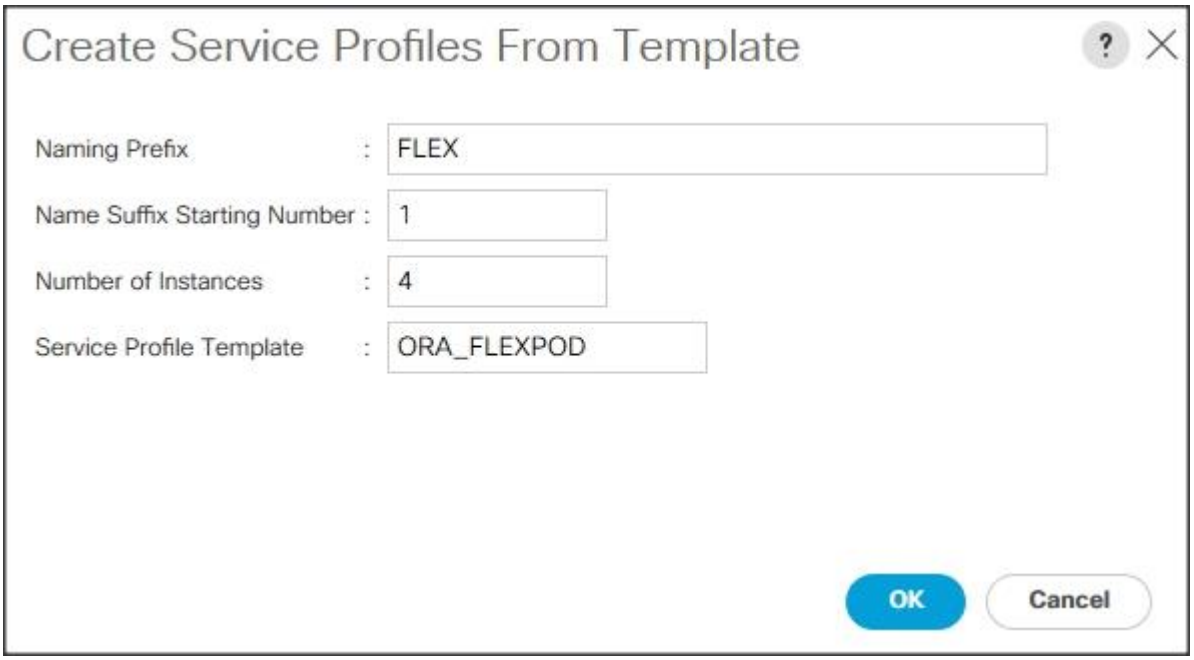


When the service profiles are created, associate them to the servers as described below.

## Associate Service Profiles to the Servers

To associate service profiles to the servers, follow these 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 the existing server that you want to assign and click OK.

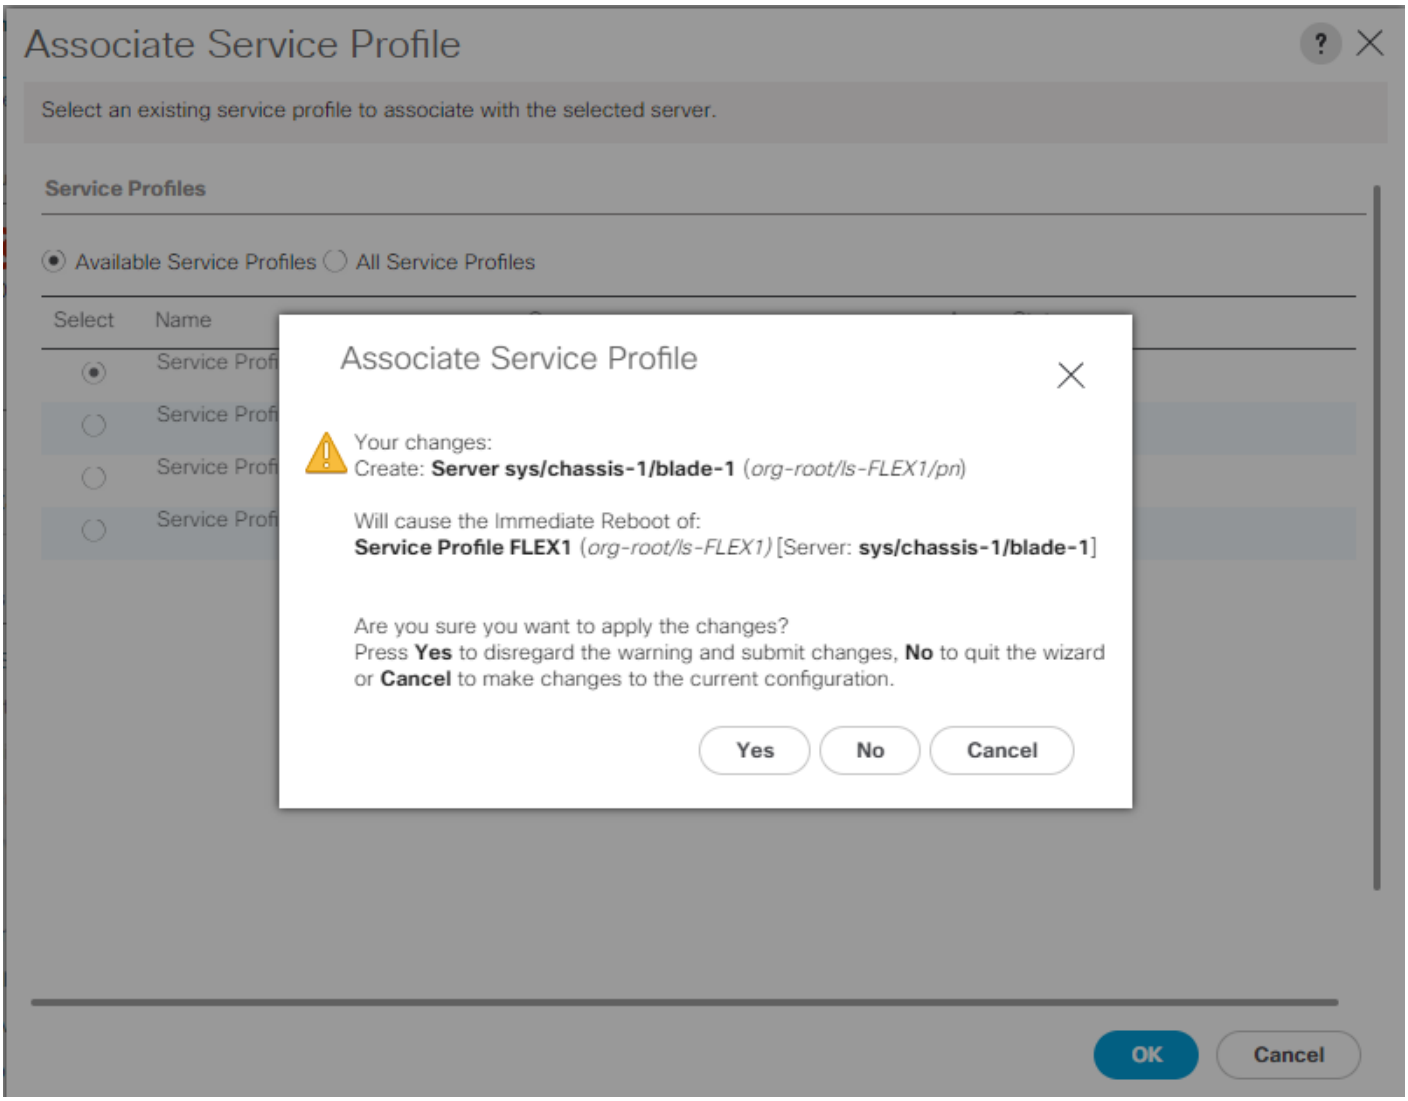


4. Assign service profiles FLEX1 to Chassis 1 Server 1 and service profile FLEX2 to Chassis 1 Server 2. Similarly, we will assign, service profiles FLEX3 to Chassis 2 Server 1 and service profile FLEX4 to Chassis 2 Server 2.

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

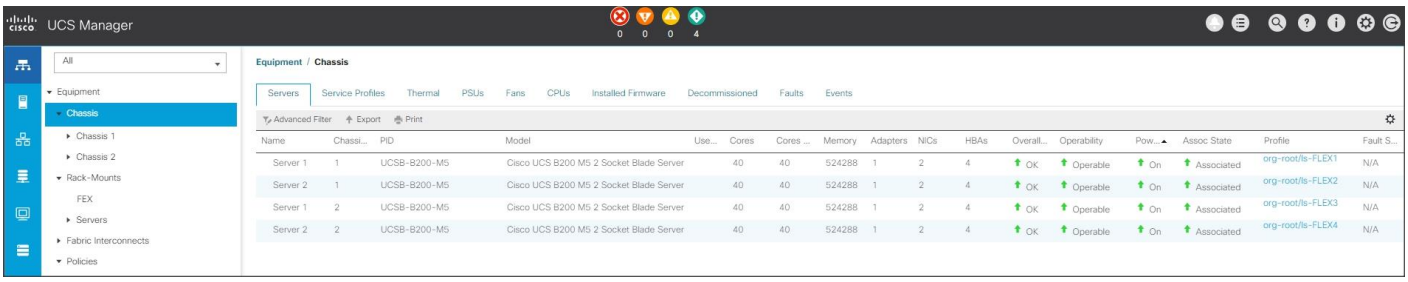


5. As shown above, make sure all the server nodes have no major or critical faults and all are in an operable state.

This completes the configuration required for Cisco UCS Manager Setup.

**Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations**

## Configure Cisco MDS 9148S Switch

This section provides a detailed procedure for configuring the Cisco MDS 9148S Switches.

**Follow these steps precisely because failure to do so could result in an improper configuration.**

We connected MDS Switches to Fabric Interconnects and NetApp AFF A700s Storage System as shown below.

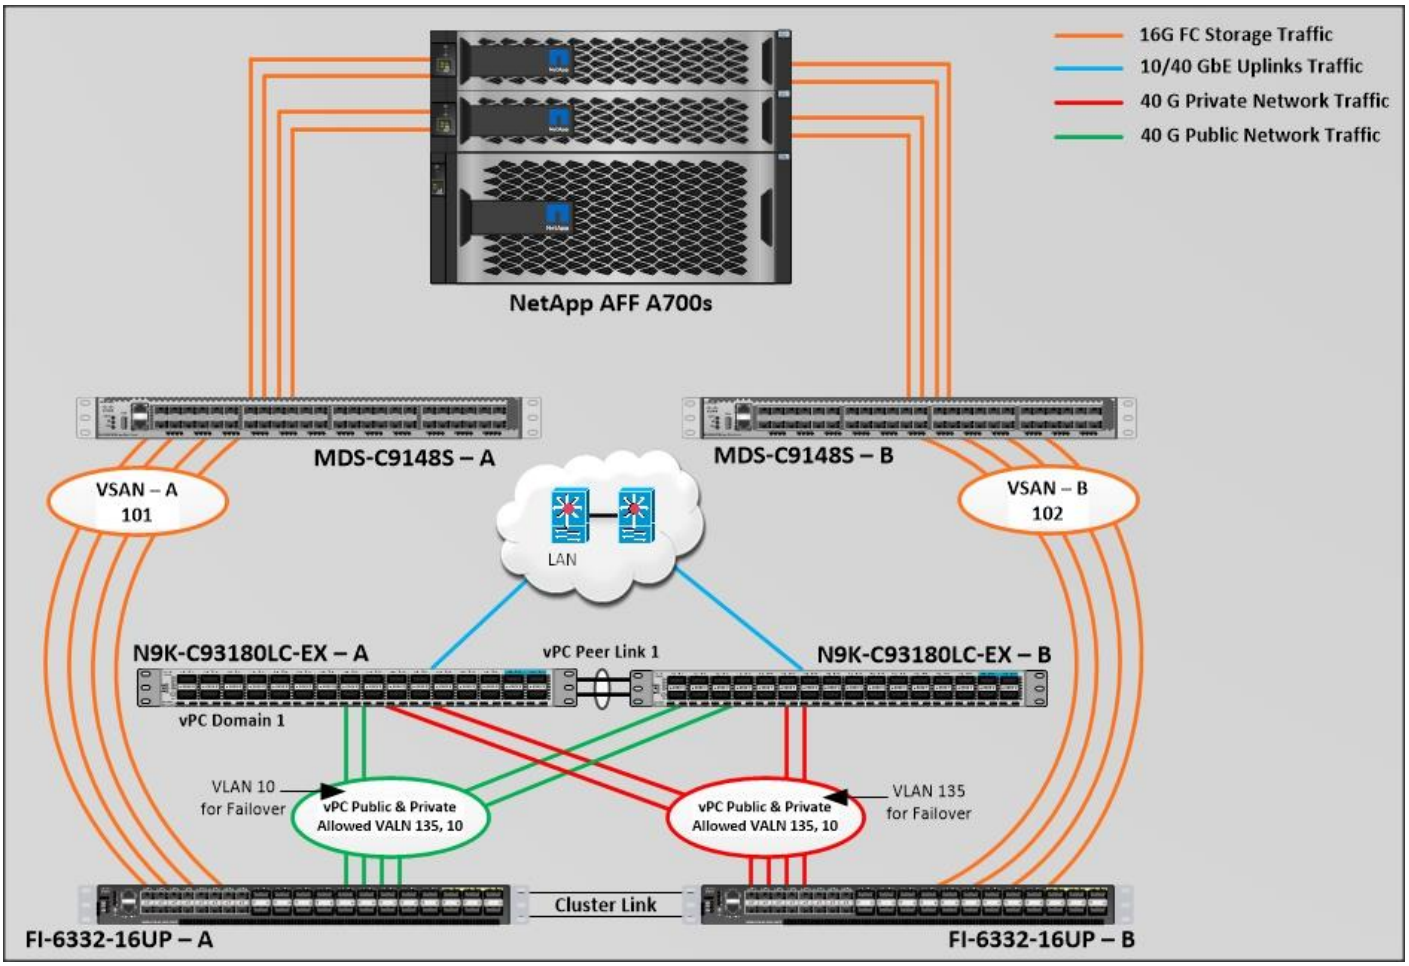


For this solution, we connected four ports (ports 1 to 4) of MDS Switch A to Fabric Interconnect A (ports 1-4). We also connected four ports (ports 1 to 4) of MDS Switch B to Fabric Interconnect B (ports 1-4). All ports carry 16 Gb/s FC Traffic. Table 8 lists port connectivity of Cisco MDS Switches to the Fabric Interconnects.

Table 8 MDS Switch Connectivity to the Fabric Interconnects

| MDS Switch | MDS Switch Port | FI Ports | Fabric Interconnect |
|---|---|---|---|
| MDS Switch A | FC Port 1/1 | FI-A Port 1/1 | Fabric Interconnect A (FI-A) |

| MDS Switch | MDS Switch Port | FI Ports | Fabric Interconnect |
|---|---|---|---|
| | FC Port 1/2 | FI-A Port 1/2 | |
| | FC Port 1/3 | FI-A Port 1/3 | |
| | FC Port 1/4 | FI-A Port 1/4 | |
| MDS Switch B | FC Port 1/1 | FI-B Port 1/1 | Fabric Interconnect B (FI-B) |
| | FC Port 1/2 | FI-B Port 1/2 | |
| | FC Port 1/3 | FI-B Port 1/3 | |
| | FC Port 1/4 | FI-B Port 1/4 | |

For these solution, we have connected four ports (ports 9 to 12) of MDS Switch A to the NetApp AFF A700s Storage controller. Similarly, we have connected four ports (ports 9 to 12) of MDS Switch B to the NetApp AFF A700s Storage controller. All ports carry 16 Gb/s FC Traffic. Table 9 lists port connectivity of Cisco MDS Switches to NetApp AFF A700s Controller.

**Table 9 MDS Switch Connectivity to the NetApp AFF A700s Storage**

| MDS Switch | MDS Switch Port | NetApp Storage Controller | NetApp Controller Ports | NetApp Port Description |
|---|---|---|---|---|
| MDS Switch A | FC Port 1/9 | NetApp AFF A700s Controller 1 | FC Port 0 – Slot 2 | NetApp-A700s-01-2a |
| | FC Port 1/10 | NetApp AFF A700s Controller 1 | FC Port 0 – Slot 3 | NetApp-A700s-01-3a |
| | FC Port 1/11 | NetApp AFF A700s Controller 2 | FC Port 0 – Slot 2 | NetApp-A700s-02-2a |
| | FC Port 1/12 | NetApp AFF A700s Controller 2 | FC Port 0 – Slot 3 | NetApp-A700s-02-3a |
| MDS Switch B | FC Port 1/9 | NetApp AFF A700s Controller 1 | FC Port 1 – Slot 2 | NetApp-A700s-01-2b |
| | FC Port 1/10 | NetApp AFF A700s Controller 1 | FC Port 1 – Slot 3 | NetApp-A700s-01-3b |
| | FC Port 1/11 | NetApp AFF A700s Controller 2 | FC Port 1 – Slot 2 | NetApp-A700s-02-2b |
| | FC Port 1/12 | NetApp AFF A700s Controller 2 | FC Port 1 – Slot 3 | NetApp-A700s-02-3b |

## Configure feature for MDS Switch A and MDS Switch B

To configure feature on MDS Switches, follow these steps:

1. Log in as admin user into MDS Switch A:

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

2. Login as admin user into MDS Switch B:

```
FLEXPOD-MDS-B# config terminal
FLEXPOD-MDS-B(config)# feature npiv
FLEXPOD-MDS-B(config)# feature telnet
```

```
FLEXPOD-MDS-B(config)# switchname FLEXPOD-MDS-B
FLEXPOD-MDS-B(config)# copy running-config startup-config
```

## Configure VSANs for MDS Switch A and MDS Switch B

To create VSANs, follow these steps:

1. Log in as admin user into MDS Switch A.
2. Create VSAN 101 for Storage Traffic:

```
FLEXPOD-MDS-A # config terminal
FLEXPOD-MDS-A(config)# VSAN database
FLEXPOD-MDS-A(config-vsan-db)# vsan 101
FLEXPOD-MDS-A(config-vsan-db)# vsan 101 interface fc 1/1-12
FLEXPOD-MDS-A(config-vsan-db)# exit

FLEXPOD-MDS-A(config)# interface fc 1/1-12
FLEXPOD-MDS-A(config-if)# switchport trunk allowed vsan 101
FLEXPOD-MDS-A(config-if)# switchport trunk mode off
FLEXPOD-MDS-A(config-if)# port-license acquire
FLEXPOD-MDS-A(config-if)# no shutdown
FLEXPOD-MDS-A(config-if)# exit
FLEXPOD-MDS-A(config)# copy running-config startup-config
```

3. Login as admin user into MDS Switch B.
4. Create VSAN 102 for Storage Traffic:

```
FLEXPOD-MDS-B # config terminal
FLEXPOD-MDS-B(config)# VSAN database
FLEXPOD-MDS-B(config-vsan-db)# vsan 102
FLEXPOD-MDS-B(config-vsan-db)# vsan 102 interface fc 1/1-12
FLEXPOD-MDS-B(config-vsan-db)# exit

FLEXPOD-MDS-B(config)# interface fc 1/1-12
FLEXPOD-MDS-B(config-if)# switchport trunk allowed vsan 102
FLEXPOD-MDS-B(config-if)# switchport trunk mode off
FLEXPOD-MDS-B(config-if)# port-license acquire
```

```
FLEXPOD-MDS-B(config-if)# no shutdown

FLEXPOD-MDS-B(config-if)# exit

FLEXPOD-MDS-B(config)# copy running-config startup-config
```

## Create and Configure Fibre Channel Zoning

This procedure sets up the Fibre Channel connections between the Cisco MDS 9148S switches, the Cisco UCS Fabric Interconnects, and the NetApp AFF Storage 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 Fibre channel zoning, follow these steps:

1. Log into 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 below.

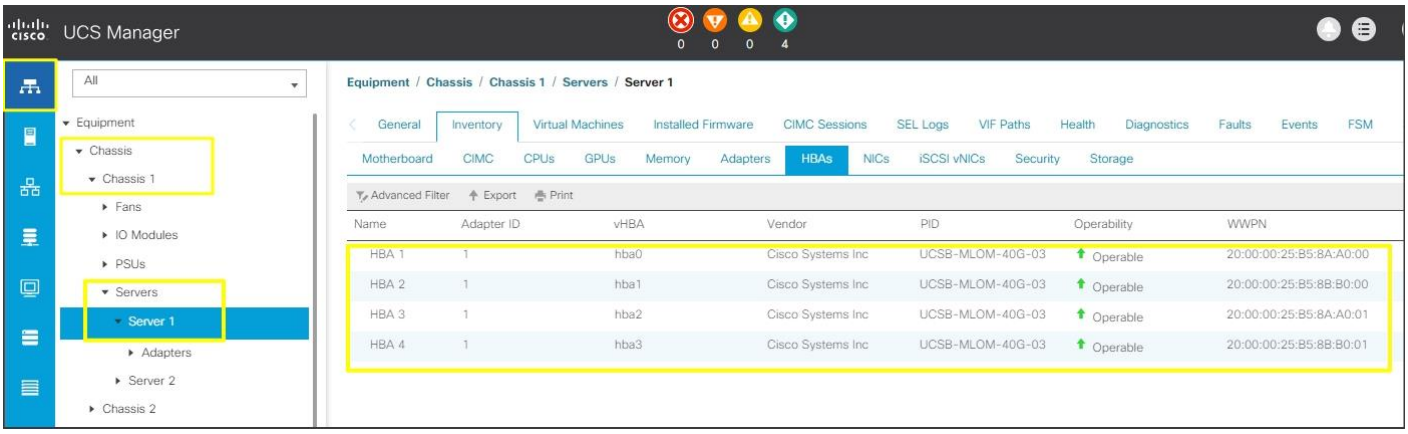


2. Log into NetApp storage controller and extract the WWPN of FC LIFs configured and verify all the port information is correct. This information can be found in the NetApp Storage GUI under Network > Network Interfaces.

The screenshot below shows the network interface, WWPN and ports connectivity configured for NetApp AFF A700s controller. Four Fibre Channel logical interfaces (LIFs) are created on storage controller cluster node 1 (node1_lif02a, node1_lif02b, node1_lif03a and node1_lif03b) and four Fibre Channel LIFs are created on storage controller cluster node 2 (node2_lif02a, node2_lif02b, node2_lif03a and node2_lif03b).

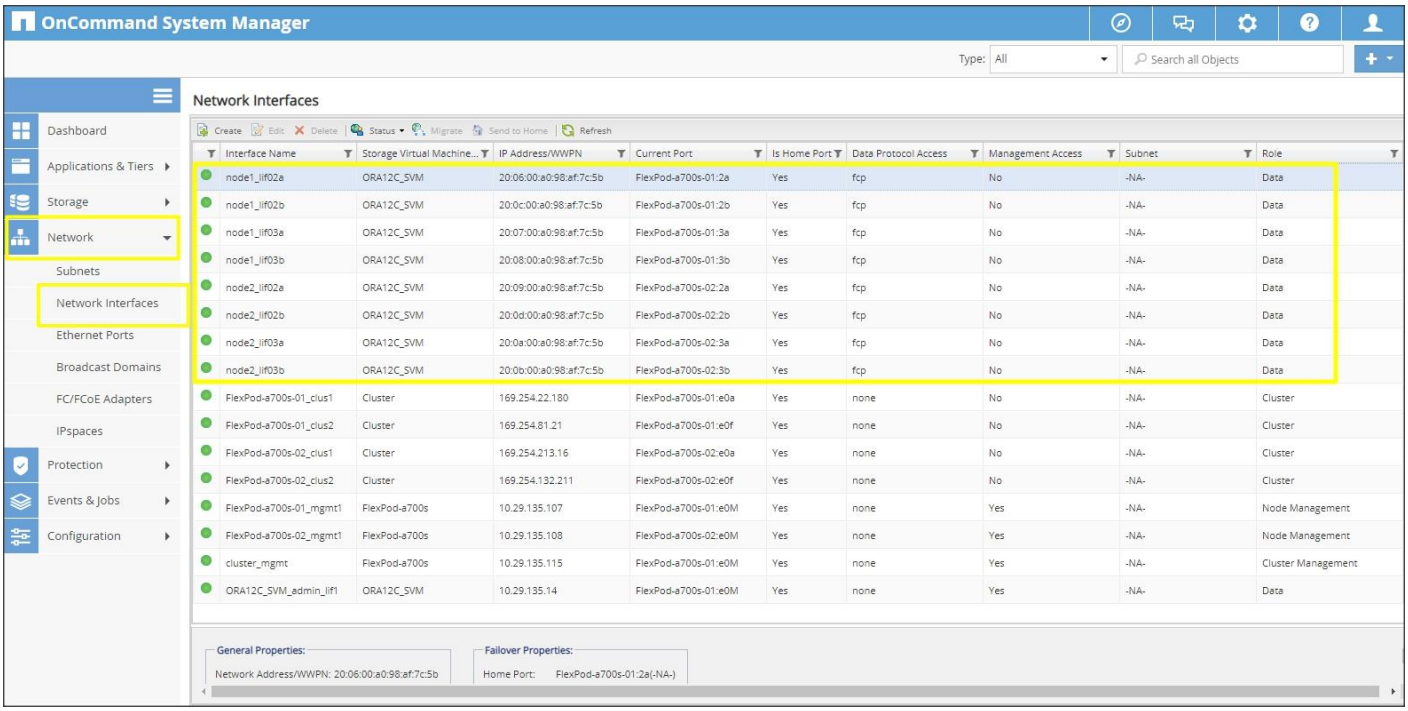


**You can also obtain this information by login to the storage cluster and run the network interface show command**

The NetApp Storage AFF A700s have eight active FC connection goes to the Cisco MDS switches. Four FC ports are connected to Cisco MDS-A, and other four FC ports are connected to Cisco MDS-B Switches. The SAN Ports FC-Port-0-Slot-2 and FC-Port-0-Slot-3 of NetApp AFF A700s Controller – 1 are connected to Cisco MDS Switch A and FC-Port-1-Slot-2 and FC-Port-1-Slot-3 are connected to Cisco MDS Switch B. Similarly, the SAN Ports FC-Port-0-Slot-2 and FC-Port-0-Slot-3 of NetApp AFF A700s Controller – 2 are connected to Cisco MDS Switch A and FC-Port-1-Slot-2 and FC-Port-1-Slot-3 are connected to Cisco MDS Switch B.

## Create Device Aliases for Fibre Channel Zoning on Cisco MDS Switch A

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

1. Log in as admin user and run the following commands:

```
configure terminal

device-alias database

  device-alias name flex1-hba0 pwwn 20:00:00:25:b5:8a:a0:00

  device-alias name flex1-hba2 pwwn 20:00:00:25:b5:8a:a0:01

  device-alias name flex2-hba0 pwwn 20:00:00:25:b5:8a:a0:02

  device-alias name flex2-hba2 pwwn 20:00:00:25:b5:8a:a0:03

  device-alias name flex3-hba0 pwwn 20:00:00:25:b5:8a:a0:04

  device-alias name flex3-hba2 pwwn 20:00:00:25:b5:8a:a0:05
```

```
  device-alias name flex4-hba0 pwwn 20:00:00:25:b5:8a:a0:06
  device-alias name flex4-hba2 pwwn 20:00:00:25:b5:8a:a0:07
  device-alias name NetApp-A700s-01-2A pwwn 20:06:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-01-3A pwwn 20:07:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-2A pwwn 20:09:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-3A pwwn 20:0a:00:a0:98:af:7c:5b
  device-alias commit
copy run start
```

### Create Device Aliases for Fibre Channel Zoning on Cisco MDS Switch B

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

1. Log in as admin user and run the following commands:

```
configure terminal
device-alias database
  device-alias name flex1-hba1 pwwn 20:00:00:25:b5:8b:b0:00
  device-alias name flex1-hba3 pwwn 20:00:00:25:b5:8b:b0:01
  device-alias name flex2-hba1 pwwn 20:00:00:25:b5:8b:b0:02
  device-alias name flex2-hba3 pwwn 20:00:00:25:b5:8b:b0:03
  device-alias name flex3-hba1 pwwn 20:00:00:25:b5:8b:b0:04
  device-alias name flex3-hba3 pwwn 20:00:00:25:b5:8b:b0:05
  device-alias name flex4-hba1 pwwn 20:00:00:25:b5:8b:b0:06
  device-alias name flex4-hba3 pwwn 20:00:00:25:b5:8b:b0:07
  device-alias name NetApp-A700s-01-2B pwwn 20:0c:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-01-3B pwwn 20:08:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-2B pwwn 20:0d:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-3B pwwn 20:0b:00:a0:98:af:7c:5b
device-alias commit
copy run start
```

## Create Zoning

### Cisco MDS Switch A

To configure zones for the Cisco MDS Switch A, follow these steps:

1. Create a zone for each service profile as shown below.

2. Log in as admin user into MDS Switch A and run these commands to create the zone:

```
configure terminal
zone name flex1 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:00
    member pwwn 20:00:00:25:b5:8a:a0:01
    member pwwn 20:06:00:a0:98:af:7c:5b
    member pwwn 20:07:00:a0:98:af:7c:5b
    member pwwn 20:09:00:a0:98:af:7c:5b
    member pwwn 20:0a:00:a0:98:af:7c:5b

zone name flex2 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:02
    member pwwn 20:00:00:25:b5:8a:a0:03
    member pwwn 20:06:00:a0:98:af:7c:5b
    member pwwn 20:07:00:a0:98:af:7c:5b
    member pwwn 20:09:00:a0:98:af:7c:5b
    member pwwn 20:0a:00:a0:98:af:7c:5b

zone name flex3 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:04
    member pwwn 20:00:00:25:b5:8a:a0:05
    member pwwn 20:06:00:a0:98:af:7c:5b
    member pwwn 20:07:00:a0:98:af:7c:5b
    member pwwn 20:09:00:a0:98:af:7c:5b
    member pwwn 20:0a:00:a0:98:af:7c:5b

zone name flex4 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:06
    member pwwn 20:00:00:25:b5:8a:a0:07
    member pwwn 20:06:00:a0:98:af:7c:5b
    member pwwn 20:07:00:a0:98:af:7c:5b
    member pwwn 20:09:00:a0:98:af:7c:5b
```

```
    member pwwn 20:0a:00:a0:98:af:7c:5b
```

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

```
zoneset name flex vsan 101
    member flex1
    member flex2
    member flex3
    member flex4
```

4. Activate the zone set by running following commands

```
zoneset activate name flex vsan 101
copy run start
```

### Cisco MDS Switch B

To configure zones for the Cisco MDS Switch B, follow these steps:

1. Create a zone for each service profile.
2. Log in as admin user into MDS Switch B and run the below commands to create the zone:

```
configure terminal
zone name flex1 vsan 102
    member pwwn 20:00:00:25:b5:8b:b0:00
    member pwwn 20:00:00:25:b5:8b:b0:01
    member pwwn 20:0c:00:a0:98:af:7c:5b
    member pwwn 20:08:00:a0:98:af:7c:5b
    member pwwn 20:0d:00:a0:98:af:7c:5b
    member pwwn 20:0b:00:a0:98:af:7c:5b

zone name flex2 vsan 102
    member pwwn 20:00:00:25:b5:8b:b0:02
    member pwwn 20:00:00:25:b5:8b:b0:03
    member pwwn 20:0c:00:a0:98:af:7c:5b
    member pwwn 20:08:00:a0:98:af:7c:5b
    member pwwn 20:0d:00:a0:98:af:7c:5b
    member pwwn 20:0b:00:a0:98:af:7c:5b
```

```
zone name flex3 vsan 102
    member pwwn 20:00:00:25:b5:8b:b0:04
    member pwwn 20:00:00:25:b5:8b:b0:05
    member pwwn 20:0c:00:a0:98:af:7c:5b
    member pwwn 20:08:00:a0:98:af:7c:5b
    member pwwn 20:0d:00:a0:98:af:7c:5b
    member pwwn 20:0b:00:a0:98:af:7c:5b

zone name flex4 vsan 102
    member pwwn 20:00:00:25:b5:8b:b0:06
    member pwwn 20:00:00:25:b5:8b:b0:07
    member pwwn 20:0c:00:a0:98:af:7c:5b
    member pwwn 20:08:00:a0:98:af:7c:5b
    member pwwn 20:0d:00:a0:98:af:7c:5b
    member pwwn 20:0b:00:a0:98:af:7c:5b
```

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

```
zoneset name flex vsan 102
    member flex1
    member flex2
    member flex3
    member flex4
```

4. Activate the zone set by running following commands:

```
zoneset activate name flex vsan 102
copy run start
```

## Verify FC Ports on MDS Switch

To verify FC ports on the MDS switch, follow these steps:

1. Log in as admin user into MDS Switch A and run the “show flogi database vsan 101” to verify all FC ports.

```
FLEXPOD-MDS-A#
FLEXPOD-MDS-A# show flogi database vsan 101
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/1            101   0xa20400  20:01:00:de:fb:92:99:00 20:65:00:de:fb:92:99:01
fc1/1            101   0xa20401  20:00:00:25:b5:8a:a0:00 20:00:00:25:b5:7a:00:00
                           [flex1-hba0]
fc1/1            101   0xa20402  20:00:00:25:b5:8a:a0:01 20:00:00:25:b5:7a:00:00
                           [flex1-hba2]
fc1/2            101   0xa20500  20:02:00:de:fb:92:99:00 20:65:00:de:fb:92:99:01
fc1/2            101   0xa20501  20:00:00:25:b5:8a:a0:03 20:00:00:25:b5:7a:00:01
                           [flex2-hba2]
fc1/2            101   0xa20503  20:00:00:25:b5:8a:a0:02 20:00:00:25:b5:7a:00:01
                           [flex2-hba0]
fc1/3            101   0xa20600  20:03:00:de:fb:92:99:00 20:65:00:de:fb:92:99:01
fc1/3            101   0xa20601  20:00:00:25:b5:8a:a0:06 20:00:00:25:b5:7a:00:03
                           [flex4-hba0]
fc1/3            101   0xa20602  20:00:00:25:b5:8a:a0:07 20:00:00:25:b5:7a:00:03
                           [flex4-hba2]
fc1/4            101   0xa20700  20:04:00:de:fb:92:99:00 20:65:00:de:fb:92:99:01
fc1/4            101   0xa20701  20:00:00:25:b5:8a:a0:05 20:00:00:25:b5:7a:00:02
                           [flex3-hba2]
fc1/4            101   0xa20702  20:00:00:25:b5:8a:a0:04 20:00:00:25:b5:7a:00:02
                           [flex3-hba0]
fc1/9            101   0xa20000  50:0a:09:81:80:12:f8:47 50:0a:09:80:80:12:f8:47
fc1/9            101   0xa20002  20:06:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-01-2A]
fc1/10           101   0xa20100  50:0a:09:83:80:12:f8:47 50:0a:09:80:80:12:f8:47
fc1/10           101   0xa20102  20:07:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-01-3A]
fc1/11           101   0xa20200  50:0a:09:81:80:12:f8:3d 50:0a:09:80:80:12:f8:3d
fc1/11           101   0xa20201  20:09:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-02-2A]
fc1/12           101   0xa20300  50:0a:09:83:80:12:f8:3d 50:0a:09:80:80:12:f8:3d
fc1/12           101   0xa20301  20:0a:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-02-3A]

Total number of flogi = 20.
```

2. Log in as admin user into MDS Switch B and run the “show flogi database vsan 101” to verify all FC ports.

```
FLEXPOD-MDS-B#
FLEXPOD-MDS-B# show flogi database vsan 102
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/1            102   0x8e0400  20:01:00:de:fb:92:95:40 20:66:00:de:fb:92:95:41
fc1/1            102   0x8e0401  20:00:00:25:b5:8b:b0:00 20:00:00:25:b5:7a:00:00
                           [flex1-hba1]
fc1/1            102   0x8e0402  20:00:00:25:b5:8b:b0:01 20:00:00:25:b5:7a:00:00
                           [flex1-hba3]
fc1/2            102   0x8e0500  20:02:00:de:fb:92:95:40 20:66:00:de:fb:92:95:41
fc1/2            102   0x8e0501  20:00:00:25:b5:8b:b0:03 20:00:00:25:b5:7a:00:01
                           [flex2-hba3]
fc1/2            102   0x8e0502  20:00:00:25:b5:8b:b0:02 20:00:00:25:b5:7a:00:01
                           [flex2-hba1]
fc1/3            102   0x8e0600  20:03:00:de:fb:92:95:40 20:66:00:de:fb:92:95:41
fc1/3            102   0x8e0601  20:00:00:25:b5:8b:b0:06 20:00:00:25:b5:7a:00:03
                           [flex4-hba1]
fc1/3            102   0x8e0603  20:00:00:25:b5:8b:b0:07 20:00:00:25:b5:7a:00:03
                           [flex4-hba3]
fc1/4            102   0x8e0700  20:04:00:de:fb:92:95:40 20:66:00:de:fb:92:95:41
fc1/4            102   0x8e0701  20:00:00:25:b5:8b:b0:05 20:00:00:25:b5:7a:00:02
                           [flex3-hba3]
fc1/4            102   0x8e0702  20:00:00:25:b5:8b:b0:04 20:00:00:25:b5:7a:00:02
                           [flex3-hba1]
fc1/9            102   0x8e0000  50:0a:09:82:80:12:f8:47 50:0a:09:80:80:12:f8:47
fc1/9            102   0x8e0002  20:0c:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-01-2B]
fc1/10           102   0x8e0100  50:0a:09:84:80:12:f8:47 50:0a:09:80:80:12:f8:47
fc1/10           102   0x8e0102  20:08:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-01-3B]
fc1/11           102   0x8e0200  50:0a:09:82:80:12:f8:3d 50:0a:09:80:80:12:f8:3d
fc1/11           102   0x8e0201  20:0d:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-02-2B]
fc1/12           102   0x8e0300  50:0a:09:84:80:12:f8:3d 50:0a:09:80:80:12:f8:3d
fc1/12           102   0x8e0301  20:0b:00:a0:98:af:7c:5b 20:01:00:a0:98:af:7c:5b
                           [NetApp-A700s-02-3B]

Total number of flogi = 20.
```

# Configure NetApp AFF A700s Storage

## NetApp Storage Connectivity

It is beyond the scope of this document to describe all the detailed information about NetApp storage connectivity and infrastructure configuration. Follow the link below for the installation and setup instructions for the NetApp AFF A700s System:

https://library.netapp.com/ecm/ecm_download_file/ECMLP2619982

For more technical information, refer to the following Cisco site:

https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html

Please refer to the following technical reports of “Oracle Databases on ONTAP” and “Database Data Protection: Backup, Recovery, Replication, and DR” from NetApp for the best practices:

https://www.netapp.com/us/media/tr-3633.pdf

https://www.netapp.com/us/media/tr-4591.pdf

This section describes the storage layout and design considerations for the database deployment.

A NetApp ONTAP cluster serves data through at least one and possibly multiple storage virtual machines (SVMs; formerly called Vservers). An SVM is a logical abstraction that represents the set of physical resources of the cluster. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and might reside on any node in the cluster to which the SVM has been given access. An SVM might own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node in the storage cluster to another. For example, a NetApp FlexVol® flexible volume can be non-disruptively moved to a new node, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware, and thus it is not tied to any specific physical hardware.

**The screenshot below shows the SVM and FC Interfaces configuration. For this solution, we configured one SVM as "ORA12C_SVM."**

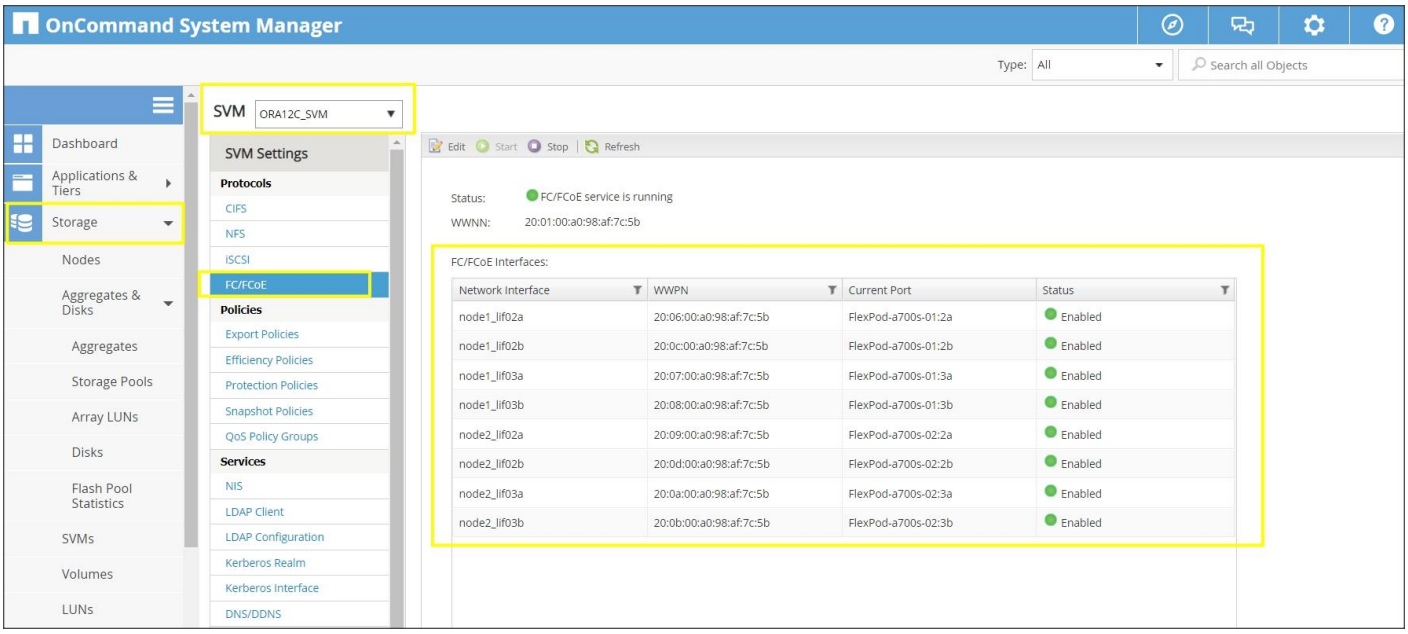


As shown above, for both the storage controller nodes (FlexPod-a700s-01 and FlexPod-a700s-02), we used ports 2a, 2b and 3a, 3b to configure LIFs. WWPN of these LIFs are used for zoning into the MDS switches for storage to MDS connectivity. The screenshot below shows the network interface configuration for this solution.

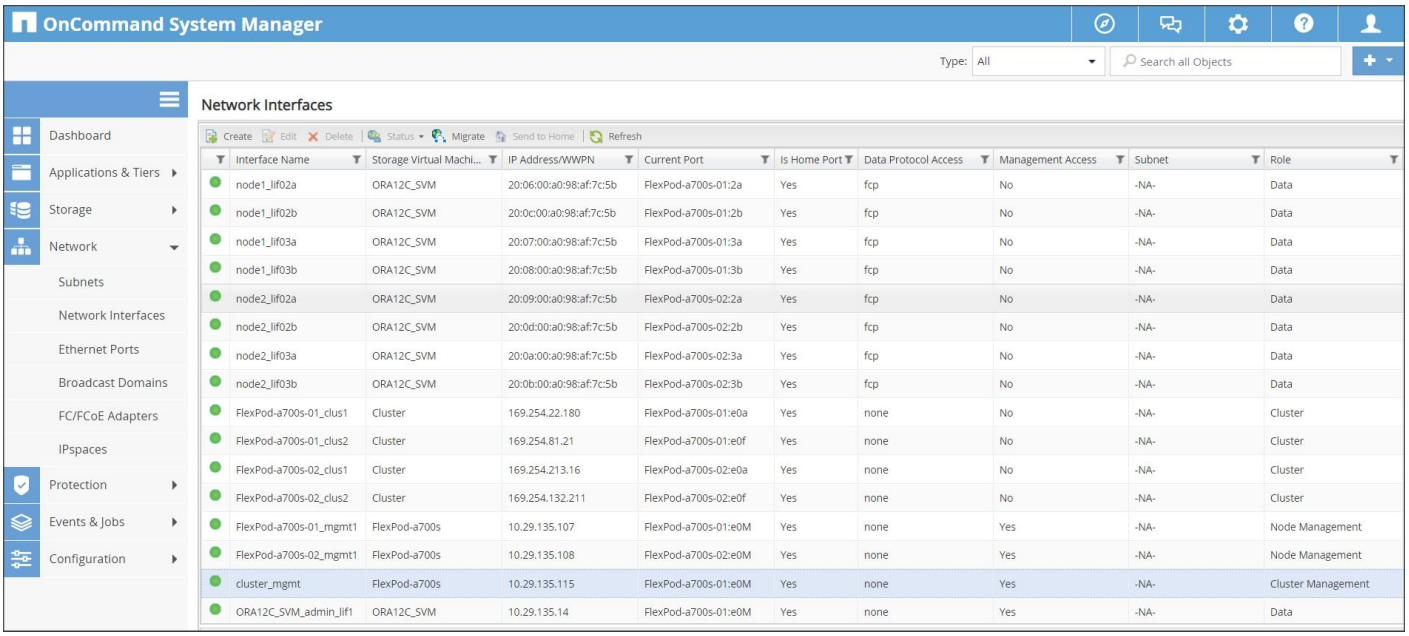


| Interface Name | Storage Virtual Machi... | IP Address/WWPN | Current Port | Is Home Port | Data Protocol Access | Management Access | Subnet | Role |
|---|---|---|---|---|---|---|---|---|
| node1_lif02a | ORA12C_SVM | 20:06:00:a0:98:af:7c:5b | FlexPod-a700s-01:2a | Yes | fcp | No | -NA- | Data |
| node1_lif02b | ORA12C_SVM | 20:0c:00:a0:98:af:7c:5b | FlexPod-a700s-01:2b | Yes | fcp | No | -NA- | Data |
| node1_lif03a | ORA12C_SVM | 20:07:00:a0:98:af:7c:5b | FlexPod-a700s-01:3a | Yes | fcp | No | -NA- | Data |
| node1_lif03b | ORA12C_SVM | 20:08:00:a0:98:af:7c:5b | FlexPod-a700s-01:3b | Yes | fcp | No | -NA- | Data |
| node2_lif02a | ORA12C_SVM | 20:09:00:a0:98:af:7c:5b | FlexPod-a700s-02:2a | Yes | fcp | No | -NA- | Data |
| node2_lif02b | ORA12C_SVM | 20:0d:00:a0:98:af:7c:5b | FlexPod-a700s-02:2b | Yes | fcp | No | -NA- | Data |
| node2_lif03a | ORA12C_SVM | 20:0a:00:a0:98:af:7c:5b | FlexPod-a700s-02:3a | Yes | fcp | No | -NA- | Data |
| node2_lif03b | ORA12C_SVM | 20:0b:00:a0:98:af:7c:5b | FlexPod-a700s-02:3b | Yes | fcp | No | -NA- | Data |
| FlexPod-a700s-01_clus1 | Cluster | 169.254.22.180 | FlexPod-a700s-01:e0a | Yes | none | No | -NA- | Cluster |
| FlexPod-a700s-01_clus2 | Cluster | 169.254.81.21 | FlexPod-a700s-01:e0f | Yes | none | No | -NA- | Cluster |
| FlexPod-a700s-02_clus1 | Cluster | 169.254.213.16 | FlexPod-a700s-02:e0a | Yes | none | No | -NA- | Cluster |
| FlexPod-a700s-02_clus2 | Cluster | 169.254.132.211 | FlexPod-a700s-02:e0f | Yes | none | No | -NA- | Cluster |
| FlexPod-a700s-01_mgmt1 | FlexPod-a700s | 10.29.135.107 | FlexPod-a700s-01:e0M | Yes | none | Yes | -NA- | Node Management |
| FlexPod-a700s-02_mgmt1 | FlexPod-a700s | 10.29.135.108 | FlexPod-a700s-02:e0M | Yes | none | Yes | -NA- | Node Management |
| cluster_mgmt | FlexPod-a700s | 10.29.135.115 | FlexPod-a700s-01:e0M | Yes | none | Yes | -NA- | Cluster Management |
| ORA12C_SVM_admin_lif1 | ORA12C_SVM | 10.29.135.14 | FlexPod-a700s-01:e0M | Yes | none | Yes | -NA- | Data |

For all the database deployment, we configured two aggregates (one aggregate on each storage node) into a single SVM (ORA12C_SVM) and each aggregate contains 35 SSD (960GB Each) drives that were subdivided into RAID DP groups, plus one spare drive as shown below.

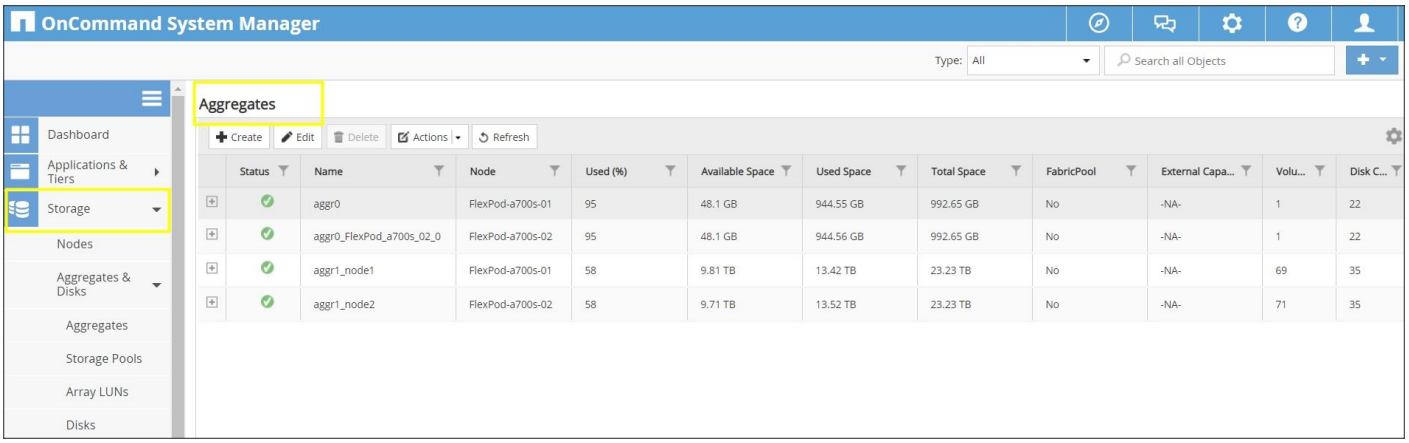


| Status | Name | Node | Used (%) | Available Space | Used Space | Total Space | FabricPool | External Capa... | Volu... | Disk C... |
|---|---|---|---|---|---|---|---|---|---|---|
| ✓ | aggr0 | FlexPod-a700s-01 | 95 | 48.1 GB | 944.55 GB | 992.65 GB | No | -NA- | 1 | 22 |
| ✓ | aggr0_FlexPod_a700s_02_0 | FlexPod-a700s-02 | 95 | 48.1 GB | 944.56 GB | 992.65 GB | No | -NA- | 1 | 22 |
| ✓ | aggr1_node1 | FlexPod-a700s-01 | 58 | 9.81 TB | 13.42 TB | 23.23 TB | No | -NA- | 69 | 35 |
| ✓ | aggr1_node2 | FlexPod-a700s-02 | 58 | 9.71 TB | 13.52 TB | 23.23 TB | No | -NA- | 71 | 35 |

For each RAC database deployment, we distributed equal number of volumes and LUNs on both the storage node by placing those volumes and LUNs into both the aggregate.

# Operating System and Database Deployment

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 NetApp AFF Storage A700s 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 4 servers quickly with a standard configuration. SAN boot volumes for these servers were hosted on the NetApp AFF Storage A700s Cluster. Once the stateless servers were provisioned, following process was performed to enable Rapid deployment of 4 RAC nodes.

Each Server node has dedicated single LUN to install operating system and all the four server node was booted off SAN. For this solution, we have installed Oracle Linux 7.5(4.1.12-124.16.4.el7uek.x86_64) on this LUNs and performed all the pre-requisite packages for Oracle Database 12cR2 to create four node Oracle RAC database solution.

## Operating System Configuration

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

1. Download Oracle Linux 7.5 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.

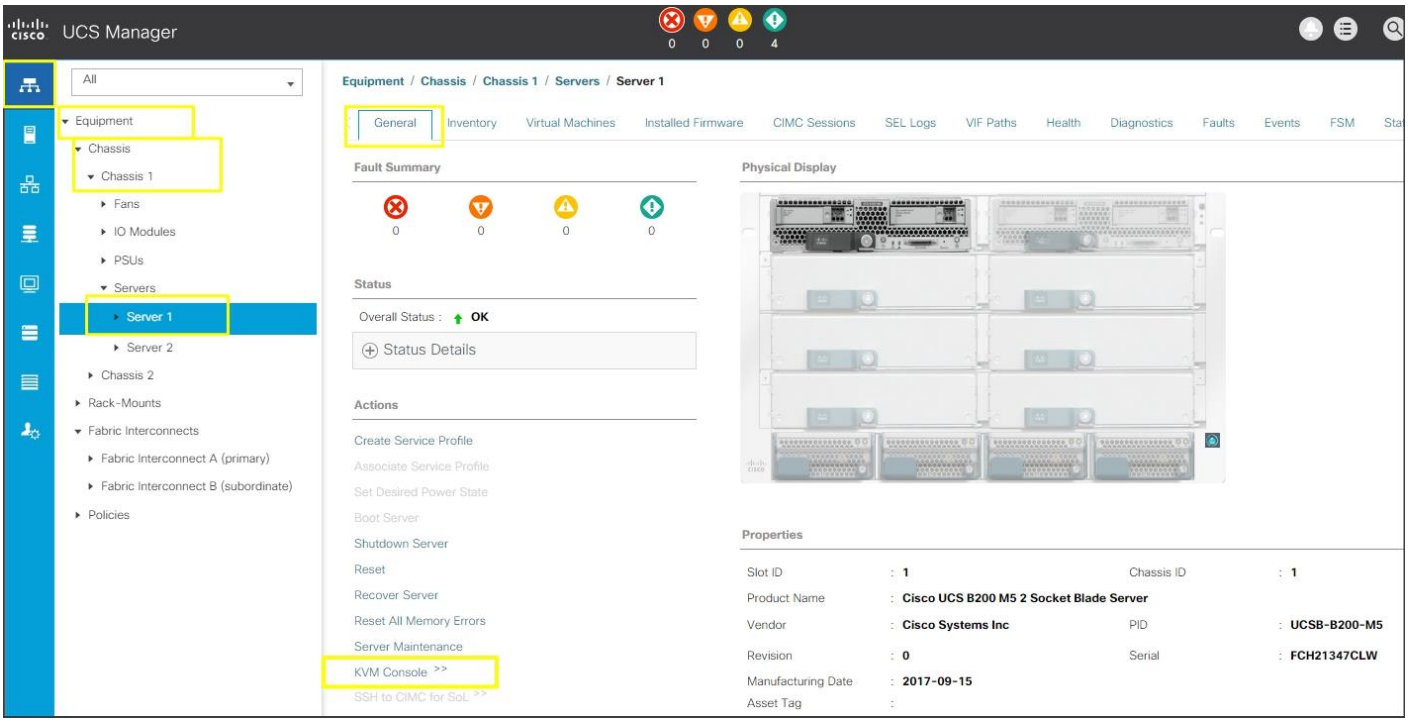


3. Click Accept security and open KVM. Enable virtual media, map the Oracle Linux ISO image and reset the server.

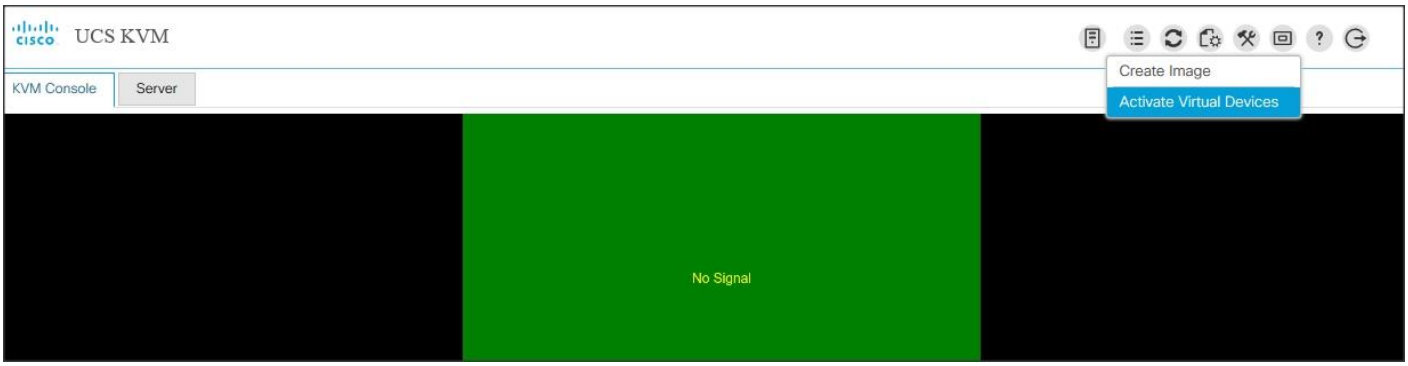


4. When the Server starts booting, it will detect the NetApp Storage active FC paths as shown below. If you see the following message in the KVM console while the server is rebooting along with the target WWPNs, it confirms the setup is done correctly and boot from SAN will be successful.

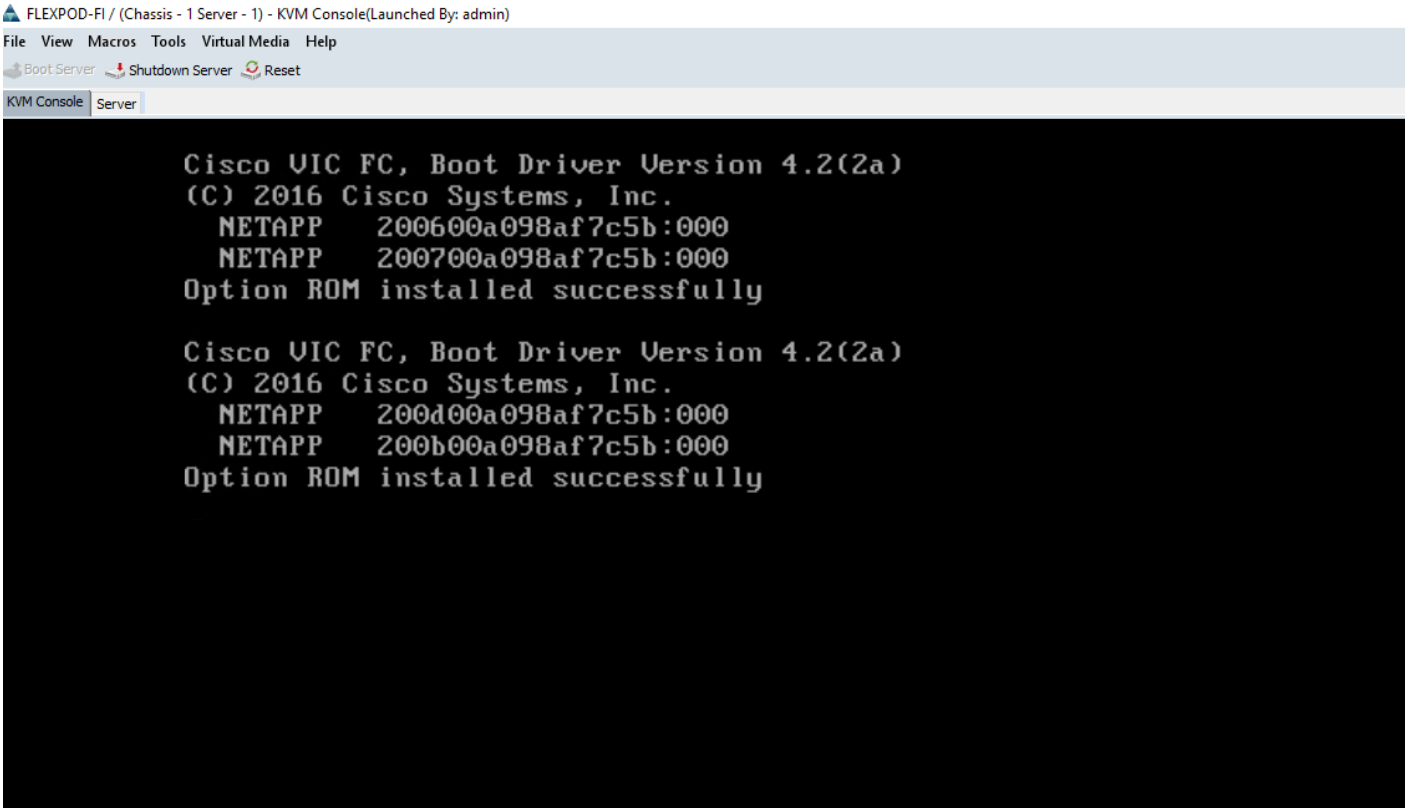


5. During server boot order, it will detect the virtual media connected as Oracle Linux cd. It should launch the Oracle Linux installer. Select language and assign the Installation destination as NetApp Storage LUN. Apply hostname and click “Configure Network” to configure all network interfaces. Alternatively,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

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

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

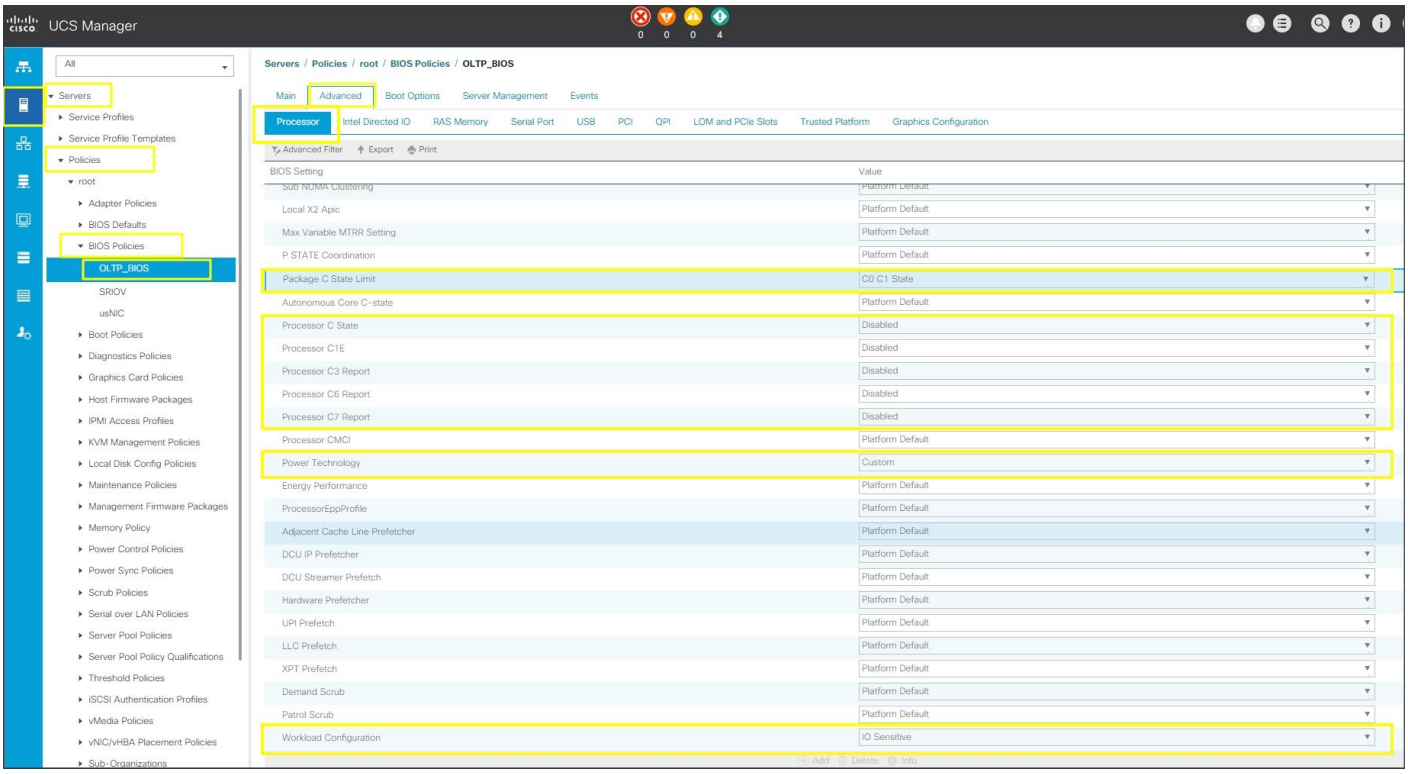


For more information about BIOS settings, refer to:
https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/whitepaper_c11-740098.pdf

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@flex1 ~]# 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 biosdevname=0 net.ifnames=0
intel_idle.max_cstate=0 processor.max_cstate=0 rdloaddriver=scsi_dh_alua"

GRUB_DISABLE_RECOVERY="true"
```

## Prerequisites Automatic Installation

After installing Oracle Linux 7.5(`4.1.12-124.16.4.el7uek.x86_64`) on all the server nodes (flex1, flex2, flex3 and flex4), you have to configure operating system pre-requisites on all the four nodes to successfully install Oracle RAC Database 12cR2.

To configure operating system pre-requisite for Oracle 12cR2 software on all four nodes, follow these 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 four nodes, follow these 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 have 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 log in as root user and issue the following command:

```
[root@flex1 ~]# 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 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 database software installations on all the four nodes as described in the following sections.

### Disable SELinux

Since most Organizations might 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 this command below 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

For DM-Multipath Configuration and best practice, refer to the NetApp Support https://library.netapp.com/ecmdocs/ECMP1217221/html/GUID-34FA2578-0A83-4ED3-B4B3-8401703D65A6.html

**Follow the steps below on all the four oracle RAC nodes.**

You can configure DM-Multipath for use in multipathing in environments that use native Linux solutions. With DM-Multipath, you can configure multiple I/O paths between a host and storage controllers into a single device. If one path fails, DM-Multipath reroutes I/Os to the remaining paths. Configure multipaths to access the LUNs presented from NetApp Storage to the nodes as shown below.

Add or modify `"/etc/multipath.conf"` file accordingly to give the alias name of each LUN id presented from NetApp 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 {
        find_multipaths yes
        user_friendly_names no
}
multipaths {
        multipath {
                wwid            3600a098038304173475d4c766a49744e
                alias           node1_os
        }
        multipath {
                wwid            3600a098038304173475d4c766a49754a
                alias           ocrvote_1
        }
}
```

**Make sure the LUNs wwid address reflects the correct value for all four nodes in "`/etc/multi-path.conf`"**

**We made sure the multipathing packages were installed and enabled for automatic restart across re-boots. We will add more LUNs and associated wwid into "`/etc/multipath.conf`" file later on as we add more LUNs for Databases.**

## Configure UDEV Rules

You need to configure UDEV rules to assign permission in all the Oracle RAC nodes to access NetApp Storage LUNs. This includes the device details along with required permissions to enable grid and oracle user to have read/write privileges on these devices. Configure UDEV rules on all the Oracle Nodes as shown below:

Create a new file named "`/etc/udev/rules.d/99-oracleasm.rules`" with the following entries on all nodes:

```
 [root@flex1 ~]# cat /etc/udev/rules.d/99-oracle-asmdevices.rules
#All LUNs which starts with dg_orarac_* #
ENV{DM_NAME}=="ocrvote_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"
#All LUNs which starts with dg_oradata_* #
ENV{DM_NAME}=="oradata_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"
```

```
#All LUNs which starts with dg_oraredo_* #
ENV{DM_NAME}=="oraredo_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"
#All LUNs which starts with dg_oraarchive_* #
ENV{DM_NAME}=="oraarchive_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"
```

**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**

## 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 (NICs), or network adapters. One adapter is for the public network traffic and the other adapter is for the private network traffic (the node interconnects).

Log in as a root user on 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

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

```
[root@flex1 ~]# cat /etc/hosts
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4


##### Public IP #####
10.29.135.121   flex1.cisco.com flex1
10.29.135.122   flex2.cisco.com flex2
10.29.135.123   flex3.cisco.com flex3
10.29.135.124   flex4.cisco.com flex4
##### Virtual IP #####
10.29.135.125   flex1-vip       flex1-vip.cisco.com
10.29.135.126   flex2-vip       flex2-vip.cisco.com
10.29.135.127   flex3-vip       flex3-vip.cisco.com
10.29.135.128   flex4-vip       flex4-vip.cisco.com
##### Private IP #####
10.10.10.121  flex1-priv        flex1-priv.cisco.com
10.10.10.122  flex2-priv        flex2-priv.cisco.com
```

```
10.10.10.123  flex3-priv        flex3-priv.cisco.com
10.10.10.124  flex4-priv        flex4-priv.cisco.com
#SCAN IP
10.29.135.129   flex-cluster    flex-cluster.cisco.com
10.29.135.130   flex-cluster    flex-cluster.cisco.com
10.29.135.131   flex-cluster    flex-cluster.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 database cluster

These steps complete the prerequisite for Oracle Database 12cR2 Installation at OS level on Oracle RAC Nodes.

**For this Solution, we used 4 identical Cisco UCS B-Series B200 M5 blade servers for hosting the four node Oracle RAC databases. All of the steps described above were also performed on all the four nodes to create 4 node Oracle RAC solution.**

**NetApp Host Utilities Kit: Please refer to this link for the best practice component for all host operating systems connected to NetApp storage via SAN protocols according to your environment: https://mysupport.netapp.com//documentation/productlibrary/index.html?productID=61343**

**Customers should download and install the host utilities kit to provide data collection tools used for NetApp support. However for this solution deployment, we did not used host utility tool kit.**

When the O.S level prerequisites are completed, you are ready to install the Oracle Grid Infrastructure as a 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 Oracle Database FlexPod solution, you will install the Oracle Grid and Database software on all four nodes (flex1, flex2, flex3 and flex4). The installation guides you through the gathering of all node information and configuring the 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 alternative, you can install Oracle validated RPM that will make sure all prerequisites are meet before Oracle grid install.

For the Oracle Database 12c Release 2 install and upgrade guide, click this link: https://docs.oracle.com/en/database/oracle/oracle-database/12.2/install-and-upgrade.html

## Create Directory Structure

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

**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 the other nodes in the cluster.**

1. Log in as a root user and create the following directory structure.

```
mkdir -p /u01/app/grid

mkdir -p /u01/app/12.2.0/grid

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


chown -R grid:oinstall /u01/app/grid/

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

chown -R oracle:oinstall /u01/app/oracle/
```

2. Log in as a grid user, download the Oracle Grid Infrastructure image files and unzip the files into the Grid home:

```
cd /u01/app/12.2.0/grid

unzip -q download_location/linuxx64_12201_grid_home
```

## Run Cluster Verification Utility

This step verifies that all the 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, log in 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 flex1,flex2,flex3,flex4 –verbose
```

## Configure HugePages

HugePages is a method to have a 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

## 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.**

**For this solution, we installed Oracle binaries on the boot LUN of the nodes. The OCR and Voting Disk files resides in the Oracle ASM disk group created from CRS LUNs.**

To install Oracle Database Grid Infrastructure Software, follow these 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.

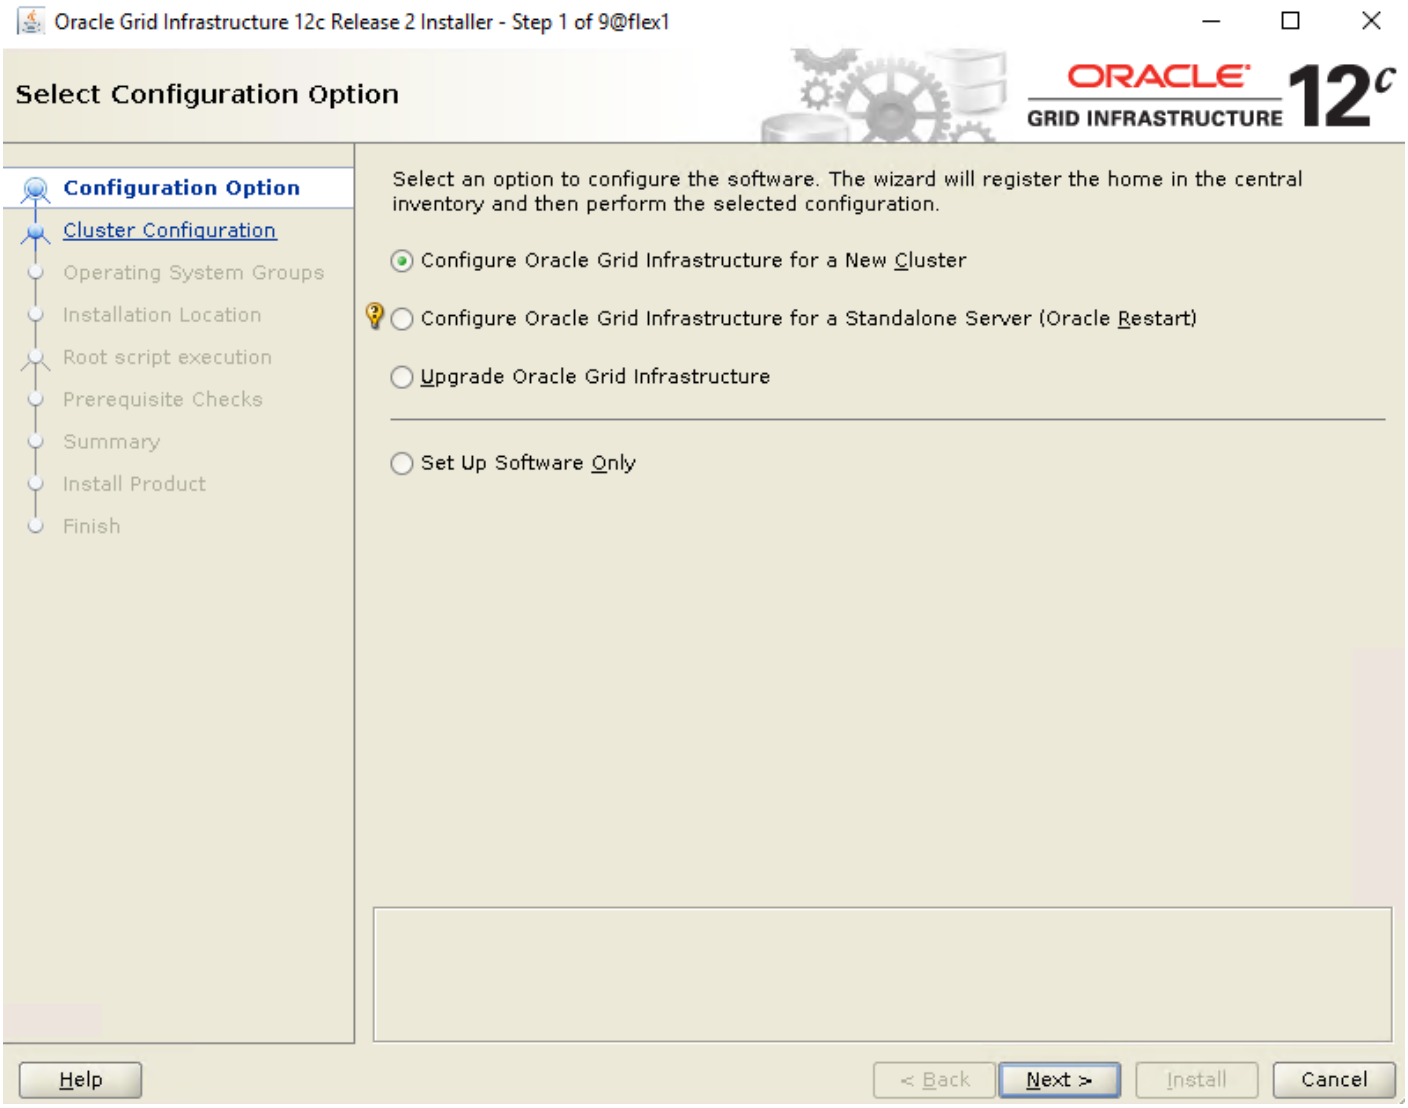


4. Select cluster configuration options “Configure an Oracle Standalone Cluster”, then click Next.

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 Cluster node information window, click the "Add" button to add all four nodes Public Hostname and Virtual Hostname as shown below:

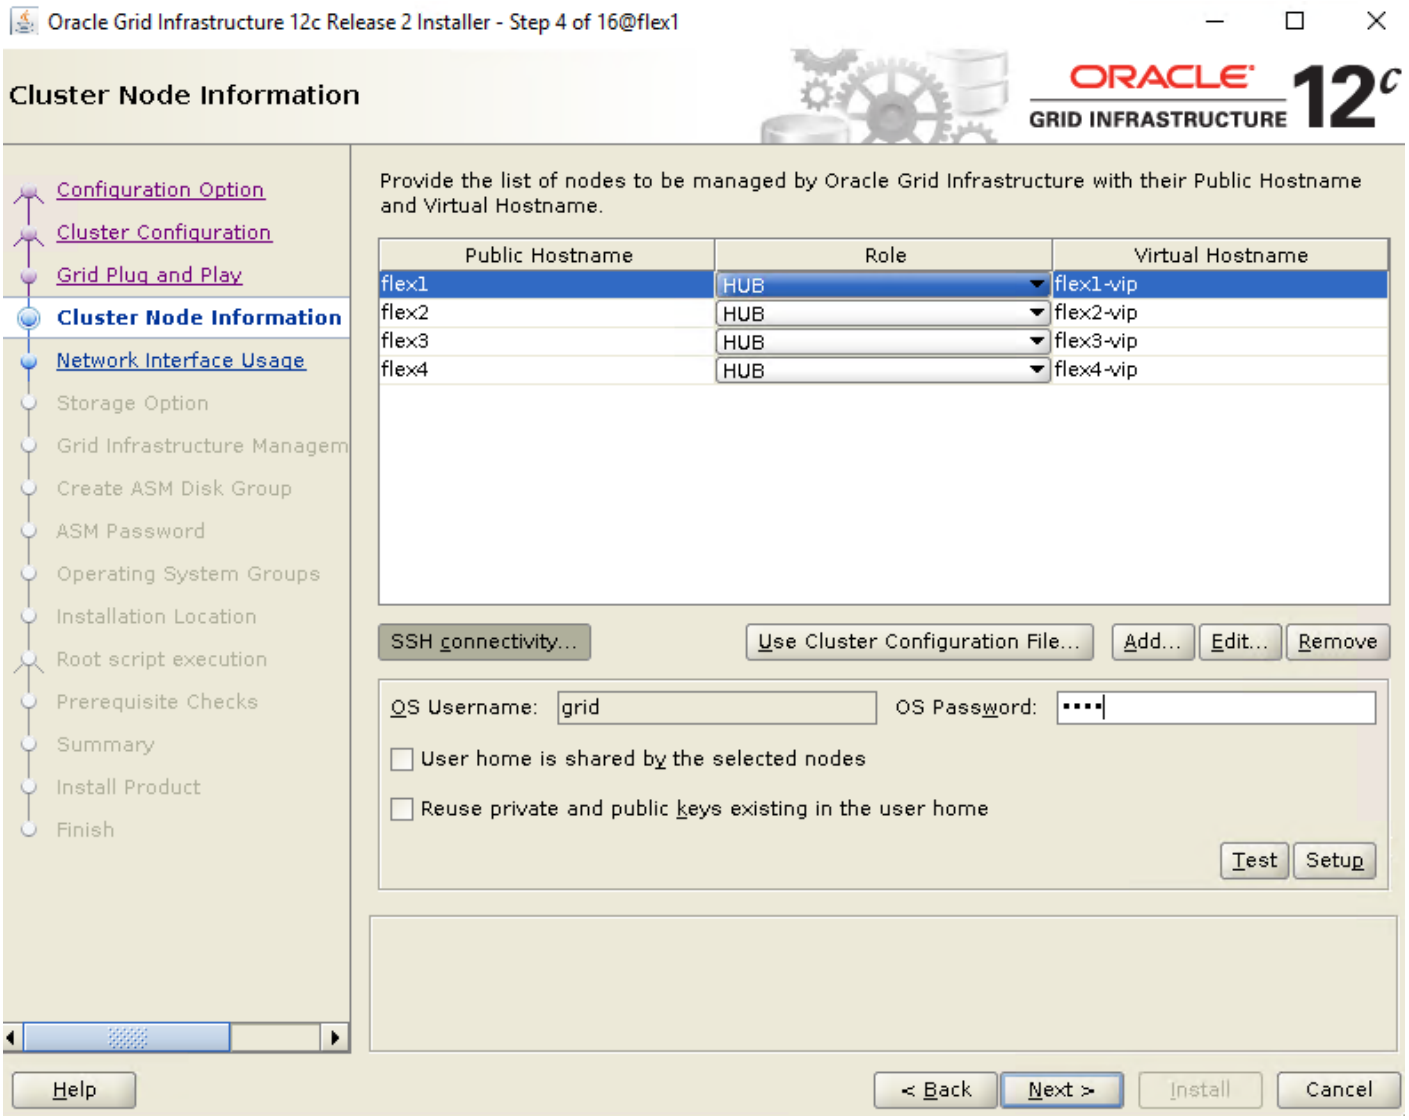


7. As shown above, you will see all nodes listed in the table of cluster nodes. Make sure the Role column is set to HUB for all four nodes. Click the SSH Connectivity button at the bottom of the window. Enter the operating system user name and password for the Oracle software owner (grid). Click Setup.

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.

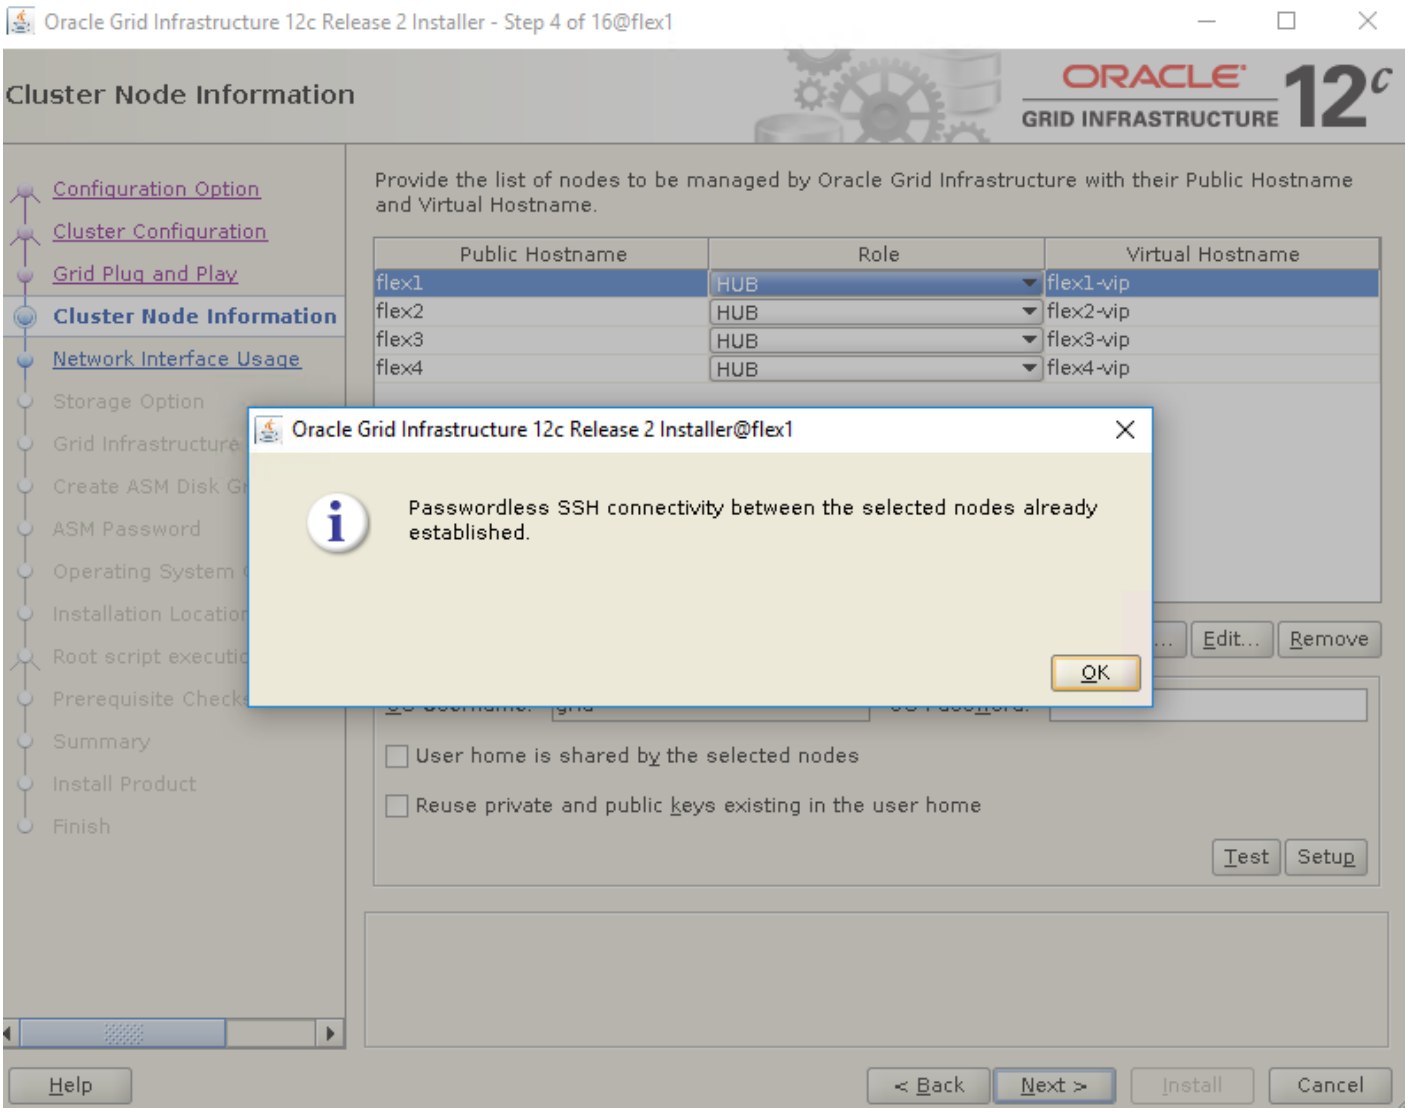


9. In Network Interface Usage screen, select the usage type for each network interface displayed as shown below:

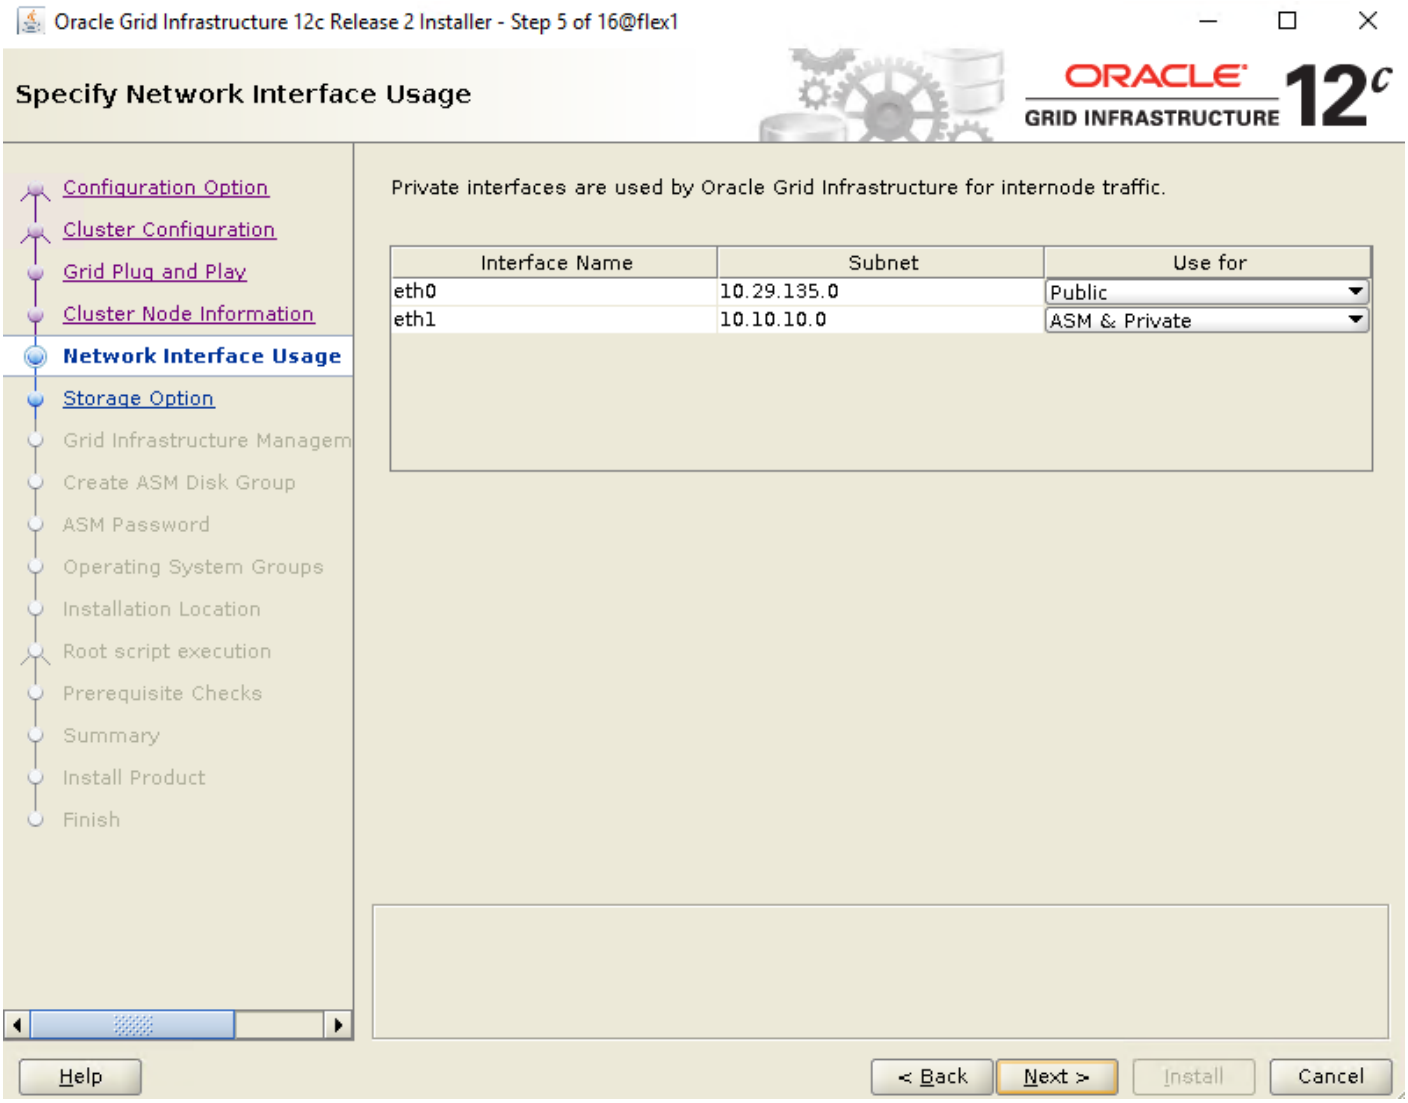


10. Select the Oracle ASM storage configuration option as “Configure ASM using block devices.” Click Next and Choose “No” option into separate ASM disk group for the Grid Infrastructure Management Repository data, and then click Next.

11. In the Create ASM Disk Group window, select the OCRVOTE LUNs assigned from NetApp Storage to store OCR and Voting disk files. Enter the name of disk group as OCRVOTE and select appropriate redundancy options as show below.

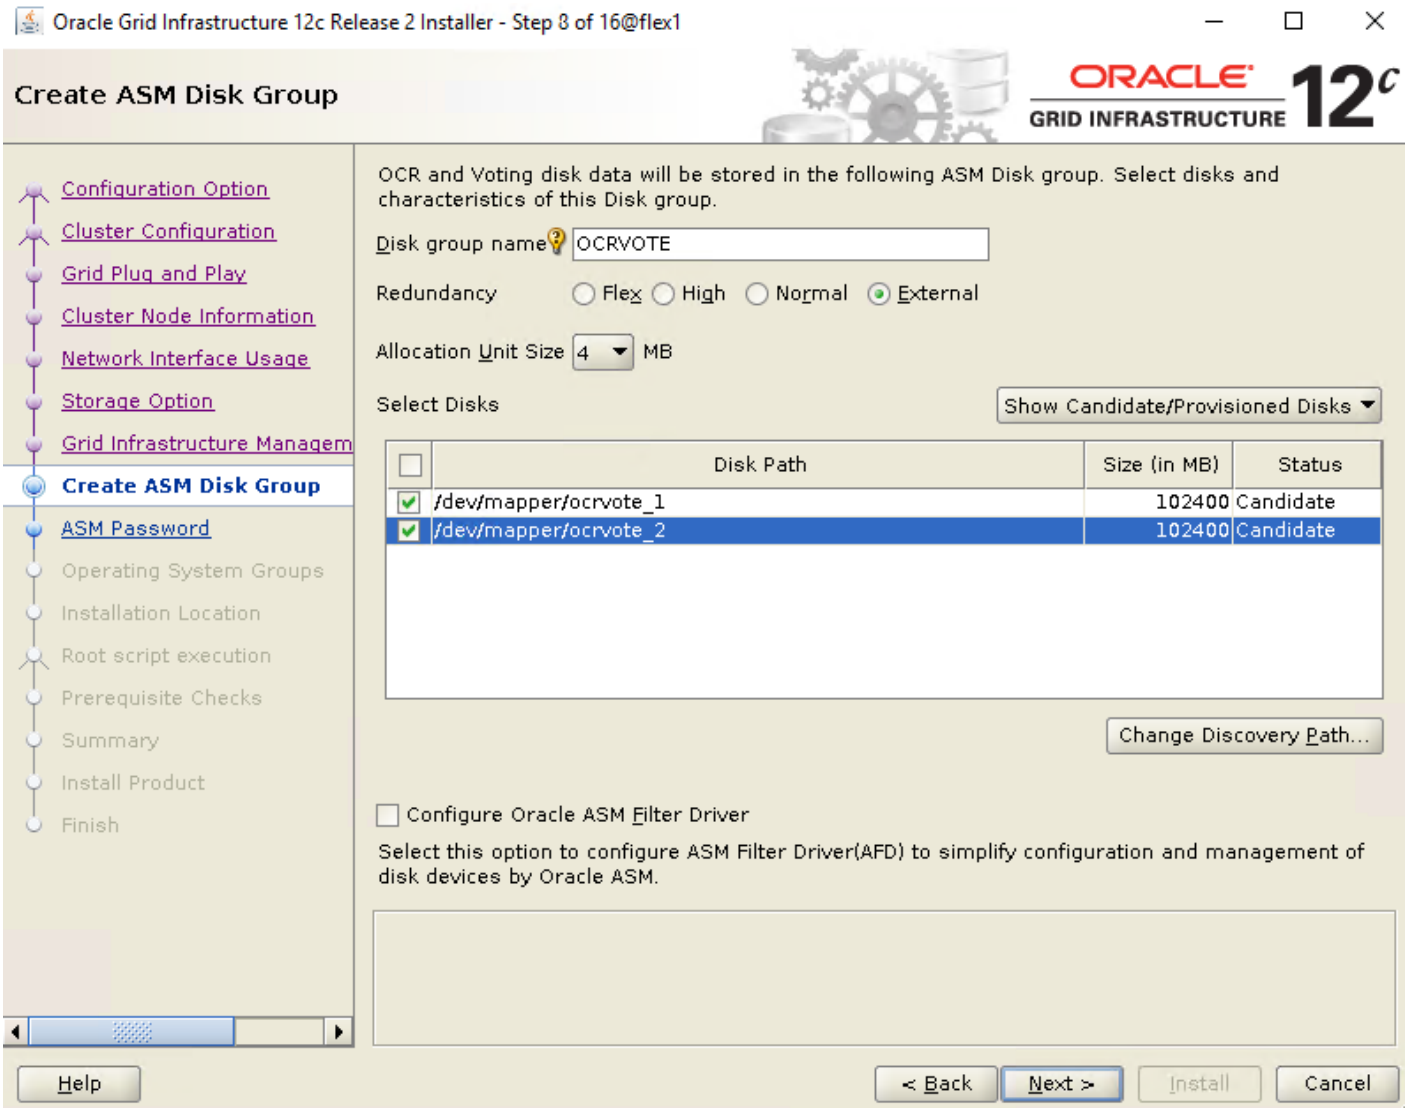


For maximum resiliency, we recommend that the cluster control ASM disk group redundancy be configured with a redundancy level of "High" rather than "External" (https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwlin/oracle-clusterware-storage-space-requirements.html#GUID-97FD5D40-A65B-4575-AD12-06C491AF3F41), preferably with the copies distributed across both controllers. This will increase the resiliency of the cluster in case there is interrupted access to the volume, aggregate, or controller hosting the cluster control ASM disk group (in the case of a volume running out of space, for example). However, for this solution, we configured redundancy as "External."

12. Choose the password for the Oracle ASM SYS and ASMSNMP account, then click Next.

13. Select the option "Do not use Intelligent Platform Management Interface (IPMI)", then click Next. You can configure to have this instance of Oracle Grid Infrastructure and Oracle Automatic Storage Management to be managed by Enterprise Manager Cloud Control. Specify the details of the Cloud Control configuration to perform and 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. Click Next and choose Inventory Directory according to your setup.

If you copied the Oracle Grid Infrastructure installation files into the Oracle Grid home directory as instructed 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.

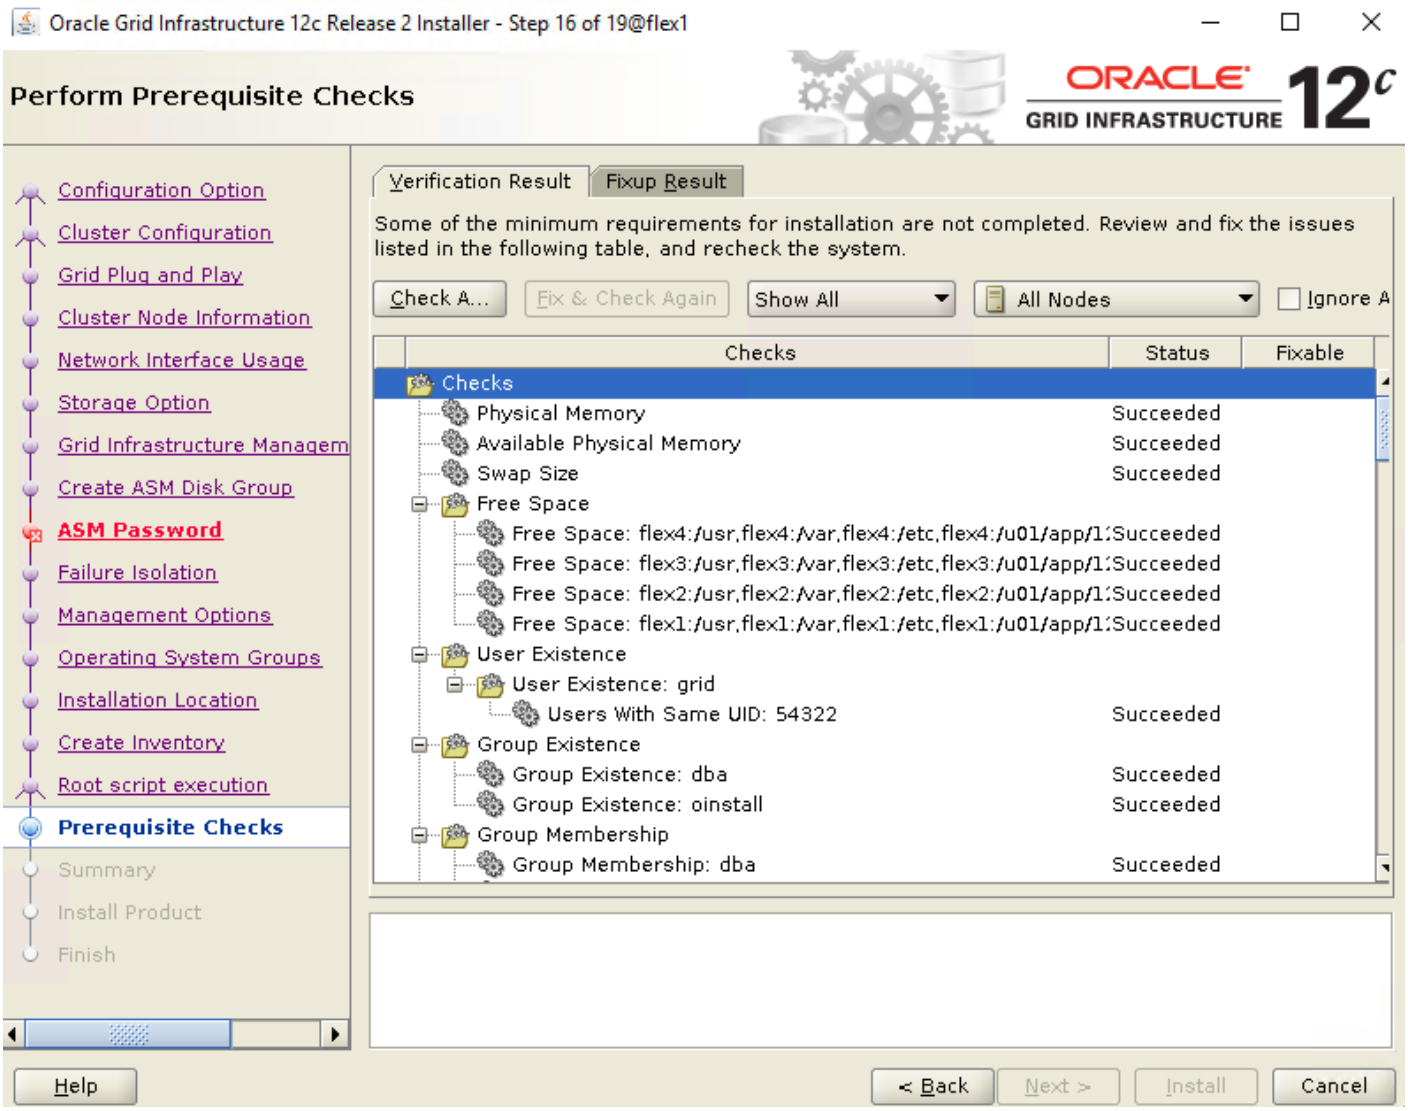


18. Review the contents of the Summary window and then click Install. The installer displays a progress indicator enabling you to monitor the installation process.

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

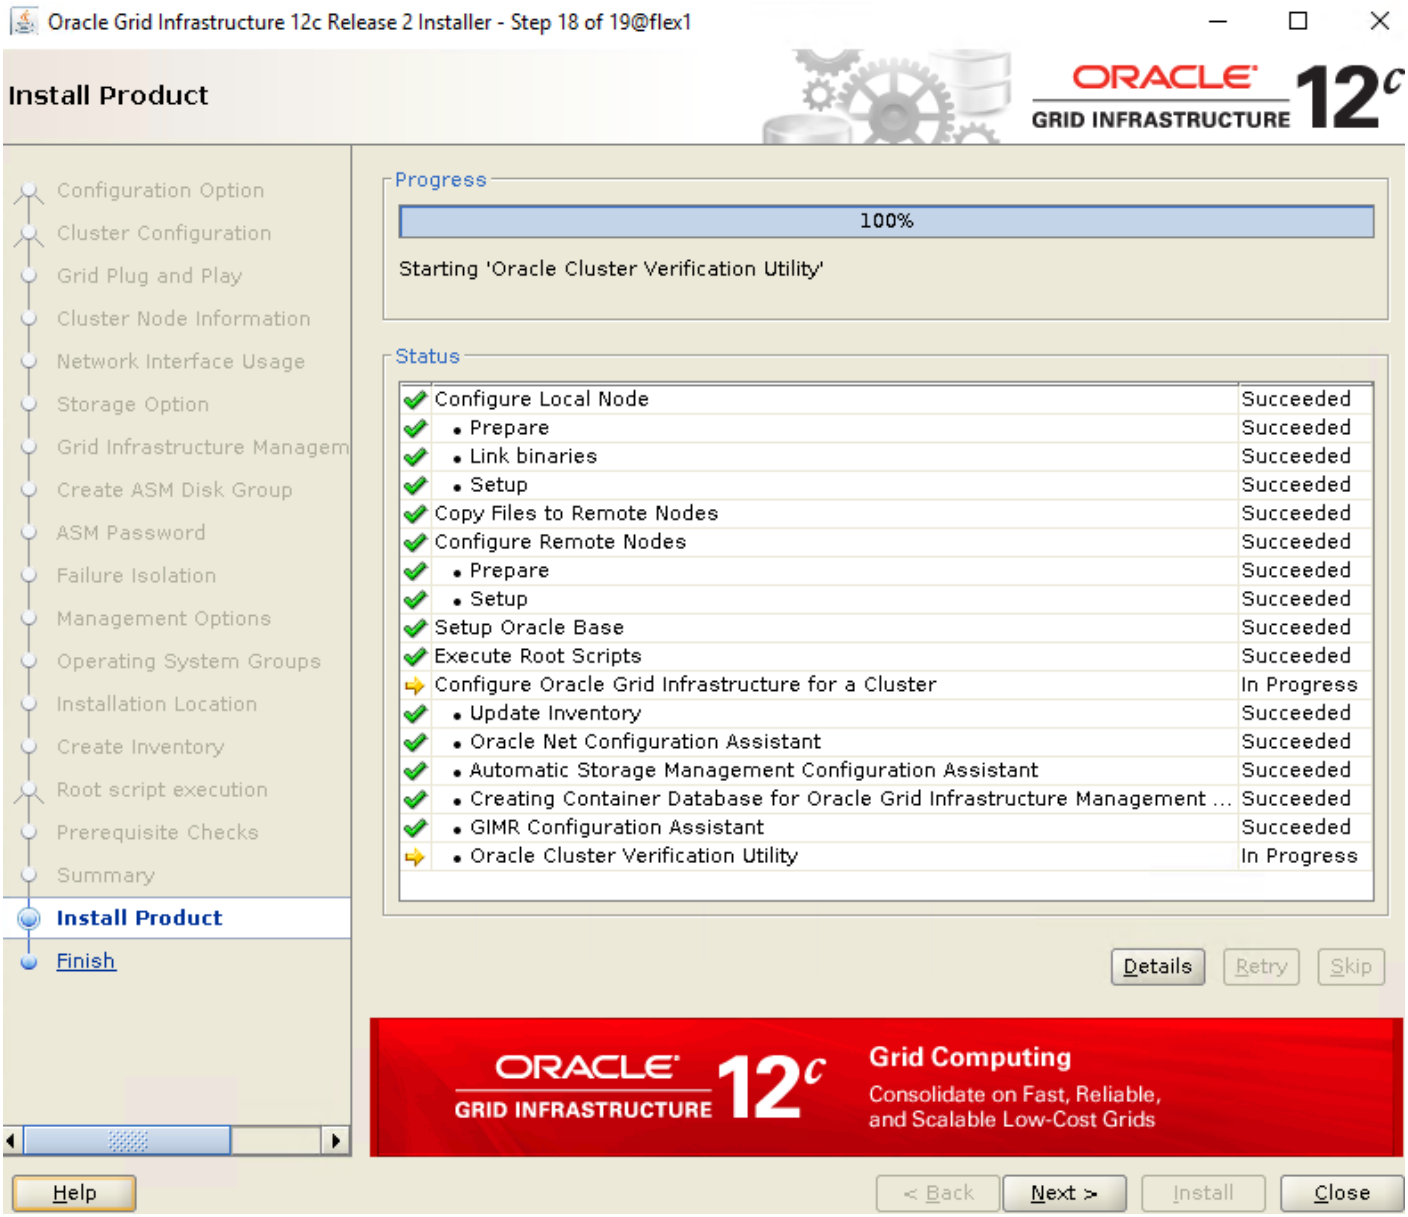


20. When the configuration complete successfully, click Close to finish and exit the grid installer.

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 a 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 provide a 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, follow these 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.

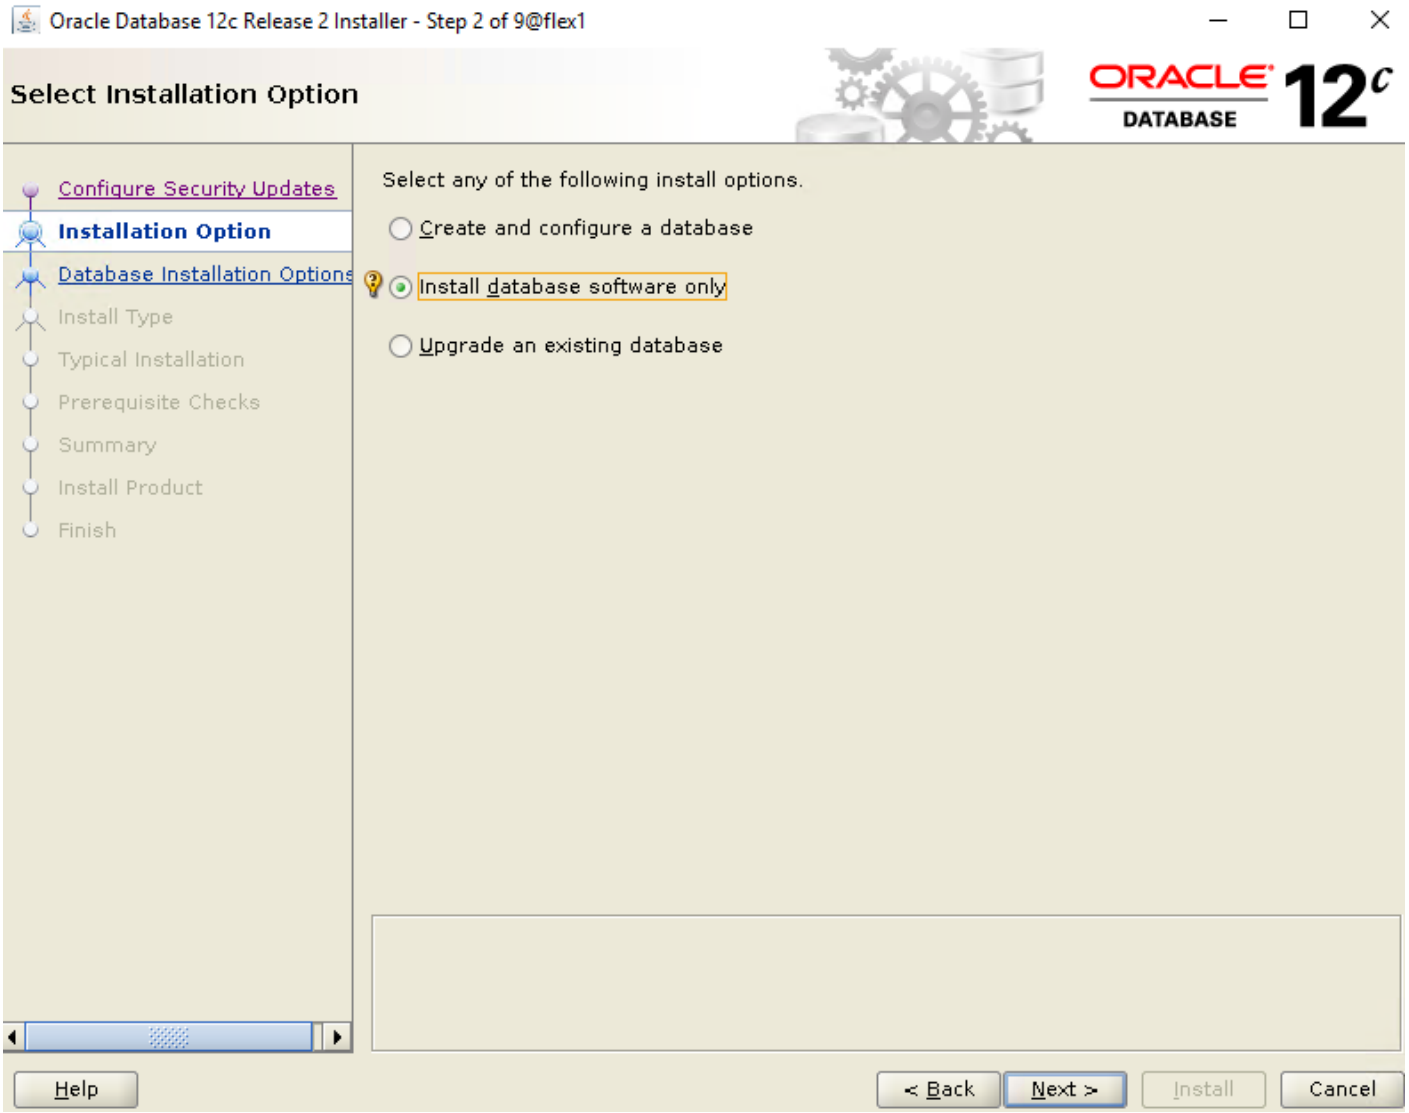


3. Select option "Oracle Real Application Clusters database installation" and click Next.

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.

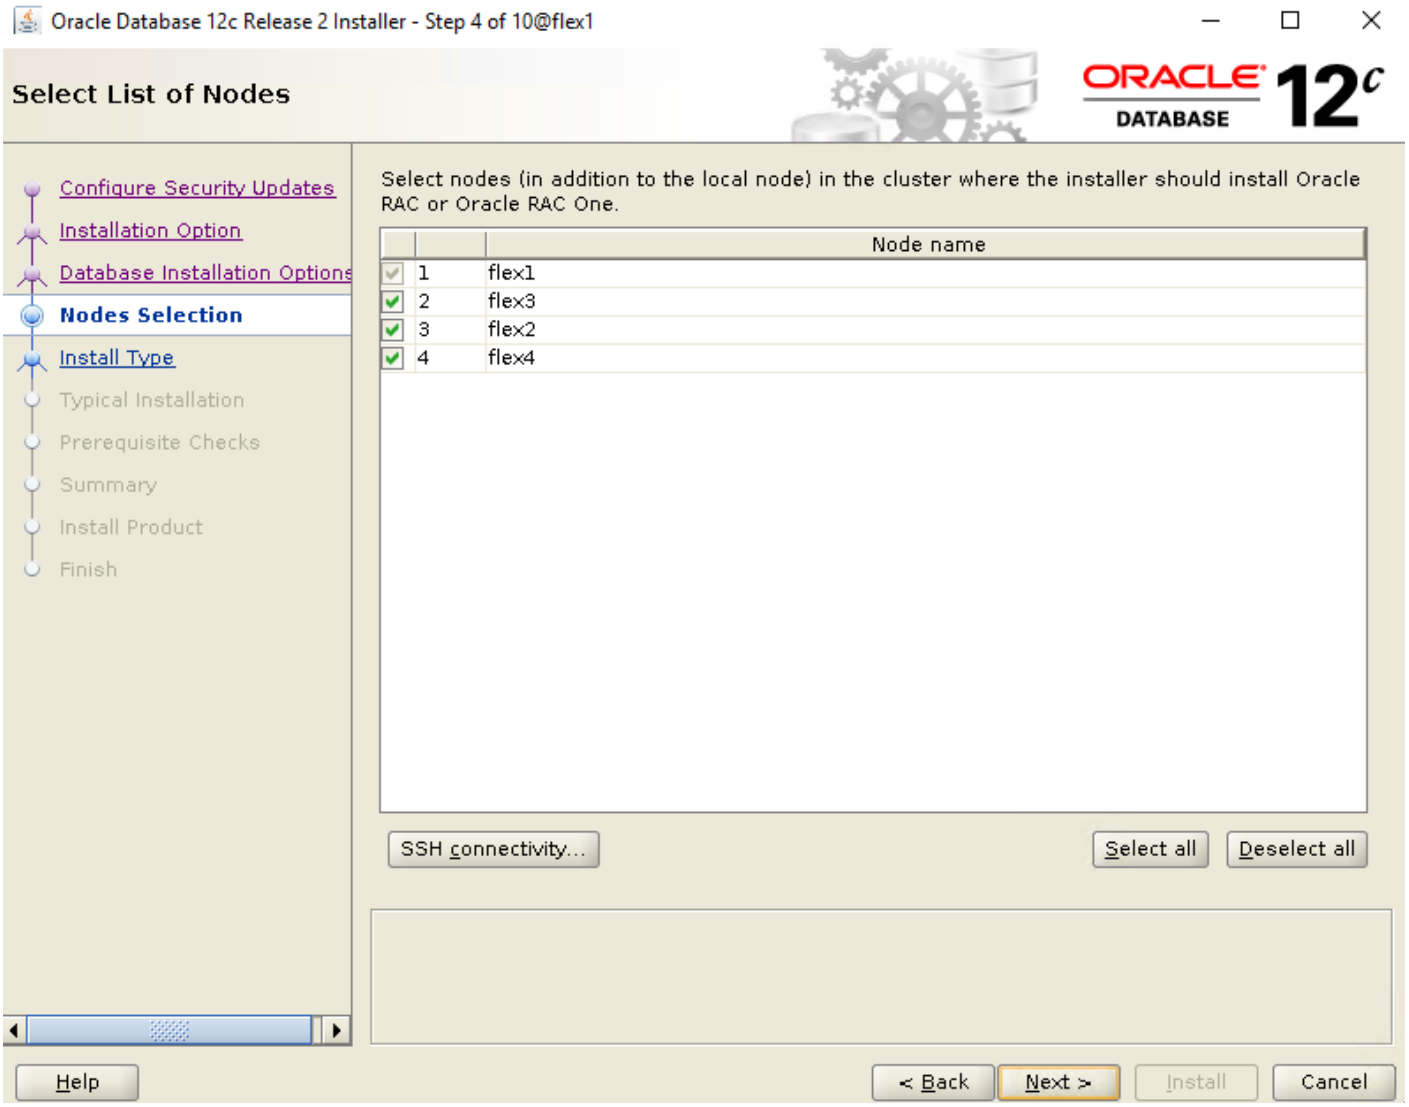


5. Click the "SSH Connectivity..." button and enter the password for the "oracle" user. Click the "Setup" button to configure passwordless SSH connectivity, and the "Test" button to test it when it is complete. When the test is complete, click Next.

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.

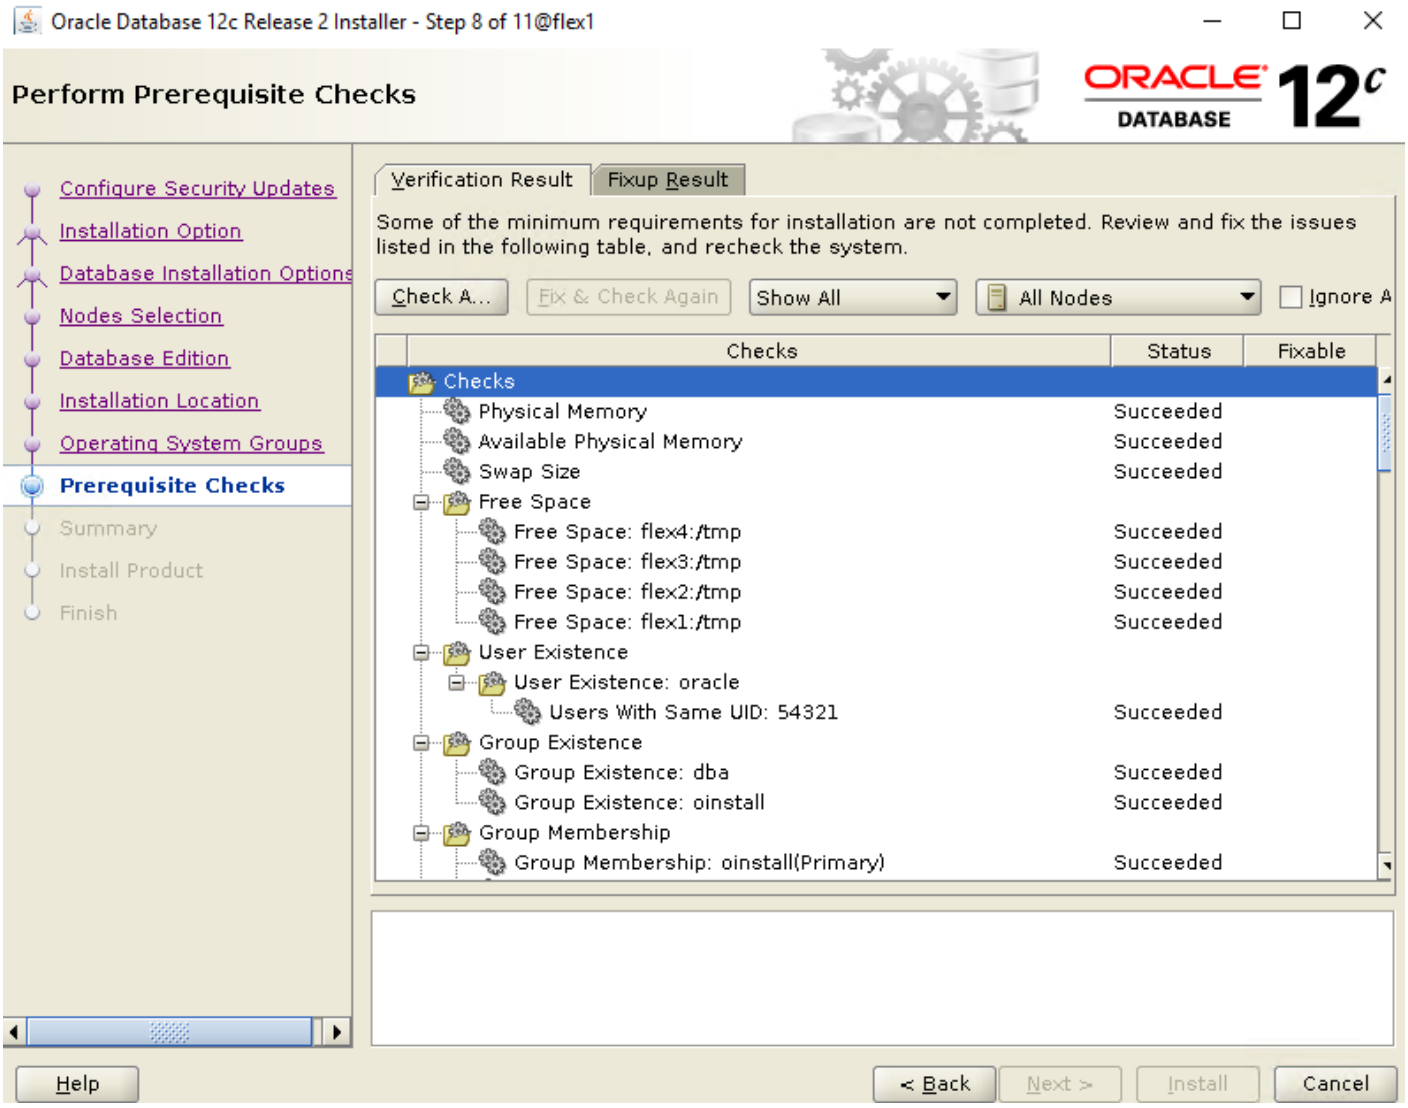


10. Verify the Oracle Database summary information, click Install.

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

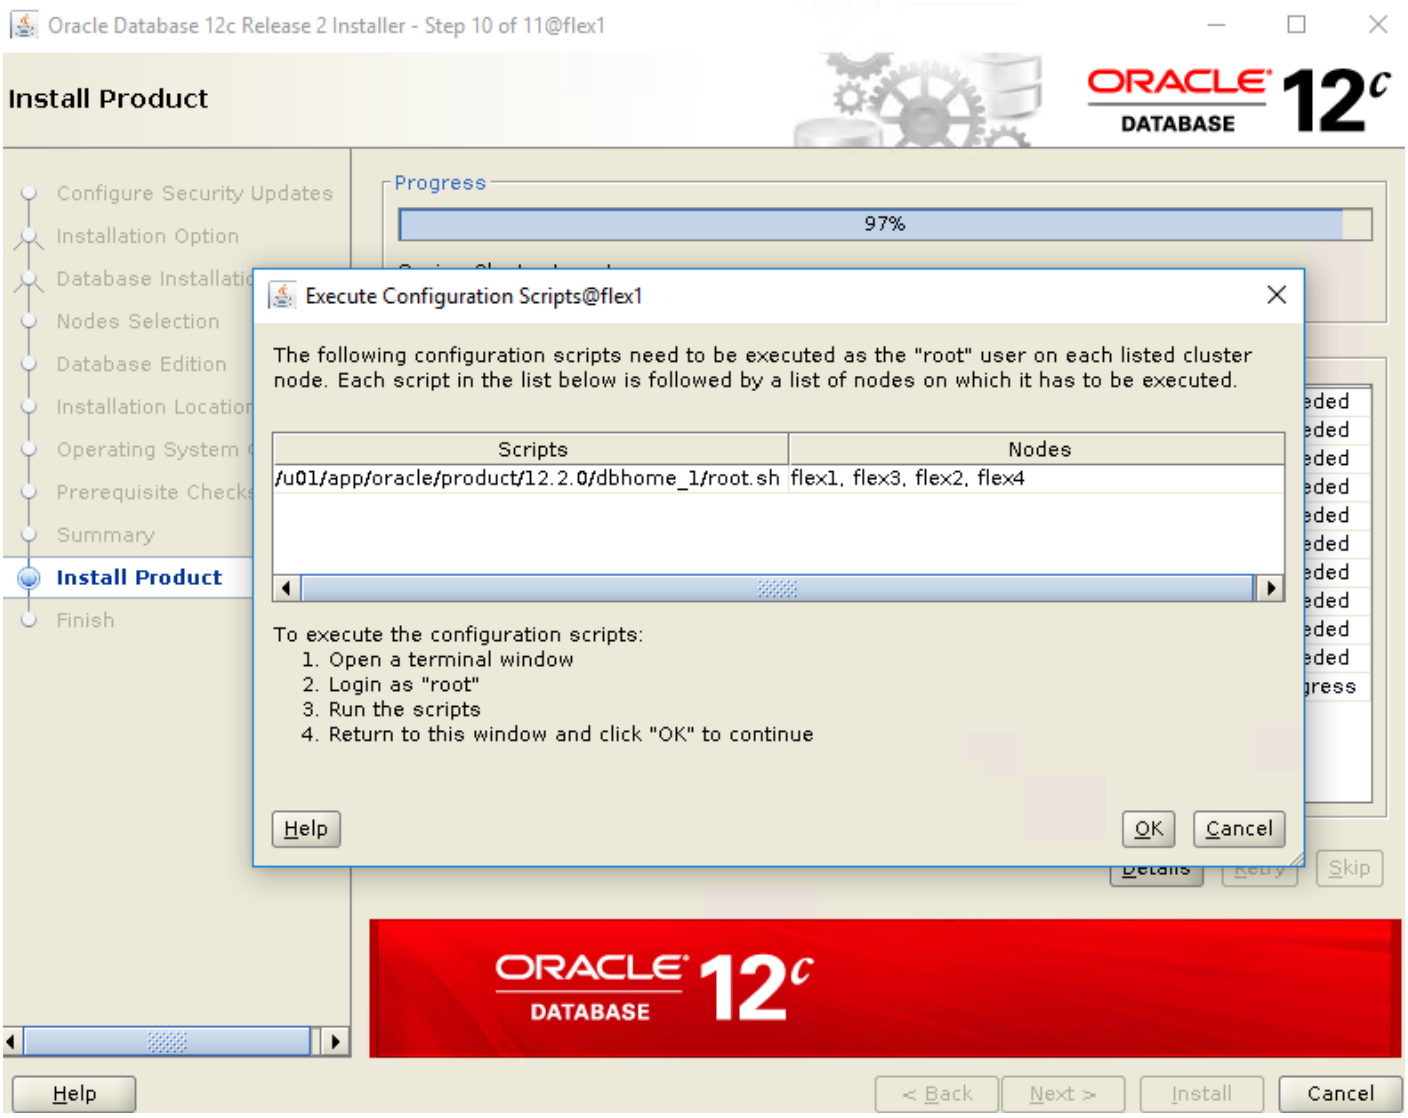


12. After the installation of Oracle Database finish successful, click Close to exit of the installer.

## Oracle Flex ASM

Oracle Flex ASM enables an Oracle ASM instance to run on a separate physical server from the database servers. With this deployment, larger clusters of Oracle ASM instances can support more database clients while reducing the Oracle ASM footprint for the overall system.

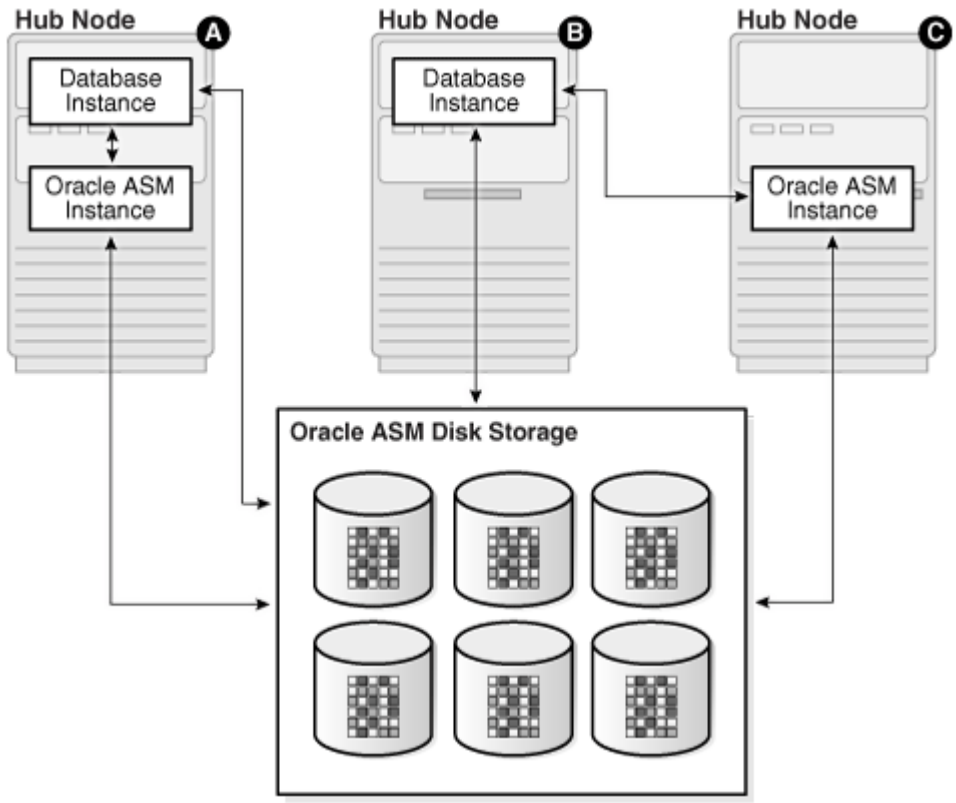


When using Oracle Flex ASM, Oracle ASM clients are configured with direct access to storage. With Oracle Flex ASM, you can consolidate all the storage requirements into a single set of disk groups. All these disk groups are mounted by and managed by a small set of Oracle ASM instances running in a single cluster. You can specify the number of Oracle ASM instances with a cardinality setting. The default is three instances.

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

```
SQL> select INST_ID,GROUP_NUMBER, INSTANCE_NAME, DB_NAME, INSTANCE_NAME||':'||DB_NAME client_id, STATUS from gv$asm_client;

   INST_ID GROUP_NUMBER INSTANCE_NAME                 DB_NAME  CLIENT_ID                      STATUS
---------- ------------ ------------------------------ -------- ------------------------------ ------------
         1           14 +ASM1                          +ASM     +ASM1:+ASM                     CONNECTED
         1           14 -MGMTDB                        _mgmtdb  -MGMTDB:_mgmtdb                CONNECTED
         1           14 flex1.cisco.com                _OCR     flex1.cisco.com:_OCR           CONNECTED
         2           14 +ASM2                          +ASM     +ASM2:+ASM                     CONNECTED
         2           14 flex2.cisco.com                _OCR     flex2.cisco.com:_OCR           CONNECTED
         4           14 +ASM4                          +ASM     +ASM4:+ASM                     CONNECTED
         4           14 flex3.cisco.com                _OCR     flex3.cisco.com:_OCR           CONNECTED
         4           14 flex4.cisco.com                _OCR     flex4.cisco.com:_OCR           CONNECTED

8 rows selected.
```

As you can see from the above query, instance1 (flex1), instance2 (flex2) and instance4 (flex4) are connected to +ASM. Also, the screenshot below shows a few more commands to check the cluster and FLEX ASM details.

```
[grid@flex1 ~]$ ps -ef | grep pmon
grid      35795      1  0 13:48 ?        00:00:00 asm_pmon_+ASM1
grid      40278      1  0 13:49 ?        00:00:00 mdb_pmon_-MGMTDB
grid     174061 117814  0 14:13 pts/1    00:00:00 grep --color=auto pmon
[grid@flex1 ~]$
[grid@flex1 ~]$ crsctl check cluster
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
[grid@flex1 ~]$
[grid@flex1 ~]$ srvctl status asm -detail
ASM is running on flex4,flex2,flex1
ASM is enabled.
ASM instance +ASM1 is running on node flex1
Number of connected clients: 2
Client names: -MGMTDB:_mgmtdb:flex-cluster flex1.cisco.com:_OCR:flex-cluster
ASM instance +ASM2 is running on node flex2
Number of connected clients: 1
Client names: flex2.cisco.com:_OCR:flex-cluster
ASM instance +ASM4 is running on node flex4
Number of connected clients: 2
Client names: flex3.cisco.com:_OCR:flex-cluster flex4.cisco.com:_OCR:flex-cluster
[grid@flex1 ~]$
[grid@flex1 ~]$ srvctl config asm -detail
ASM home: <CRS home>
Password file: +OCRVOTE/orapwASM
Backup of Password file:
ASM listener: LISTENER
ASM is enabled.
ASM is individually enabled on nodes:
ASM is individually disabled on nodes:
ASM instance count: 3
Cluster ASM listener: ASMNET1LSNR_ASM
[grid@flex1 ~]$ asmcmd
ASMCMD> showclustermode
ASM cluster : Flex mode enabled
ASMCMD> showclusterstate
Normal
```

# 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 solution, you will test and validate node and user scalability on all 4 node Oracle RAC Databases with various database benchmarking tools as explained below.

## Hardware Calibration Test using FIO

FIO is short for Flexible IO, a versatile IO workload generator. FIO is a tool that will spawn a number of threads or processes doing a particular type of I/O action as specified by the user. For our solution, we use FIO to measure the performance of a NetApp storage device over a given period of time. For the FIO Tests, we created 16 volumes and each volume has one LUN. These 16 LUNs were shared across all the four nodes for read/write IO operations.

We run various FIO tests for measuring IOPS, Latency and Throughput performance of this solution by changing block size parameter into the FIO test. For each FIO test, we also changed read/write ratio as 0/100 percent read/write, 50/50 percent read/write, 70/30 percent read/write, 90/10 percent read/write and 100/0 percent read/write to scale the performance of the system. We also ran the tests for at least 4 hours to help ensure that this configuration is capable of sustaining this type of load for longer period of time.

### 4k and 8k Random Read/Write IOPs Tests

The following chart shows results for the random read/write FIO test for the 4k block size.

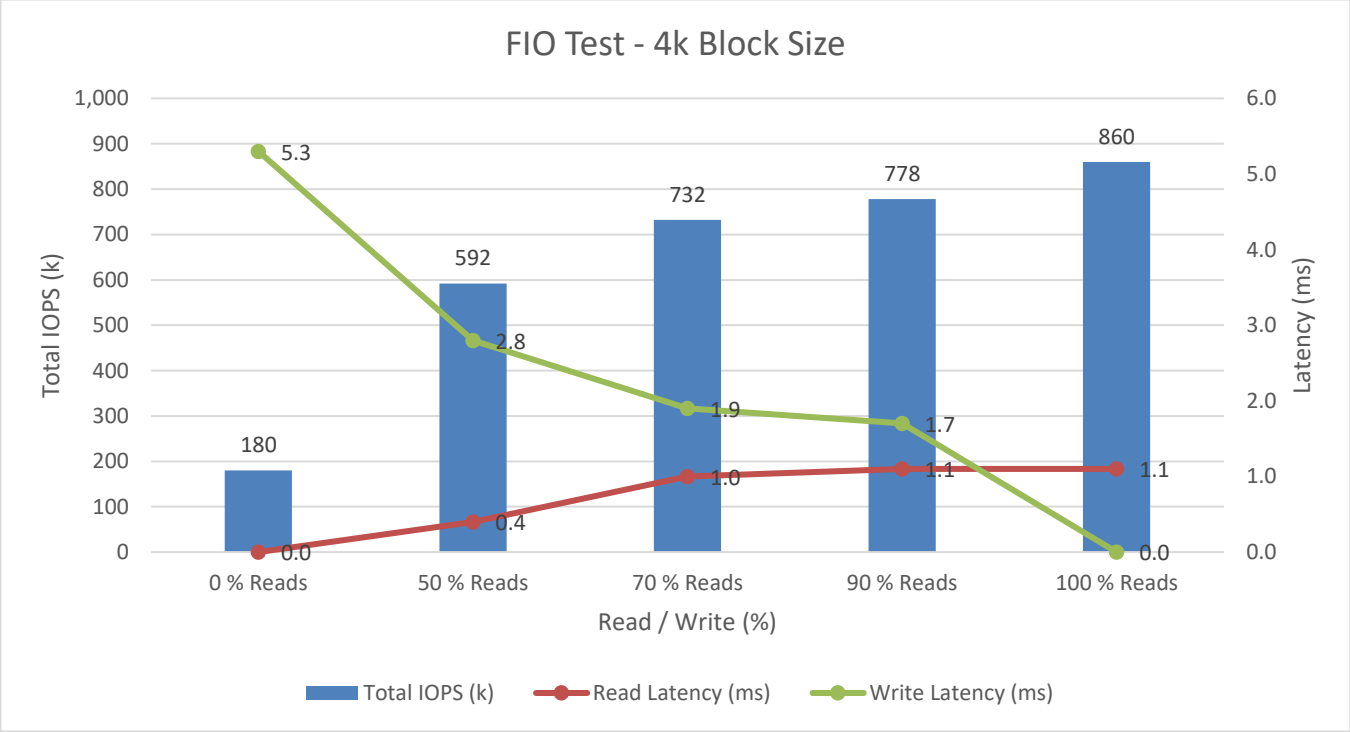


For the 100/0 percent read/write test, we achieved about 860k IOPS with the read latency about 1.1 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 778k IOPS with the read latency about 1.1 millisecond and the write latency about 1.7 millisecond. For the 70/30 percent read/write test, we achieved 732k IOPS with the read latency about 1 millisecond and the write latency about 1.9 millisecond. For the 50/50 percent

read/write test, we achieved about 592k IOPS with the read latency about 0.4 millisecond and the write latency about 2.8 millisecond. For the 0/100 percent read/write test, we achieved 180k IOPS with the write latency about 5.3 millisecond.

The chart below shows results for the random read/write FIO test for the 8k block size representing OLTP type of workloads.

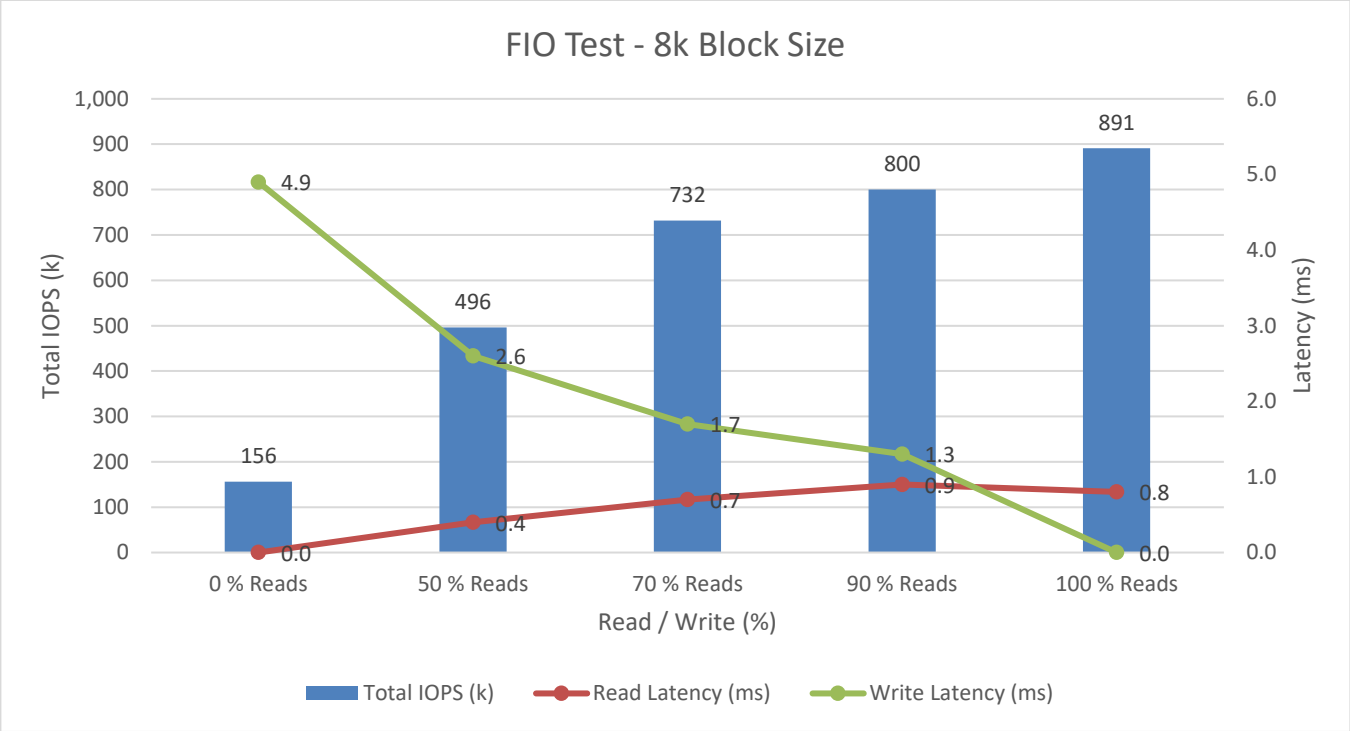


For the 100/0 percent read/write test, we achieved about 891k IOPS with the read latency about 0.8 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 800k IOPS with the read latency about 0.9 millisecond and the write latency about 1.3 millisecond. For the 70/30 percent read/write test, we achieved about 732k IOPS with the read latency about 0.7 millisecond and the write latency about 1.7 millisecond. For the 50/50 percent read/write test, we achieved about 496k IOPS with the read latency about 0.4 millisecond and the write latency about 2.6 millisecond. For the 0/100 percent read/write test, we achieved about 156k IOPS with the write latency about 4.9 millisecond.

## Bandwidth Tests

The bandwidth tests are carried out with 512k and 1MB IO Size and represents the DSS database type workloads. The chart below shows results for the sequential read/write FIO test for the 512k block size.

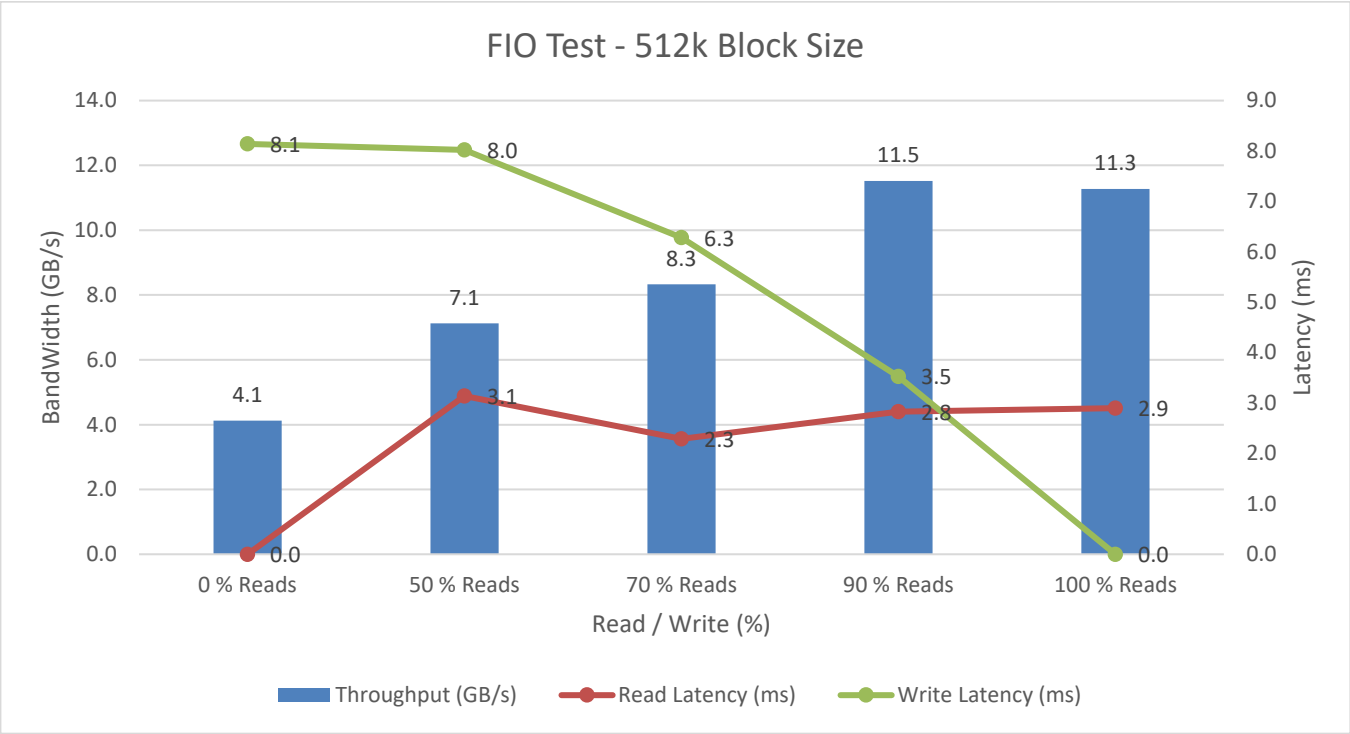


For the 100/0 percent read/write test, we achieved about 11.3 GB/s throughput with the read latency about 2.9 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 11.5 GB/s throughput with the read latency about 2.8 millisecond and the write latency about 3.5 millisecond. For the 70/30 percent read/write test, we achieved about 8.3 GB/s throughput with the read latency about 2.3 millisecond and the write latency about 6.3 millisecond. For the 50/50 percent read/write test, we achieved about 7.1 GB/s throughput with the read latency about 3.1 millisecond and the write latency about 8 millisecond. For the 0/100 percent read/write test, we achieved about 4.1 GB/s throughput with the write latency about 8.1 millisecond.

The chart below shows results for the sequential read/write FIO test for the 1MB block size.

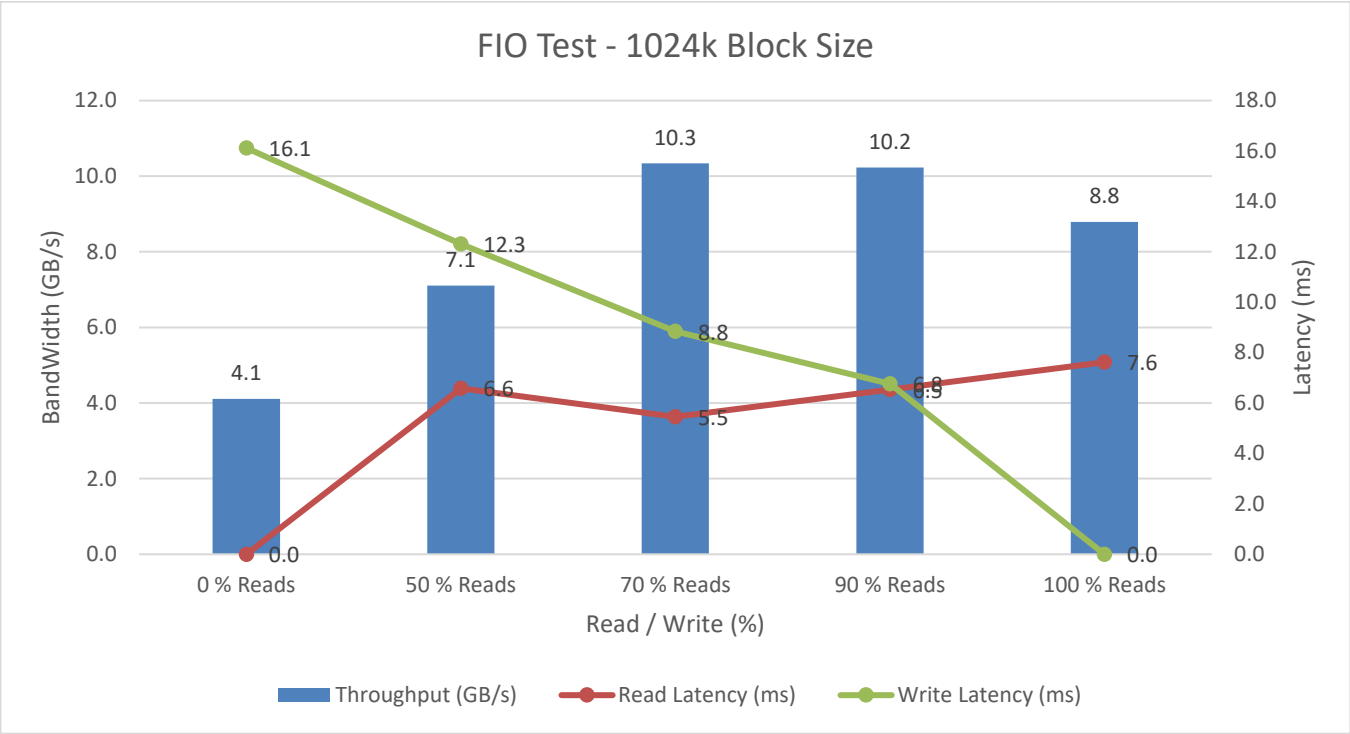


For the 100/0 percent read/write test, we achieved around 8.8 GB/s throughput with the read latency around 7.6 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 10.2 GB/s throughput with the read latency about 6.5 millisecond and the write latency about 6.7 millisecond. For the 70/30 percent read/write test, we achieved about 10.3 GB/s throughput with the read latency about 5.5 millisecond and the write latency about 8.8 millisecond. For the 50/50 percent read/write test, we achieved about 7.1 GB/s throughput with the read latency about 6.6 millisecond and the write latency about 7.1 millisecond. For the 0/100 percent read/write test, we achieved about 4.1 GB/s throughput with the write latency about 16.1 millisecond.

The slight decrease in bandwidth between the 90 percent read workload and 100 percent read workload is a result of ONTAP's extremely efficient write processing. Inbound write data is journaled into mirrored NVRAM. At this point, the write operation is complete and on durable media, and the IO is then acknowledged to the client. NVRAM is not cache, it is a journal that is only used for recovery of an interrupted write operation. The actual write operation to the SSD drives is a direct memory-to-drive transfer. The bandwidth and latency characteristics of NVRAM exceeds even that of SSD drives, so any tests with write IO will result in an increase in overall bandwidth compared to 100percent read tests where all IO must be drawn from slower Flash-based media.

We did not see any performance dips or degradation over the period of run time. It is also important to note that this is not a benchmarking exercise and the numbers presented are not the peak numbers where there is hardware resource saturation. These are practical and out of box test numbers that can be easily reproduced by any one. At this time, we are ready to create OLTP database(s) and continue with database tests.

## Database Creation with DBCA

We used Oracle Database Configuration Assistant (DBCA) to create three OLTP (SLOB, SOE and OLTP) and one DSS (DSS) databases for SLOB and Swingbench test calibration. Alternatively, you can use Database creation scripts to create the databases as well.

For all the database deployment, we have configured two aggregates (one aggregate on each storage node) into a single SVM and each aggregate contains 35 SSD (960GB Each) drives that were subdivided into RAID DP groups, plus one spare drive as explained earlier in the storage configuration section.

For each RAC database, we have created total number of 18 volumes and each volume contains one LUN. We distributed equal number of volumes and LUNs on both the storage node by placing those volumes and LUNs into both the aggregates. All the databases files were also spread evenly across the 2 nodes of the storage system so that each storage node served data for the databases. The screenshot below shows the storage layout of all the volumes and LUN configuration for all the databases.

| | Node: FlexPod-a700s-01 | | | | | Node: FlexPod-a700s-02 | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | Aggregate: aggr1_node1 (RAID Group - raid_dp, 23.23 TB) | | | | | Aggregate: aggr1_node2 (RAID Group - raid_dp, 23.23 TB) | | | | |
| | Volume | Size (GB) | Lun | Size (GB) | Notes | Volume | Size (GB) | Lun | Size (GB) | Notes |
| | | | | | | | | | | |
| ASM | crs_1 | 150 | ocrvote_1 | 100 | OCR & Voting Disks | crs_2 | 150 | ocrvote_2 | 100 | OCR & Voting Disks |
| | | | | | | | | | | |
| SLOB Database | slobvol_1 | 400 | oradata_slob_1 | 300 | Data Files | slobvol_2 | 400 | oradata_slob_2 | 300 | Data Files |
| | slobvol_3 | 400 | oradata_slob_3 | 300 | Data Files | slobvol_4 | 400 | oradata_slob_4 | 300 | Data Files |
| | slobvol_5 | 400 | oradata_slob_5 | 300 | Data Files | slobvol_6 | 400 | oradata_slob_6 | 300 | Data Files |
| | slobvol_7 | 400 | oradata_slob_7 | 300 | Data Files | slobvol_8 | 400 | oradata_slob_8 | 300 | Data Files |
| | slobvol_9 | 400 | oradata_slob_9 | 300 | Data Files | slobvol_10 | 400 | oradata_slob_10 | 300 | Data Files |
| | slobvol_11 | 400 | oradata_slob_11 | 300 | Data Files | slobvol_12 | 400 | oradata_slob_12 | 300 | Data Files |
| | slobvol_13 | 400 | oradata_slob_13 | 300 | Data Files | slobvol_14 | 400 | oradata_slob_14 | 300 | Data Files |
| | slobvol_15 | 400 | oradata_slob_15 | 300 | Data Files | slobvol_16 | 400 | oradata_slob_16 | 300 | Data Files |
| | slobvol_17 | 150 | oraredo_slob_1 | 100 | Control Files + Redo Log Files | slobvol_18 | 150 | oraredo_slob_2 | 100 | Control Files + Redo Log Files |
| | | | | | | | | | | |
| SOE Database | soevol_1 | 1000 | oradata_soe_1 | 625 | Data Files | soevol_2 | 1000 | oradata_soe_2 | 625 | Data Files |
| | soevol_3 | 1000 | oradata_soe_3 | 625 | Data Files | soevol_4 | 1000 | oradata_soe_4 | 625 | Data Files |
| | soevol_5 | 1000 | oradata_soe_5 | 625 | Data Files | soevol_6 | 1000 | oradata_soe_6 | 625 | Data Files |
| | soevol_7 | 1000 | oradata_soe_7 | 625 | Data Files | soevol_8 | 1000 | oradata_soe_8 | 625 | Data Files |
| | soevol_9 | 1000 | oradata_soe_9 | 625 | Data Files | soevol_10 | 1000 | oradata_soe_10 | 625 | Data Files |
| | soevol_11 | 1000 | oradata_soe_11 | 625 | Data Files | soevol_12 | 1000 | oradata_soe_12 | 625 | Data Files |
| | soevol_13 | 1000 | oradata_soe_13 | 625 | Data Files | soevol_14 | 1000 | oradata_soe_14 | 625 | Data Files |
| | soevol_15 | 1000 | oradata_soe_15 | 625 | Data Files | soevol_16 | 1000 | oradata_soe_16 | 625 | Data Files |
| | soevol_17 | 200 | oraredo_soe_1 | 100 | Control Files + Redo Log Files | soevol_18 | 200 | oraredo_soe_2 | 100 | Control Files + Redo Log Files |
| | | | | | | | | | | |
| OLTP Database | oltpvol_1 | 900 | oradata_oltp_1 | 500 | Data Files | oltpvol_2 | 900 | oradata_oltp_2 | 500 | Data Files |
| | oltpvol_3 | 900 | oradata_oltp_3 | 500 | Data Files | oltpvol_4 | 900 | oradata_oltp_4 | 500 | Data Files |
| | oltpvol_5 | 900 | oradata_oltp_5 | 500 | Data Files | oltpvol_6 | 900 | oradata_oltp_6 | 500 | Data Files |
| | oltpvol_7 | 900 | oradata_oltp_7 | 500 | Data Files | oltpvol_8 | 900 | oradata_oltp_8 | 500 | Data Files |
| | oltpvol_9 | 900 | oradata_oltp_9 | 500 | Data Files | oltpvol_10 | 900 | oradata_oltp_10 | 500 | Data Files |
| | oltpvol_11 | 900 | oradata_oltp_11 | 500 | Data Files | oltpvol_12 | 900 | oradata_oltp_12 | 500 | Data Files |
| | oltpvol_13 | 900 | oradata_oltp_13 | 500 | Data Files | oltpvol_14 | 900 | oradata_oltp_14 | 500 | Data Files |
| | oltpvol_15 | 900 | oradata_oltp_15 | 500 | Data Files | oltpvol_16 | 900 | oradata_oltp_16 | 500 | Data Files |
| | oltpvol_17 | 200 | oraredo_oltp_1 | 100 | Control Files + Redo Log Files | oltpvol_18 | 200 | oraredo_oltp_2 | 100 | Control Files + Redo Log Files |
| | | | | | | | | | | |
| DSS Database | dssvol_1 | 1500 | oradata_dss_1 | 1000 | Data Files | dssvol_2 | 1500 | oradata_dss_2 | 1000 | Data Files |
| | dssvol_3 | 1500 | oradata_dss_3 | 1000 | Data Files | dssvol_4 | 1500 | oradata_dss_4 | 1000 | Data Files |
| | dssvol_5 | 1500 | oradata_dss_5 | 1000 | Data Files | dssvol_6 | 1500 | oradata_dss_6 | 1000 | Data Files |
| | dssvol_7 | 1500 | oradata_dss_7 | 1000 | Data Files | dssvol_8 | 1500 | oradata_dss_8 | 1000 | Data Files |
| | dssvol_9 | 1500 | oradata_dss_9 | 1000 | Data Files | dssvol_10 | 1500 | oradata_dss_10 | 1000 | Data Files |
| | dssvol_11 | 1500 | oradata_dss_11 | 1000 | Data Files | dssvol_12 | 1500 | oradata_dss_12 | 1000 | Data Files |
| | dssvol_13 | 1500 | oradata_dss_13 | 1000 | Data Files | dssvol_14 | 1500 | oradata_dss_14 | 1000 | Data Files |
| | dssvol_15 | 1500 | oradata_dss_15 | 1000 | Data Files | dssvol_16 | 1500 | oradata_dss_16 | 1000 | Data Files |
| | dssvol_17 | 200 | oraredo_dss_1 | 100 | Control Files + Redo Log Files | dssvol_18 | 200 | oraredo_dss_2 | 100 | Control Files + Redo Log Files |
| | | | | | | | | | | |

As shown above, for each database, we created total 18 volumes and each volume contains one LUN. On these 18 LUNs, we created two disk groups to store the data and redolog files for the database. We used 16 LUNs to create Oracle ASM "Data" disk group and 2 LUNs to create Oracle ASM "redolog" disk group for each database.

We used a widely adopted SLOB and Swingbench database performance test tools to test and validate throughput, IOPS, and latency for various test scenarios as explained below.

## SLOB Test

The Silly Little Oracle Benchmark (SLOB) is a toolkit for generating and testing I/O through an Oracle database. SLOB is very effective to test 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).

To test the SLOB workload, we created one database as SLOB. For SLOB database, we created total 18 volumes and each volume contains one LUN. On these 18 LUNs, we created two disk groups to store the data and redolog files for the SLOB database. First disk-group "DATASLOB" was created with 16 LUNs (300 GB each) while second disk-group "REDOSLOB" was created with two LUNs (100 GB each) as shown below.

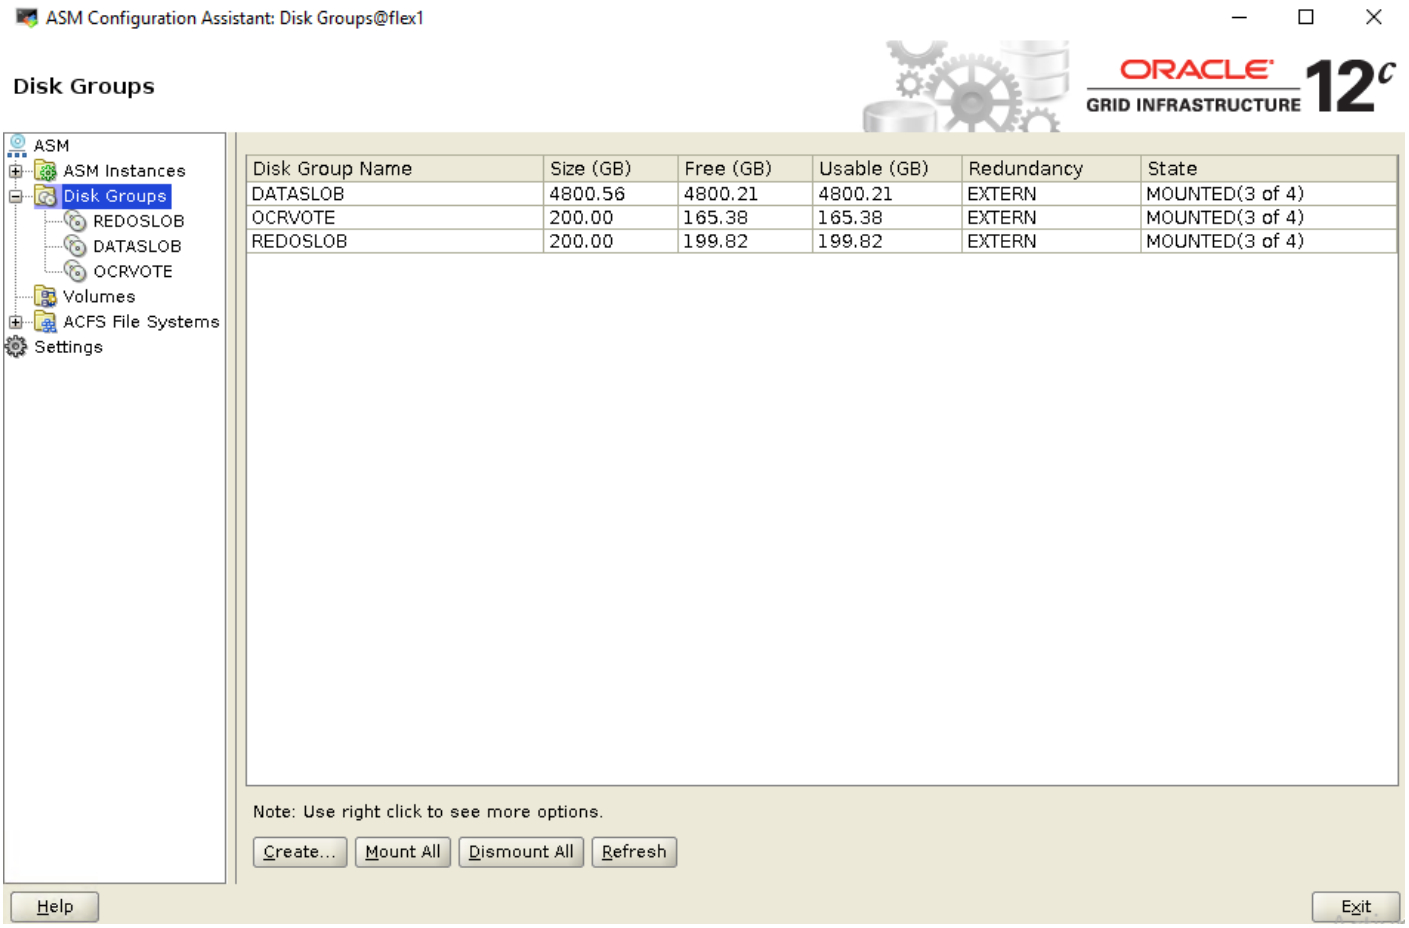


Those ASM disk groups provided the storage required to create the tablespaces for the SLOB Database. We loaded SLOB schema on "DATASLOB" disk-group of up to 3 TB in size.

We used SLOB2 to generate our OLTP workload. Each database server applied the workload to Oracle database, log, and temp files. 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 Test

SLOB2 was configured to run against all the four RAC nodes and the concurrent users were equally spread across all the nodes. We tested the environment by increasing the number of Oracle users in database from a minimum of 32 users up to a maximum of 512 users across all the four nodes. At each load point, we verified that the storage system and the server nodes could maintain steady-state behavior without failure. We also made sure that there were no bottlenecks across servers or networking systems.

We performed User Scalability test with 32, 64, 128, 192, 256 and 512 users on 4 Oracle RAC nodes by varying read/write ratio as explained below:

- Varying workloads
  - 100% read (0% update)
  - 90% read (10% update)
  - 70% read (30% update)
  - 50% read (50% update)

The following table illustrate total number of IOPS (both read and write) for user scalability test when run with 32, 64, 128, 192, 256 and 512 Users on the SLOB database.

Table 10 Total IOPS for SLOB User Scalability Tests

| Users | IOPS for Read/Write % (100/0) | IOPS for Read/Write % (90/10) | IOPS for Read/Write % (70/30) | IOPS for Read/Write % (50/50) |
|---|---|---|---|---|
| 32 | 186,386 | 188,906 | 209,264 | 220,718 |
| 64 | 316,731 | 324,739 | 333,120 | 343,007 |
| 128 | 461,265 | 452,060 | 456,660 | 469,478 |
| 192 | 518,950 | 504,624 | 503,086 | 514,103 |
| 256 | 541,849 | 532,473 | 529,033 | 532,116 |
| 512 | 603,039 | 559,831 | 547,240 | 555,170 |

The following graphs demonstrate total number of IOPS while running SLOB workload for various concurrent users for each test scenario.

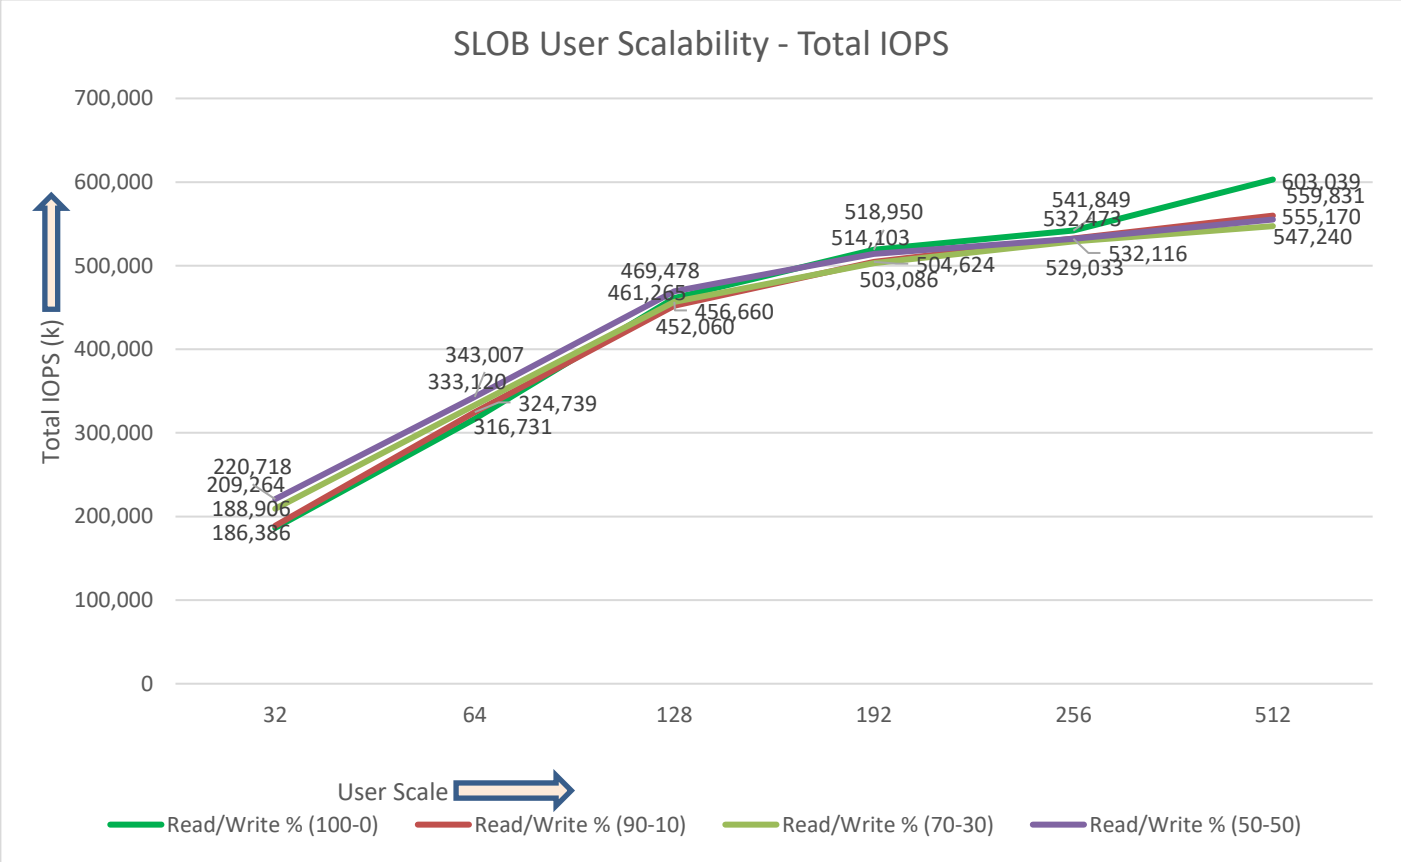


The above graph shows the linear scalability with increased users and similar IOPS from 32 users to 512 users with 100% read, 90% read, 70% read and 50% read.

The AWR screenshot below was captured from a 100% Read (0% update) Test scenario while running SLOB test for 512 users. The snapshot shows a section from the Oracle AWR report from the run that highlights Physical Reads/Sec and Physical Writes/Sec for each instance.

System Statistics - Per Second DB/Inst: SLOB/slob1 Snaps: 310-311

| 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 | 155,795.40 | 148,933.1 | 5.2 | 5.5 | 14.6 | 1.4 | 2,369.7 | 5.0 | 0.20 | 0.1 |
| 2 | 210,927.91 | 149,900.2 | 2.4 | 5.9 | 33.7 | 1.4 | 2,388.8 | 5.2 | 0.19 | 0.2 |
| 3 | 159,063.72 | 151,961.7 | 1.5 | 5.3 | 12.1 | 1.4 | 2,416.7 | 4.4 | 0.20 | 0.1 |
| 4 | 200,586.75 | 152,232.4 | 2.5 | 4.4 | 10.5 | 1.6 | 2,424.2 | 4.9 | 0.20 | 0.1 |
| Sum | 726,373.78 | 603,027.4 | 11.6 | 21.2 | 70.9 | 5.7 | 9,599.4 | 19.5 | 0.78 | 0.6 |
| Avg | 181,593.44 | 150,756.9 | 2.9 | 5.3 | 17.7 | 1.4 | 2,399.9 | 4.9 | 0.20 | 0.2 |
| Std | 28,251.16 | 1,600.9 | 1.6 | 0.7 | 10.8 | 0.1 | 25.2 | 0.4 | 0.00 | 0.0 |

The screenshot above 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 help ensure consistency in behavior and test results.

The screenshot below was captured from a 70% Read (30% update) Test scenario while running SLOB test for 512 users. The snapshot shows a section from AWR report from the run that highlights Physical Reads/Sec and Physical Writes/Sec for each instance.

System Statistics - Per Second DB/Inst: SLOB/slob1 Snaps: 334-335

| 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 | 109,377.19 | 102,785.5 | 31,050.0 | 10,864.7 | 62,872.8 | 1.1 | 1,645.7 | 5.8 | 0.15 | 488.3 |
| 2 | 114,843.30 | 107,949.3 | 32,596.1 | 11,347.6 | 66,022.1 | 1.1 | 1,726.8 | 5.6 | 0.15 | 512.8 |
| 3 | 149,685.24 | 105,692.9 | 31,895.3 | 11,109.2 | 64,587.8 | 1.3 | 1,688.8 | 5.9 | 0.15 | 501.6 |
| 4 | 110,480.18 | 103,911.3 | 31,359.7 | 10,962.4 | 63,504.7 | 1.1 | 1,664.5 | 5.9 | 0.15 | 493.3 |
| Sum | 484,385.91 | 420,339.0 | 126,901.0 | 44,283.8 | 256,987.5 | 4.7 | 6,725.8 | 23.2 | 0.59 | 1,996.0 |
| Avg | 121,096.48 | 105,084.8 | 31,725.3 | 11,071.0 | 64,246.9 | 1.2 | 1,681.5 | 5.8 | 0.15 | 499.0 |
| Std | 19,204.75 | 2,253.8 | 677.5 | 210.0 | 1,379.2 | 0.1 | 35.0 | 0.1 | 0.00 | 10.7 |

The following graph illustrates the latency exhibited by the NetApp AFF A700s Storage across different workloads. All the workloads experienced less than 1 millisecond latency and it varies based on the workloads. As expected, the 50% read (50% update) test exhibited higher latencies as the user counts increases.

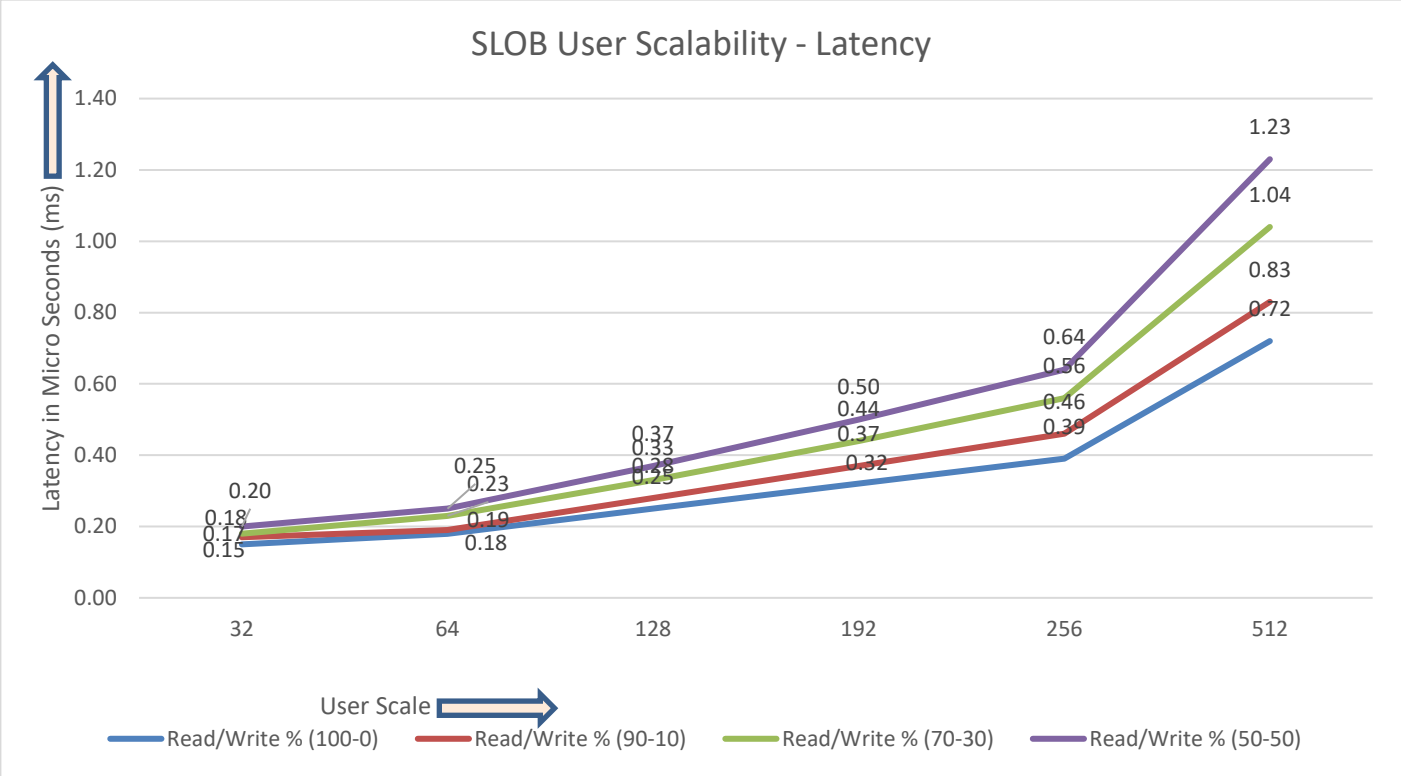


The following screenshot was captured from 100 % Read (0% Update) Test scenario while running SLOB test for 512 users. The screenshot shows a section of AWR report from the run that highlights top timed Events.

```
Top Timed Events                          DB/Inst: SLOB/slob1  Snaps: 310-311
-> 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                    1.080367E+09      0.0     788,506.70 729.85us    85.53     729.87us   727.62us   733.91us     2.78us    4
                 DB CPU                                             N/A      N/A      85,995.01              9.33                                                    4
      Cluster    gc cr grant 2-way                           524,206,578      0.0      66,697.65 127.24us     7.24     127.29us   122.65us   134.05us     5.17us    4
      Cluster    gc cr grant congested                         5,393,910      0.0      21,808.33   4.04ms     2.37       4.04ms     3.93ms     4.12ms    91.61us    4
      Applicatio enq: TX - row lock contention                       5      0.0         920.63 184.13 s     0.10     186.06 s   181.05 s   195.70 s  8358.21ms    4
      Other      gcs drm freeze in enter server mode             2,380      0.0         501.40 210.67ms     0.05     210.61ms   201.53ms   224.49ms    10.56ms    4
      Cluster    gc buffer busy release                            402      0.0         193.59 481.58ms     0.02     388.31ms   170.57us   660.45ms   279.10ms    4
      Other      latch free                                     67,725      0.0         192.59   2.84ms     0.02       2.86ms     1.69ms     4.32ms     1.09ms    4
      Other      RMA: IPC0 completion sync                       7,534      0.0         146.25  19.41ms     0.02      19.41ms    19.37ms    19.44ms    30.68us    4
      Concurrenc library cache load lock                         6,188      0.0         104.40  16.87ms     0.01      15.02ms     5.26ms    41.87ms    17.91ms    4
```

## Swingbench Test

We used Swingbench for OLTP and DSS 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 about 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 4-node Oracle RAC cluster configuration.

- OLTP database user 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 (OLTP and SOE) and one DSS (DSS) database to demonstrate database multi-tenancy capability, performance and sustainability. We created approximately 3 TB of OLTP Database, 4 TB of SOE Database and 6 TB of DSS database to perform Order Entry and Sales Entry swingbench workload testing.

The first step after the databases creation is calibration; about the number of concurrent users, nodes, throughput, IOPS and latency for database optimization. For this FlexPod solution, we have tested system performance with different databases running at a time and capture the results as explained in the following sections.

## One (SOE) Database Performance

For one OLTP database workload featuring Order Entry schema, we used one SOE database. For the SOE database (4 TB), we used 128GB size of System Global Area (SGA). We also made sure 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 SOE database and generate AWR reports for each scenario as shown below.

## User Scalability

Table 11 lists the Transaction Per Minutes (TPM), IOPS and System Utilization for SOE Database while running swingbench workloads from 100 users to 800 users.

**Table 11 TPM, TPS, IOPS and System Utilization for SOE Database User Scale Tests**

| Users | Transactions | | Storage IOPS (AWR) | | | System Utilization (%) |
|---|---|---|---|---|---|---|
| | Per Second (TPS) | Per Minute (TPM) | Reads/s | Writes/s | Total Phy. IOPS | |
| 100 | 8,280 | 496,776 | 39,645 | 21,404 | 61,049 | 9.3 |
| 200 | 15,225 | 913,470 | 74,088 | 36,921 | 111,009 | 17.8 |
| 200 | 20,273 | 1,216,356 | 98,384 | 50,057 | 148,442 | 22.8 |
| 400 | 24,777 | 1,486,590 | 118,738 | 60,882 | 179,620 | 27.4 |
| 500 | 28,084 | 1,685,010 | 134,636 | 71,005 | 205,641 | 32.5 |
| 600 | 30,590 | 1,835,376 | 146,532 | 80,368 | 226,900 | 37.1 |
| 700 | 33,105 | 1,986,270 | 154,331 | 82,346 | 236,677 | 43.1 |
| 800 | 36,467 | 2,188,002 | 162,648 | 83,012 | 245,660 | 47.4 |

The chart below shows the IOPS and Latency of SOE database while running swingbench workload for 100 users to 800 users on all 4 RAC nodes.

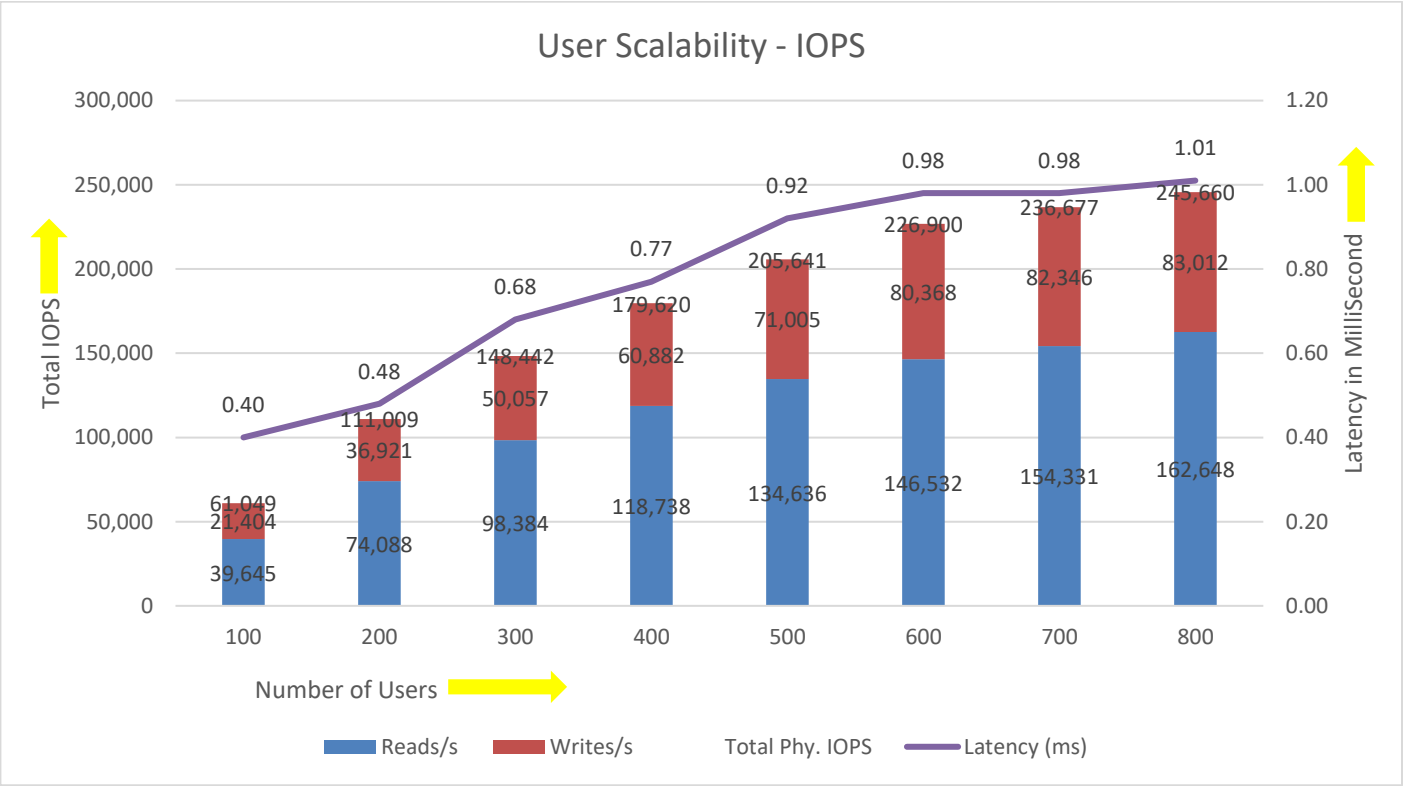


The chart below shows TPM and System Utilization of the system while running swingbench workload for 100 user to 800 users.

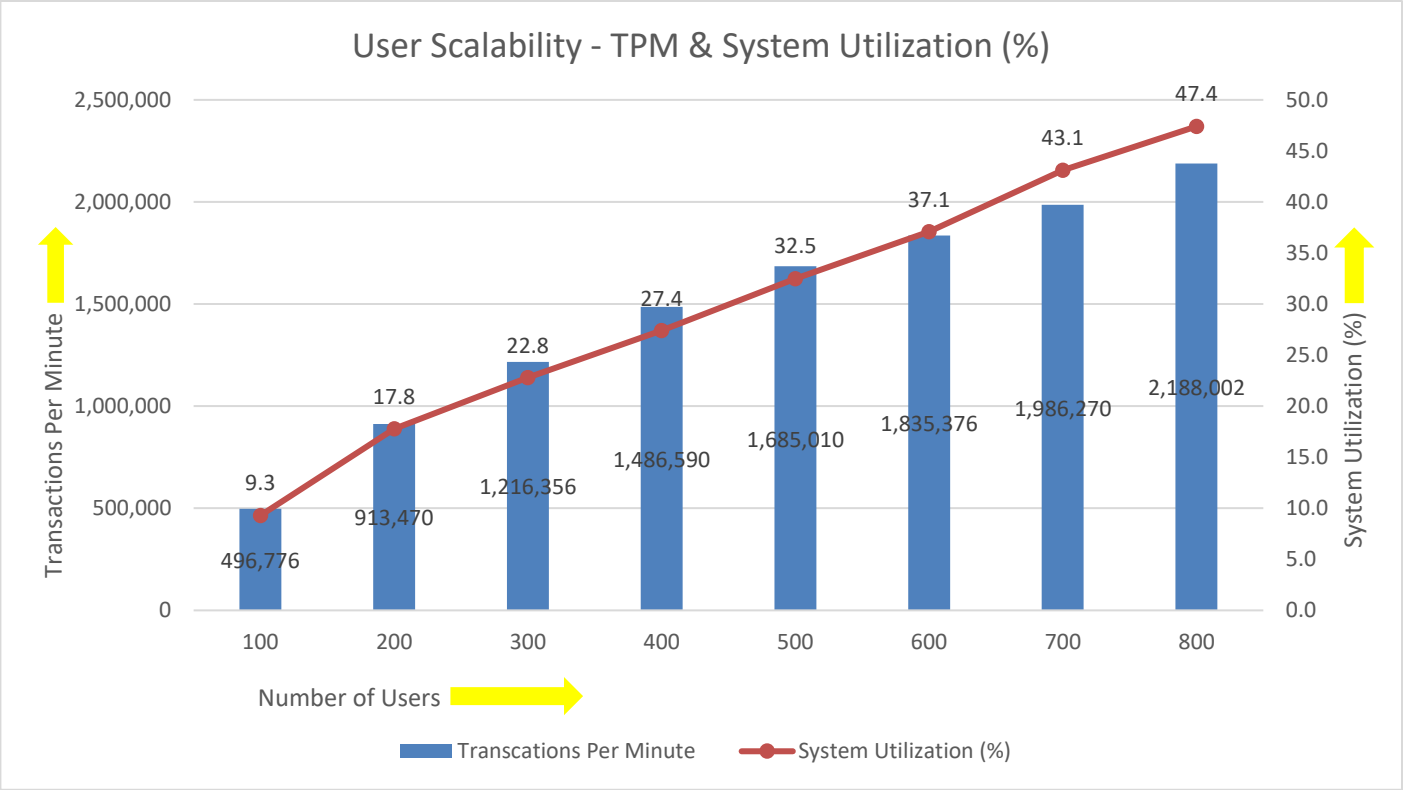

The screenshot below was captured from the 800 User Scale Test scenario while running Swingbench workload on SOE database. The snapshot shows a section from 4-hour window of AWR Global report from the run that highlights Physical Reads/Sec and Physical Writes/Sec for each instance. Notice that IO load is distributed across all the cluster nodes performing workload operations.

```
System Statistics - Per Second                DB/Inst: SOE/soe1  Snaps: 74-78

            Logical     Physical     Physical        Redo        Block        User
  I#        Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s     Calls/s      Execs/s     Parses/s   Logons/s       Txns/s
---- --------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- ------------
   1      904,296.03     38,610.9     19,668.8     29,381.9    185,892.7     25,933.3     89,726.4     10,306.8       0.45      8,643.7
   2      938,176.47     40,253.8     20,390.3     30,495.9    193,291.3     27,006.9     93,437.7     10,733.3       0.45      9,001.6
   3    1,136,462.10     39,304.7     19,951.9     29,485.2    187,806.4     26,314.6     91,051.3     10,457.7       0.44      8,770.8
   4    1,058,048.40     44,478.6     23,001.4     33,857.1    215,500.1     30,154.1    104,323.5     11,984.7       0.45     10,050.7
~~~ ~~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~
 Sum    4,036,983.00    162,648.0     83,012.4    123,220.1    782,490.4    109,409.0    378,538.8     43,482.5       1.80     36,466.7
 Avg    1,009,245.75     40,662.0     20,753.1     30,805.0    195,622.6     27,352.2     94,634.7     10,870.6       0.45      9,116.7
 Std      107,441.43      2,632.0      1,528.0      2,095.9     13,617.6      1,920.1      6,639.2        763.4       0.00        640.0
                    ----------------------------------------------------------------------------------------------------------------------
```

The AWR screenshot below shows top timed events for the same 800 User Scale Test while Swingbench test was running.

```
Top Timed Events                              DB/Inst: SOE/soe1  Snaps: 74-78
-> 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                  1.679536E+09      0.0    1,629,626.08     .97ms    29.13       .97ms      .95ms      .98ms    16.51us    4
                 DB CPU                                            N/A      N/A    1,306,367.35             23.35                                                   4
      Commit     log file sync                              391,693,265      0.0      840,605.91   2.15ms    15.03      2.15ms     1.96ms     2.44ms   215.23us    4
      Cluster    gc current block congested                  40,536,239      0.0      247,195.21   6.10ms     4.42      6.01ms     5.62ms     6.21ms   264.28us    4
      Cluster    gc current block busy                      464,121,194      0.0      246,956.51 532.09us     4.41    535.99us   438.04us   619.09us    75.63us    4
      Cluster    gc current grant 2-way                     591,737,257      0.0      220,612.94 372.82us     3.94    372.72us   358.28us   380.07us     9.99us    4
      Cluster    gc cr block 3-way                          304,639,822      0.0      206,520.00 677.92us     3.69    662.38us   485.33us   758.60us   122.42us    4
      Cluster    gc cr block congested                       30,781,329      0.0      193,042.74   6.27ms     3.45      6.17ms     5.70ms     6.39ms   318.34us    4
      Cluster    gc current grant busy                      111,064,019      0.0      171,241.91   1.54ms     3.06      1.54ms     1.51ms     1.59ms    33.34us    4
      Cluster    gc current block 2-way                     356,478,432      0.0      122,968.64 344.95us     2.20    350.50us   261.83us   396.23us    60.41us    4
```

The Oracle Enterprise Manager screenshot below shows All Wait Events, IO Requests Per Second, CPU Utilization for the same 800 User Scale Test while Swingbench test was running for 24 hours.

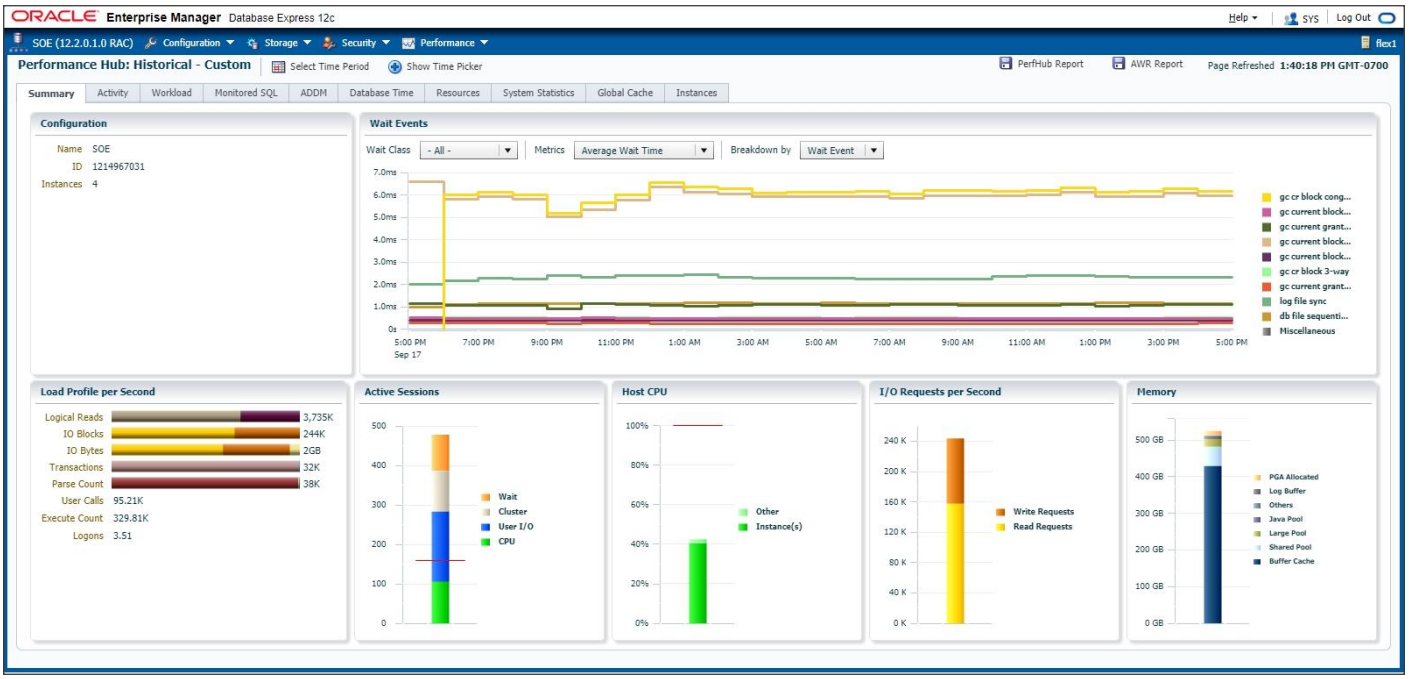

## Two (SOE and OLTP) Database Performance

For two OLTP database workload featuring Order Entry schema, we used SOE and OLTP database. For both the databases, we used 128GB of System Global Area (SGA). We also made sure that HugePages were in use all the time while databases were running. The SOE + OLTP Database scalability test were run for at least 12 hours and ensured that results are consistent for the duration of the full run.

Notice that SOE database is of approximately 4 TB in size, while OLTP database is 3 TB in size. While running below tests on both the database together, we scaled more users on the SOE database compared to the OLTP database. We ran the Swingbench scripts on each node to start SOE and OLTP database and generate AWR reports for each scenario as shown below.

## User Scalability

Table 12 lists the IOPS and System Utilization for SOE + OLTP databases running together Swingbench workloads from 200 users to 800 users across all the four RAC nodes:

**Table 12 Total IOPS for SLOB + OLTP Database for User Scalability Tests**

| Users | IOPS for SOE DB | IOPS for OLTP DB | Total IOPS | System Utilization (%) |
|---|---|---|---|---|
| 200 | 81,838 | 27,755 | 109,593 | 18.9 |
| 300 | 119,816 | 26,097 | 145,913 | 25.3 |
| 400 | 128,855 | 44,909 | 173,764 | 32.1 |
| 500 | 139,657 | 61,445 | 201,102 | 39.3 |
| 600 | 145,508 | 75,465 | 220,973 | 44.2 |
| 700 | 150,922 | 89,918 | 240,840 | 50.7 |
| 800 | 142,873 | 96,206 | 239,079 | 55.9 |

The graph below demonstrates the total IOPS for both SOE and OLTP Database while running Swingbench workload together from 200 users to 800 users.

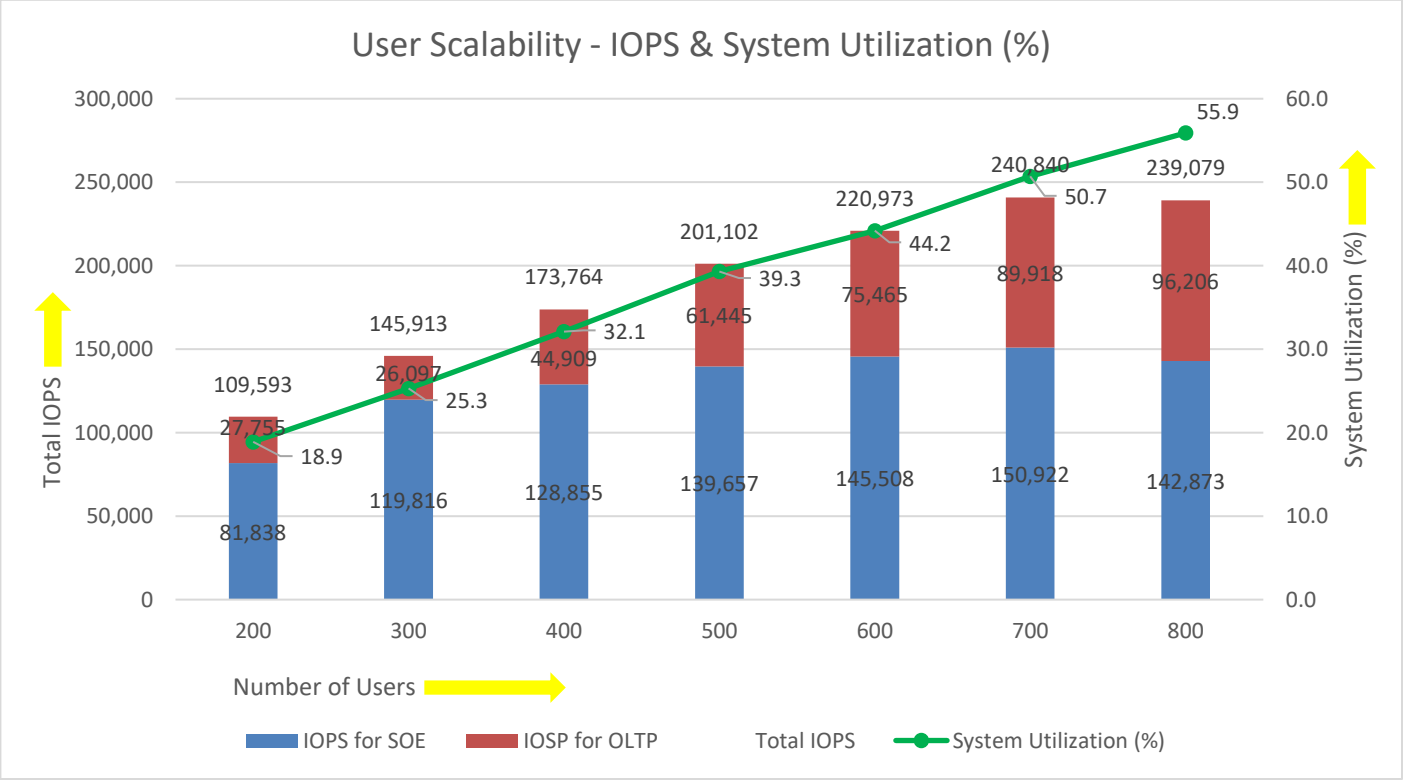

Table 13 lists the TPM and System Utilization for SOE + OLTP databases running Swingbench workloads together from 100 users to 800 users.

Table 13 Total TPM for SLOB + OLTP Database for User Scalability Tests

| Users | TPM for SOE DB | TPM for OLTP DB | Total TPM | System Utilization (%) |
|---|---|---|---|---|
| 200 | 691,092 | 232,542 | 923,634 | 18.9 |
| 300 | 1,051,458 | 217,620 | 1,269,078 | 25.3 |
| 400 | 1,129,806 | 395,370 | 1,525,176 | 32.1 |
| 500 | 1,227,612 | 552,582 | 1,780,194 | 39.3 |
| 600 | 1,298,058 | 696,966 | 1,995,024 | 44.2 |
| 700 | 1,306,896 | 821,701 | 2,128,597 | 50.7 |
| 800 | 1,265,010 | 898,944 | 2,163,954 | 55.9 |

The graph below illustrates the total IOPS for both SOE and OLTP Database while running Swingbench workload together from 200 users to 800 users.

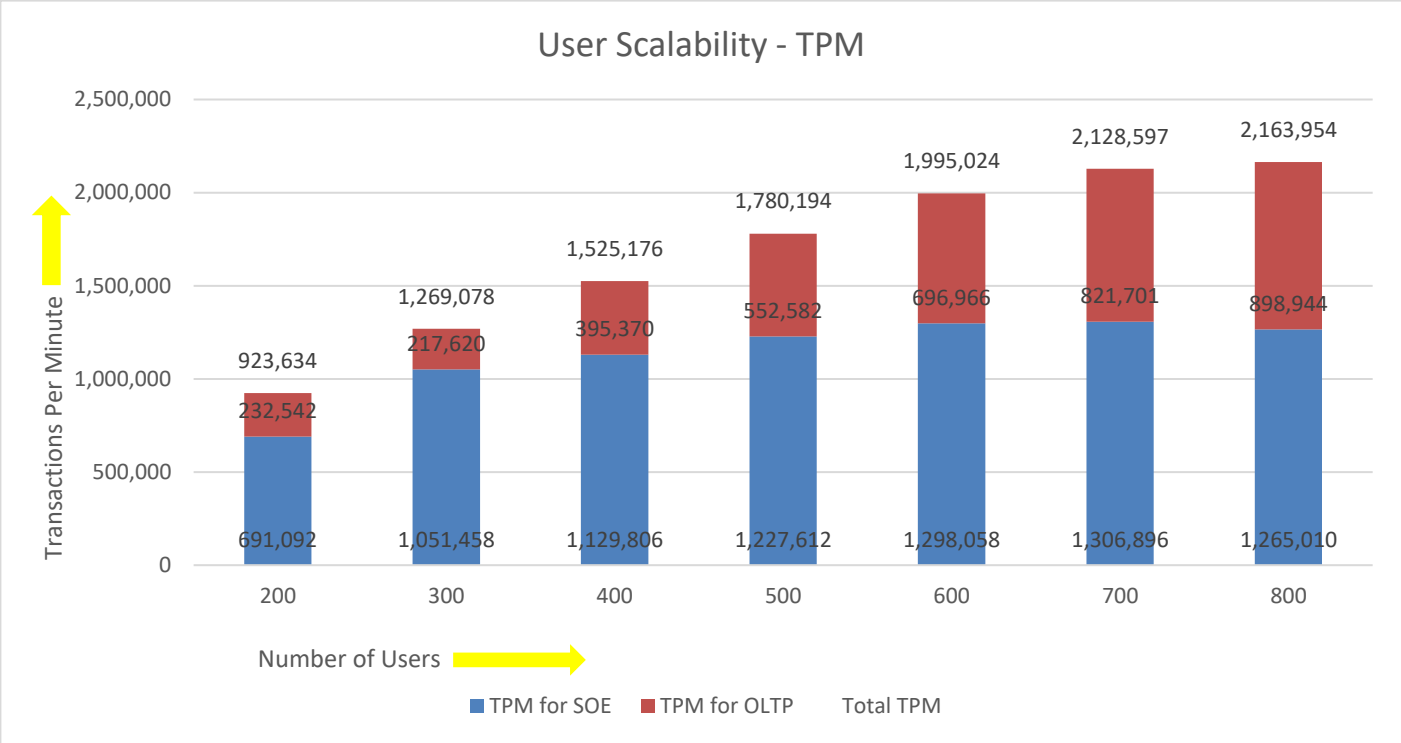


The results were in line with prior assessments where the user scalability was almost linear till 600 users and beyond 600 users the rate of IOPS increase slowed down.

The screenshot shown below was captured from the 800 Users Scale Test scenario while running Swingbench workload on two database at the same time for 24 hours. The snapshot shows a section from 24-hour window of AWR Global report from the run that highlights Physical Reads/Sec, Physical Writes/Sec and Transactions per Seconds for each instance for both the databases. Notice that IO load is distributed across all the cluster nodes performing workload operations.

```
System Statistics - Per Second              DB/Inst: SOE/soe1  Snaps: 358-382

            Logical     Physical     Physical         Redo        Block         User
  I#        Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s      Calls/s      Execs/s     Parses/s   Logons/s       Txns/s
---- --------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- ------------
   1      566,176.94     20,297.2     11,410.4     17,767.1    103,924.6     14,269.6     49,389.6      5,667.3       0.17      4,754.0
   2      567,314.38     20,247.1     11,329.0     17,636.6    103,050.3     14,148.4     48,974.7      5,619.4       0.17      4,713.9
   3      585,066.16     20,807.2     11,567.7     17,951.8    105,003.0     14,426.0     49,936.8      5,729.9       0.17      4,806.5
   4      689,140.63     23,551.8     13,308.4     20,344.8    120,754.4     16,628.9     57,557.4      6,604.3       0.16      5,540.3
 ~~~ ~~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~
 Sum    2,407,698.10     84,903.2     47,615.5     73,700.3    432,732.3     59,472.9    205,858.5     23,620.9       0.67     19,814.7
 Avg      601,924.52     21,225.8     11,903.9     18,425.1    108,183.1     14,868.2     51,464.6      5,905.2       0.17      4,953.7
 Std       58,783.80      1,571.1        941.6      1,286.3      8,418.8      1,179.3      4,080.9        468.2       0.00        392.9
```

```
System Statistics - Per Second              DB/Inst: OLTP/oltp1  Snaps: 318-342

            Logical     Physical     Physical         Redo        Block         User
  I#        Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s      Calls/s      Execs/s     Parses/s   Logons/s       Txns/s
---- --------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- ------------
   1      349,578.30     12,930.7      7,763.6     12,796.0     73,189.0      9,997.0     34,671.2      3,969.4       0.11      3,330.2
   2      402,893.75     13,027.3      7,887.3     12,871.6     73,746.0     10,082.1     34,963.6      4,003.5       0.11      3,358.5
   3      364,991.20     14,922.1      8,114.9     13,176.1     76,002.4     10,396.1     36,054.9      4,128.2       0.11      3,463.1
   4      467,315.99     15,813.3      9,518.9     15,263.8     90,031.4     12,350.9     42,832.2      4,904.0       0.11      4,114.3
 ~~~ ~~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~
 Sum    1,584,779.23     56,693.3     33,284.7     54,107.5    312,968.8     42,826.1    148,521.9     17,005.1       0.43     14,266.1
 Avg      396,194.81     14,173.3      8,321.2     13,526.9     78,242.2     10,706.5     37,130.5      4,251.3       0.11      3,566.5
 Std       52,440.00      1,426.9        811.6      1,169.5      7,953.1      1,109.6      3,847.5        440.5       0.00        369.7
```

The screenshot below shows top timed events for SOE database while running the same above Swingbench test for 24 hours.

```
Top Timed Events                            DB/Inst: SOE/soe1  Snaps: 358-382
-> 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                7.050821E+09       0.0   7,442,779.62   1.06ms    24.59       1.06ms     1.05ms     1.08ms    15.75us    4
                 DB CPU                                         N/A       N/A   6,075,235.66            20.07                                                        4
      Commit     log file sync                          1.705457E+09       0.0   4,018,701.82   2.36ms    13.28       2.36ms     1.95ms     3.15ms   551.02us    4
      Cluster    gc current grant 2-way                 2.671641E+09       0.0   1,550,173.17 580.23us     5.12     616.19us   194.39us      .96ms   317.46us    4
      Cluster    gc current block busy                  1.925001E+09       0.0   1,446,570.87 751.46us     4.78     757.45us   488.76us   944.48us   192.07us    4
      Cluster    gc current block congested               42,098,159       0.0   1,360,282.90  32.31ms     4.49      31.98ms    20.19ms    43.48ms     9.52ms    4
      Cluster    gc cr block 3-way                      1.494389E+09       0.0   1,341,880.74 897.95us     4.43     901.96us   531.05us     1.18ms   305.83us    4
      Cluster    gc cr grant 2-way                      2.210294E+09       0.0   1,191,479.42 539.06us     3.94     556.54us   188.89us   928.93us   305.73us    4
      Cluster    gc current block 2-way                 1.651133E+09       0.0   1,059,897.19 641.92us     3.50     625.57us   197.85us   789.64us   285.88us    4
      Cluster    gc cr block congested                    30,264,924       0.0   1,042,665.74  34.45ms     3.45      33.51ms    20.52ms    44.32ms     9.80ms    4
```

The screenshot below shows top timed events for OLTP database while running the same above Swingbench test for 24 hours.

```
Top Timed Events                            DB/Inst: OLTP/oltp1  Snaps: 318-342
-> 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                4.576784E+09       0.0   4,544,231.69    .99ms    27.16        .99ms      .98ms     1.02ms    17.06us    4
                 DB CPU                                         N/A       N/A   4,115,305.12            24.60                                                        4
      Commit     log file sync                          1.228804E+09       0.0   2,641,063.77   2.15ms    15.79       2.15ms     1.78ms     2.79ms   448.51us    4
      Cluster    gc current block busy                  1.496741E+09       0.0     888,689.80 593.75us     5.31     601.35us   378.92us   732.51us   154.14us    4
      System I/O db file parallel write                 1.303143E+09       0.0     858,059.42 658.45us     5.13     661.20us   614.98us   698.89us    39.17us    4
      Cluster    gc cr block 3-way                      1.146138E+09       0.0     846,483.49 738.55us     5.06     744.56us   389.49us     1.00ms   272.62us    4
      Cluster    gc current grant 2-way                 1.763565E+09       0.0     802,341.22 454.95us     4.80     485.86us   154.32us   697.67us   233.53us    4
      Cluster    gc current block 2-way                 1.250840E+09       0.0     677,436.64 541.59us     4.05     532.38us   159.00us   683.20us   250.24us    4
      Cluster    gc cr grant 2-way                      1.469241E+09       0.0     593,276.94 403.80us     3.55     418.79us   150.37us   677.93us   216.05us    4
      Cluster    gc cr block 2-way                      1.040361E+09       0.0     572,519.33 550.31us     3.42     531.39us   158.06us   668.10us   249.15us    4
```

The screenshots shown below were captured from the Oracle Enterprise Manager while running 800 users Swingbench workload test for 24 hours.

Figure 3 Wait Events, IO Requests per Second, and CPU Utilization for the SOE Database

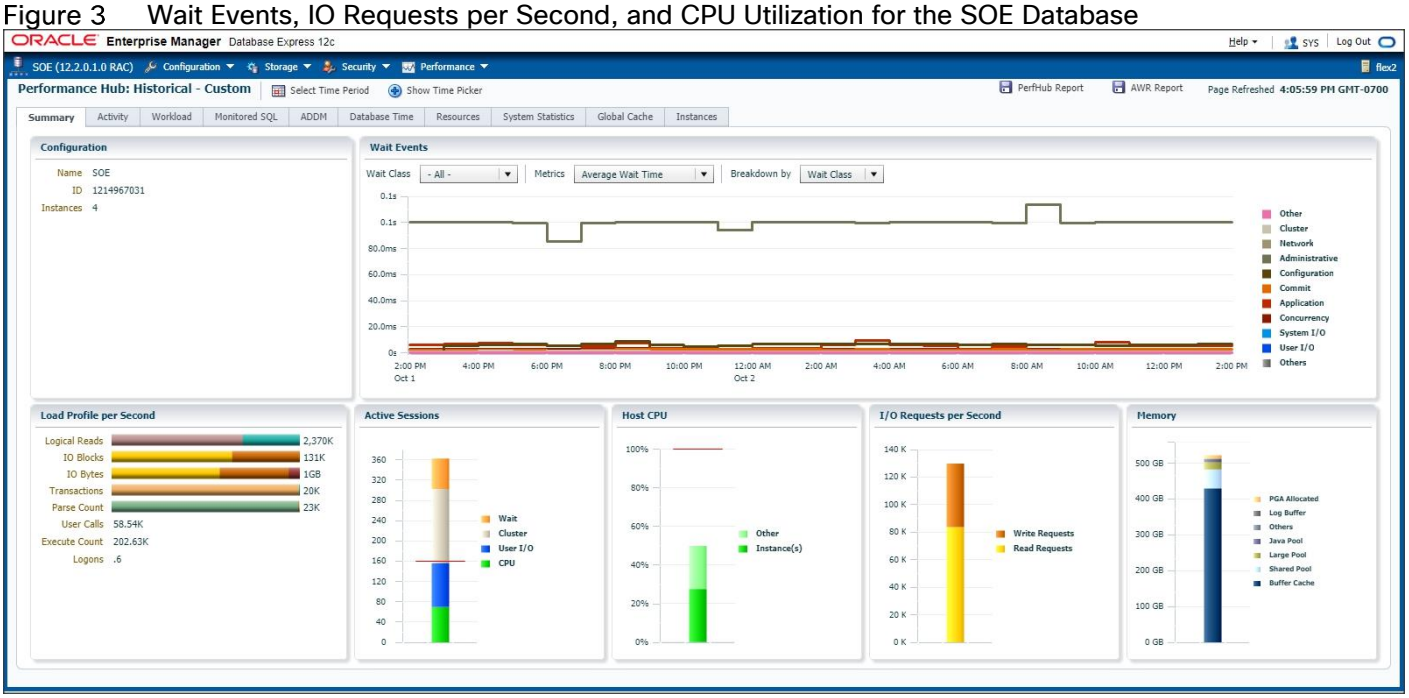


Figure 4 Wait Events, IO Requests per Second and CPU Utilization for the OLTP Database

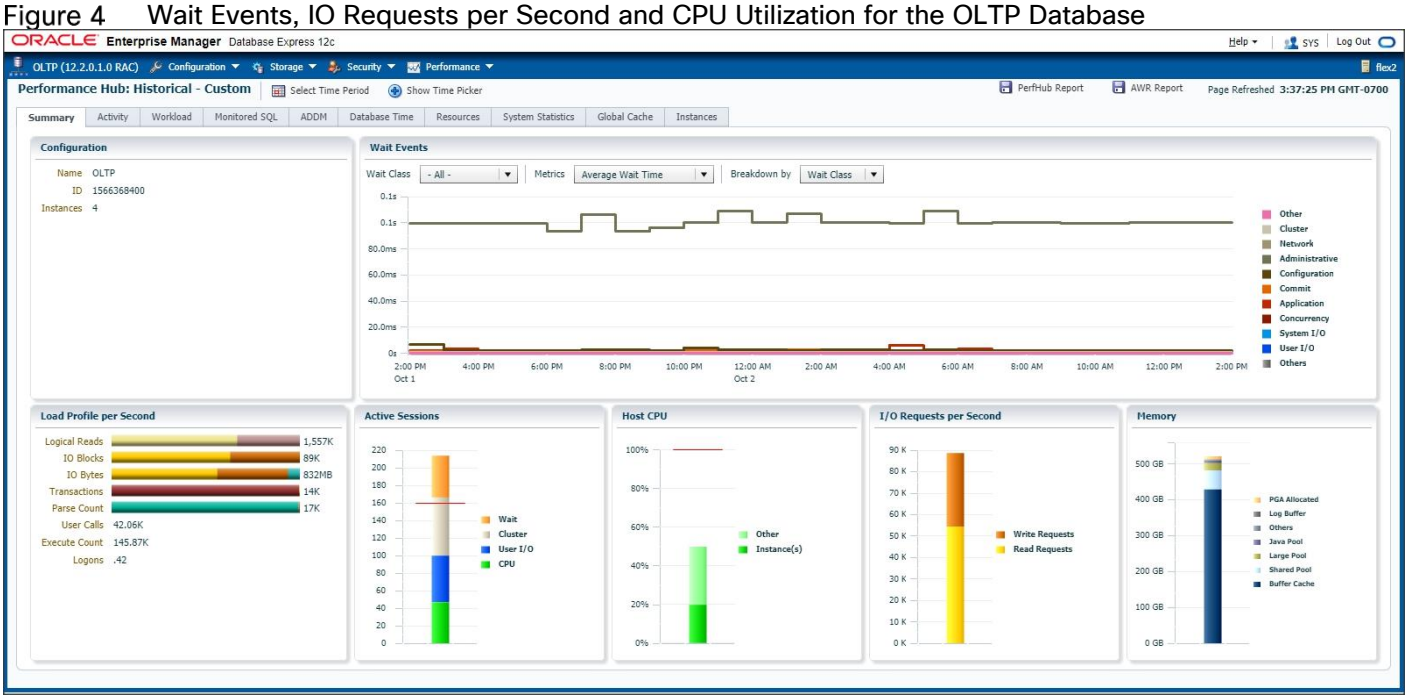


## One (DSS) Database Performance

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 6 TB of DSS database by loading Swingbench sh schema into Datafile Tablespace.

DSS Database activity is captured using Oracle Enterprise Manager for 24 hour Swingbench workload test as shown below.

For 24 hour DSS workload test, we observed the total sustained IO bandwidth average was about 10.5 GB/sec after the initial ramp up workload. As shown in above screenshot, 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 was captured from Oracle AWR report while running Swingbench SH workload on DSS database for 24 hours.

```
Top Timed Events                          DB/Inst: DSS/dss1  Snaps: 30-54
-> 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   direct path read temp                     92,645,271       0.0    1,967,690.73  21.24ms    35.36      21.26ms    20.76ms    22.13ms   618.92us    4
      User I/O   direct path read                         59,652,703       0.0    1,372,420.31  23.01ms    24.66      23.01ms    22.30ms    23.74ms   638.35us    4
                 DB CPU                                          N/A       N/A    1,371,639.33             24.65                                                      4
      User I/O   direct path write temp                   15,973,481       0.0      162,568.24  10.18ms     2.92      10.18ms     9.45ms    10.48ms   488.29us    4
      System I/O control file sequential read              2,221,964       0.0       30,841.66  13.88ms     0.55      13.88ms    13.70ms    14.09ms   167.34us    4
      Cluster    gc current block busy                       648,267       0.0        8,619.08  13.30ms     0.15      13.17ms     9.15ms    15.04ms     2.76ms    4
      Other      RMA: IPC0 completion sync                   361,383       0.0        7,009.92  19.40ms     0.13      19.40ms    19.39ms    19.40ms     4.09us    4
      Cluster    gc buffer busy acquire                       21,048       0.0        5,935.45 282.00ms     0.11     277.95ms   224.45ms   322.18ms    46.40ms    4
      Cluster    gc buffer busy release                      116,551       0.0        5,807.59  49.83ms     0.10      49.84ms    43.86ms    56.12ms     5.06ms    4
      Other      Sync ASM rebalance                           41,264       0.0        4,812.83 116.63ms     0.09     115.65ms    20.98us   155.16ms    77.09ms    4
```

The screenshot below shows Wait Class events for each Instance of DSS Database while running Swingbench workload test for 24 hours.

The screenshot below shows IO throughput and CPU utilization for each instance of DSS Database while running Swingbench workload test for 24 hours.

## All Three (SOE, OLTP and DSS) Database Performance

In this test, we ran both OLTP (SOE+OLTP) and DSS (DSS) Database Swingbench workload at the same time to measure the system performance on small random queries presented via OLTP databases as well as large and sequential transactions submitted via DSS database workload.

The screenshots shown below were captured from Oracle AWR reports while running the Swingbench workload tests on all three database at the same time for 24 hours. We achieved about 128k IOPS for SOE database, 68k IOPS for OLTP database and 4.8 GB/s throughput for DSS database.

SOE Database AWR snapshot that highlights Physical Reads/Sec, Physical Writes/Sec and Transactions per Seconds.

```
System Statistics - Per Second              DB/Inst: SOE/soe1  Snaps: 360-385

             Logical     Physical     Physical        Redo        Block        User
  I#         Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s     Calls/s      Execs/s     Parses/s   Logons/s       Txns/s
---- --------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- ------------
   1      562,479.34     19,378.5     10,856.6     17,033.1     99,294.6     13,594.0     47,056.5      5,401.4       0.12      4,530.2
   2      503,273.70     19,970.1     11,283.6     17,480.4    102,079.2     14,003.0     48,472.6      5,563.3       0.11      4,666.5
   3      638,590.27     22,721.7     12,458.0     18,984.4    111,824.1     15,393.2     53,283.1      6,115.3       0.11      5,129.6
   4      581,833.41     20,373.0     11,578.9     17,719.4    103,659.3     14,238.1     49,284.7      5,657.1       0.11      4,745.2
~~~ ~~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~
Sum    2,286,176.72     82,443.4     46,177.1     71,217.2    416,857.2     57,228.3    198,096.9     22,737.1       0.46     19,071.5
Avg      571,544.18     20,610.8     11,544.3     17,804.3    104,214.3     14,307.1     49,524.2      5,684.3       0.12      4,767.9
Std       55,809.43      1,465.3        677.5        836.5      5,384.5        771.5      2,669.7        306.1       0.00        257.0
-------------------------------------------------------------------------------------------------------------------------------------
```

OLTP Database AWR snapshot that highlights Physical Reads/Sec, Physical Writes/Sec and Transactions per Seconds.

```
System Statistics - Per Second              DB/Inst: OLTP/oltp1  Snaps: 320-345

             Logical     Physical     Physical        Redo        Block        User
  I#         Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s     Calls/s      Execs/s     Parses/s   Logons/s       Txns/s
---- --------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- ------------
   1      277,441.84      9,875.5      6,176.9     10,107.1     57,135.4      7,738.0     26,836.1      3,074.3       0.10      2,577.9
   2      303,117.35     10,150.5      6,414.6     10,454.3     59,252.4      8,049.5     27,915.3      3,196.7       0.10      2,681.5
   3      301,983.71     11,930.5      6,770.1     11,057.5     62,961.7      8,583.7     29,768.1      3,408.8       0.10      2,859.6
   4      345,490.05     10,436.7      6,505.0     10,606.5     60,200.7      8,195.5     28,417.3      3,255.3       0.10      2,730.6
~~~ ~~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~
Sum    1,228,032.95     42,393.2     25,866.6     42,225.4    239,550.1     32,566.7    112,936.7     12,935.0       0.40     10,849.6
Avg      307,008.24     10,598.3      6,466.6     10,556.3     59,887.5      8,141.7     28,234.2      3,233.8       0.10      2,712.4
Std       28,257.19        917.2        245.1        394.1      2,417.1        351.1      1,216.9        138.9       0.00        116.9
-------------------------------------------------------------------------------------------------------------------------------------
```

DSS Database AWR snapshot that highlights database MB per Seconds.

```
IO Profile (Global)                          DB/Inst: DSS/dss1  Snaps: 26-51

Statistic                      Read+Write/s        Reads/s       Writes/s
------------------------------ -------------- -------------- --------------
Total Requests                     6,154.03       5,290.44         863.59
Database Requests                  6,100.75       5,240.52         860.23
Optimized Requests                     0.00           0.00           0.00
Redo Requests                          0.44            N/A           0.44
Total (MB)                         4,814.09       4,651.36         162.72
Database (MB)                      4,813.26       4,650.58         162.67
Optimized Total (MB)                   0.00           0.00           0.00
Redo (MB)                              0.01            N/A           0.01
Database (blocks)                616,096.82     595,274.70      20,822.12
Via Buffer Cache (blocks)             95.71           2.67          93.05
Direct (blocks)                  616,001.11     595,272.03      20,729.08
-------------------------------------------------------------------------
```

The screenshots shown below were captured from the Oracle Enterprise Manager while running the Swingbench workload tests on all three database together for 24 hours.

Figure 5 SOE Database Wait Events per Instance, IO Requests per Second and CPU Utilization

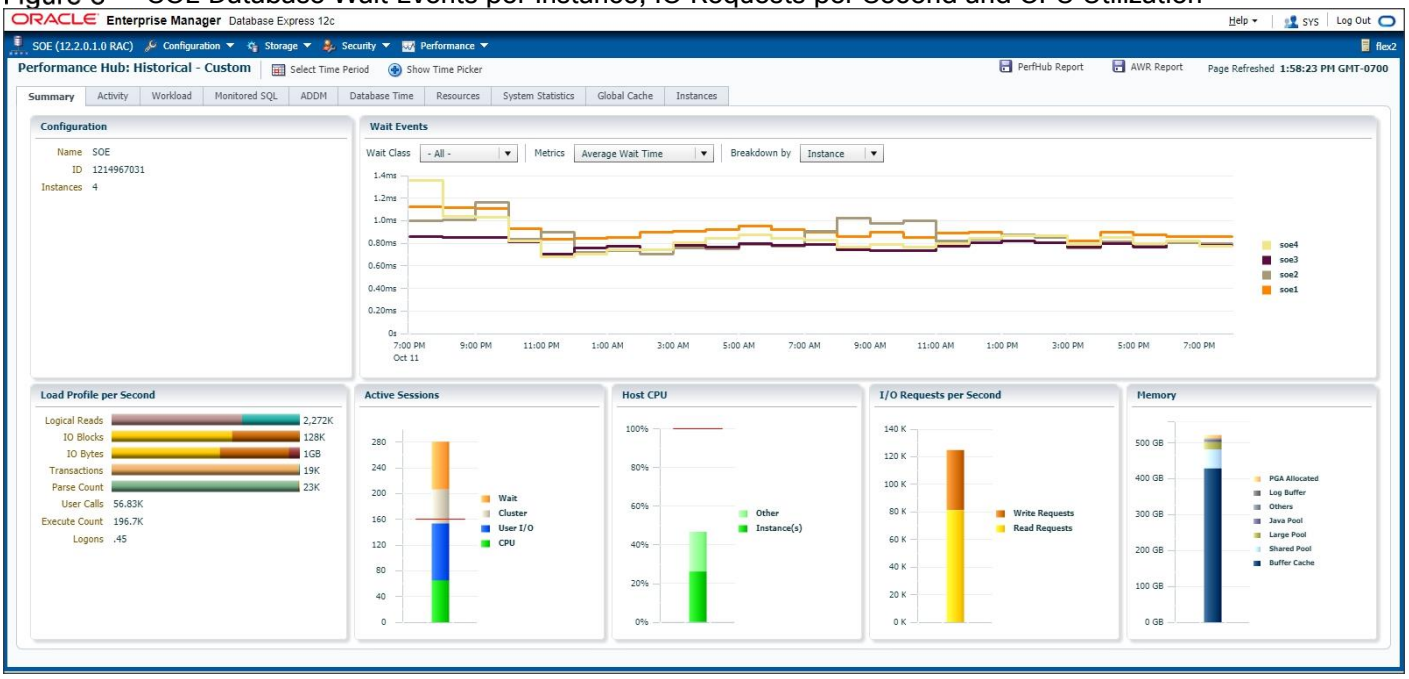

Figure 6 OLTP Database Wait Events per Instance, IO Requests per Second and CPU Utilization

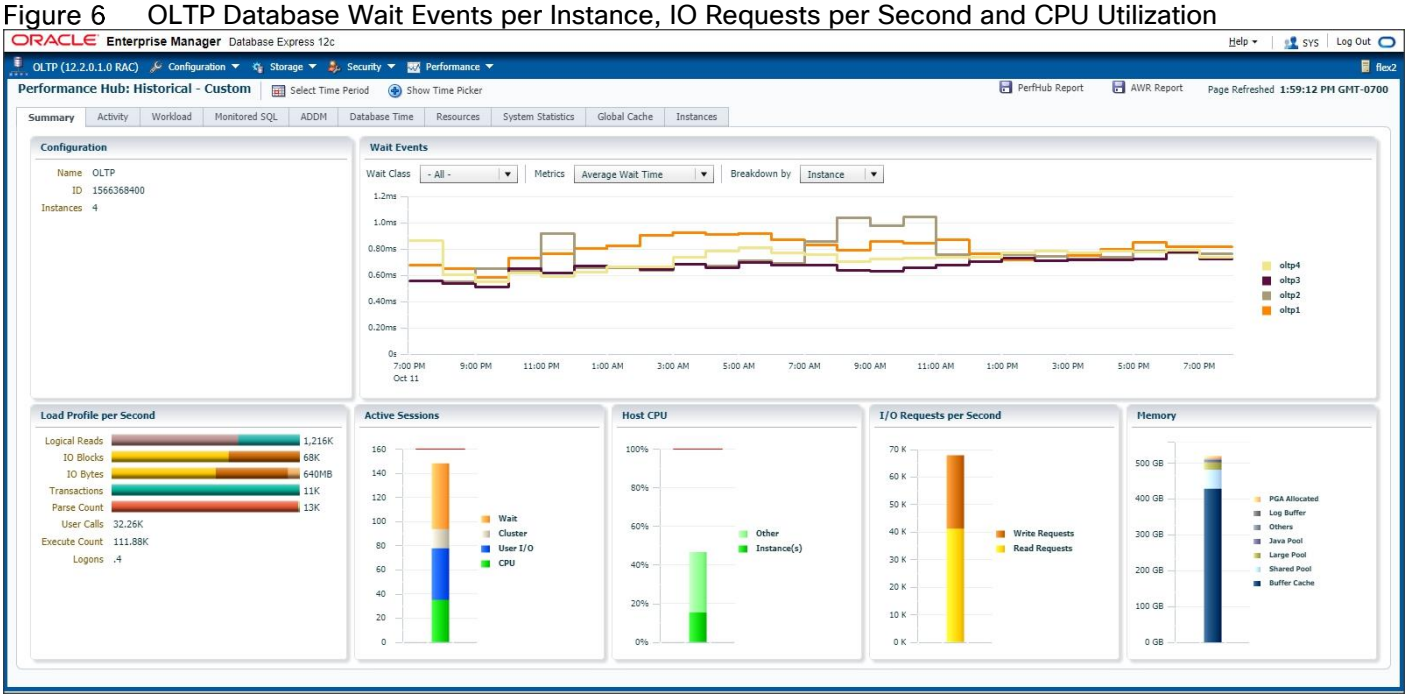


Figure 7 DSS Database Wait Events per Instance, IO Bytes and CPU Utilization

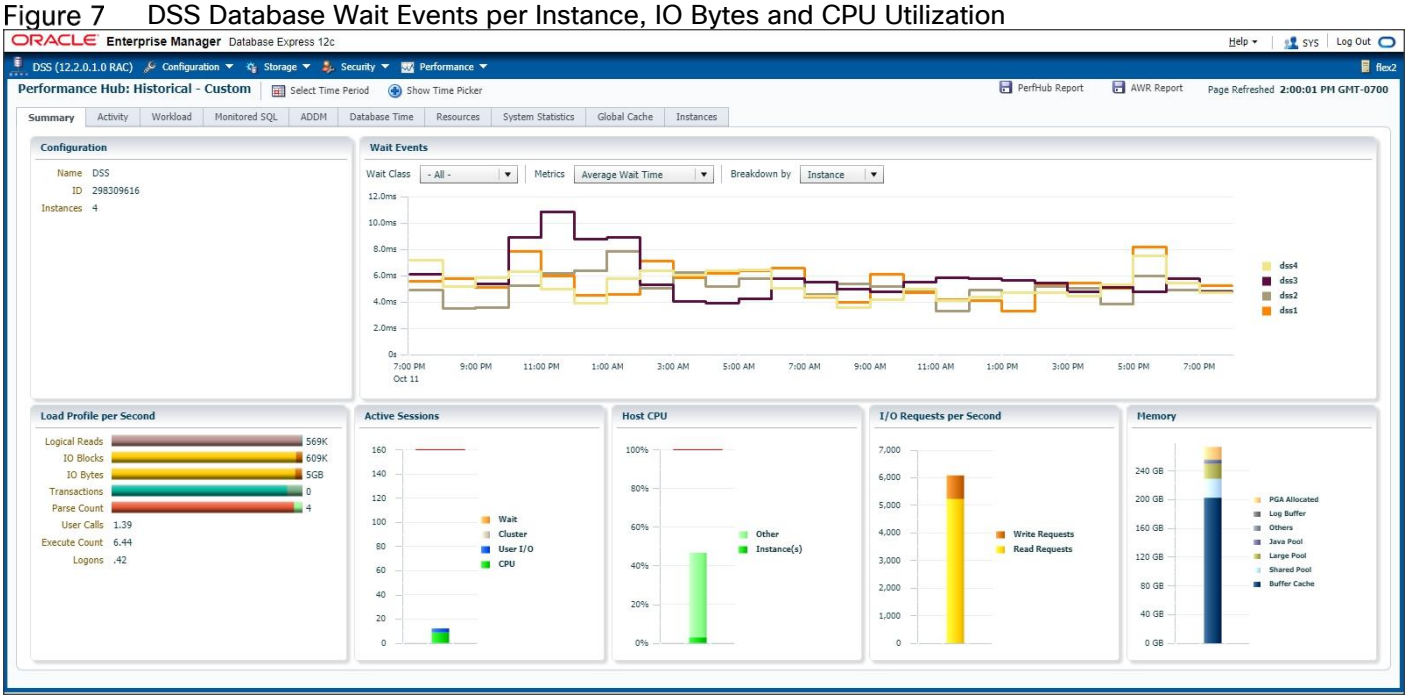

# Resiliency and Failure Tests

The goal of these tests is to help 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.

Table 14 Hardware Failover Tests

| Scenario | Tests |
|---|---|
| Test 1 – UCS FI – A Failure | **Run the system on Full Database work Load.**<br>Power Off Fabric Interconnect – A and check network traffic on Fabric Interconnect – B. |
| Test 2 – UCS FI – B Failure | **Run the system on Full Database work Load.**<br>Power Off Fabric Interconnect – B and check network traffic on Fabric Interconnect – A |
| Test 3 – UCS Nexus Switch – A Failure | **Run the system on Full Database work Load.**<br>Power Off Nexus Switch – A and check network traffic on Nexus Switch – B. |
| Test 4 – UCS Nexus Switch – B Failure | **Run the system on Full Database work Load.**<br>Power Off Nexus Switch – B and check network traffic on Nexus Switch – A. |
| Test 5 – UCS MDS Switch – A Failure | **Run the system on Full Database work Load.**<br>Power Off MDS Switch – A and check storage traffic on MDS Switch – B |
| Test 6 – UCS MDS Switch – B Failure | **Run the system on Full Database work Load.**<br>Power Off MDS Switch – B and check storage traffic on MDS Switch – A |
| Test 7 – UCS Chassis 1 and Chassis 2 IOM Links Failure | Run the system on full Database work Load. Disconnect one link from each Chassis 1 IOM and Chassis 2 IOM by pulling it out and reconnect it after 5 minutes. |

Figure 8 The FlexPod Solution Infrastructure Diagram Under Normal Operating Conditions

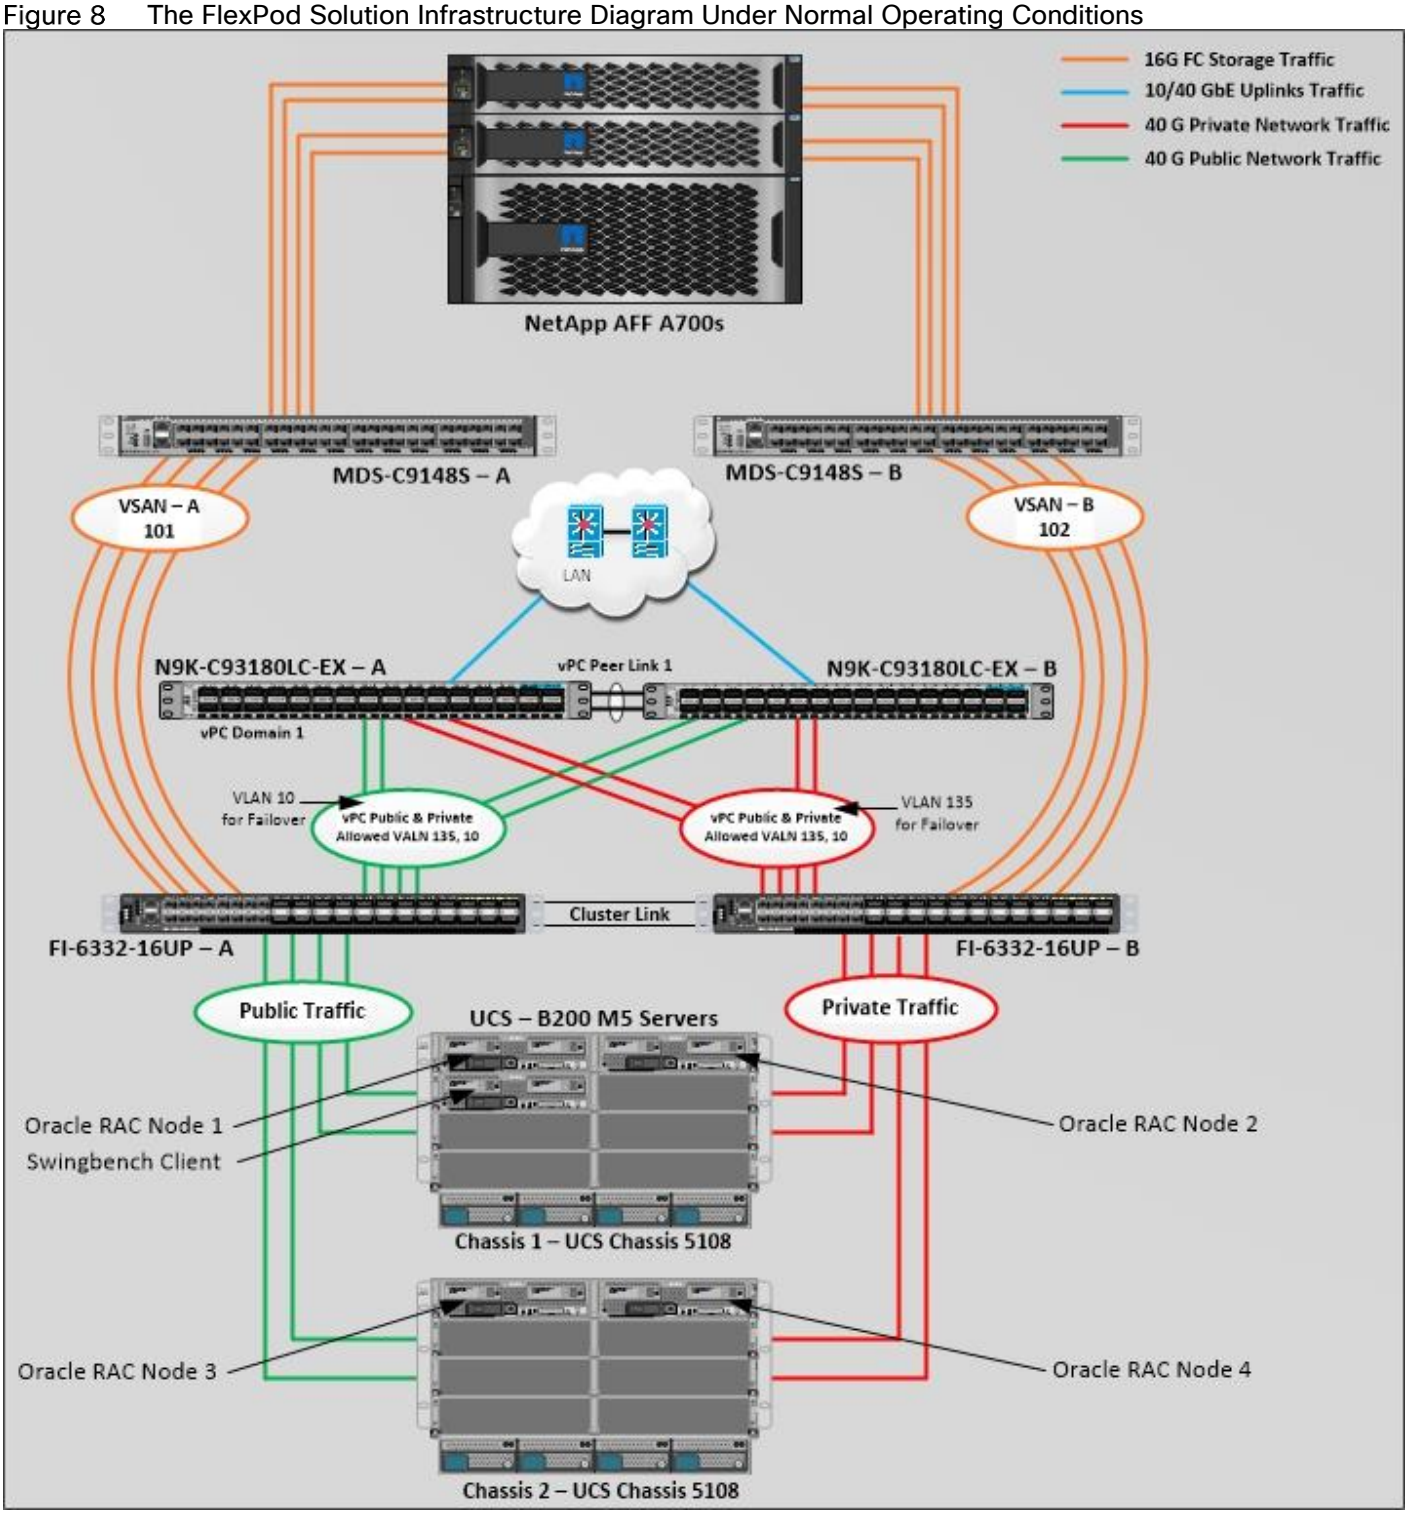


The screenshot below shows a complete infrastructure details of MAC address and VLAN information for Cisco UCS Fabric Interconnect – A switch before failover test.

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:

```
FLEXPOD-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
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F    F  Veth807
* 135      000c.293a.2afe    dynamic   0          F    F  Veth807
* 135      0025.b589.aa00    static    0          F    F  Veth773
* 135      0025.b589.aa01    static    0          F    F  Veth765
* 135      0025.b589.aa02    static    0          F    F  Veth781
* 135      0025.b589.aa03    static    0          F    F  Veth789
* 135      0025.b589.aa08    static    0          F    F  Veth807
* 4044     2c33.111c.65c3    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b365    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60d    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1f    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a1    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609c    dynamic   0          F    F  Veth32769
* 4044     cc46.d622.c9f1    dynamic   0          F    F  Eth2/1/33
```

As shown in the screenshot above, Fabric Interconnect – A carry Oracle Public Network traffic on VLAN 135 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:

```
FLEXPOD-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.b589.bb00    static    0          F    F  Veth775
* 10       0025.b589.bb01    static    0          F    F  Veth767
* 10       0025.b589.bb02    static    0          F    F  Veth783
* 10       0025.b589.bb03    static    0          F    F  Veth791
* 10       0025.b589.bb08    static    0          F    F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F    F  Veth32769
* 4044     cc46.d690.f929    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.

**All the Hardware failover tests were conducted during all the databases (SOE, OLTP and DSS) running Swingbench workloads.**

## 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 (flex1 and flex2) on chassis 1 and (flex3 and flex4) on chassis 2 will fail over the public network interface MAC addresses and its VLAN network traffic to fabric interconnect – B.

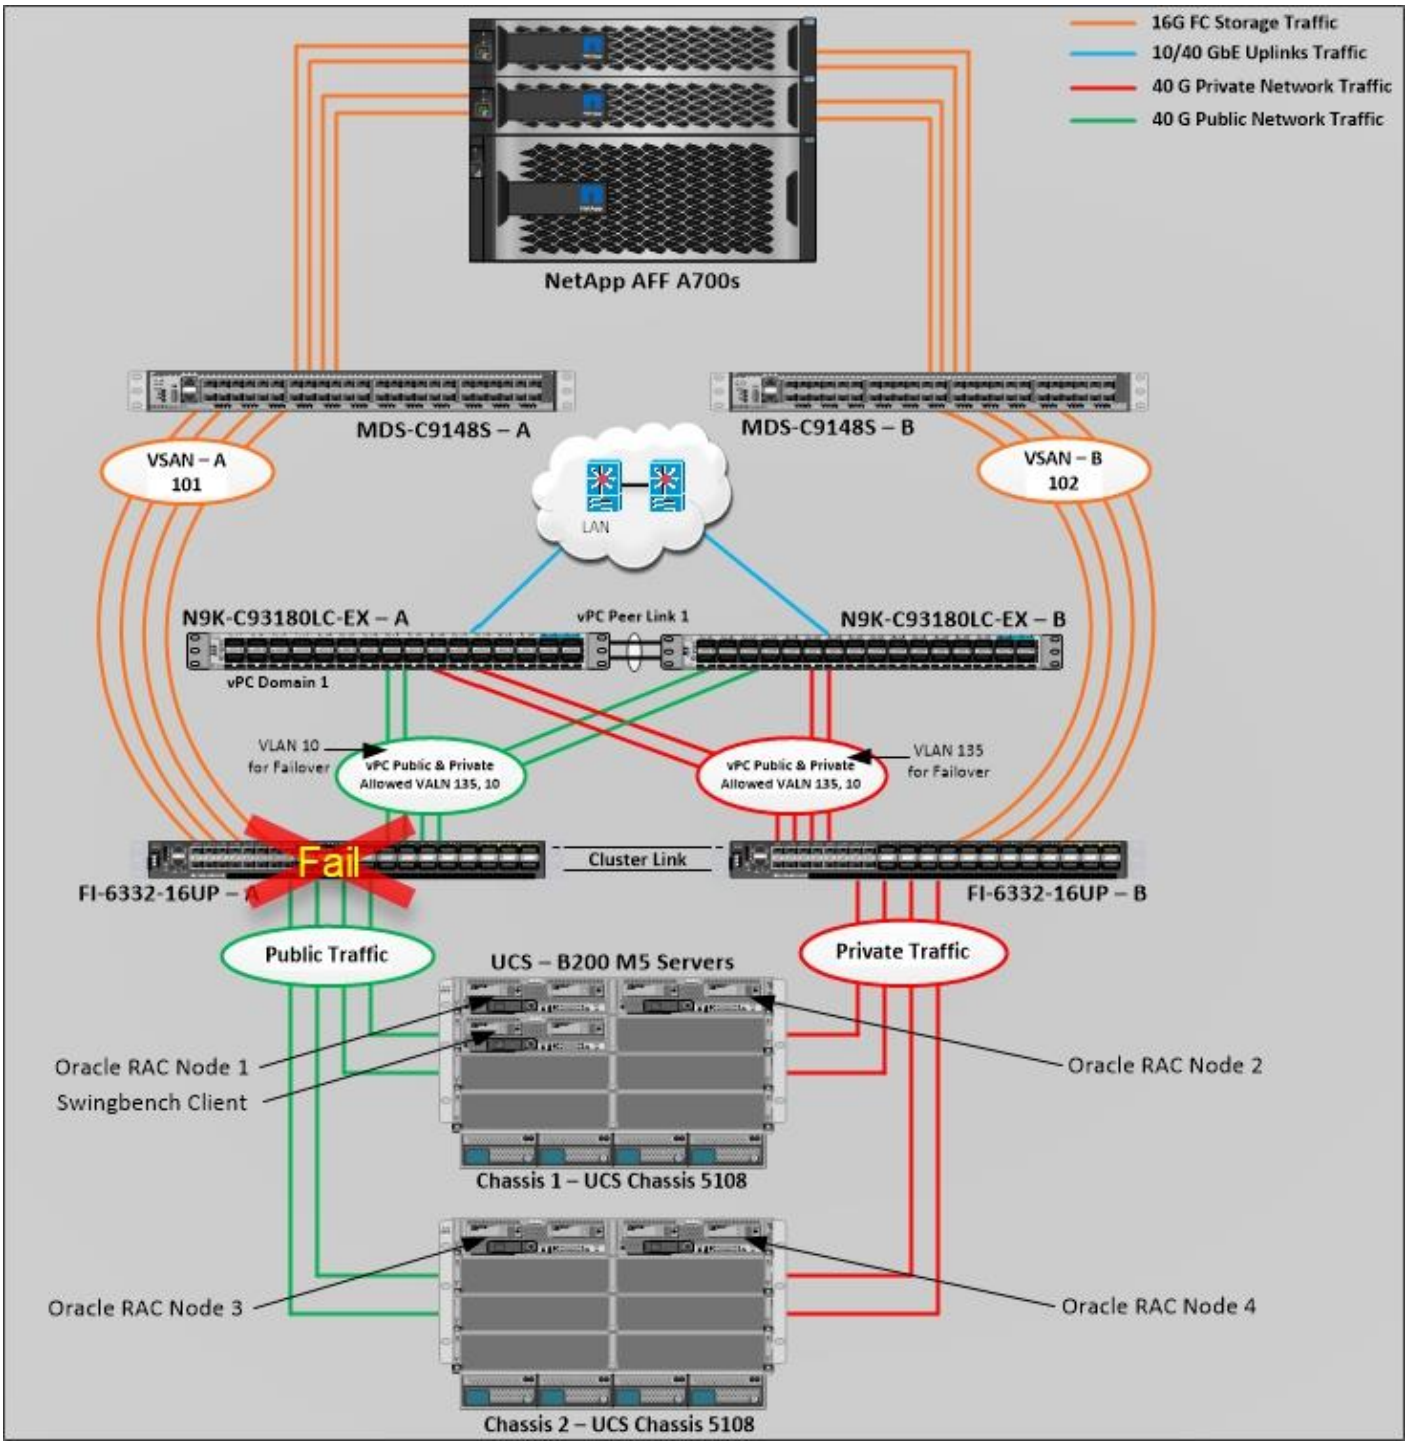


Unplug the power cable from Fabric Interconnect – A, and check the MAC address and VLAN information on Cisco UCS Fabric Interconnect – B.

```
FLEXPOD-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
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F    F  Veth808
* 135      000c.293a.2afe    dynamic   0          F    F  Veth808
* 135      0025.b589.aa00    static    0          F    F  Veth774
* 135      0025.b589.aa01    static    0          F    F  Veth766
* 135      0025.b589.aa02    static    0          F    F  Veth782
* 135      0025.b589.aa03    static    0          F    F  Veth790
* 135      0025.b589.aa08    static    0          F    F  Veth808
* 10       0025.b589.bb00    static    0          F    F  Veth775
* 10       0025.b589.bb01    static    0          F    F  Veth767
* 10       0025.b589.bb02    static    0          F    F  Veth783
* 10       0025.b589.bb03    static    0          F    F  Veth791
* 10       0025.b589.bb08    static    0          F    F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F    F  Veth32769
* 4044     cc46.d690.f929    dynamic   0          F    F  Eth2/1/33
```

We noticed (as shown in the screenshot above), when the Fabric Interconnect – A failed, it routed all the Public Network traffic of VLAN 135 to Fabric Interconnect – B. Also, we observed some performance impact on databases workload as shown in the screenshots below.

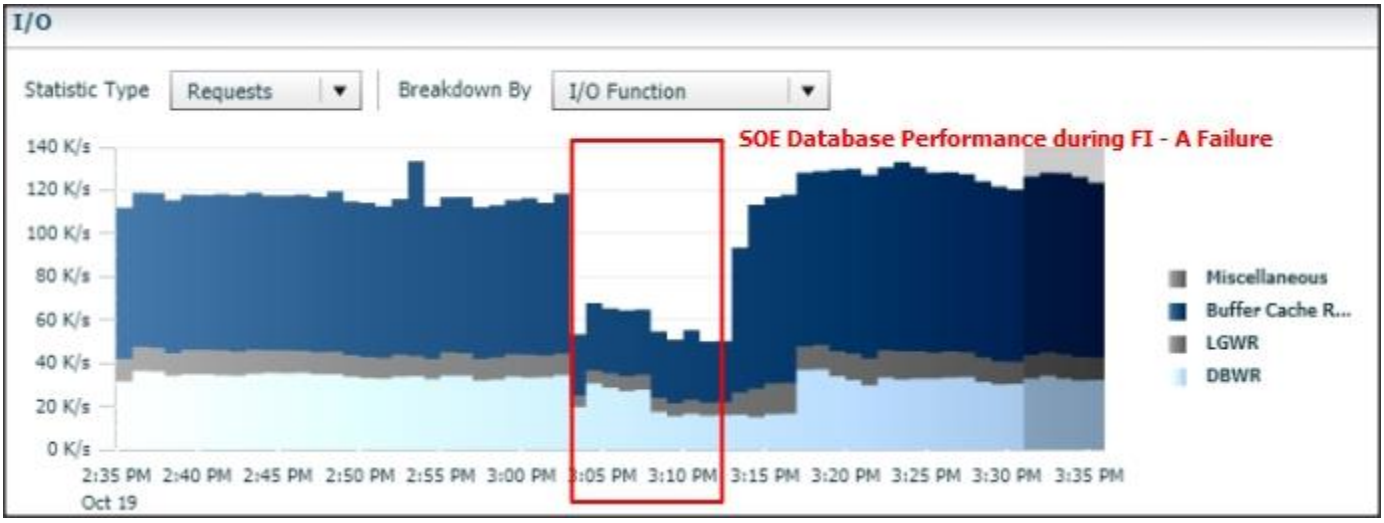

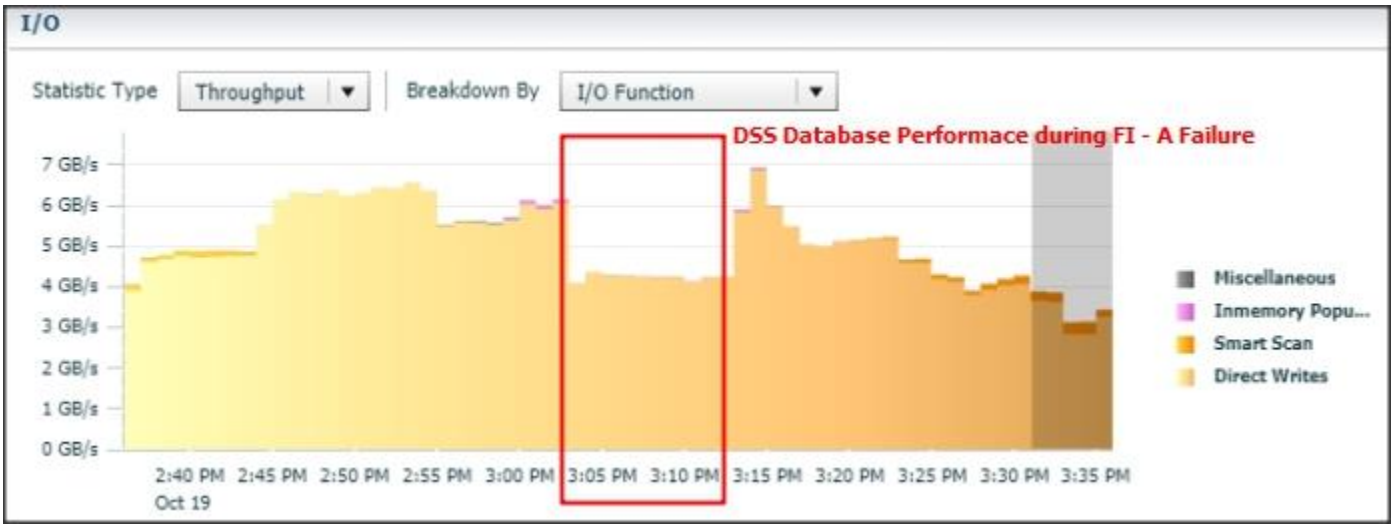


When we disconnected the power from Fabric Interconnect – A, some performance issues on the total IOPS, latency on SOE and OLTP database, as well as throughput on DSS database, as shown in the screenshots (above). However, we did not see any interruption in any Private, Public and Storage Network Traffic.

We noticed this behavior because each server node can failover vNICs from one fabric interconnect switch to another fabric interconnect switch, but there is no vHBAs failover from one fabric interconnect switch to another fabric interconnect switch. Therefore, if one fabric interconnect failure occurred, we would lose half of the number of vHBAs and consequently experience some performance impact on the databases, as shown in the screenshot (above).

After plugging the power cable into the Fabric Interconnect – A Switch, the respective blades (flex1 and flex2) on chassis 1 and (flex3 and flex4) on chassis 2 will route back the MAC addresses and its VLAN traffic to Fabric Interconnect – A. In normal operating conditions, the operating system level multipath configuration will bring back all the failed storage path back to active and database performance will resume to peak performance as previously shown.

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

```
FLEXPOD-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
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F    F  Veth807
* 135      000c.293a.2afe    dynamic   0          F    F  Veth807
* 135      0025.b589.aa00    static    0          F    F  Veth773
* 135      0025.b589.aa01    static    0          F    F  Veth765
* 135      0025.b589.aa02    static    0          F    F  Veth781
* 135      0025.b589.aa03    static    0          F    F  Veth789
* 135      0025.b589.aa08    static    0          F    F  Veth807
* 4044     2c33.111c.65c3    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b365    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60d    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1f    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a1    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609c    dynamic   0          F    F  Veth32769
* 4044     cc46.d622.c9f1    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.

```
FLEXPOD-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.b589.bb00    static    0          F    F  Veth775
* 10       0025.b589.bb01    static    0          F    F  Veth767
* 10       0025.b589.bb02    static    0          F    F  Veth783
* 10       0025.b589.bb03    static    0          F    F  Veth791
* 10       0025.b589.bb08    static    0          F    F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F    F  Veth32769
* 4044     cc46.d690.f929    dynamic   0          F    F  Eth2/1/33
```

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

As similar to the test above, we conducted a hardware failure test on Fabric Interconnect – B by disconnecting power cable to the switch.

The figure below illustrates how during Fabric Interconnect – B switch failure, the respective blades (flex1 and flex2) on chassis 1 and (flex3 and flex4) on chassis 2 will fail over the private network interface MAC addresses and its VLAN network traffic to fabric interconnect – A.

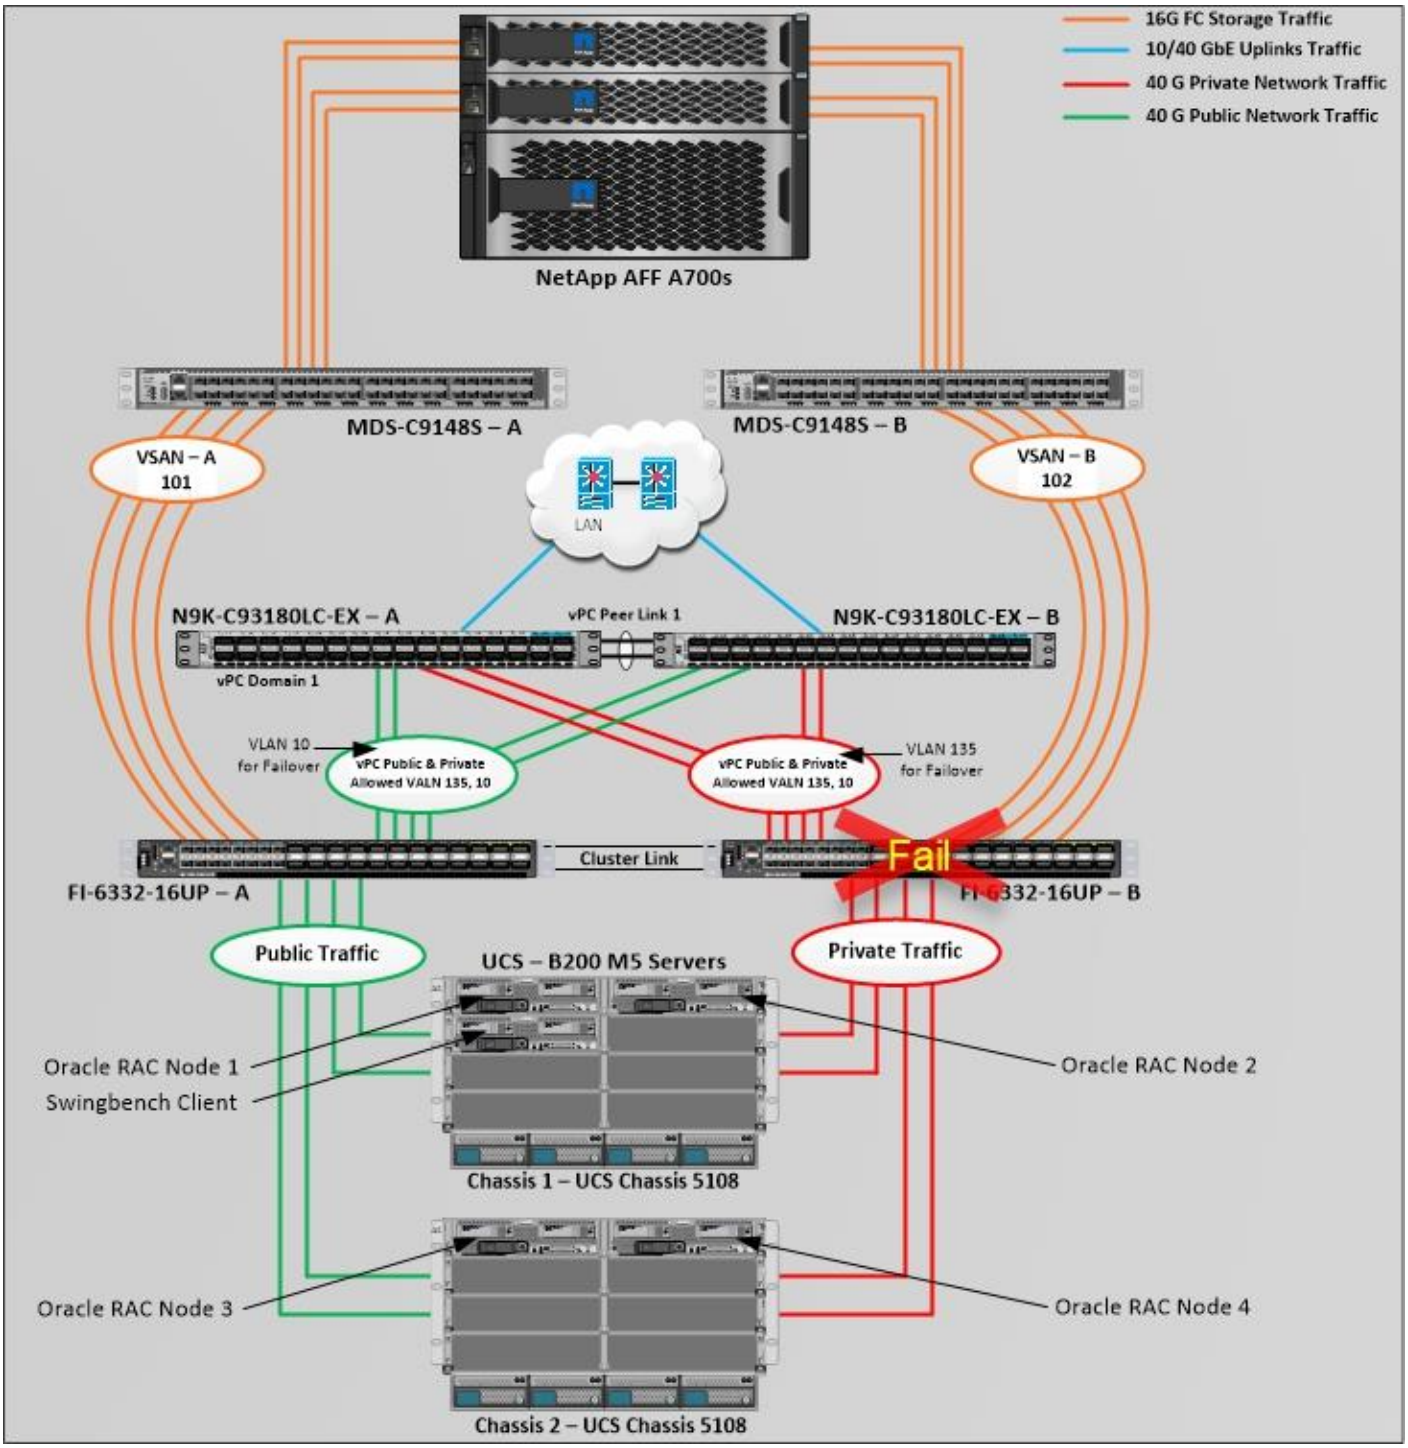


Unplug power cable from Fabric Interconnect – B, and check the MAC address and VLAN information on Cisco UCS Fabric Interconnect – A.

```
FLEXPOD-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
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F    F  Veth807
* 135      000c.293a.2afe    dynamic   0          F    F  Veth807
* 135      0025.b589.aa00    static    0          F    F  Veth773
* 135      0025.b589.aa01    static    0          F    F  Veth765
* 135      0025.b589.aa02    static    0          F    F  Veth781
* 135      0025.b589.aa03    static    0          F    F  Veth789
* 135      0025.b589.aa08    static    0          F    F  Veth807
* 10       0025.b589.bb00    static    0          F    F  Veth776
* 10       0025.b589.bb01    static    0          F    F  Veth768
* 10       0025.b589.bb02    static    0          F    F  Veth784
* 10       0025.b589.bb03    static    0          F    F  Veth792
* 10       0025.b589.bb08    static    0          F    F  Veth810
* 4044     2c33.111c.65c3    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b365    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60d    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1f    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a1    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609c    dynamic   0          F    F  Veth32769
* 4044     cc46.d622.c9f1    dynamic   0          F    F  Eth2/1/33
```

As shown in the figure above, when the Fabric Interconnect – B failed, it routed all the Private Network traffic of VLAN 10 to Fabric Interconnect – A.

When we disconnected power from Fabric Interconnect – B, some performance issues occurred on total IOPS, latency on SOE and OLTP Database as well as throughput on DSS database similar as Fabric Interconnect A failure test. Notice that, we did not see any interruption in any Private, Public and Storage Network Traffic.

We noticed this behavior because each server node can failover vNICs from one fabric interconnect switch to another fabric interconnect switch but there is no vHBAs failover from one fabric interconnect switch to another fabric interconnect switch. Therefore, in case of any one fabric interconnect failure, we would lose half of the number of vHBAs and consequently some performance impact on the databases.

After plug back power cable to Fabric Interconnect – B Switch, the respective blades (flex1 and flex2) on chassis 1 and (flex3 and flex4) on chassis 2 will route back the MAC addresses and its VLAN traffic to Fabric Interconnect – B. In normal operating conditions, the operating system level multipath configuration will bring back all the failed storage path back to active and database performance will resume to peak performance as previously shown.

## Test 3 and 4 – Cisco Nexus Switch Failure Test

We conducted a hardware failure test on Nexus Switch – A by disconnecting power cable to the switch and checking the MAC address and VLAN information on Cisco UCS Nexus Switch – B. During Nexus Switch – A failure, it routed 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.

Similarly, we conducted a hardware failure test on Nexus Switch – B by disconnecting power cable to the switch and checking the MAC address and VLAN information on Cisco UCS Nexus Switch – A. During Nexus Switch – B failure, it routed all the Private Network and Public Network Traffic of VLAN 10 and VLAN 134 to Nexus Switch – A.

## Test 5 and 6 – Cisco MDS Switch Failure Test

We conducted hardware failure test on MDS Switch – A by disconnecting power cable to the Switch and checking the storage network traffic on MDS Switch – B. As we expected, we observed some impact on all three databases performance as we lost half of the VSAN (VSAN-A 101) traffic. While VSAN-A (101) stays locally into

the switch and only carry storage traffic through the MDS switch A, VSAN-A doesn't failover to MDS Switch B therefore we reduced server to storage connectivity into half during MDS Switch A failure.

As a result, during one MDS Switch failure test, we observed similar performance impact on all the databases as Fabric Interconnect Switch failure tests.

## 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 one of the server port link cable from the Chassis as explained below.

Unplug one server port cable 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 one link from Chassis 1 and one link from Chassis 2 IOM Failed.

```
FLEXPOD-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
---------+----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F    F  Veth807
* 135      000c.293a.2afe    dynamic   0          F    F  Veth807
* 135      0025.b589.aa00    static    0          F    F  Veth773
* 135      0025.b589.aa01    static    0          F    F  Veth765
* 135      0025.b589.aa02    static    0          F    F  Veth781
* 135      0025.b589.aa03    static    0          F    F  Veth789
* 135      0025.b589.aa08    static    0          F    F  Veth807
* 4044     2c33.111c.65c3    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b365    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60d    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1f    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a1    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609c    dynamic   0          F    F  Veth32769
* 4044     cc46.d622.c9f1    dynamic   0          F    F  Eth2/1/33
```

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

```
FLEXPOD-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.b589.bb00    static    0          F    F  Veth775
* 10       0025.b589.bb01    static    0          F    F  Veth767
* 10       0025.b589.bb02    static    0          F    F  Veth783
* 10       0025.b589.bb03    static    0          F    F  Veth791
* 10       0025.b589.bb08    static    0          F    F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F    F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F    F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F    F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F    F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F    F  Veth32769
* 4044     cc46.d690.f929    dynamic   0          F    F  Eth2/1/33
```

We noticed no disruption in public and private network traffic, even though one failed traffic link from both the Chassis because of the port-channel feature.

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

# 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 such as Oracle RAC. The FlexPod Datacenter with NetApp All Flash AFF system is a converged infrastructure platform that combines best-of breed technologies from Cisco and NetApp into a powerful converged platform for enterprise applications. The pre-validated FlexPod architecture delivers proven value, agility, and performance that drive higher productivity, faster decision making, and greater opportunities for growth.

An essential feature for Oracle databases deployed on shared enterprise system is the ability to deliver consistent and dependable high performance. High performance must be coupled with non-disruptive operations, high availability, scalability, and storage efficiency. Customers can depend on Cisco UCS and NetApp Clustered Data ONTAP Storage to provide these essential elements. Built on clustered Data ONTAP unified scale-out architecture, AFF consistently meets or exceeds the high performance demands of Oracle databases. It also provides rich data management capabilities, such as integrated data protection and non-disruptive upgrades and data migration. These features allow customers to eliminate performance silos and seamlessly integrate AFF into a shared infrastructure.

Clustered Data ONTAP 9.3 delivers an enhanced inline compression capability that significantly reduces the amount of flash storage required and carries near-zero effects on system performance. The combination of Cisco UCS, NetApp and Oracle Real Application Cluster Database architecture can provide the following benefits to accelerate your IT transformation:

- 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.
- A single platform built from unified compute, fabric, and storage technologies, allowing you to scale to large-scale data centers without architectural changes.
- Enabling faster deployments, greater flexibility of choice, efficiency, high availability and lower risk.

# Appendix

## Cisco Nexus 93180 LC-EX Configuration

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

!Command: show running-config

!Time: Tue Nov 20 21:01:25 2018

version 7.0(3)I6(1)

switchname FLEXPOD-NEXUS-A

policy-map type network-qos jumbo

  class type network-qos class-default

    mtu 9216

vdc FLEXPOD-NEXUS-A id 1

  limit-resource vlan minimum 16 maximum 4094

  limit-resource vrf minimum 2 maximum 4096

  limit-resource port-channel minimum 0 maximum 511

  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 interface-vlan

feature hsrp

feature lacp

feature vpc


no password strength-check

username admin password 5 $5$FWM47Q8G$RxxMd920jZrjYX5GXLRCkfWgDtka29dQV1TiP/k4VmD
role network-admin

ip domain-lookup

system default switchport

system qos

  service-policy type network-qos jumbo
```

```
copp profile strict

snmp-server user admin network-admin auth md5 0x54fc9b6e68ba6a06acf4165853e18078
priv 0x54fc9b6e68ba6a06acf4165853e18078 localizedkey

rmon event 1 description FATAL(1) owner PMON@FATAL

rmon event 2 description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 description ERROR(3) owner PMON@ERROR

rmon event 4 description WARNING(4) owner PMON@WARNING

rmon event 5 description INFORMATION(5) owner PMON@INFO

ntp server 72.163.32.44 use-vrf default


vlan 1,10,135

vlan 10

  name Oracle_RAC_Private_Traffic

vlan 135

  name Oracle_RAC_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 10.29.135.1

port-channel load-balance src-dst l4port

vpc domain 1

  peer-keepalive destination 10.29.135.104 source 10.29.135.103

  auto-recovery


interface Vlan1


interface port-channel1

  description vPC peer-link

  switchport mode trunk

  switchport trunk allowed vlan 1,10,135

  spanning-tree port type network

  vpc peer-link
```

```
interface port-channel51
  description Port-Channel FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  vpc 51

interface port-channel52
  description Port-Channel FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  vpc 52

interface Ethernet1/1
  description Peer link connected to N9K-B-Eth1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  channel-group 1 mode active

interface Ethernet1/2
  description Peer link connected to N9K-B-Eth1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  channel-group 1 mode active

interface Ethernet1/3
...........
...........
```

```
interface Ethernet1/20

interface Ethernet1/21
  description Fabric-Interconnect-A-31
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active

interface Ethernet1/22
  description Fabric-Interconnect-A-32
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active

interface Ethernet1/23
  description Fabric-Interconnect-B-31
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 52 mode active

interface Ethernet1/24
  description Fabric-Interconnect-B-32
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
```

```
  channel-group 52 mode active

interface Ethernet1/25
  description NetApp Controller 1 e5e
  switchport access vlan 135

interface Ethernet1/26
...........
interface Ethernet1/28

interface Ethernet1/29
  description connect to uplink switch
  switchport access vlan 135
  speed 1000

interface Ethernet1/30
...........
...........
interface Ethernet1/32

interface mgmt0
  vrf member management
  ip address 10.29.135.103/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I6.1.bin
no system default switchport shutdown
MDS 9148S FC Switch Configuration
FLEXPOD-MDS-A# show running-config
!Command: show running-config
!Time: Tue Nov 20 21:19:36 2018
version 7.3(0)D1(1)
```

```
power redundancy-mode redundant

feature npiv

no feature http-server

role name default-role

  description This is a system defined role and applies to all users.

  rule 5 permit show feature environment

  rule 4 permit show feature hardware

  rule 3 permit show feature module

  rule 2 permit show feature snmp

  rule 1 permit show feature system

username admin password 5 $5$lC5R4cI9$rfRLLbkccOM1lwxALFaHFjx1XtCse5pPjAk3i/HJi0/
role network-admin

no password strength-check

ip domain-lookup

ip host FLEXPOD-MDS-A  10.29.135.105

aaa group server radius radius

snmp-server user admin network-admin auth md5 0xe2b9f6bd292aa43dc651e566873540cf
priv 0xe2b9f6bd292aa43dc651e566873540cf localizedkey

rmon event 1 log description FATAL(1) owner PMON@FATAL

rmon event 2 log description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 log description ERROR(3) owner PMON@ERROR

rmon event 4 log description WARNING(4) owner PMON@WARNING

rmon event 5 log description INFORMATION(5) owner PMON@INFO

ntp server 72.163.32.44

vsan database

  vsan 101

device-alias database

  device-alias name flex1-hba0 pwwn 20:00:00:25:b5:8a:a0:00

  device-alias name flex1-hba2 pwwn 20:00:00:25:b5:8a:a0:01

  device-alias name flex2-hba0 pwwn 20:00:00:25:b5:8a:a0:02

  device-alias name flex2-hba2 pwwn 20:00:00:25:b5:8a:a0:03

  device-alias name flex3-hba0 pwwn 20:00:00:25:b5:8a:a0:04

  device-alias name flex3-hba2 pwwn 20:00:00:25:b5:8a:a0:05
```

```
  device-alias name flex4-hba0 pwwn 20:00:00:25:b5:8a:a0:06
  device-alias name flex4-hba2 pwwn 20:00:00:25:b5:8a:a0:07
  device-alias name NetApp-A700s-01-2A pwwn 20:06:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-01-3A pwwn 20:07:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-2A pwwn 20:09:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-3A pwwn 20:0a:00:a0:98:af:7c:5b
device-alias commit
fcdomain fcid database
  vsan 1 wwn 20:06:8c:60:4f:bd:31:80 fcid 0x940000 dynamic
  vsan 1 wwn 20:05:8c:60:4f:bd:31:80 fcid 0x940100 dynamic
  vsan 1 wwn 20:01:00:de:fb:92:99:00 fcid 0x940300 dynamic
  vsan 1 wwn 20:02:00:de:fb:92:99:00 fcid 0x940400 dynamic
  vsan 1 wwn 20:03:00:de:fb:92:99:00 fcid 0x940500 dynamic
  vsan 1 wwn 20:04:00:de:fb:92:99:00 fcid 0x940600 dynamic
  vsan 1 wwn 50:0a:09:82:80:12:f8:47 fcid 0x940700 dynamic
  vsan 1 wwn 50:0a:09:84:80:12:f8:47 fcid 0x940800 dynamic
  vsan 1 wwn 50:0a:09:81:80:12:f8:47 fcid 0x940900 dynamic
  vsan 1 wwn 50:0a:09:83:80:12:f8:47 fcid 0x940a00 dynamic
  vsan 1 wwn 50:0a:09:81:80:12:f8:3d fcid 0x940b00 dynamic
  vsan 1 wwn 50:0a:09:83:80:12:f8:3d fcid 0x940c00 dynamic
  vsan 101 wwn 50:0a:09:81:80:12:f8:47 fcid 0xa20000 dynamic
  vsan 101 wwn 50:0a:09:83:80:12:f8:47 fcid 0xa20100 dynamic
  vsan 101 wwn 50:0a:09:81:80:12:f8:3d fcid 0xa20200 dynamic
  vsan 101 wwn 50:0a:09:83:80:12:f8:3d fcid 0xa20300 dynamic
  vsan 101 wwn 20:01:00:de:fb:92:99:00 fcid 0xa20400 dynamic
  vsan 101 wwn 20:00:00:25:b5:8a:a0:00 fcid 0xa20702 dynamic
!              [flex1-hba0]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:04 fcid 0xa20501 dynamic
!              [flex3-hba0]
  vsan 1 wwn 20:4d:54:7f:ee:76:cd:80 fcid 0x940200 dynamic
  vsan 101 wwn 20:00:00:25:b5:8a:a0:06 fcid 0xa20502 dynamic
!              [flex4-hba0]
```

```
  vsan 101 wwn 20:00:00:25:b5:8a:a0:02 fcid 0xa20701 dynamic
!              [flex2-hba0]
  vsan 101 wwn 20:02:00:de:fb:92:99:00 fcid 0xa20500 dynamic
  vsan 101 wwn 20:03:00:de:fb:92:99:00 fcid 0xa20600 dynamic
  vsan 101 wwn 20:04:00:de:fb:92:99:00 fcid 0xa20700 dynamic
  vsan 101 wwn 20:02:00:a0:98:af:7c:5b fcid 0xa20001 dynamic
  vsan 101 wwn 20:04:00:a0:98:af:7c:5b fcid 0xa20101 dynamic
  vsan 101 wwn 20:06:00:a0:98:af:7c:5b fcid 0xa20002 dynamic
!              [NetApp-A700s-01-2A]
  vsan 101 wwn 20:07:00:a0:98:af:7c:5b fcid 0xa20102 dynamic
!              [NetApp-A700s-01-3A]
  vsan 101 wwn 20:09:00:a0:98:af:7c:5b fcid 0xa20201 dynamic
!              [NetApp-A700s-02-2A]
  vsan 101 wwn 20:0a:00:a0:98:af:7c:5b fcid 0xa20301 dynamic
!              [NetApp-A700s-02-3A]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:01 fcid 0xa20601 dynamic
!              [flex1-hba2]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:03 fcid 0xa20404 dynamic
!              [flex2-hba2]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:05 fcid 0xa20401 dynamic
!              [flex3-hba2]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:07 fcid 0xa20603 dynamic
!              [flex4-hba2]
!Active Zone Database Section for vsan 101
zone name flex1 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:00
!               [flex1-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:01
!               [flex1-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
```

```
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zone name flex2 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:02
!               [flex2-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:03
!               [flex2-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zone name flex3 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:04
!               [flex3-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:05
!               [flex3-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
```

```
!               [NetApp-A700s-02-3A]

zone name flex4 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:06
!               [flex4-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:07
!               [flex4-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zoneset name flex vsan 101
    member flex1
    member flex2
    member flex3
    member flex4

zoneset activate name flex vsan 101
do clear zone database vsan 101
!Full Zone Database Section for vsan 101
zone name flex1 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:00
!               [flex1-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:01
!               [flex1-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
```

```
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zone name flex2 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:02
!               [flex2-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:03
!               [flex2-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zone name flex3 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:04
!               [flex3-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:05
!               [flex3-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
```

```
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zone name flex4 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:06
!               [flex4-hba0]
    member pwwn 20:00:00:25:b5:8a:a0:07
!               [flex4-hba2]
    member pwwn 20:06:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-2A]
    member pwwn 20:07:00:a0:98:af:7c:5b
!               [NetApp-A700s-01-3A]
    member pwwn 20:09:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-2A]
    member pwwn 20:0a:00:a0:98:af:7c:5b
!               [NetApp-A700s-02-3A]

zoneset name flex vsan 101
    member flex1
    member flex2
    member flex3
    member flex4

interface mgmt0
  ip address 10.29.135.105 255.255.255.0
vsan database
  vsan 101 interface fc1/1
  vsan 101 interface fc1/2
  vsan 101 interface fc1/3
  vsan 101 interface fc1/4
  vsan 101 interface fc1/5
  vsan 101 interface fc1/6
```

```
  vsan 101 interface fc1/7
  vsan 101 interface fc1/8
  vsan 101 interface fc1/9
  vsan 101 interface fc1/10
  vsan 101 interface fc1/11
  vsan 101 interface fc1/12
switchname FLEXPOD-MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz-npe.7.3.0.D1.1.bin
boot system bootflash:/m9100-s5ek9-mz-npe.7.3.0.D1.1.bin
interface fc1/1
……————————
……————————
interface fc1/48
interface fc1/1
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/2
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/3
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/4
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/5
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/6
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/7
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/8
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/9
```

```
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/10
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/11
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/12
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/13
  port-license acquire
……————
……————
interface fc1/48
  port-license acquire

no system default switchport shutdown
ip default-gateway 10.29.135.1
```

## Multipath Configuration "/etc/multipath.conf"

```
defaults {
        find_multipaths yes
        user_friendly_names no
}
multipaths {
        multipath {
                wwid            3600a098038304173475d4c766a49744e
                alias           node1_os
        }
        multipath {
                wwid            3600a098038304173475d4c766a49754a
                alias           ocrvote_1
        }
        multipath {
                wwid            3600a098038304173452b4c7959614634
                alias           ocrvote_2
        }
        multipath {
                wwid            3600a098038304173475d4c766a49754b
                alias           oradata_slob_1
        }
        multipath {
                wwid            3600a098038304173452b4c7959614635
                alias           oradata_slob_2
        }
        multipath {
                wwid            3600a098038304173475d4c766a49754c
                alias           oradata_slob_3
        }
        multipath {
                wwid            3600a098038304173452b4c7959614636
```

```
        alias           oradata_slob_4
}
multipath {
        wwid            3600a098038304173475d4c766a49754d
        alias           oradata_slob_5
}
multipath {
        wwid            3600a098038304173452b4c7959614637
        alias           oradata_slob_6
}
multipath {
        wwid            3600a098038304173475d4c766a49754e
        alias           oradata_slob_7
}
multipath {
        wwid            3600a098038304173452b4c7959614638
        alias           oradata_slob_8
}
multipath {
        wwid            3600a098038304173475d4c766a49754f
        alias           oradata_slob_9
}
multipath {
        wwid            3600a098038304173452b4c7959614639
        alias           oradata_slob_10
}
multipath {
        wwid            3600a098038304173475d4c766a497550
        alias           oradata_slob_11
}
multipath {
        wwid            3600a098038304173452b4c795961462d
```

```
        alias           oradata_slob_12
    }
    multipath {
        wwid            3600a098038304173475d4c766a497551
        alias           oradata_slob_13
    }
    multipath {
        wwid            3600a098038304173452b4c7959614641
        alias           oradata_slob_14
    }
    multipath {
        wwid            3600a098038304173475d4c766a497552
        alias           oradata_slob_15
    }
    multipath {
        wwid            3600a098038304173452b4c7959614642
        alias           oradata_slob_16
    }
    multipath {
        wwid            3600a098038304173475d4c766a497553
        alias           oraredo_slob_1
    }
    multipath {
        wwid            3600a098038304173452b4c7959614643
        alias           oraredo_slob_2
    }
    multipath {
        wwid            3600a098038304173475d4c766a497563
        alias           oradata_oltp_1
    }
  multipath {
        wwid            3600a098038304173452b4c795961464e
```

```
        alias           oradata_oltp_2
 }
multipath {
        wwid            3600a098038304173475d4c766a497564
        alias           oradata_oltp_3
 }
multipath {
        wwid            3600a098038304173452b4c795961464f
        alias           oradata_oltp_4
 }
multipath {
        wwid            3600a098038304173475d4c766a497565
        alias           oradata_oltp_5
 }
multipath {
        wwid            3600a098038304173452b4c7959614650
        alias           oradata_oltp_6
 }
multipath {
        wwid            3600a098038304173475d4c766a497566
        alias           oradata_oltp_7
 }
multipath {
        wwid            3600a098038304173452b4c7959614651
        alias           oradata_oltp_8
 }
multipath {
        wwid            3600a098038304173475d4c766a497567
        alias           oradata_oltp_9
 }
multipath {
        wwid            3600a098038304173452b4c7959614652
```

```
        alias           oradata_oltp_10
 }
multipath {
        wwid            3600a098038304173475d4c766a497568
        alias           oradata_oltp_11
 }
multipath {
        wwid            3600a098038304173452b4c7959614653
        alias           oradata_oltp_12
 }
multipath {
        wwid            3600a098038304173475d4c766a497569
        alias           oradata_oltp_13
 }
multipath {
        wwid            3600a098038304173452b4c7959614654
        alias           oradata_oltp_14
 }
multipath {
        wwid            3600a098038304173475d4c766a49756a
        alias           oradata_oltp_15
 }
multipath {
        wwid            3600a098038304173452b4c7959614655
        alias           oradata_oltp_16
 }
 multipath {
        wwid            3600a098038304173475d4c766a49756b
        alias           oraredo_oltp_1
 }
 multipath {
        wwid            3600a098038304173452b4c7959614656
```

```
        alias           oraredo_oltp_2
}
multipath {
        wwid            3600a098038304173475d4c766a49756c
        alias           oradata_soe_1
}
multipath {
        wwid            3600a098038304173452b4c7959614657
        alias           oradata_soe_2
}
multipath {
        wwid            3600a098038304173475d4c766a49756d
        alias           oradata_soe_3
}
multipath {
        wwid            3600a098038304173452b4c7959614658
        alias           oradata_soe_4
}
multipath {
        wwid            3600a098038304173475d4c766a49756e
        alias           oradata_soe_5
}
multipath {
        wwid            3600a098038304173452b4c7959614659
        alias           oradata_soe_6
}
multipath {
        wwid            3600a098038304173475d4c766a49756f
        alias           oradata_soe_7
}
multipath {
        wwid            3600a098038304173452b4c795961465a
```

```
        alias           oradata_soe_8
}
multipath {
        wwid            3600a098038304173475d4c766a497570
        alias           oradata_soe_9
}
multipath {
        wwid            3600a098038304173452b4c795961462f
        alias           oradata_soe_10
}
multipath {
        wwid            3600a098038304173475d4c766a497571
        alias           oradata_soe_11
}
multipath {
        wwid            3600a098038304173452b4c7959614661
        alias           oradata_soe_12
}
multipath {
        wwid            3600a098038304173475d4c766a497572
        alias           oradata_soe_13
}
multipath {
        wwid            3600a098038304173452b4c7959614662
        alias           oradata_soe_14
}
multipath {
        wwid            3600a098038304173475d4c766a497573
        alias           oradata_soe_15
}
multipath {
        wwid            3600a098038304173452b4c7959614663
```

```
        alias           oradata_soe_16
}
multipath {
        wwid            3600a098038304173475d4c766a497574
        alias           oraredo_soe_1
}
multipath {
        wwid            3600a098038304173452b4c7959614664
        alias           oraredo_soe_2
}
multipath {
        wwid            3600a098038304173475d4c766a497575
        alias           oradata_dss_1
}
multipath {
        wwid            3600a098038304173452b4c7959614665
        alias           oradata_dss_2
}
multipath {
        wwid            3600a098038304173475d4c766a497576
        alias           oradata_dss_3
}
multipath {
        wwid            3600a098038304173452b4c7959614666
        alias           oradata_dss_4
}
multipath {
        wwid            3600a098038304173475d4c766a497577
        alias           oradata_dss_5
}
multipath {
        wwid            3600a098038304173452b4c7959614667
```

```
        alias           oradata_dss_6
}
multipath {
        wwid            3600a098038304173475d4c766a497578
        alias           oradata_dss_7
}
multipath {
        wwid            3600a098038304173452b4c7959614668
        alias           oradata_dss_8
}
multipath {
        wwid            3600a098038304173475d4c766a497579
        alias           oradata_dss_9
}
multipath {
        wwid            3600a098038304173452b4c7959614669
        alias           oradata_dss_10
}
multipath {
        wwid            3600a098038304173475d4c766a49757a
        alias           oradata_dss_11
}
multipath {
        wwid            3600a098038304173452b4c795961466a
        alias           oradata_dss_12
}
multipath {
        wwid            3600a098038304173475d4c766a497630
        alias           oradata_dss_13
}
multipath {
        wwid            3600a098038304173452b4c795961466b
```

```
                alias           oradata_dss_14
        }
        multipath {
                wwid            3600a098038304173475d4c766a497631
                alias           oradata_dss_15
        }
        multipath {
                wwid            3600a098038304173452b4c795961466c
                alias           oradata_dss_16
        }
        multipath {
                wwid            3600a098038304173475d4c766a497632
                alias           oraredo_dss_1
        }
        multipath {
                wwid            3600a098038304173452b4c795961466d
                alias           oraredo_dss_2
        }
}
```

## Configuration of "/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

vm.nr_hugepages=180000
```

### Configuration of "/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    473985890

# 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    473985890
```

### Configuration of "/etc/udev/rules.d/99-oracle-asmdevices.rules"

```
#All LUNs which starts with dg_orarac_* #

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

#All LUNs which starts with dg_oradata_* #

ENV{DM_NAME}=="oradata_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"

#All LUNs which starts with dg_oraredo_* #

ENV{DM_NAME}=="oraredo_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"

#All LUNs which starts with dg_oraarchive_* #

ENV{DM_NAME}=="oraarchive_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"
```

# References

The following references were used in preparing this document.

## Cisco Unified Computing System

https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html

## Cisco UCS Data Center Design Guides

https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/data-center-design-guides-all.html#Hyperconverged

## FlexPod Converged Infrastructure

https://www.cisco.com/c/en/us/solutions/data-center-virtualization/flexpod/index.html#~tab-resources

https://www.netapp.com/us/products/converged-systems/flexpod-converged-infrastructure.aspx

## Cisco UCS B200 M5 Servers

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/datasheet-c78-739296.html

## Oracle Database 12c Release 2

https://docs.oracle.com/en/database/oracle/oracle-database/12.2/index.html

## NetApp AFF A-Series All Flash Storage

https://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx

https://www.netapp.com/us/media/ds-3582_tcm10-121337.pdf

https://hwu.netapp.com/Controller/Index

https://www.netapp.com/us/products/storage-systems/disk-shelves-and-storage-media/index.aspx#tech-specs

## NetApp Support

https://mysupport.netapp.com/

For a comprehensive list of storage hardware specifications, including supported media type and capacities, see the NetApp Hardware Universe.

# 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 the 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 the Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group for configuring, implementing and validating infrastructure best practices for highly available Oracle RAC databases solutions on Cisco UCS Servers, Cisco Nexus Products and various storage technologies. Hardikkumar Vyas holds a Master's degree in Electrical Engineering and has over 5 years of experience working with Oracle RAC Databases and associated applications. Hardikkumar Vyas's main focus is developing database solutions on different platforms, perform benchmarks, prepare reference architectures and write technical documents for Oracle RAC Databases on Cisco UCS Platforms.

## Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

James Bradshaw, NetApp