

# FlashStack Data Center with Oracle RAC 12c on Oracle Linux

Deployment Guide for Oracle Database 12c RAC with Oracle Linux 7.2 on Cisco UCS and Pure Storage FlashArray//m Series

**Last Updated:** May 26, 2017



## About Cisco Validated Designs

---

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

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

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

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

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)

© 2017 Cisco Systems, Inc. All rights reserved.

# Table of Contents

About Cisco Validated Designs .....	2
Executive Summary .....	6
Solution Overview .....	7
Introduction .....	7
Audience .....	7
Goals and Objectives of this Document .....	7
FlashStack System Overview .....	7
FlashStack Solution Benefits .....	8
Solution Summary .....	10
Oracle Database 12c R1 RAC on FlashStack .....	10
FlashStack Oracle RAC Solution Design Principles .....	11
Technology Overview .....	12
Cisco Unified Computing System .....	12
Cisco Unified Computing System Components .....	13
Cisco UCS 6300 Series Fabric Interconnects .....	13
Cisco UCS 5100 Series Blade Server Chassis .....	13
Cisco UCS 2300 Series Fabric Extender .....	13
Cisco UCS B-Series Blade Servers .....	14
Cisco UCS Adapters .....	14
Cisco UCS Manager .....	14
Cisco Nexus 9000 Series Switches .....	14
Cisco MDS 9100 Series Multilayer Fabric Switches .....	14
Cisco UCS Manager .....	14
Cisco UCS Service Profile .....	15
Cisco UCS 6300 Unified Fabric Interconnects .....	18
Cisco UCS 5108 Blade Server Chassis .....	19
Cisco UCS 2304XP Fabric Extenders .....	20
Cisco UCS Blade Servers .....	22
Cisco UCS Virtual Interface Card 1340 .....	24
Cisco UCS Virtual Interface Card 1380 .....	25
Cisco Nexus 9372PX-E Switches .....	25
Cisco MDS 9148S 16G FC Switches .....	26
Pure Storage FlashArray //m .....	27
Oracle 12c (12.1.0.2) Database .....	32
Oracle Linux 7.2 .....	33
Solution Design .....	34
Design Topology .....	34
Physical Topology .....	35

Deployment Hardware and Software .....	38
Cisco UCS Configuration Overview.....	38
High Level Steps to Configure Cisco Unified Computing System .....	38
Configure Fabric Interconnects for a Cluster Setup .....	38
Configure Fabric Interconnects for Chassis and Blade Discovery .....	42
Configure LAN and SAN on Cisco UCS Manager.....	45
Configure UUID, MAC, WWNN and WWPN Pools .....	48
Set Jumbo Frames in both the Cisco Fabric Interconnect .....	52
Configure vNIC and vHBA Template .....	52
Configure Ethernet Uplink Port-Channel.....	55
Create Server Boot Policy for SAN Boot.....	56
Configure and Create a Service Profile Template.....	66
Create Service Profiles from Template and Associate to Servers .....	79
Configure Cisco Nexus 9372PX-E Switches .....	82
Configure Global Settings for Cisco Nexus A and Cisco Nexus B .....	82
Configure VLANs for Cisco Nexus A and Cisco Nexus B .....	83
Virtual Port Channel (vPC) Summary for Data and Storage Network .....	84
Create vPC Peer-Link Between Two Nexus Switches .....	84
Create vPC Configuration between Nexus 9372PX-E and Fabric Interconnects .....	86
Verify vPC Status on both Cisco Nexus 9372PX-E Switches .....	89
Configure Cisco MDS 9148S Switches .....	91
Configure Features for MDS Switch A and MDS Switch B .....	93
Configure VSANs for MDS Switch A and MDS Switch B .....	94
Create and Configure Fiber Channel Zoning .....	95
Create Device Aliases for Fiber Channel Zoning .....	96
Create Zoning .....	97
Configure Pure Storage.....	100
OS and Database Deployment .....	105
Operating System Configuration .....	105
Miscellaneous Post-Installation Steps.....	108
Cloning a Linux Host with SAN Boot on Pure Storage .....	111
Pre-clone Steps on the Source Linux Host .....	111
Clone the boot LUN on the Pure Storage .....	113
Post-clone Steps on the Source Linux Host.....	114
Gather vHBA Details for Setup .....	114
Volume Setup on Pure Storage for the Clone .....	115
Target Host Setup.....	118
Configure Host Groups .....	120
Configure CRS, Data and Redo Log Volumes .....	122
Configure UDEV Rule .....	124



Oracle Database 12c GRID Infrastructure Deployment.....	124
Create the Local Directory Structure and Ownership on each RAC Node .....	125
Configure Public and Private NICs on each RAC Node .....	125
Configure “/etc/hosts” on each RAC Node .....	125
Establish SSH Connectivity.....	126
Run Cluster Verification Utility.....	127
Configure HugePages.....	127
Install and Configure Oracle Database Grid Infrastructure Software .....	127
Install Oracle Database Software.....	136
Scalability Test and Results.....	142
Hardware Calibration using FIO on FlashArray //m20 .....	142
Database Creation with DBCA .....	142
Database Workload Configuration .....	143
SLOB Performance on FlashArray //m20 .....	143
User Scalability Performance on FlashArray //m20.....	143
Node Scalability Performance on FlashArray //m20.....	145
SwingBench Performance on FlashArray //m20.....	146
Scalability Performance on FlashArray //m20 .....	146
FlashStack Scalability .....	154
Hardware Calibration using FIO on FlashArray //m70 .....	156
SLOB Performance on FlashArray //m70 .....	157
User Scalability Performance on FlashArray //m20.....	157
SwingBench Performance on FlashArray //m70.....	159
Scalability Performance on FlashArray //m70 .....	159
Resiliency and Failure Tests .....	166
Summary .....	183
Appendix.....	184
About the Authors .....	199
Acknowledgements .....	199



## Executive Summary

---

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

This Cisco Validated Design (CVD) describes a FlashStack reference architecture for deploying a highly available Oracle Database environment on Pure Storage FlashArray//m using Cisco UCS compute servers, Cisco MDS Switches, Cisco Nexus Switches and Oracle Linux. Cisco and Pure Storage have validated the reference architecture with OLTP and Data Warehouse workload in Cisco's lab. This document presents the hardware and software configuration of the components involved, results of various tests and offers implementation and best practices guidance.

FlashStack is designed to increase IT responsiveness to business demands while reducing the overall cost of computing. FlashStack components are integrated and standardized to help you achieve timely, repeatable, consistent deployments.

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

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

- Leverage a pre-validated platform to minimize business disruption and improve IT agility and reduce deployment time from months to weeks.
- Guarantee customer success with prebuilt, pre-tested drivers, Oracle database software
- A cohesive, integrated system that is managed, serviced and tested as a whole
- Slash administration time and total cost of ownership (TCO) by up to 50 percent

# Solution Overview

---

## Introduction

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

This CVD describes how Cisco UCS can be used in conjunction with Pure Storage FlashArray//m systems to implement an Oracle Real Application Clusters (RAC) 12c solution.

## Audience

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

## Goals and Objectives of this Document

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

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

Here are the objectives we would like to accomplish in this reference architecture document:

1. Build, validate and predict performance of Server, Network and Storage platform on a per workload basis
2. Seamless scalability of performance and capacity to meet growth needs of Oracle Database
3. High availability of DB instances without performance compromise through software and hardware upgrades

We will demonstrate the scalability and performance aspects by running Swingbench and SLOB (Silly Little Oracle Benchmark) on OLTP (Online Transaction Processing) and DSS (Decision-support System) like benchmark with varying nodes, users and read/update workload characteristics.

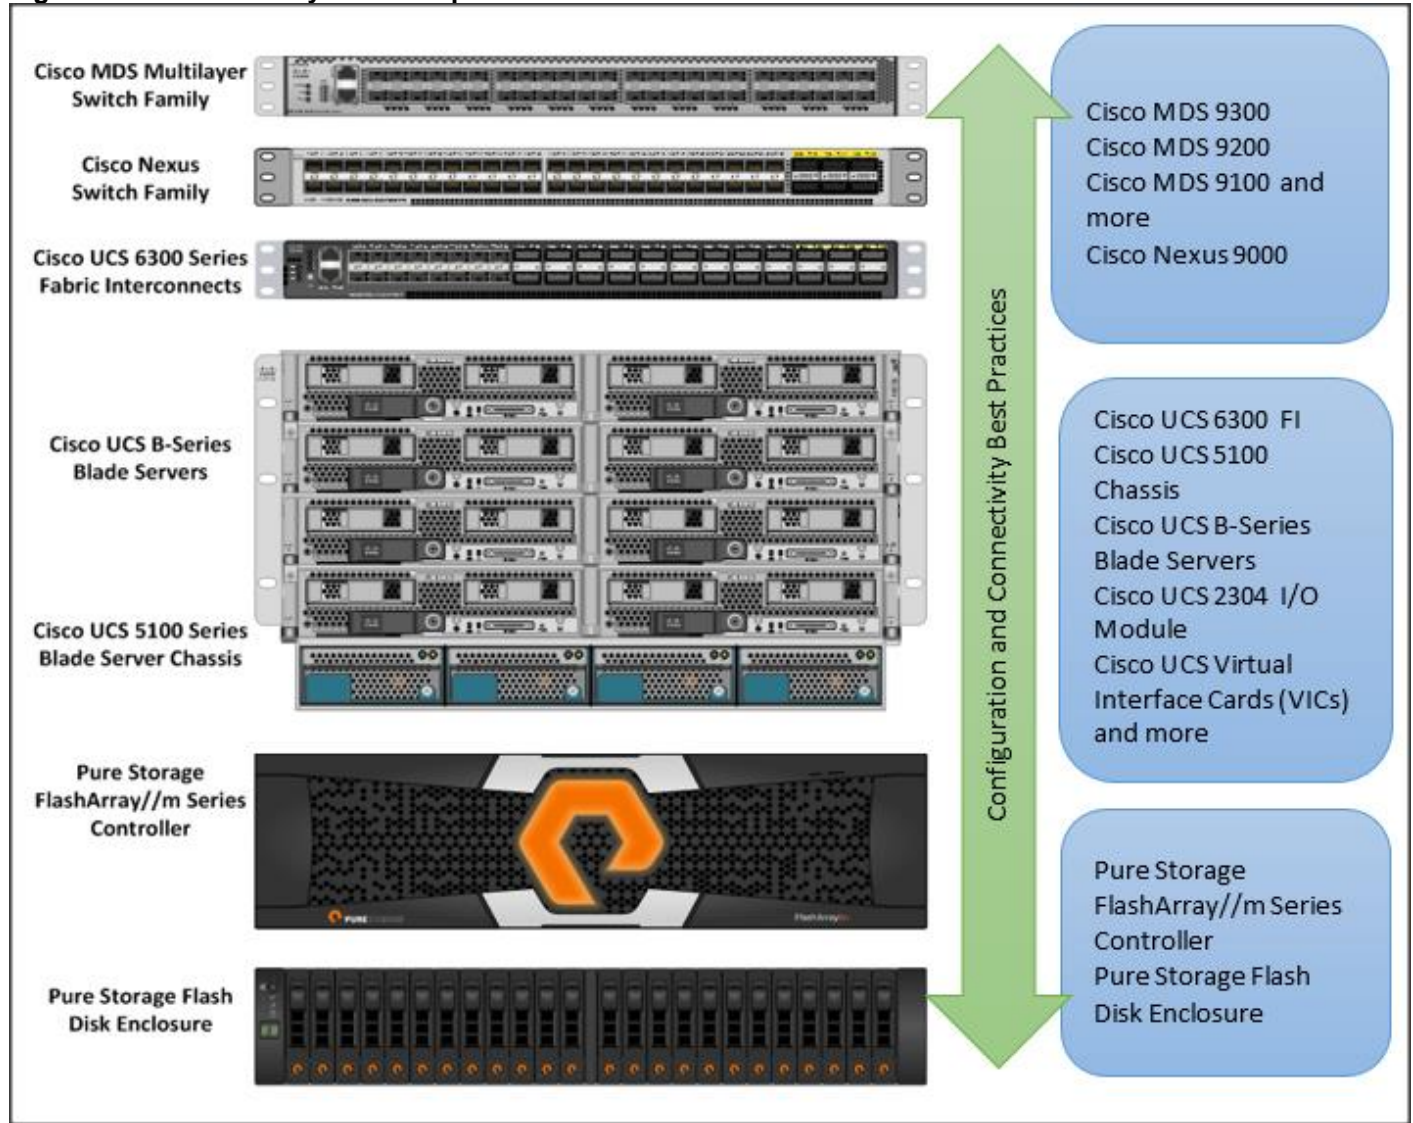
## FlashStack System Overview

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

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

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

**Figure 1 FlashStack System Components**



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

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

## FlashStack Solution Benefits

Key Benefits of the FlashStack solution are:

1. Consistent Performance and Scalability

- Consistent sub-millisecond latency with 100% flash storage. Consolidate 100s of enterprise-class applications in a single rack.
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines, and the capability to scale I/O bandwidth to match demand without disruption
- Repeatable growth through multiple FlashStack CI deployments.

2. Operational Simplicity

- Fully tested, validated, and documented for rapid deployment
- Reduced management complexity
- No storage tuning or tiers necessary
- Auto-aligned 512b architecture eliminates storage alignment headaches

3. Lowest TCO

- Dramatic savings in power, cooling and space with Cisco UCS and 100% Flash Industry leading data reduction

4. Enterprise Grade Resiliency

- Highly available architecture and redundant components
- Non-disruptive operations
- Upgrade and expand without downtime or performance loss
- Native data protection: snapshots and replication

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

## Solution Summary

### Oracle Database 12c R1 RAC on FlashStack

The FlashStack Data Center with Oracle RAC on Oracle Linux solution provides an end-to-end architecture with Cisco UCS, Oracle, and Pure Storage technologies and demonstrates the FlashStack configuration benefits for running Oracle Database 12c RAC with Cisco VICs (Virtual Interface Cards).

This section describes the design considerations for the Oracle Database 12c RAC on FlashStack deployment. The following table lists the inventory of the components used in the FlashStack solution. In this solution design, we used two chassis with 8 identical Intel CPU based Cisco UCS B-series B200 M4 blade servers for hosting the 8-node Oracle RAC database.

**Table 1 Inventory and Bill of Material**

Vendor	Name	Version / Model	Description	Qty
Cisco	Cisco Nexus 9372PX-E Switch	N9K-C9372PX-E	Cisco Nexus 9000 Series Switches	2
Cisco	Cisco MDS 9148S 16G Fabric Switch	DS-C9148S-12PK9	Cisco MDS 9100 Series Multilayer Fabric Switches	2
Cisco	Cisco UCS 6332-16UP Fabric Interconnect	UCS-FI-6332-16UP	Cisco 6300 Series Fabric Interconnects	2
Cisco	Cisco UCS Fabric Extender	UCS-IOM-2304	UCS 2304XP I/O Module (4 External, 8 Internal 40Gb Ports)	4
Cisco	Cisco UCS 5108 Blade Server Chassis	UCSB-5108-AC2-UPG	Cisco UCS 5100 Series Blade Server Chassis	2
Cisco	Cisco UCS B200 M4 Blade Servers	UCSB-B200-M4	UCS B-Series Blade Servers	8
Cisco	Cisco UCS VIC 1340	UCSB-MLOM-40G-03	Cisco UCS VIC 1340 modular LOM for blade servers	8
Cisco	Cisco UCS VIC 1380	UCSB-VIC-M83-8P	Cisco UCS VIC 1380 mezzanine adapter for blade servers	8
Pure Storage	Pure FlashArray //m20 Controller	Purity 4.5.12	Pure Storage FlashArray	1
Pure Storage	Pure FlashArray //m70 Controller	Purity 4.5.12	Pure Storage FlashArray	1



Vendor	Name	Version / Model	Description	Qty
Pure Storage	Pure FlashArray Disk Enclosure	Purity 4.5.12	Flash Storage	2

## FlashStack Oracle RAC Solution Design Principles

The FlashStack for Oracle RAC solution addresses the following primary design principles:


- **Repeatable:** Create a scalable building block that can be easily replicated at any customer site. Publish the version of various firmware under test and weed out any issues in the lab before customers deploy this solution.
- **Available:** Create a design that is resilient and not prone to failure of a single component. For example, we include best practices to enforce multiple paths to storage, multiple NICs for connectivity, and high availability (HA) clustering with the use of Oracle RAC.
- **Efficient:** Take advantage of inline data reduction, higher bandwidth and low latency of the Pure Storage FlashArray//m used in the FlashStack solution.
- **Simple:** Avoid unnecessary or complex tweaks to make the results look better than a normal out-of-box environment.
- **Scalable:** By reporting the linear scaling of Oracle databases within the FlashStack architecture and best-in-class flash storage performance.

## Technology Overview

This section provides a technical overview of products used in this solution.

### Cisco Unified Computing System

**Figure 2 Cisco UCS System**



The Cisco Unified Computing System™ is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 and 40 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Cisco UCS is a next-generation solution for blade and rack server computing.

The Cisco UCS unites the following main components:

- Computing

The system is based on an entirely new class of computing system that incorporates rack mount and blade servers based on Intel Xeon Processor E5 and E7.

- Network

The system is integrated onto a low-latency, lossless, 10 and 40-Gbps unified network fabric. This network foundation consolidates LAN, SAN and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.

- Virtualization

The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.

- Storage Access

The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access, the Cisco Unified Computing System can access storage over Ethernet (NFS or iSCSI), Fibre Channel over Ethernet (FCoE) and Fibre Channel (FC). This provides customers with choice for storage access and investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.

- Management

The system uniquely integrates all system components to enable the entire solution to be managed as a single entity by the Cisco UCS Manager. The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a powerful scripting library module for Microsoft PowerShell built on a robust application programming interface (API) to manage all system configuration and operations.

Cisco UCS fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability. Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and non-virtualized systems.

## Cisco Unified Computing System Components

### Cisco UCS 6300 Series Fabric Interconnects

(<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6300-series-fabric-interconnects/index.html>)

Cisco UCS 6300 Series Fabric Interconnects provides line-rate, low-latency, lossless, 10 and 40-Gigabit Ethernet (varies by model), Fibre Channel over Ethernet (FCoE) and Fibre Channel (FC). Cisco UCS 6300 Series Fabric provides management and communication backbone for Cisco UCS B-Series Blade Servers, Cisco UCS 5100 Series Blade Server Chassis, and Cisco UCS C-Series Rack Servers.

### Cisco UCS 5100 Series Blade Server Chassis

(<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-5100-series-blade-server-chassis/index.html>)

Cisco UCS 5108 Blade Server Chassis is a 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.

### Cisco UCS 2300 Series Fabric Extender

(<http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-675243.html>)

Cisco UCS 2300 series Fabric Extender brings the unified fabric into the blade server enclosure, providing multiple 10 and 40 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

## Cisco UCS B-Series Blade Servers

(<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>)

Based on Intel® Xeon® processor E7 and E5 product families, Cisco UCS B-Series Blade Servers work with virtualized and non-virtualized applications to increase:

- Performance
- Energy efficiency
- Flexibility
- Administrator productivity

## Cisco UCS Adapters

(<http://www.cisco.com/c/en/us/products/interfaces-modules/unified-computing-system-adapters/index.html>)

The Cisco Unified Computing System supports Converged Network Adapters (CNAs) obviate the need for multiple network interface cards (NICs) and host bus adapters (HBAs) by converging LAN and SAN traffic in a single interface.

## Cisco UCS Manager

(<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>) Streamline many of your most time-consuming daily activities, including configuration, provisioning, monitoring, and problem resolution with Cisco UCS Manager. It reduces TCO and simplifies daily operations to generate significant savings.

## Cisco Nexus 9000 Series Switches

(<http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>)

The 9000 Series offers modular 9500 switches and fixed 9300 and 9200 switches with 1/10/25/50/40/100 Gigabit Ethernet switch configurations. 9200 switches are optimized for high-performance and density in NX-OS mode operations.

## Cisco MDS 9100 Series Multilayer Fabric Switches

(<http://www.cisco.com/c/en/us/products/storage-networking/mds-9100-series-multilayer-fabric-switches/index.html>)

The Cisco MDS 9100 Series Multilayer Fabric Switches consists of Cisco MDS 9148S, a 48-port, 16 Gbps Fibre Channel switch.

## Cisco UCS Manager

Cisco UCS Manager provides unified, centralized, embedded management of all Cisco UCS software and hardware components across multiple chassis and thousands of virtual machines. Administrators use the software to manage the entire Cisco UCS as a single logical entity through an intuitive GUI, a command-line interface (CLI), or an XML API.

Cisco UCS Manager manages Cisco UCS systems through an intuitive HTML 5 or Java user interface and a command-line interface (CLI) enabling centralized management of distributed systems scaling to thousands of servers. Cisco UCS Manager is embedded on a pair of Cisco UCS 6300 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The manager gives administrators a single interface for performing server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection.

Cisco UCS management software provides a model-based foundation for streamlining the day-to-day processes of updating, monitoring, and managing computing resources, local storage, storage connections, and network connections. By enabling better automation of processes, Cisco UCS Manager allows IT organizations to achieve greater agility and scale in their infrastructure operations while reducing complexity and risk.

Cisco UCS Manager provides an easier, faster, more flexible, and unified solution for managing firmware across the entire hardware stack than traditional approaches to server firmware provisioning. Using service profiles, administrators can associate any compatible firmware with any component of the hardware stack. After the firmware versions are downloaded from Cisco, they can be provisioned within minutes on components in the server, fabric interconnect, and fabric extender based on the required network, server, and storage policies for each application and operating system. The firmware's auto-installation capability simplifies the upgrade process by automatically sequencing and applying upgrades to individual system elements.

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

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

Cisco UCS Manager provides end-to-end management of all the devices in the Cisco UCS domain it manages. Devices that are uplinked from the fabric interconnect must be managed by their respective management applications.

Cisco UCS Manager is provided at no additional charge with every Cisco UCS platform.

For more information on Cisco UCS Manager, visit:

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

## Cisco UCS Service Profile

Service profiles are essential to the automation functions in Cisco UCS Manager. They provision and manage Cisco UCS systems and their I/O properties within a Cisco UCS domain. Infrastructure policies are created by server, network, and storage administrators and are stored in the Cisco UCS Fabric Interconnects. The infrastructure policies needed to deploy applications are encapsulated in the service profiles templates, which are collections of policies needed for the specific applications. The service profile templates are then used to create one or more service profiles, which provide the complete definition of the server. The policies coordinate and automate element management at every layer of the hardware stack, including RAID levels, BIOS settings, firmware revisions and settings, server identities, adapter settings, VLAN and VSAN network settings, network quality of service (QoS), and data center connectivity.

A server's identity is made up of many properties such as UUID, boot order, IPMI settings, BIOS firmware, BIOS settings, RAID settings, disk scrub settings, number of NICs, NIC speed, NIC firmware, MAC and IP addresses, number of HBAs, HBA WWNs, HBA firmware, FC fabric assignments, QoS settings, VLAN assignments, remote keyboard/video/monitor etc. I think you get the idea. It's a LONG list of "points of configuration" that need to be configured to give this server its identity and make it unique from every other server within your data center. Some of these parameters are kept in the hardware of the server itself (like BIOS firmware version, BIOS settings, boot order, FC boot settings, etc.) while some settings are kept on your network and storage switches (like VLAN

assignments, FC fabric assignments, QoS settings, ACLs, etc.). This results in following server deployment challenges:

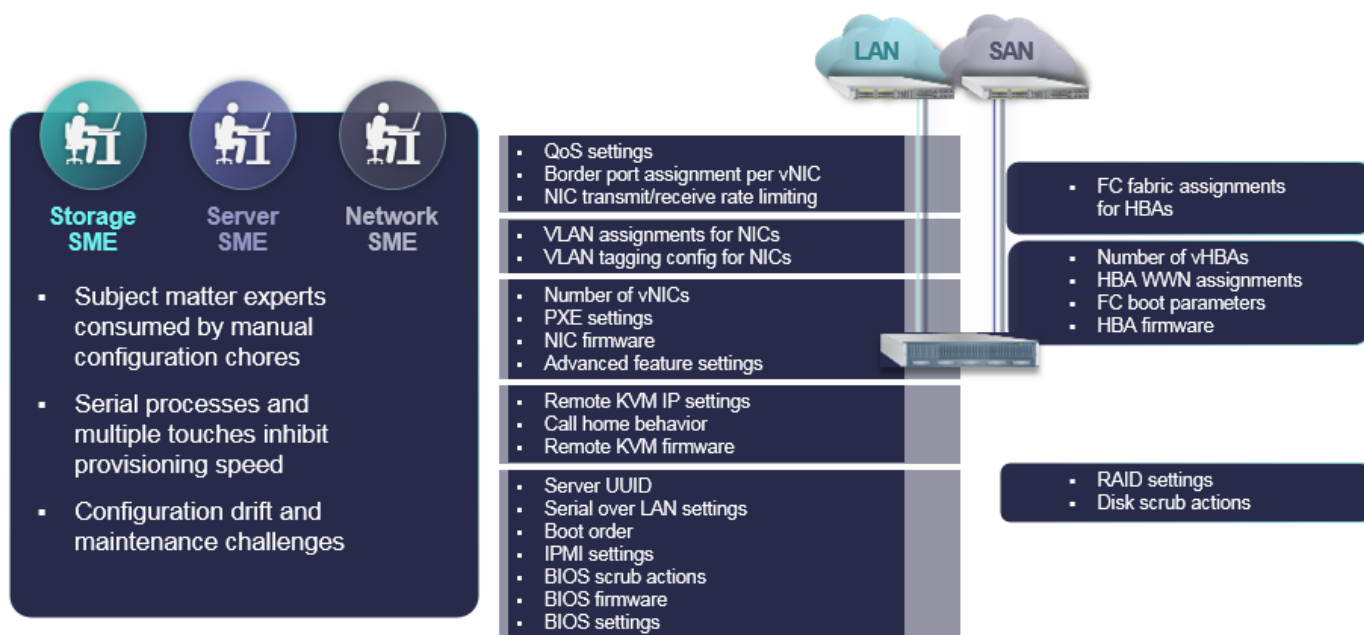
- Every deployment requires coordination among server, storage, and network teams
- Need to ensure correct firmware and settings for hardware components
- Need appropriate LAN and SAN connectivity

The service profile consists of a software definition of a server and the associated LAN and SAN connectivity that the server requires. When a service profile is associated with a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile. Service profiles improve IT productivity and business agility because they establish the best practices of your subject-matter experts in software. With service profiles, infrastructure can be provisioned in minutes instead of days, shifting the focus of IT staff from maintenance to strategic initiatives. Service profiles enable pre-provisioning of servers, enabling organizations to configure new servers and associated LAN and SAN access settings even before the servers are physically deployed.

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. By abstracting these settings from the physical server into a Cisco Service Profile, the Service Profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, Service Profiles can, at any time, be migrated from one physical server to another. This logical abstraction of the server personality separates the dependency of the hardware type or model and is a result of Cisco's unified fabric model (rather than overlaying software tools on top).

**Figure 3 Traditional Components**

## Traditional Element Configuration



### Compute, LAN, SAN Seamlessly Through Software

Service profiles benefit both virtualized and non-virtualized environments. Workloads may need to be moved from one server to another to change the hardware resources assigned to a workload or to take a server offline for maintenance. Service profiles can be used to increase the mobility of non-virtualized servers. They also can be used in conjunction with virtual clusters to bring new resources online easily, complementing existing virtual



machine mobility. Service profiles are also used to enable Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) capabilities for servers that will run hypervisors enabled for VM-FEX.

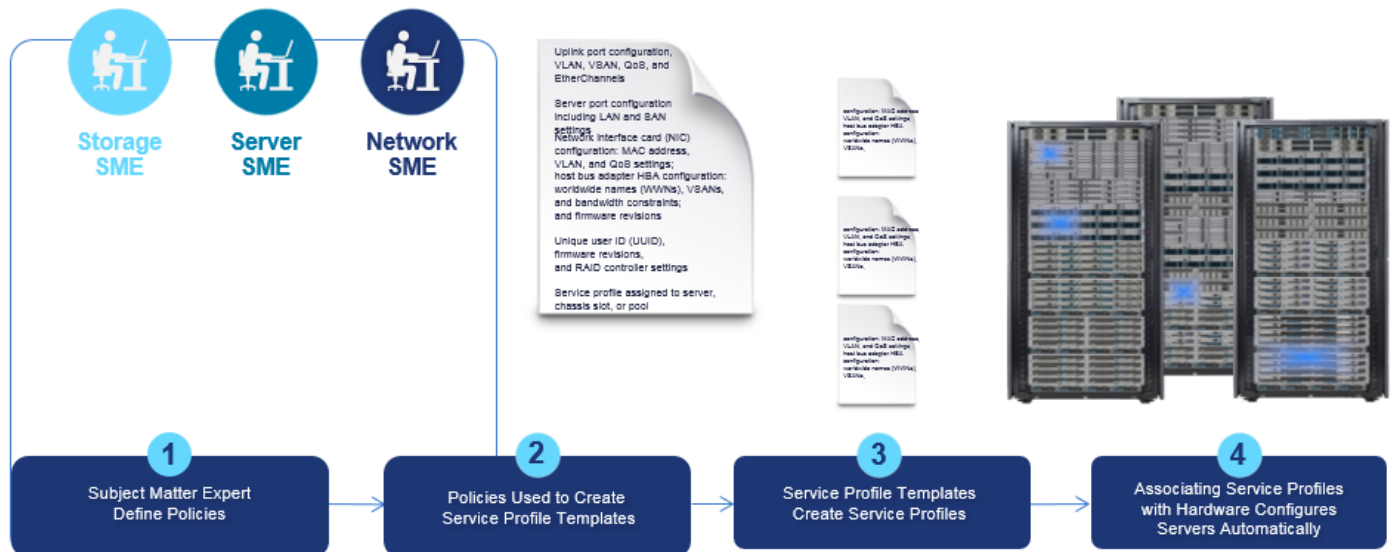
Cisco UCS has uniquely addressed these challenges with the introduction of service profiles that enables integrated, policy based infrastructure management. Cisco UCS Service Profiles hold the DNA for nearly all configurable parameters required to set up a physical server. A set of user defined policies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

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

**Figure 4 Cisco UCS Management**

## UCS: Embedded Automation

Integrated, Policy-Based Infrastructure Management



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

- **Service profiles and templates.** Service profile templates are stored in the Cisco UCS 6300 Series Fabric Interconnects for reuse by server, network, and storage administrators. Service profile templates consist of server requirements and the associated LAN and SAN connectivity. Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools.

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

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

Service profile templates are used to simplify the creation of new service profiles, helping ensure consistent policies within the system for a given service or application. Whereas a service profile is a description of a logical server and there is a one-to-one relationship between the profile and the physical server, a service

profile template can be used to define multiple servers. The template approach enables you to configure hundreds of servers with thousands of virtual machines as easily as you can configure one server. This automation reduces the number of manual steps needed, helping reduce the opportunities for human error, improve consistency, and further reducing server and network deployment times.

- **Programmatically deploying server resources.** Cisco UCS Manager provides centralized management capabilities, creates a unified management domain, and serves as the central nervous system of the Cisco UCS. Cisco UCS Manager is embedded device management software that manages the system from end-to-end as a single logical entity through an intuitive GUI, CLI, or XML API. Cisco UCS Manager implements role- and policy-based management using service profiles and templates. This construct improves IT productivity and business agility. Now infrastructure can be provisioned in minutes instead of days, shifting IT's focus from maintenance to strategic initiatives.
- **Dynamic provisioning.** Cisco UCS resources are abstract in the sense that their identity, I/O configuration, MAC addresses and WWNs, firmware versions, BIOS boot order, and network attributes (including QoS settings, ACLs, pin groups, and threshold policies) all are programmable using a just-in-time deployment model. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, enabling flexibility in the use of system resources. Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools.

## Cisco UCS 6300 Unified Fabric Interconnects

The Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6300 Series offers line-rate, low-latency, lossless 10 and 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

**Figure 5 Cisco UCS 6300 Series Fabric Interconnect**



The Cisco UCS 6300 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers, 5100 Series Blade Server Chassis, and C-Series Rack Servers managed by Cisco UCS. All servers attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6300 Series provides both LAN and SAN connectivity for all servers within its domain.

From a networking perspective, the Cisco UCS 6300 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 and 40 Gigabit Ethernet ports, switching capacity of 2.56 terabits per second (Tbps), and 320 Gbps of bandwidth per chassis, independent of packet size and enabled services. The product family supports Cisco® low-latency, lossless 10 and 40 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the server through the fabric interconnect. Significant TCO savings can be achieved with an FCoE optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

### Lower Total Cost of Ownership


The 6300 Series offers several key features and benefits that can lower TCO. Some examples include:

- Bandwidth up to 2.56 Tbps
- Centralized unified management with Cisco UCS Manager software

### Highly Scalable Architecture

Cisco Fabric Extender technology scales up to 20 chassis in just one unified system without additional complexity. The result is that customers can eliminate dedicated chassis management and blade switches, as well as reduce cabling.

**Figure 6 Cisco UCS 6332-16UP Fabric Interconnect**





For this Oracle 12c RAC solution we used FI 6332-16UP. As shown in **Error! Reference source not found.**, FI 6332-16UP is a one-rack-unit (1RU) 40 Gigabit Ethernet/FCoE switch and 1/10 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 2.24 Tbps throughput and up to 40 ports. The switch has 24 40Gbps fixed Ethernet/FCoE ports and 16 1/10Gbps Ethernet/FCoE or 4/8/16G Fiber Channel ports. This Fabric Interconnect is targeted for FC storage deployments requiring high performance 16G FC connectivity to MDS switches.

## 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.

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

**Figure 7 Cisco UCS Blade Server Chassis Front View and Rear View**




Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2200 and 2300 Series Fabric Extenders.

## Cisco UCS 2304XP Fabric Extenders

Cisco UCS 2304XP 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. It is a third-generation I/O module (IOM) that shares the same form factor as the second-generation Cisco UCS 2200 Series Fabric Extenders and is backward compatible with the shipping Cisco UCS 5108 Blade Server Chassis.




**Figure 8 Cisco UCS 2304XP Fabric Extender**



The Cisco UCS 2304 Fabric Extender has four 40 Gigabit Ethernet, FCoE-capable, Quad Small Form-Factor Pluggable (QSFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2304 has four 40 Gigabit Ethernet ports connected through the mid-plane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 320 Gbps of I/O to the chassis.

**Figure 9 Rear of Cisco UCS 5108 Blade Server Chassis with Two Cisco UCS 2304 Fabric Extenders Inserted**



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. Because the fabric extender is similar to a distributed line card, it does not perform any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity and enabling Cisco UCS to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

## Cisco UCS Blade Servers

Cisco UCS offers a variety of x86-based compute portfolio to address the needs of today’s workloads. Based on Intel® Xeon® processor E7 and E5 product families, Cisco UCS B-Series Blade Servers work with virtualized and non-virtualized applications to increase:









- Performance
- Energy efficiency
- Flexibility
- Administrator productivity

Figure 10 illustrates a complete summary of fourth generation Cisco UCS compute portfolio featuring Blade and Rack-Mount Servers.

**Figure 10 Cisco UCS Compute Portfolio**

### UCS Compute Portfolio

Performance Optimized for Bare Metal, Virtualized, and Cloud Applications

Cloud Scale	Enterprise Performance	Intensive/Mission Critical
 <p><b>UCS C3000 Series</b> Ideal Capacity-Optimized Platform for Large Object Storage at Scale</p>	 <p><b>UCS C240 M4</b> Ideal Platform for Big Data, ERP, and Database Applications</p>  <p><b>UCS C220 M4</b> Versatile, General Purpose Enterprise Infrastructure, and Application Server</p>  <p><b>UCS B200 M4</b> Optimal Choice for VDI, Private Cloud, or Dense Virtualization/Consolidation Workloads</p>	 <p><b>UCS C460 M4</b> Mission-Critical, 4-Socket Server for Large, CPU-Intensive Applications</p> <hr/> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;">  <p><b>UCS B420 M4</b> Enterprise Class, 4-Socket Blade for Large, Memory-Intensive Bare Metal and Virtualized Applications</p> </div> <div style="text-align: center;">  <p><b>UCS B260 M4</b> Mission-Critical, 2-Socket Blade for Large, CPU-Intensive Bare Metal and Virtualized Applications</p> </div> <div style="text-align: center;">  <p><b>UCS B460 M4</b> Mission-Critical, 4-Socket Blade for Large, CPU-Intensive Bare Metal and Virtualized Applications</p> </div> </div>

Rack
Blade

For this Oracle 12c RAC solution, we used enterprise-class, Cisco UCS B200 M4 Blade Servers. The Cisco UCS B200 M4 Blade Server delivers record performance, expandability, and configurability for workloads ranging from web infrastructure to distributed databases. Optimized for data center or cloud, the Cisco UCS B200 M4 can quickly deploy stateless physical and virtual workloads, with the programmability of the Cisco UCS Manager.

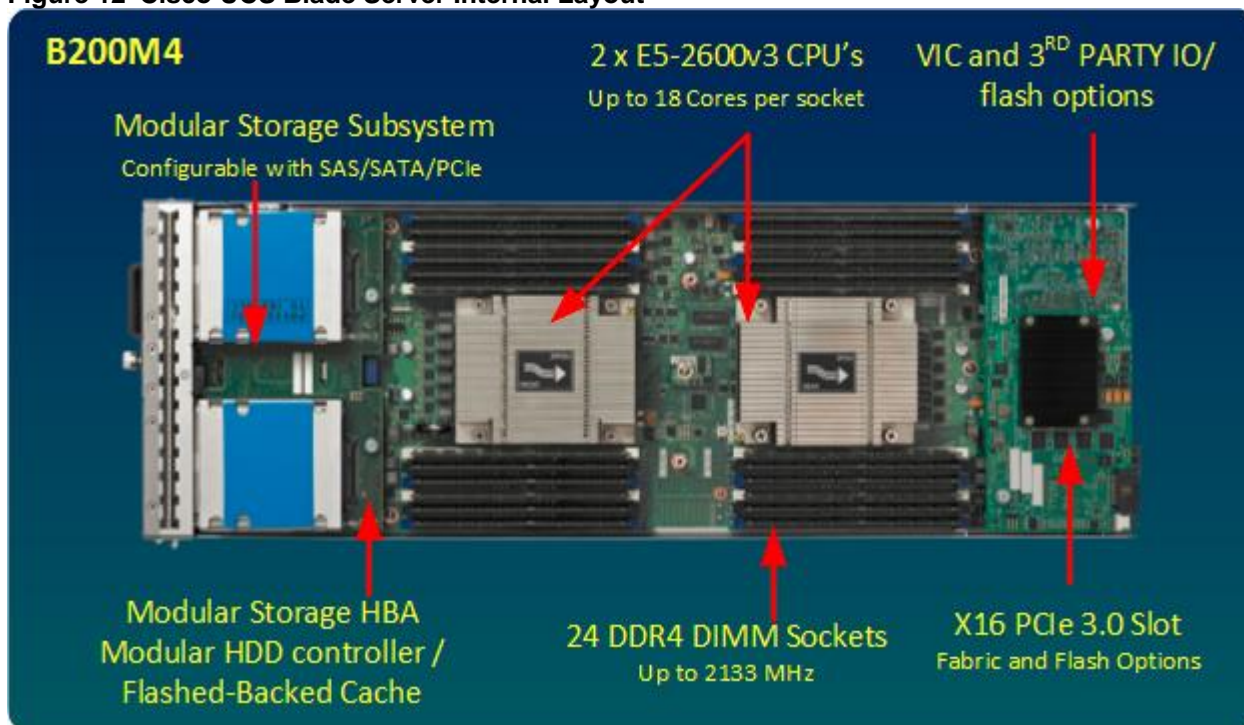
**Figure 11 Cisco UCS Blade Server B200 M4**





The Cisco UCS B200 M4 is built with the Intel® Xeon® E5-2600 v4 and v3 processor family, up to 1.5 TB of memory (with 64 GB DIMMs), up to two drives, and up to 80 Gbps total bandwidth. It offers exceptional levels of performance, flexibility, and I/O throughput to run the most demanding applications. In addition, Cisco UCS has the architectural advantage of not having to power and cool switches in each blade chassis. Having a larger power budget available for blades allows Cisco to design uncompromised expandability and capabilities in its blade servers.

**Figure 12 Cisco UCS Blade Server Internal Layout**



**Table 2 Cisco UCS B200 M4 Blade Feature**  
**Cisco UCS B200 M4 Key Features**

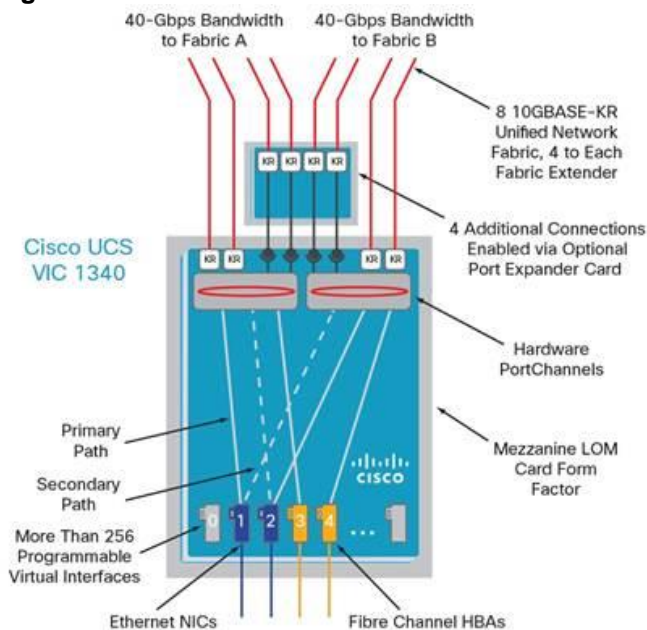
Item	Specification
------	---------------

<b>Form factor</b>	Half-width blade form factor
<b>Processors</b>	Either 1 or 2 Intel® Xeon® processor E5-2600 v4 and v3 product family CPUs
<b>Chipset</b>	Intel C610 series
<b>Memory</b>	Up to 24 double-data-rate 4 (DDR4) dual in-line memory (DIMMs) at 2400 and 2133 MHz speeds
<b>Mezzanine adapter slots</b>	2
<b>Hard drives</b>	Two optional, hot-pluggable, SAS, SATA hard disk drives (HDDs) or solid-state drives (SSDs)
<b>Maximum internal storage</b>	Up to 3.2 TB

## Cisco UCS Virtual Interface Card 1340

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 13), is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet.

**Figure 13 Cisco Virtual Interface Card 1340**



The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1340 supports Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.

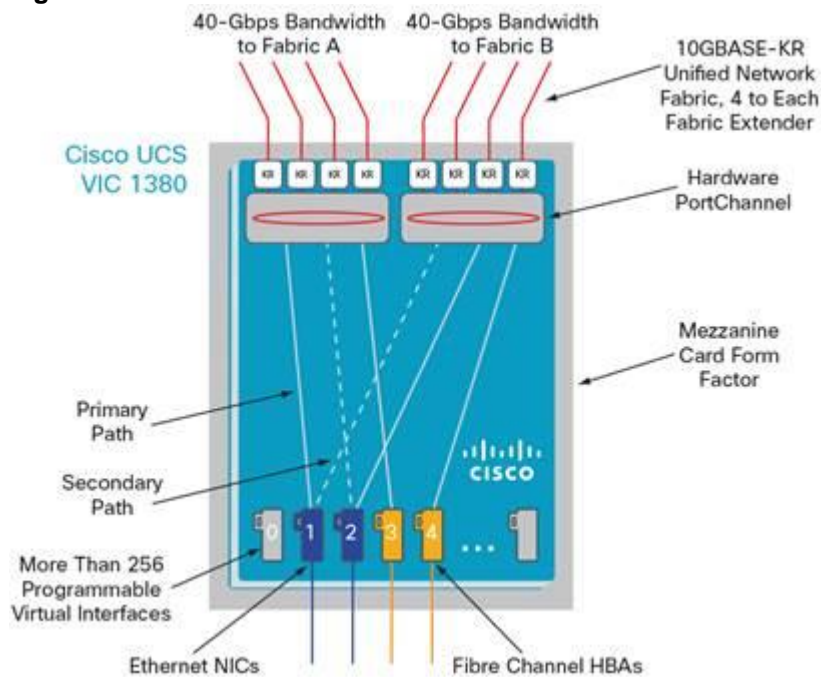
The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile.

The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure (Figure 13)

## Cisco UCS Virtual Interface Card 1380

The Cisco UCS Virtual Interface Card (VIC) 1380 is a dual-port 40-Gbps Ethernet, or dual 4 x 10 Fibre Channel over Ethernet (FCoE)-capable mezzanine card designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. The card enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1380 supports Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.

**Figure 14 Cisco Virtual Interface Card 1380**



The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure.

## Cisco Nexus 9372PX-E Switches

The Cisco Nexus 9372PX/9372PX-E Switches has 48 1/10-Gbps Small Form Pluggable Plus (SFP+) ports and 6 Quad SFP+ (QSFP+) uplink ports. All the ports are line rate, delivering 1.44 Tbps of throughput in a 1-rack-unit (1RU) form factor.

**Figure 15 Cisco Nexus 9372PX-E Switch**



The Cisco Nexus 9372PX switch benefits are listed below.

### Architectural Flexibility

- Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
- Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support

### Feature Rich

- Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
- ACI-ready infrastructure helps users take advantage of automated policy-based systems management
- Virtual Extensible LAN (VXLAN) routing provides network services
- Nexus 9372PX-E supports IP-based endpoint group (EPG) classification in ACI mode

### Simplified Operations

- An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
- Python Scripting for programmatic access to the switch command-line interface (CLI)

### Investment Protection

- A Cisco 40 Gb bidirectional transceiver allows for reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet
- Support for 1 Gb and 10 Gb access connectivity for data centers migrating access switching infrastructure to faster speeds

## Cisco MDS 9148S 16G FC Switches

The Cisco® MDS 9148S 16G Multilayer Fabric Switch (Figure 16) is the next generation of the highly reliable, flexible, and low-cost Cisco MDS 9100 Series switches. It combines high performance with exceptional flexibility and cost effectiveness. This powerful, compact one rack-unit (1RU) switch scales from 12 to 48 line-rate 16 Gbps Fibre Channel ports.

**Figure 16 Cisco MDS 9148S 16G FC Switch**



The Cisco MDS 9148S is excellent for:

- A standalone SAN in small departmental storage environments
- A top-of-the-rack switch in medium-sized redundant fabrics
- An edge switch in enterprise data center core-edge topologies

The Cisco MDS 9148S is powered by Cisco NX-OS and Cisco Prime™ Data Center Network Manager (DCNM) software. It delivers advanced storage networking features and functions with ease of management and compatibility with the entire Cisco MDS 9000 Family portfolio for reliable end-to-end connectivity.

The Cisco MDS 9148S features and benefits are as follows:

- Port speed: 2/4/8/16-Gbps autosensing with 16 Gbps of dedicated bandwidth per port
- Enhance reliability, speed problem resolution, and reduce service costs by using Fibre Channel ping and traceroute to identify exact path and timing of flows, as well as Cisco Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) and Cisco Fabric Analyzer to capture and analyze network traffic.
- Automate deployment and upgrade of software images.
- Reduce consumption of hardware resources and administrative time needed to create and manage zones.
- Optimize bandwidth utilization by aggregating up to 16 physical ISLs into a single logical Port-Channel bundle with multipath load balancing.

## Pure Storage FlashArray //m

Who knew that moving to all-flash storage could help reduce the cost of IT? FlashArray//m makes server and workload investments more productive, while also lowering storage spend. With FlashArray//m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient, accelerating your journey to the cloud.




**Figure 17 Pure Storage FlashArray //m**

- FlashArray//m's performance can also make your business smarter by unleashing the power of real-time analytics, driving customer loyalty, and creating new, innovative customer experiences that simply weren't possible with disk. All by Transforming Your Storage with FlashArray//m.
- FlashArray//m enables you to transform your data center, cloud, or entire business with an affordable all-flash array capable of consolidating and accelerating all your key applications.
- FlashArray//m features and benefits are as follows:
  - Mini Size—Reduce power, space and complexity by 90 percent
  - 3U base chassis with 15-136+ TBs usable
  - ~1kW of power
  - 6 cables
  - Mighty Performance—Transform your datacenter, cloud, or entire business
  - Up to 300,000 32K IOPS
  - Up to 9 GB/s bandwidth
  - <1ms average latency
  - Modular Scale—Scale FlashArray//m inside and outside of the chassis for generations
  - Expandable to ~½ PB usable via expansion shelves
  - Upgrade controllers and drives to expand performance and/or capacity
  - Meaningful Simplicity—Appliance-like deployment with worry-free operations
  - Plug-and-go deployment that takes minutes, not days
  - Non-disruptive upgrades and hot-swap everything
  - Less parts = more reliability
- The FlashArray//m expands upon the FlashArray's modular, stateless architecture, designed to enable expandability and upgradability for generations. The FlashArray//m leverages a chassis-based design with customizable modules, enabling both capacity and performance to be independently improved over time with advances in compute and flash, to meet your businesses needs today and tomorrow.

**Figure 18 Pure Storage FlashArray//m**



## Modular Upgradability for Generations



The Pure Storage FlashArray is ideal for:

- **Accelerating Databases and Applications** Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.
- **Virtualizing and Consolidating Workloads** Easily accommodate the most IO-hungry Tier-1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.
- **Delivering the Ultimate Virtual Desktop Experience** Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance for under \$100/desktop.
- **Protecting and Recovering Vital Data Assets** Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with Flash Recover.
- **Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:**
  - **Consistent Performance** FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.
  - **Less Cost** than Disk Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure.
  - **Mission-Critical Resiliency** FlashArray delivers >99.999 percent proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.
- **Disaster Recovery Built-In** FlashArray offers native, fully integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.
- **Simplicity Built-In** FlashArray offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray the way you like it: Web-based GUI, CLI, VMware vCenter, Rest API, or OpenStack.

Figure 19 Pure Storage FlashArray//m



Table 3 Pure Storage FlashArray//m

	//M10	//M20	//M50	//M70
<b>Capacity</b>	Up to 25 TBs effective capacity* 5 – 10TBs raw capacity	Up to 120+ TBs effective capacity* 5 – 40TBs raw capacity	Up to 250+ TBs effective capacity* 30 – 88TBs raw capacity	Up to 450+ TBs effective capacity* 44 – 136TBs raw capacity
<b>Performance</b>	Up to 100,000 32K IOPS** <1ms average latency Up to 3 GB/s bandwidth	Up to 150,000 32K IOPS** <1ms average latency Up to 5 GB/s bandwidth	Up to 220,000 32K IOPS** <1ms average latency Up to 7 GB/s bandwidth	Up to 300,000 32K IOPS** <1ms average latency Up to 9 GB/s bandwidth
<b>Connectivity</b>	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 1 Gb/s Management & Replication ports	8 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 10 Gb/s Replication ports 1 Gb/s Management ports	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 10 Gb/s Replication ports 1 Gb/s Management ports	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 10 Gb/s Replication ports 1 Gb/s Management ports
<b>Physical</b>	3U 610 Watts (nominal) 105 lbs (47.6 kg) 5.12" x 18.94" x 29.72" chassis	3U*** 742 Watts (nominal) 110 lbs (49.9 kg) fully loaded 5.12" x 18.94" x 29.72" chassis	3U – 7U 1007 - 1447 Watts (nominal) 110 lbs (49.9 kg) fully loaded+ 44 lbs per expansion shelf 5.12" x 18.94" x 29.72" chassis	5U – 11U 1439 – 2099 Watts (nominal) 110 lbs (49.9 kg) fully loaded+ 44 lbs per expansion shelf 5.12" x 18.94" x 29.72" chassis

Purity Operating Environment

Purity implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashArray//m.

- Storage Software Built for Flash — The FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity to both extend the life and ensure the maximum performance of consumer-grade MLC flash.
- Granular and Adaptive — Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity’s data and flash management services to operate at the highest efficiency.
- Best Data Reduction Available — FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte

aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.

- Highly Available and Resilient — FlashProtect implements high availability, dual-parity RAID-3D, non-disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.
- Backup and Disaster Recovery Built In — FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully integrated in the FlashArray and leverage the native data reduction capabilities.



- Pure1 Manage — By combining local web-based management with cloud-based monitoring, Pure1 Manage allows you to manage your FlashArray wherever you are – with just a web browser.
- Pure1 Connect — A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashArray//m to all your data center and cloud monitoring, management, and orchestration tools.
- Pure1 Support — FlashArray//m is constantly cloud- connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.
- Pure1 Collaborate — Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.

### Experience Evergreen Storage

Tired of the 3 to 5 year array replacement merry-go-round? The move to FlashArray//m can be your last data migration. Purchase and deploy storage once and once only – then expand capacity and performance incrementally in conjunction with your business needs and without downtime. Pure Storage’s vision for Evergreen Storage is delivered by a combination of the FlashArray’s stateless, modular architecture and the Forever Flash business model, enabling you to extend the lifecycle of storage from 3-5 years to a decade or more.

### Non-Disruptive Upgrades

In general, upgrading a storage device will involve some form of data migration driven from the host. In many cases, an upgrade may result in downtime to the application environment, adversely affecting business operations.

With Pure Storage, Non Disruptive Upgrades (NDU) are supported from any Pure Storage FA-400-series system to any Pure Storage FlashArray//m, and of course between models in the same family. Additionally, all of our customers running FA-300 systems can also follow the same process and upgrade directly to a FlashArray//m non-disruptively.

Compared to the migration based NDU method we’ve had to deal with in the past, the Pure Storage NDU process is simply some multipath failover and path recovery events for the hosts. The value in terms of efficiency and use of resources this returns to large organizations with many arrays to manage is immeasurable.

### Migration-based NDU method

Allocate new rack space  
 Install new equipment  
 Allocate more switch ports  
 Allocate more cables  
 Allocate more power  
 Configure zoning/VLANs  
 Provision new LUNs  
 Migrate data (in stages)



### Pure Storage Method

Hosts experience path failover  
 and path recovery



It is a paradigm shift in the way of thinking: that upgrading hardware to get better performance is easier to do without shutting anything down.

## Oracle 12c (12.1.0.2) Database

Oracle revolutionized the field of enterprise database management systems with most extensive self-management capabilities in the industry, ranging from zero-overhead instrumentation to integrate self-healing and business-driven management. Oracle Database 12c, the next generation of the world's most popular database, make DBA lives easier by providing a various feature like change and configuration management, patching, provisioning, testing, performance management and automatic tuning. Oracle Database high-availability (HA) technologies, collectively referred to as Oracle Maximum Availability Architecture (MAA), provide complete resiliency against all types of outages – from component failures to natural disasters. Industry-leading Oracle HA technology such as Oracle Real Application Clusters (Oracle RAC) provides the highest levels of server HA which Oracle Active Data Guard protects data and applications against site wide outages.

Oracle Multitenant is a new feature of Oracle Database 12c, and allows each database plugged into the new multitenant architecture to look and feel like a standard Oracle Database to applications; so existing applications can run unchanged. Oracle Database 12c multitenant architecture makes it easy to consolidate many databases quickly and manage them as a cloud service. The Oracle Database 12c release 12.0.1.2 features the Oracle Database 12c In-Memory, an optional add-on that provides in-memory capabilities. The in-memory option makes Oracle Database 12c the first Oracle database to offer real-time analytics. Additional database innovations deliver new levels of efficiency, performance, security, and availability.

Oracle Database 12c introduces a rich set of new or enhanced features. It's beyond the scope to document all the features of 12c in this solution but we will showcase some of the features such as multi-tenancy, consolidation and rapid provisioning in this solution.

Oracle Real Application Clusters (RAC) harnesses the processing power of multiple, interconnected servers on a cluster; allows access to a single database from multiple servers on a cluster, insulating both applications and database users from server failures, while providing performance that scales out on-demand and is a vital component of grid computing that allows multiple servers to access a single database at one time.

The FlashStack solution for Oracle includes the following Oracle 12c components:

- Oracle Database 12c Release 1 (12.1.0.2) Enterprise Edition
- Oracle Grid Infrastructure 12c (12.1.0.2)
- Oracle Automatic Storage Management (ASM) & ASM Cluster File System (ACFS)

## Oracle Linux 7.2

Oracle Linux, formerly known as Oracle Enterprise Linux, is a Linux distribution based on Red Hat Enterprise Linux (RHEL), repackaged and freely distributed by Oracle, available under the GNU General Public License (GPL) since late 2006. Oracle Linux can be downloaded through Oracle's E-delivery service or from a variety of mirror sites, and can be deployed and distributed freely.. Commercial technical support is available through Oracle's Oracle Linux Support program, which supports Oracle Linux, and existing RHEL or CentOS Installation.

Oracle Corporation distributes Oracle Linux with two alternative kernels:

- Red Hat Compatible Kernel (RHCK) – identical to the kernel shipped in Red Hat Enterprise Linux
- Unbreakable Enterprise Kernel (UEK) – based on newer mainline Linux kernel versions, with Oracle's own enhancements for OLTP, InfiniBand, and SSD disk access, NUMA-optimizations, Reliable Datagram Sockets (RDS), async I/O, OCFS2, and networking.

Oracle Linux Support Program provides support for KVM components as part of Oracle Linux 5, Oracle Linux 6, Oracle Linux 7, RHEL5, RHEL6 and RHEL7. This does not include Oracle Product support on KVM offerings.

Cisco submitted 2 TPC-C benchmark results that run Oracle Linux with the Unbreakable Enterprise Kernel R2 on Cisco Unified Computing System. For this FlashStack solution, Oracle Linux Server release 7.2 was used.

## Solution Design

### Design Topology

This section describes the design considerations for the Oracle Database 12c RAC on FlashStack deployment. In this solution design, we have used two chassis with 8 identical Intel CPU based Cisco UCS B-series B200 M4 blade servers for hosting the 8-node Oracle RAC database.

The server has Cisco UCS VIC 1340 and VIC 1380 cards and they were connected four ports from each Cisco Fabric extender of the Cisco UCS chassis to the Cisco Fabric Interconnect, which were in turn connected to the Cisco MDS 9148S for upstream connectivity to access the Pure Storage FlashArray//m LUNs. The server configuration is described in Table 4 .

**Table 4 Cisco UCS Blade Server Configuration**

Server Configuration	
<b>Processor</b>	2 x Intel Xeon E5-2697 V3 2.6 GHz (2 CPUs with 14 cores each)
<b>Memory</b>	256GB @ 2.1GHz (16 x 16GB) 16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v
<b>VIC (Virtual Interface Card) 1340</b>	2x ports - 40 Gbps Unified I/O ports on Cisco UCS VIC 1380 Delivers 80 Gbps to the server
<b>VIC (Virtual Interface Card) 1380</b>	2x ports - 40 Gbps Unified I/O ports on Cisco UCS VIC 1340 160 Gbps Delivers 80 Gbps to the server

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

**Table 5 Network Configuration**

LAN and SAN Configuration	
<b>VLANs:</b>	
• <b>Public VLAN</b>	134
• <b>Private VLAN (RAC Interconnect)</b>	10
<b>VSANs</b>	
• <b>VSAN – A</b>	101
• <b>VSAN – B</b>	102

The FlashStack design comprises of //m70 FlashArray for increased scalability and throughput. The table below shows the components of the array.



**Table 6 Pure Storage FlashArray Configuration**

Storage Components	Description
FlashArray	// m70
Capacity	80TB
Connectivity	8 x 16 Gb/s redundant Fibre Channel 1 Gb/s redundant Ethernet ( Management port )
Physical	3U

**Table 7 OS and Software**

Operating System and Software	
Oracle Linux Server 7.2 (64 bit)	Linux orarac1 3.8.13-98.7.1.el7uek.x86_64
Oracle 12c Release 1 GRID	12.1.0.2
Oracle 12c Release 1 Database Enterprise Edition	12.1.0.2
Cisco Nexus 9372PX-E NXOS Version	6.1(2) I2 (2a)
Cisco MDS 9148S System Version	6.2 (9)
Cisco UCS Manager	3.1 (1g)
Pure Storage Purity Version	4.5.12
Oracle Swingbench	2.5.971

## Physical Topology

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

- **Network:** Cisco Nexus 9372PX-E, Cisco MDS 9148S and Cisco UCS Fabric Interconnect 6332-16UP for external and internal connectivity of IP and FC network.
- **Storage:** Pure Storage FlashArray//m (base shelves comes with 2 NVRAM with //m20 model and 4 NVRAM with //m70 model) with 16Gb Fibre Channel connectivity
- **Compute:** Cisco UCS B200 M4 Blade Server

Figure 20 illustrates the FlashStack solution physical infrastructure.

Figure 20 Physical Topology

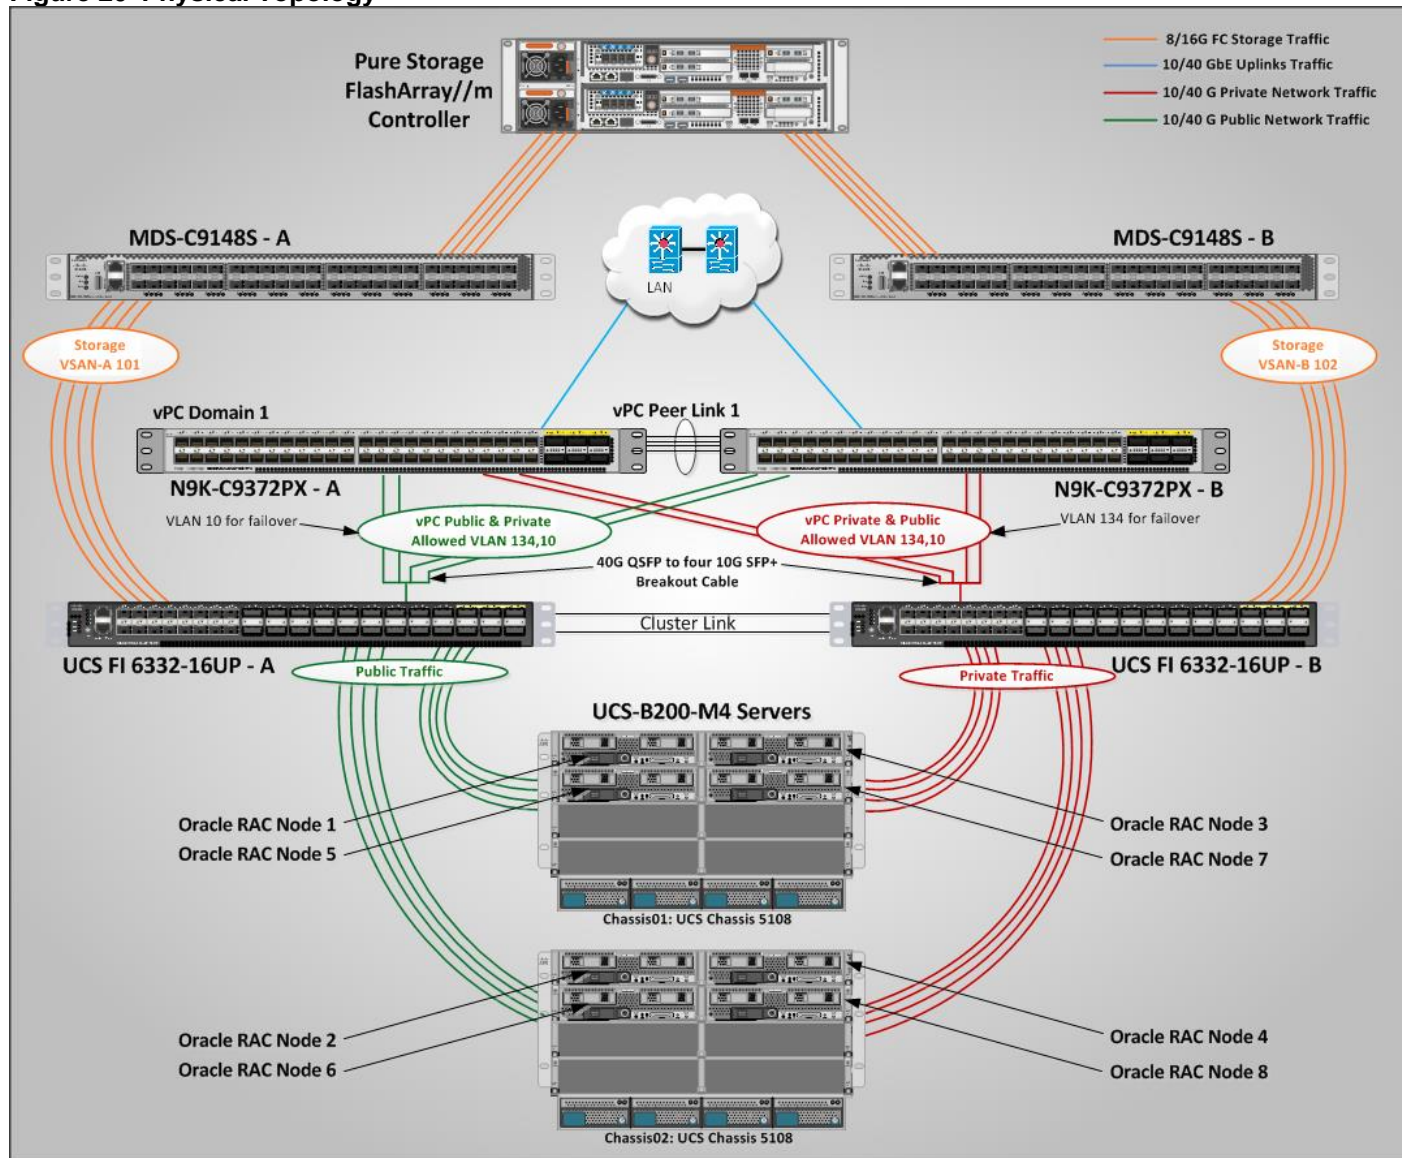


Figure 20 is a typical network configuration that can be deployed in a customer's environment. The best practices and setup recommendations are described later in this document.

As shown in Figure 20, a pair of Cisco UCS 6332-16UP fabric interconnects carries both storage and network traffic from the blades with the help of Cisco Nexus 9372PX-E and Cisco MDS 9148S switches. Both the fabric interconnect and the Cisco Nexus switch are clustered with the peer link between them to provide high availability. Two virtual Port-Channels (vPCs) are configured to provide public network and private network paths for the blades to northbound switches. Each vPC has VLANs created for application network data and management data paths.

As illustrated in Figure 20, eight (four per chassis) links go to Fabric Interconnect "A". Similarly, eight links go to Fabric Interconnect B. Fabric Interconnect-A links are used for Oracle Public network traffic shown as green lines. Fabric Interconnect-B links are used for Oracle private interconnect traffic shown as red lines. FC Storage access from Fabric Interconnect-A and Fabric Interconnect-B show as orange line.



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

terconnect and will not be routed via northbound network switch. In other words, all inter blade (or RAC node private) communication will be resolved locally at the fabric interconnect and this significantly reduces latency for Oracle Cache Fusion traffic.

---

# Deployment Hardware and Software

---

## Cisco UCS Configuration Overview

It is beyond the scope of this document to cover detailed information about UCS infrastructure setup and connectivity. The documentation guides and examples are available at [http://www.cisco.com/en/US/products/ps10281/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html).



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

---

## High Level Steps to Configure Cisco Unified Computing System

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

1. Configure Fabric Interconnects for a Cluster Setup
2. Configure Fabric Interconnects for Chassis and Blade Discovery
  - a. Configure Global Policies
  - b. Configure Server Ports
3. Configure LAN and SAN on Cisco UCS Manager
  - a. Configure Ethernet LAN Uplink Ports
  - b. Configure FC SAN Uplink Ports
  - c. Configure VLAN
  - d. Configure VSAN
4. Configure UUID, IP, MAC, WWNN and WWPN Pools
  - a. UUID Pool Creation
  - b. IP and MAC Pool Creation
  - c. WWNN and WWPN Pool Creation
5. Configure vNIC and vHBA Template
  - a. Create Public vNIC Template
  - b. Create Private vNIC Template
  - c. Create Storage vHBA Template
6. Configure Ethernet Uplink Port-Channels
7. Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

## Configure Fabric Interconnects for a Cluster Setup

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

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

**Figure 21 Fabric Interconnect A Setup**

```

COM4 - PuTTY

Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup

You have chosen to setup a new Fabric interconnect. Continue? (y/n): y

Enforce strong password? (y/n) [y]: n

Enter the password for "admin":
Confirm the password for "admin":

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes

Enter the switch fabric (A/B) []: A

Enter the system name:  PURESTG-FI

Physical Switch Mgmt0 IP address : 10.29.134.151

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.29.134.1

Cluster IPv4 address : 10.29.134.150

Configure the DNS Server IP address? (yes/no) [n]: n

Configure the default domain name? (yes/no) [n]: n

Join centralized management environment (UCS Central)? (yes/no) [n]: n

Following configurations will be applied:

Switch Fabric=A
System Name=PURESTG-FI
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=10.29.134.151
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.29.134.1
Ipv6 value=0

Cluster Enabled=yes
Cluster IP Address=10.29.134.150
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized.
      UCSM will be functional only after peer FI is configured in clustering mode.

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
  
```

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. Now, connect console port on the second Fabric Interconnect and do as following.

**Figure 22 Fabric Interconnect B Setup**

```

COM1 - PUTTY
Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? yes

Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.134.151
Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
Cluster IPv4 address      : 10.29.134.150

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address

Physical Switch Mgmt0 IP address : 10.29.134.152

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.

Mon Mar 21 20:54:11 UTC 2016
Configuration file - Ok

Cisco UCS 6300 Series Fabric Interconnect
PURESTG-FI-B login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2009, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
PURESTG-FI-B# █

```

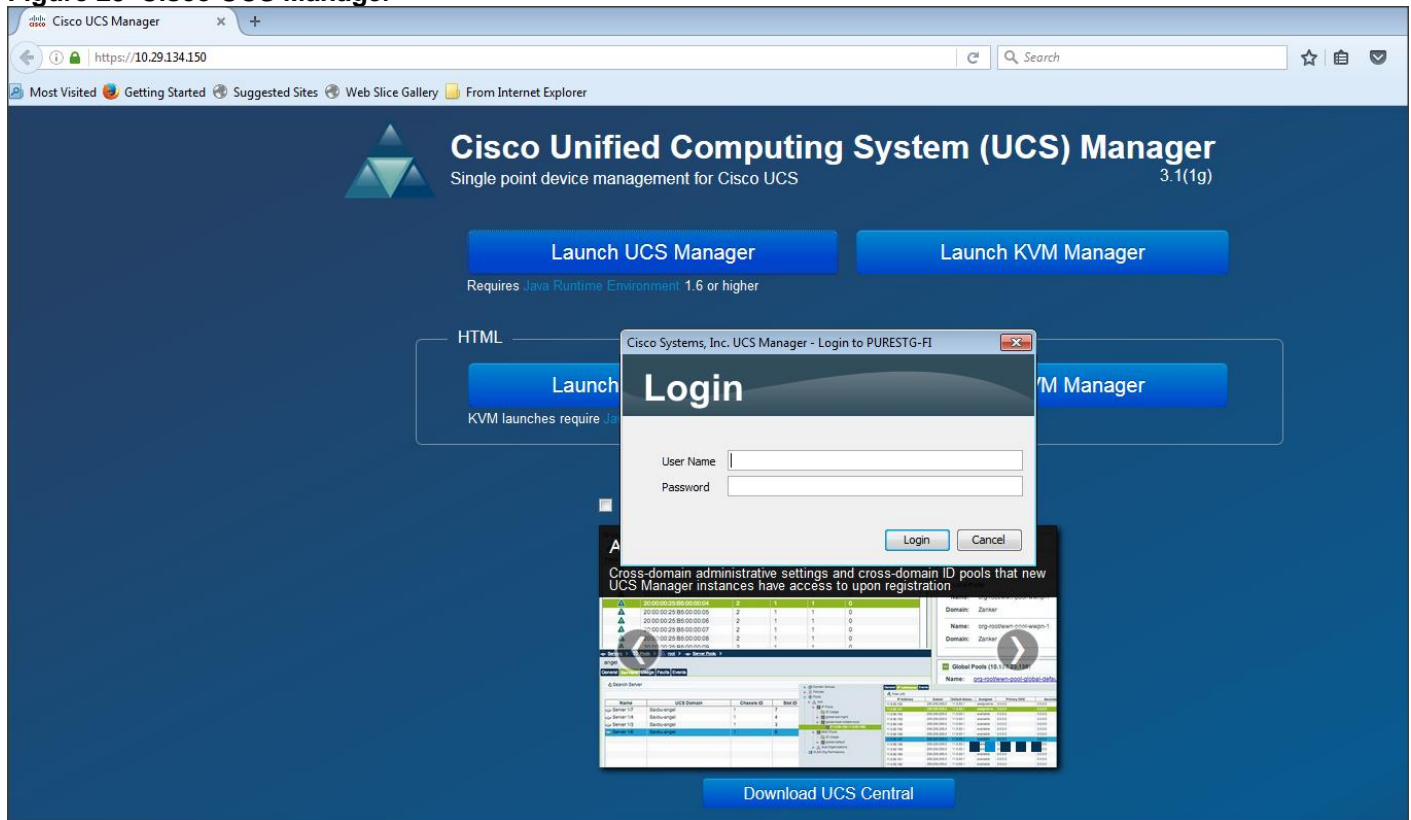
6. Review the settings printed to the console. Answer yes to apply and save the configuration.
7. Wait for the login prompt to make the configuration has been saved to Fabric Interconnect B.

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.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.




**Figure 23 Cisco UCS Manager**



4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

Figure 24 Cisco UCS System Manager



## Configure Fabric Interconnects for Chassis and Blade Discovery

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

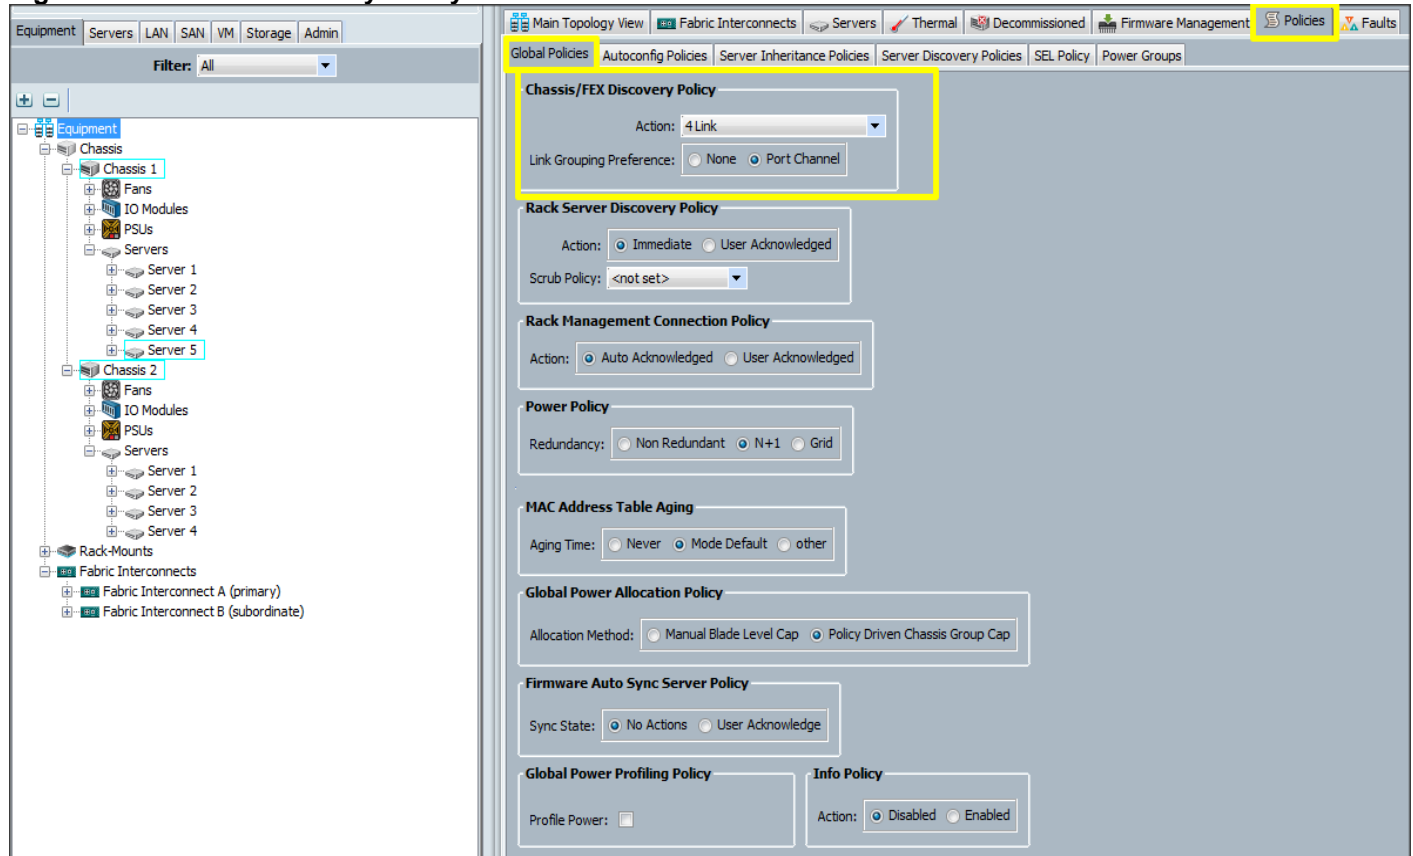
## Configure Global Policies

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, complete the following steps:

1. Go to Equipment > Policies (right pane) > Global Policies > Chassis/FEX Discovery Policies. As shown in the figure, select Action as “4 Link” from the drop down list and select Link Grouping Preference as “Port Channel”.

**Figure 25 Chassis Discovery Policy**



**Configure Server Ports**

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

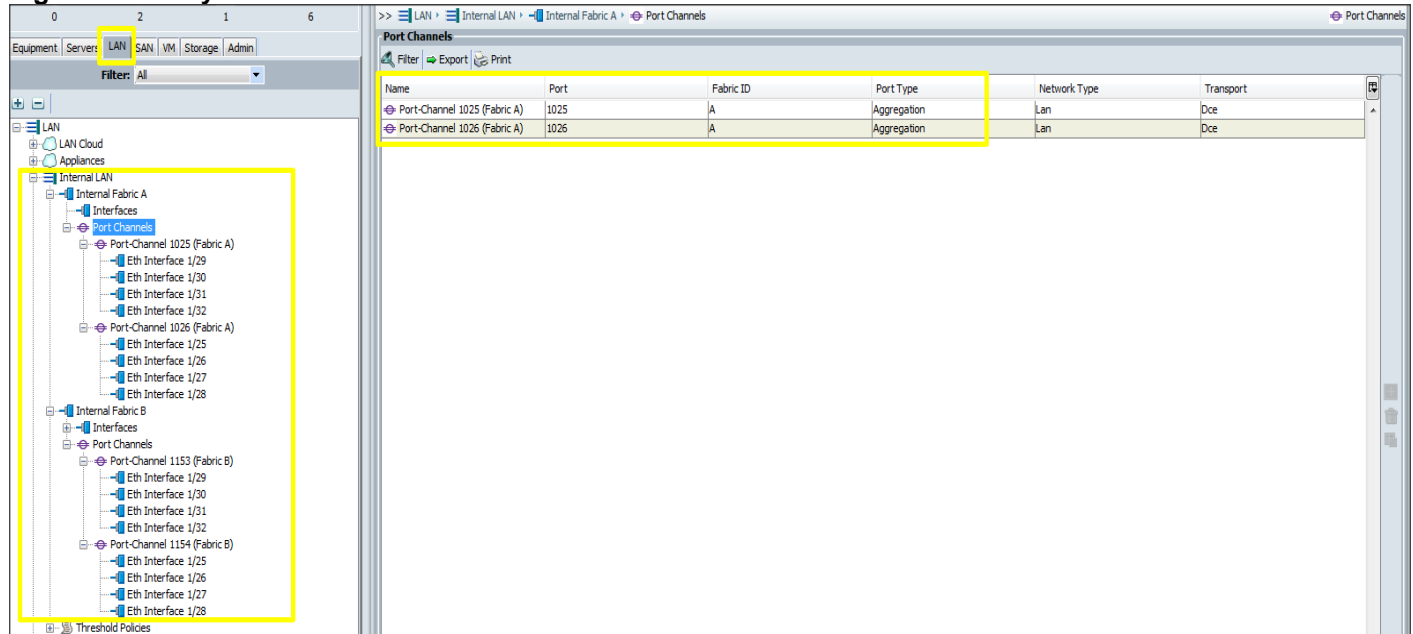
1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > Ethernet Ports.
2. Select the ports that are connected to chassis IOM.
3. Right-click “Configure as Server Port”.

Figure 26 Configure Server Ports

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 7	1	7	8C:60:4F:BD:31:A2	Unconfigured	Physical	Sfp Not Present	Disabled
Port 8	1	8	8C:60:4F:BD:31:A3	Unconfigured	Physical	Sfp Not Present	Disabled
Port 9	1	9	8C:60:4F:BD:31:A4	Unconfigured	Physical	Sfp Not Present	Disabled
Port 10	1	10	8C:60:4F:BD:31:A5	Unconfigured	Physical	Sfp Not Present	Disabled
Port 11	1	11	8C:60:4F:BD:31:A6	Unconfigured	Physical	Sfp Not Present	Disabled
Port 12	1	12	8C:60:4F:BD:31:A7	Unconfigured	Physical	Sfp Not Present	Disabled
Port 13	1	13	8C:60:4F:BD:31:A8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 14	1	14	8C:60:4F:BD:31:A9	Unconfigured	Physical	Sfp Not Present	Disabled
Port 15	1	15	8C:60:4F:BD:31:AA	Unconfigured	Physical	Sfp Not Present	Disabled
Port 16	1	16	8C:60:4F:BD:31:AB	Unconfigured	Physical	Sfp Not Present	Disabled
Port 17	1	17	8C:60:4F:BD:31:AC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 18	1	18	8C:60:4F:BD:31:BD	Unconfigured	Physical	Sfp Not Present	Disabled
Port 20	1	20	8C:60:4F:BD:31:B8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 21	1	21	8C:60:4F:BD:31:BC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 22	1	22	8C:60:4F:BD:31:C0	Unconfigured	Physical	Sfp Not Present	Disabled
Port 23	1	23	8C:60:4F:BD:31:C4	Unconfigured	Physical	Sfp Not Present	Disabled
Port 24	1	24	8C:60:4F:BD:31:C8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 25	1	25	8C:60:4F:BD:31:CC	Server	Physical	Up	Enabled
Port 26	1	26	8C:60:4F:BD:31:D0	Server	Physical	Up	Enabled
Port 27	1	27	8C:60:4F:BD:31:D4	Server	Physical	Up	Enabled
Port 28	1	28	8C:60:4F:BD:31:D8	Server	Physical	Up	Enabled
Port 29	1	29	8C:60:4F:BD:31:DC	Server	Physical	Up	Enabled
Port 30	1	30	8C:60:4F:BD:31:E0	Server	Physical	Up	Enabled
Port 31	1	31	8C:60:4F:BD:31:E4	Server	Physical	Up	Enabled
Port 32	1	32	8C:60:4F:BD:31:E8	Server	Physical	Up	Enabled
Port 33	1	33	8C:60:4F:BD:31:EC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 34	1	34	8C:60:4F:BD:31:F0	Unconfigured	Physical	Sfp Not Present	Disabled
Port 35	1	35	8C:60:4F:BD:31:F4	Unconfigured	Physical	Sfp Not Present	Disabled
Port 36	1	36	8C:60:4F:BD:31:F8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 37	1	37	8C:60:4F:BD:31:FC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 38	1	38	8C:60:4F:BD:31:00	Unconfigured	Physical	Sfp Not Present	Disabled
Port 39	1	39	8C:60:4F:BD:31:04	Unconfigured	Physical	Sfp Not Present	Disabled

4. Repeat same task for Fabric Interconnect B.
5. After configuring Server Ports, acknowledge both the Chassis. Go to Equipment > Chassis > Chassis 1 > General > Action > select "Acknowledge Chassis". Same way, acknowledge chassis 2.
6. 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 on OK. Same way, repeat the process to Re-acknowledge all the Servers.
7. Once 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 in below figure. Verify the same for Internal Fabric B.

**Figure 27 Verify Internal LAN Port-Channel**



### Configure LAN and SAN on Cisco UCS Manager

Configure Ethernet Uplink Ports to initiate Chassis and Blade discovery.

#### Configure Ethernet LAN Uplink Ports

To configure server ports, complete the following steps:

1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > Ethernet Ports.
2. Select the ports that are connected to chassis IOM.
3. Right-click “Configure as Uplink Port”.
4. Repeat these steps for Fabric Interconnect B.



Configure and enable the Ethernet LAN uplink Ports on Fabric Interconnect A and B. Here, we have created four uplink ports on each Fabric Interconnect as shown below. These ports will be used to create Virtual Port Channel in later sections.

**Figure 28 Configure Ethernet Uplinks Ports**

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 12	1	12	8C:60:4F:BD:31:A7	Unconfigured	Physical	Sfp Not Present	Disabled
Port 13	1	13	8C:60:4F:BD:31:A8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 14	1	14	8C:60:4F:BD:31:A9	Unconfigured	Physical	Sfp Not Present	Disabled
Port 15	1	15	8C:60:4F:BD:31:AA	Unconfigured	Physical	Sfp Not Present	Disabled
Port 16	1	16	8C:60:4F:BD:31:AB	Unconfigured	Physical	Sfp Not Present	Disabled
Port 17	1	17	8C:60:4F:BD:31:AC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 18	1	18	8C:60:4F:BD:31:BD	Unconfigured	Physical	Sfp Not Present	Disabled
Port 20	1	20	8C:60:4F:BD:31:B8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 21	1	21	8C:60:4F:BD:31:BC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 22	1	22	8C:60:4F:BD:31:C0	Unconfigured	Physical	Sfp Not Present	Disabled
Port 23	1	23	8C:60:4F:BD:31:C4	Unconfigured	Physical	Sfp Not Present	Disabled
Port 24	1	24	8C:60:4F:BD:31:C8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 25	1	25	8C:60:4F:BD:31:CC	Server	Physical	Up	Enabled
Port 26	1	26	8C:60:4F:BD:31:D0	Server	Physical	Up	Enabled
Port 27	1	27	8C:60:4F:BD:31:D4	Server	Physical	Up	Enabled
Port 28	1	28	8C:60:4F:BD:31:D8	Server	Physical	Up	Enabled
Port 29	1	29	8C:60:4F:BD:31:DC	Server	Physical	Up	Enabled
Port 30	1	30	8C:60:4F:BD:31:E0	Server	Physical	Up	Enabled
Port 31	1	31	8C:60:4F:BD:31:E4	Server	Physical	Up	Enabled
Port 32	1	32	8C:60:4F:BD:31:E8	Server	Physical	Up	Enabled
Port 33	1	33	8C:60:4F:BD:31:EC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 34	1	34	8C:60:4F:BD:31:F0	Unconfigured	Physical	Sfp Not Present	Disabled
Port 35	1	35	8C:60:4F:BD:31:F4	Unconfigured	Physical	Sfp Not Present	Disabled
Port 36	1	36	8C:60:4F:BD:31:F8	Unconfigured	Physical	Sfp Not Present	Disabled
Port 37	1	37	8C:60:4F:BD:31:FC	Unconfigured	Physical	Sfp Not Present	Disabled
Port 38	1	38	8C:60:4F:BD:31:00	Unconfigured	Physical	Sfp Not Present	Disabled
Port 39	1	39	8C:60:4F:BD:31:04	Unconfigured	Physical	Sfp Not Present	Disabled
Port 40	1	40	8C:60:4F:BD:31:08	Unconfigured	Physical	Sfp Not Present	Disabled
<b>Scalability Port 19</b>							
Port 1	1	1	8C:60:4F:BD:31:B4	Network	Physical	Up	Enabled
Port 2	1	2	8C:60:4F:BD:31:B5	Network	Physical	Up	Enabled
Port 3	1	3	8C:60:4F:BD:31:B6	Network	Physical	Up	Enabled
Port 4	1	4	8C:60:4F:BD:31:B7	Network	Physical	Up	Enabled

**Configure FC SAN Uplink Ports**

To configure FC SAN Uplink ports, complete the following steps:

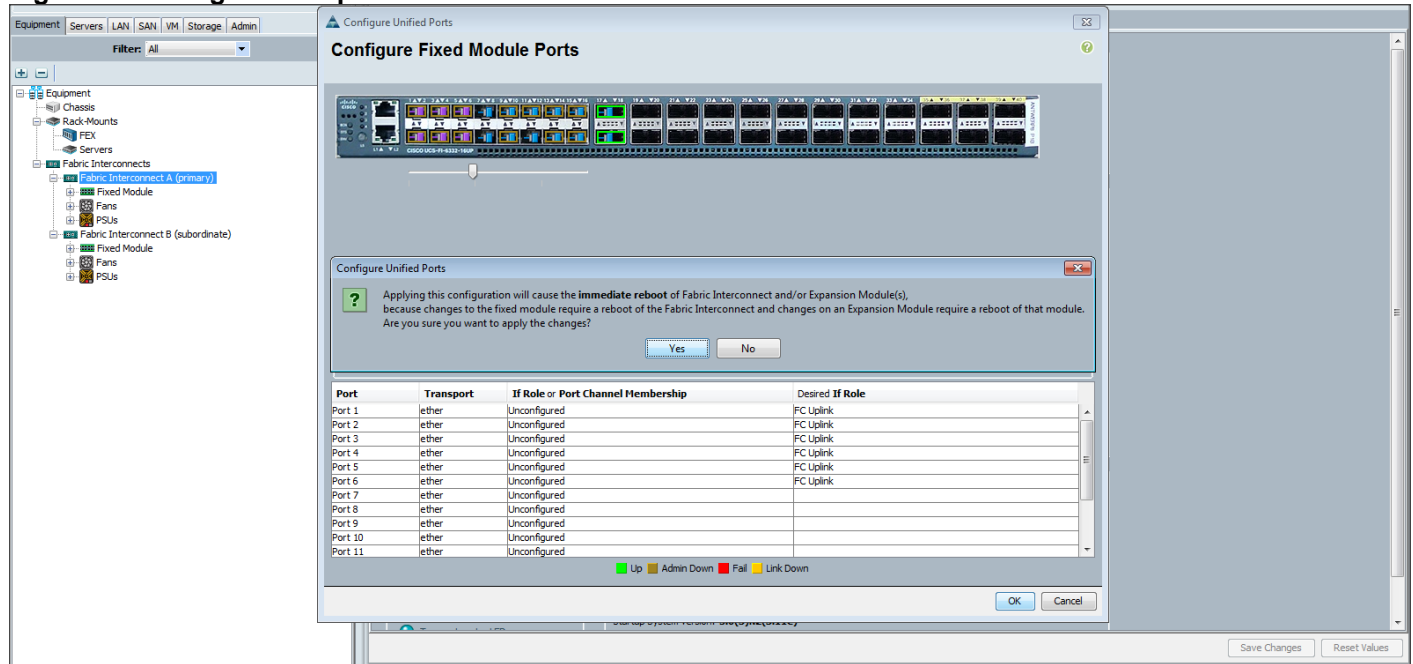
1. From Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > FC Ports menu.
2. Select the desired ports and enable those ports. The figure below shows the configuration of FC Ports.
3. Repeat these steps for Fabric Interconnect B.



Configure and enable the FC Uplink Ports on Fabric Interconnect A and B. We created four uplink ports on each Fabric Interconnect as shown below.



**Figure 29 Configure FC Uplink Ports**



### Configure VLAN

To configure VLAN, complete the following steps:

1. In Cisco UCS Manager, click LAN > LAN Cloud > VLANs.
2. Right-click Create VLANs.

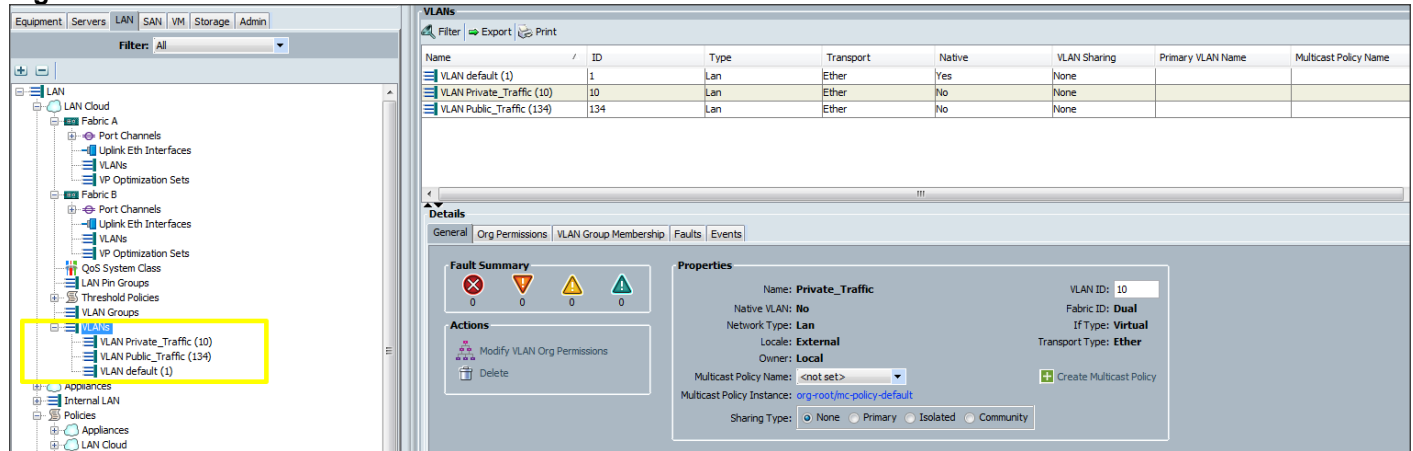


In this solution, we created 2 VLANs: one for private network (VLAN 10) and one for public network (VLAN 134) traffic. These two VLANs will be used in the vNIC templates that are discussed later.



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

**Figure 30 VLANs**



**Configure VSAN**

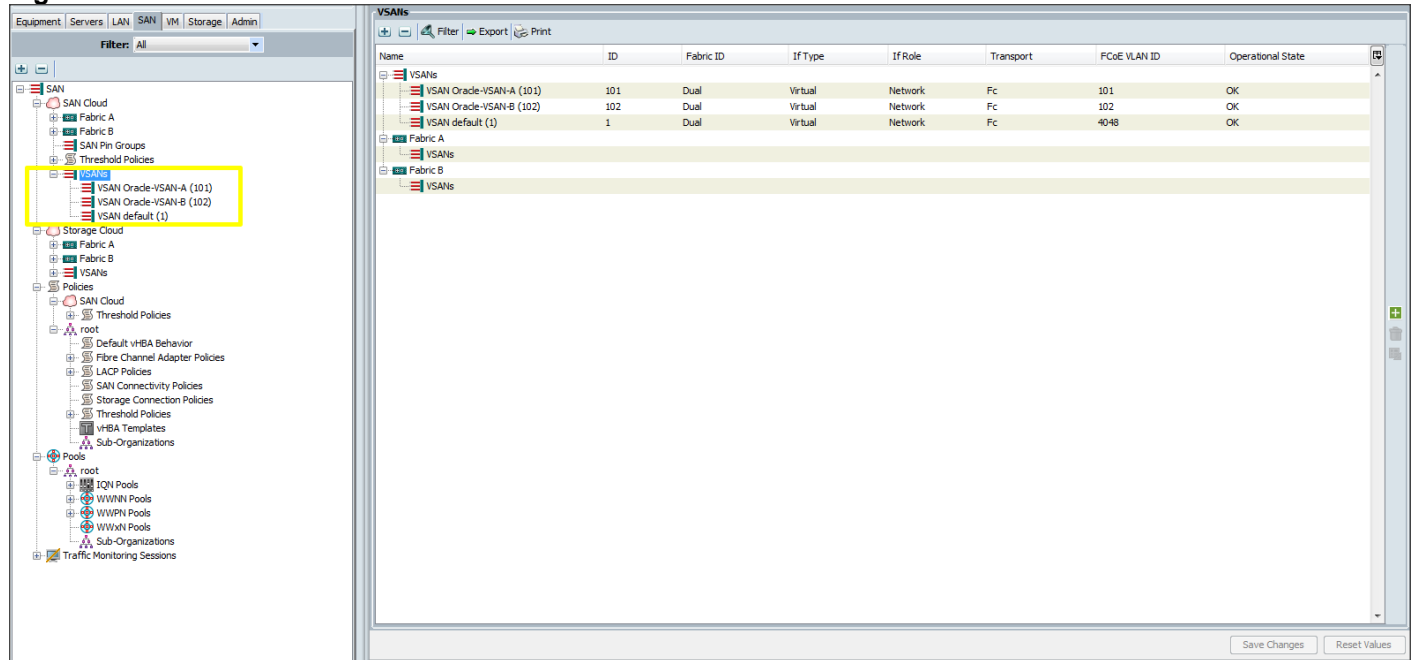
To configure VSAN, complete the following steps:

1. In Cisco UCS Manager, click SAN > SAN Cloud > VSANs and right-click Create VSANs.



In this solution, we created 2 VSANs: VSAN-A 101 and VSAN-B 102 for SAN Boot and Storage Access.

**Figure 31 VSANs**



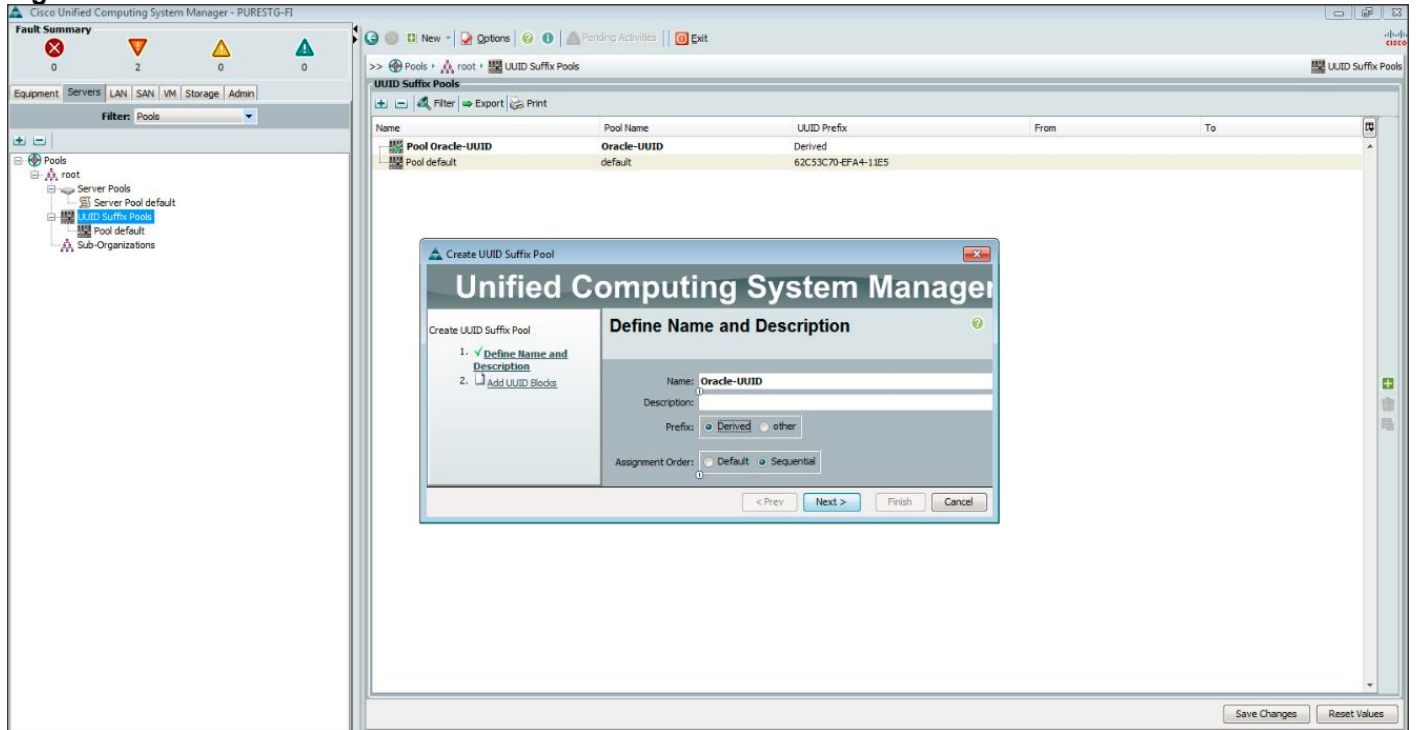
**Configure UUID, MAC, WWNN and WWPN Pools**

**Create UUID Pool**

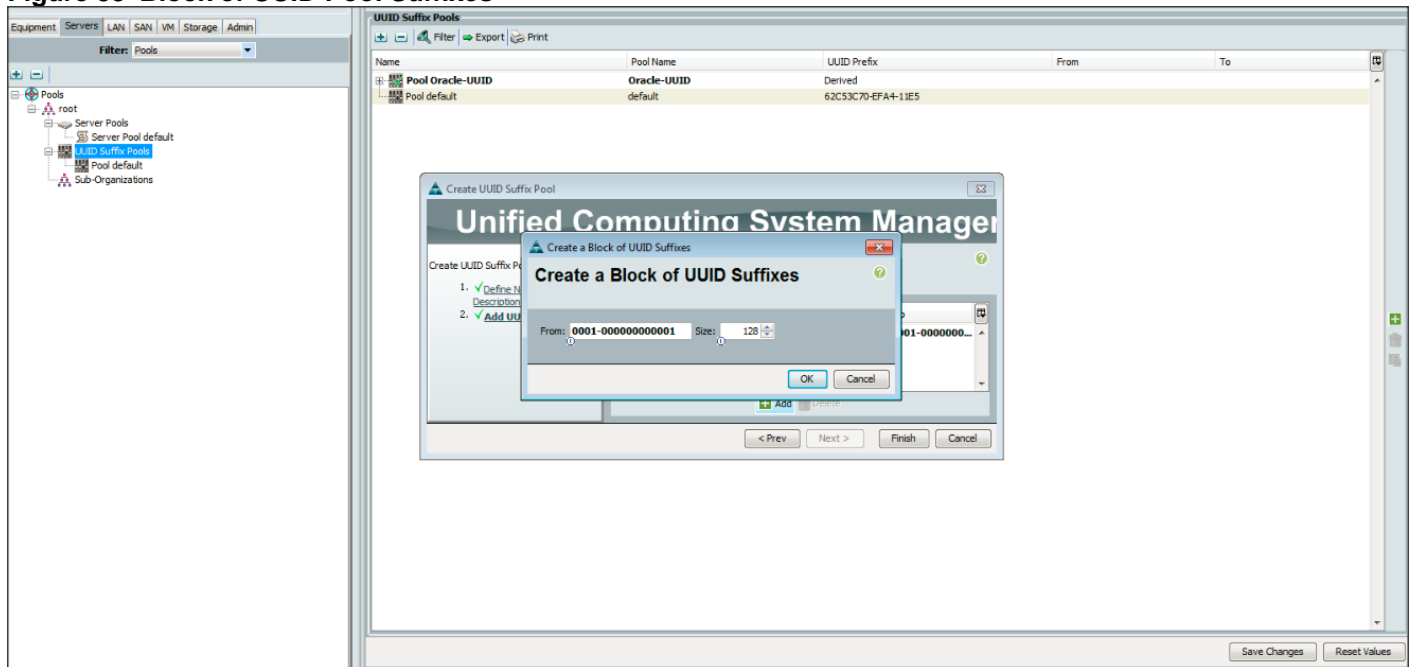
To create the UUID Pool, complete the following steps:

1. In Cisco UCS Manager, click tab Servers > Pools > root > UUID Suffix Pools.
2. Right-click "Create UUID Suffix Pool" and create a new pool as shown below.

**Figure 32 Create UUID Pool**



**Figure 33 Block of UUID Pool Suffixes**



**Create IP and MAC Pools**

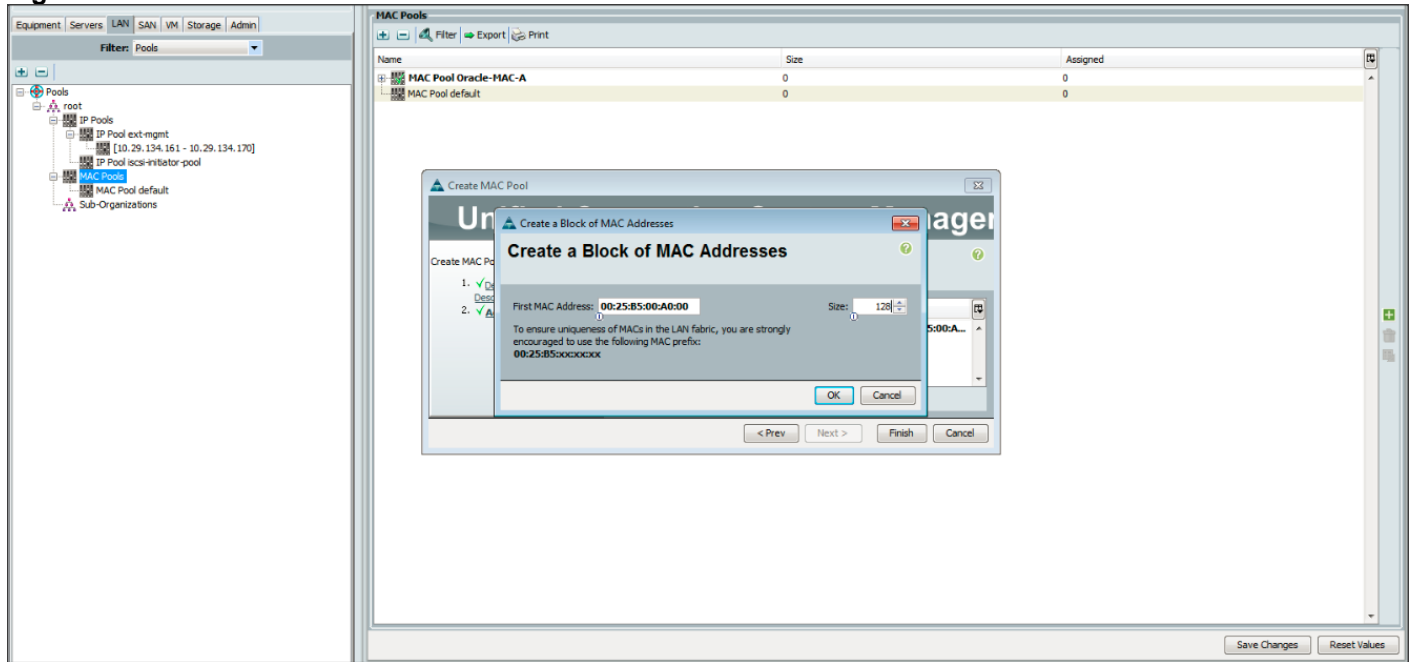
To create the IP and MAC Pools, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab > Pools > root > MAC Pools to "Create MAC Pool".

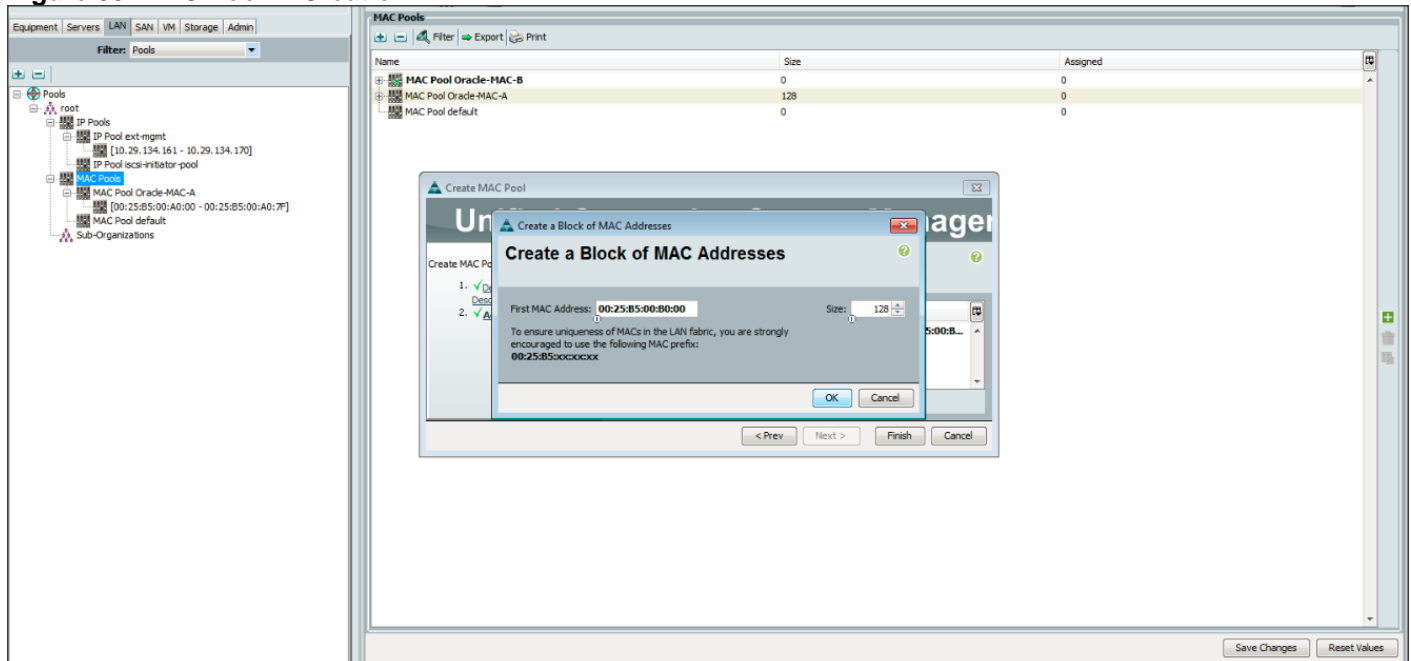


We created Oracle-MAC-A and Oracle-MAC-B as shown below for all the vNIC MAC Addresses.

**Figure 34 MAC Pool-A Creation**



**Figure 35 MAC Pool-B Creation**

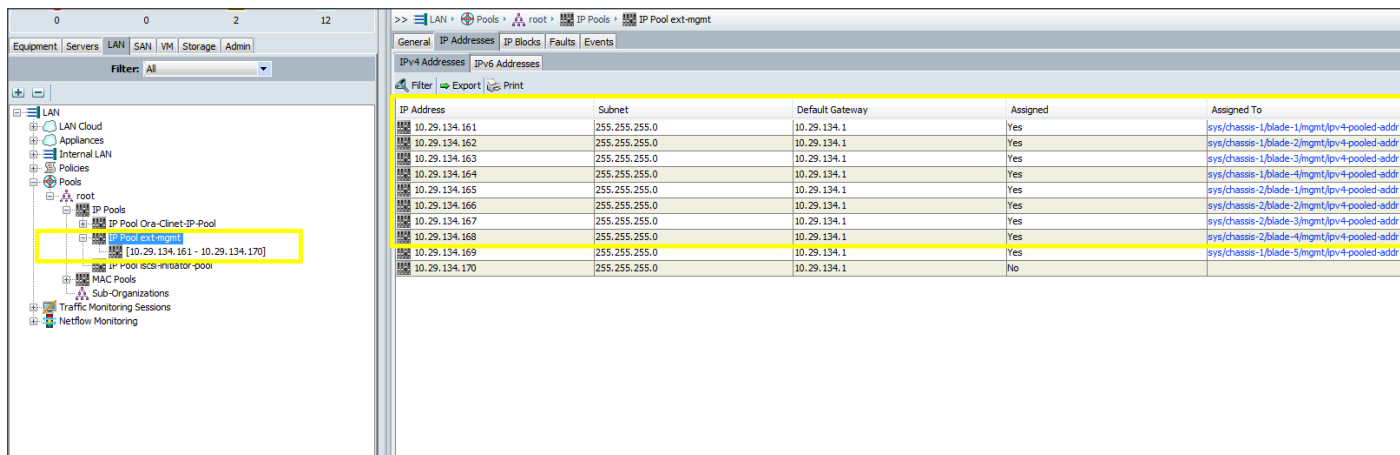


2. Click the LAN tab > Pools > root > IP Pool > IP Pool ext-mgmt to "Create IP Pool".



We created IP Pool to assign Server CIMC Address as shown below.

**Figure 36 Create IP Pool**



 For all 8 Oracle Nodes, we assigned CIMC IP as shown in the screenshot.

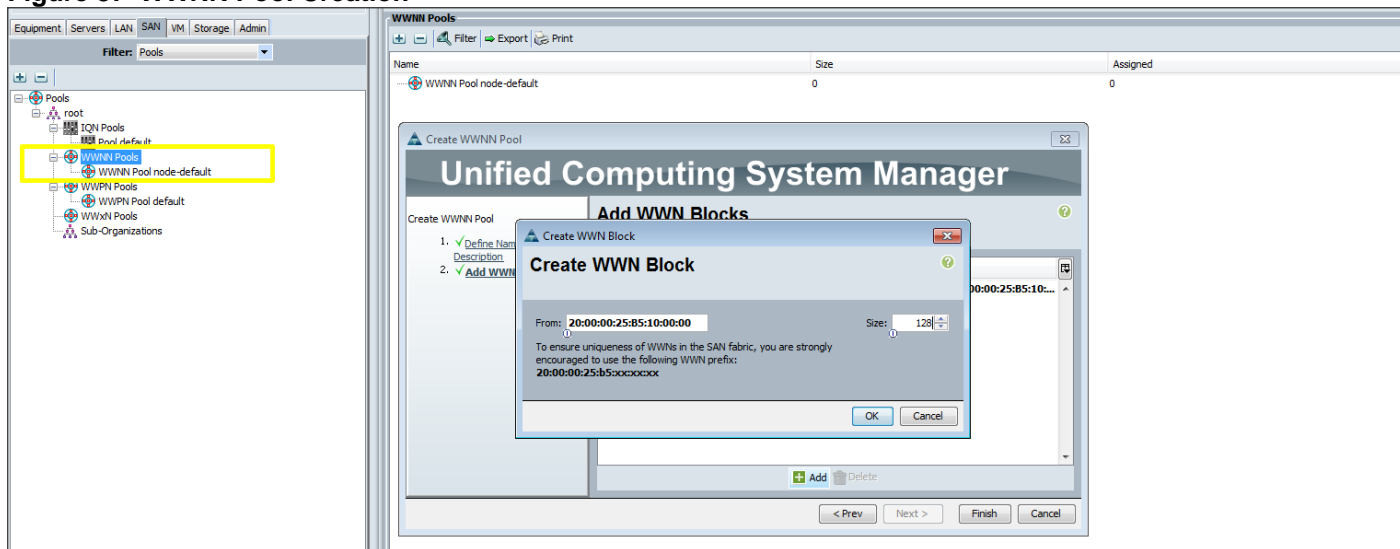
### WWNN and WWPN Pools

To create WWNN and WWPN Pools, complete the following steps:


1. In Cisco UCS Manager, click the SAN tab > Pools > root > WWNN Pools and right-click "Create WWNN Pools" as show below.

 We created Oracle-WWNN Pool as World Wide Node Name.

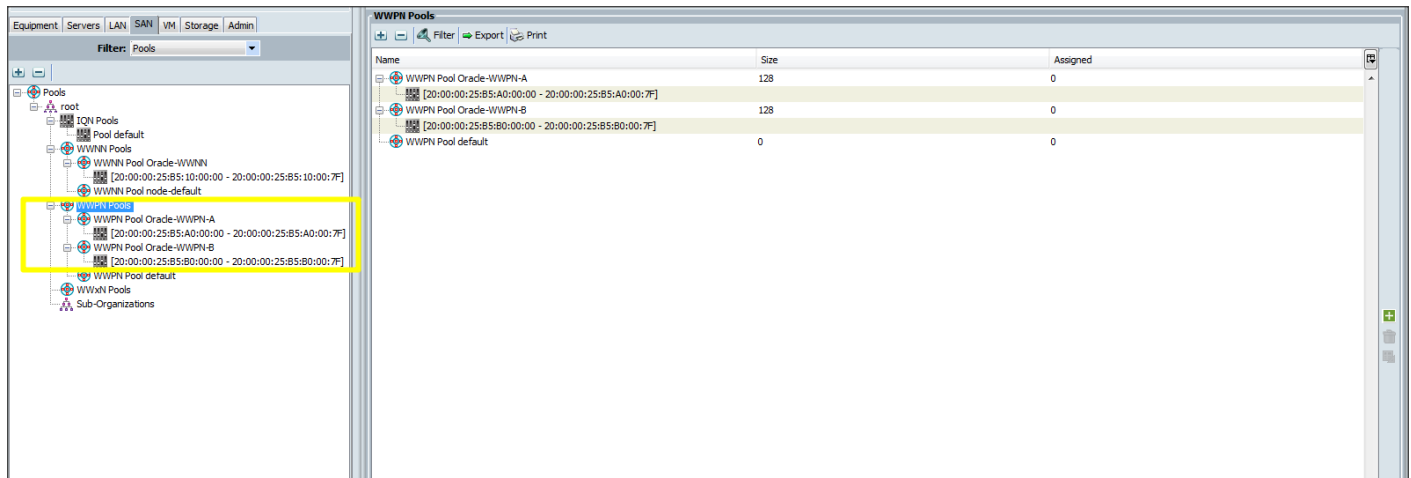
**Figure 37 WWNN Pool Creation**



2. Click WWPN Pools "Create WWPN Pools".

 We created two WWPN as Oracle-WWPN-A Pool and Oracle-WWPN-B as World Wide Port Name as shown below. These WWNN and WWPN entries will be used for Boot from SAN configuration.

**Figure 38 WWPN Pool Creation**

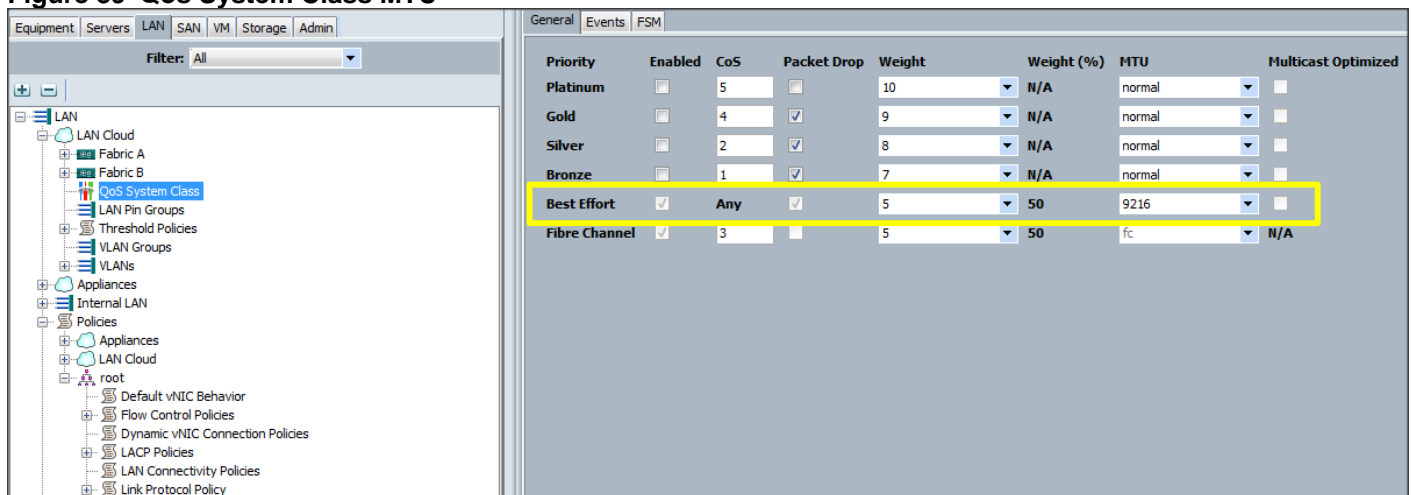


### Set Jumbo Frames in both the Cisco Fabric Interconnect

To configure jumbo frames, complete the following steps:

1. Go to the LAN tab > LAN Cloud > QoS System Class and in the right pane click the General tab.
2. On the Best Efforts row, enter 9216 in the box under the MTU column. Click Save Changes in the bottom window to save the configuration.

**Figure 39 Qos System Class MTU**



### Configure vNIC and vHBA Template

You will use the two vNIC templates for Public Network and Private Network Traffic to create the Service Profiles.

#### Create Public vNIC Template

To create a public vNIC template, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab > Policies > vNIC templates and right-click "Create vNIC Template" as shown below.

**Figure 40 vNIC Template for Public Network Traffic**



The screenshot displays a network management interface with a sidebar on the left and a main configuration window titled "Create vNIC Template".

**Sidebar (Left):** Shows a tree view of network configurations under "LAN". The "vNIC Templates" folder is expanded, showing a list of templates including "MAC Pool Oracle-MAC-A", "MAC Pool Oracle-MAC-B", and "MAC Pool Oracle-MAC-C".

**Main Window (Right):** The "Create vNIC Template" dialog box is open. It contains the following fields and options:

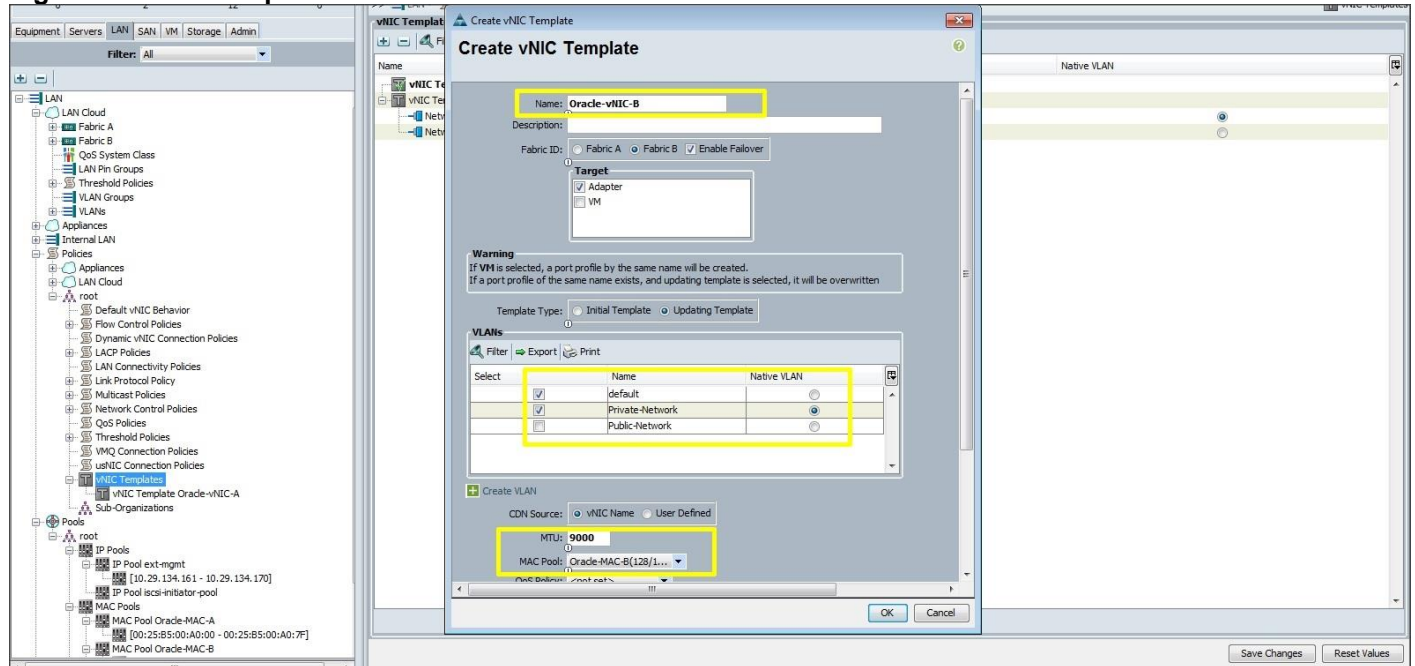
- Name:** Oracle-vNIC-A
- Description:** (empty)
- Fabric ID:** Fabric A, Fabric B, Enable Fallover (selected)
- Target:** Adapter (checked), VM (unchecked)
- Warning:** If VM is selected, a port profile by the same name will be created. If a port profile of the same name exists, and updating template is selected, it will be overwritten.
- Template Type:** Initial Template (selected), Updating Template
- VLANs:** A table with columns "Name" and "Native VLAN".
- Create VLAN:** A section with "CDN Sources" (vNIC Name, User Defined) and "MTU: 1500".
- MAC Pool:** Oracle-MAC-A(128/1500)

The "VLANs" table is highlighted with a yellow box and contains the following data:

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	Private-Network	<input type="radio"/>
<input checked="" type="checkbox"/>	Public-Network	<input checked="" type="radio"/>

Buttons for "OK" and "Cancel" are located at the bottom right of the dialog box. "Save Changes" and "Reset Values" buttons are located at the bottom right of the main window.

**Figure 41 vNIC Template for Private Network Traffic**



### Create Storage vHBA Template

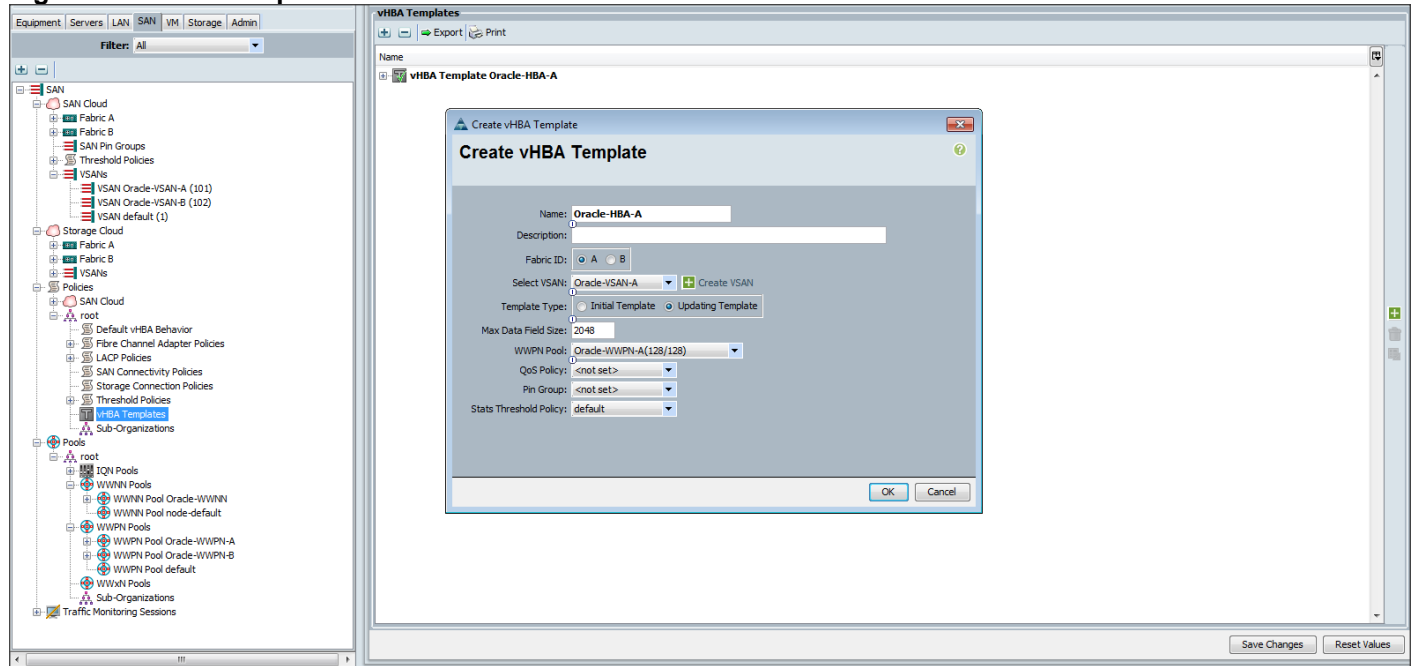
To create a storage vHBA template, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab SAN > Policies > vHBA templates and right-click “Create vHBA Template” to create vHBAs.



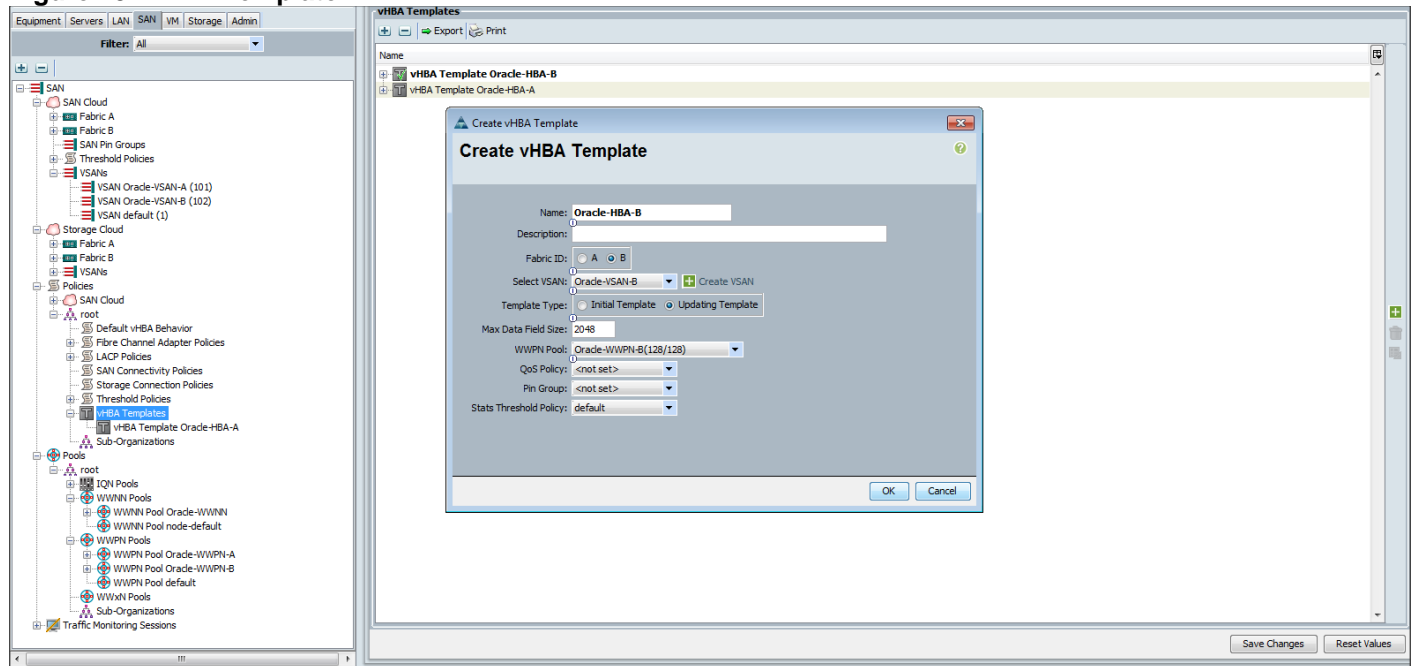
We created two vHBA as Oracle-HBA-A and Oracle-HBA-B as shown below.

**Figure 42 vHBA Template**



2. Select WWPN Pool for Oracle-HBA-A as "Oracle-WWPN-A" and Oracle-HBA-B as "Oracle-WWPN-B" created earlier.

**Figure 43 vHBA Template**



## Configure Ethernet Uplink Port-Channel

To configure the ethernet uplink Port-Channel, complete the following steps:

1. Click the LAN tab > LAN Cloud > Fabric A > Port Channels and right-click "Create Port-Channel".
2. Select the desired Ethernet Uplink ports configured earlier.

- Repeat the same steps to create Port-Channel on Fabric B. In the current setup, we used ports on Fabric A as shown below to create port-channel 19.

**Figure 44 Port Channels for Fabric A**

Name	Fabric ID	Aggr. Port ID	If Type	If Role	Transport
Port-Channel 19 (oracle-public)	A		Aggregation	Network	Ether
Eth Interface 1/19/1	A	19	Physical	Network	Ether
Eth Interface 1/19/2	A	19	Physical	Network	Ether
Eth Interface 1/19/3	A	19	Physical	Network	Ether
Eth Interface 1/19/4	A	19	Physical	Network	Ether

- Configure the ports on Fabric B to create port channel 20.

**Figure 45 Port Channels for Fabric B**

Name	Fabric ID	Aggr. Port ID	If Type	If Role	Transport
Port-Channel 20 (oracle-private)	B		Aggregation	Network	Ether
Eth Interface 1/19/1	B	19	Physical	Network	Ether
Eth Interface 1/19/2	B	19	Physical	Network	Ether
Eth Interface 1/19/3	B	19	Physical	Network	Ether
Eth Interface 1/19/4	B	19	Physical	Network	Ether

## Create Server Boot Policy for SAN Boot



We strongly recommend you 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 in the following steps.

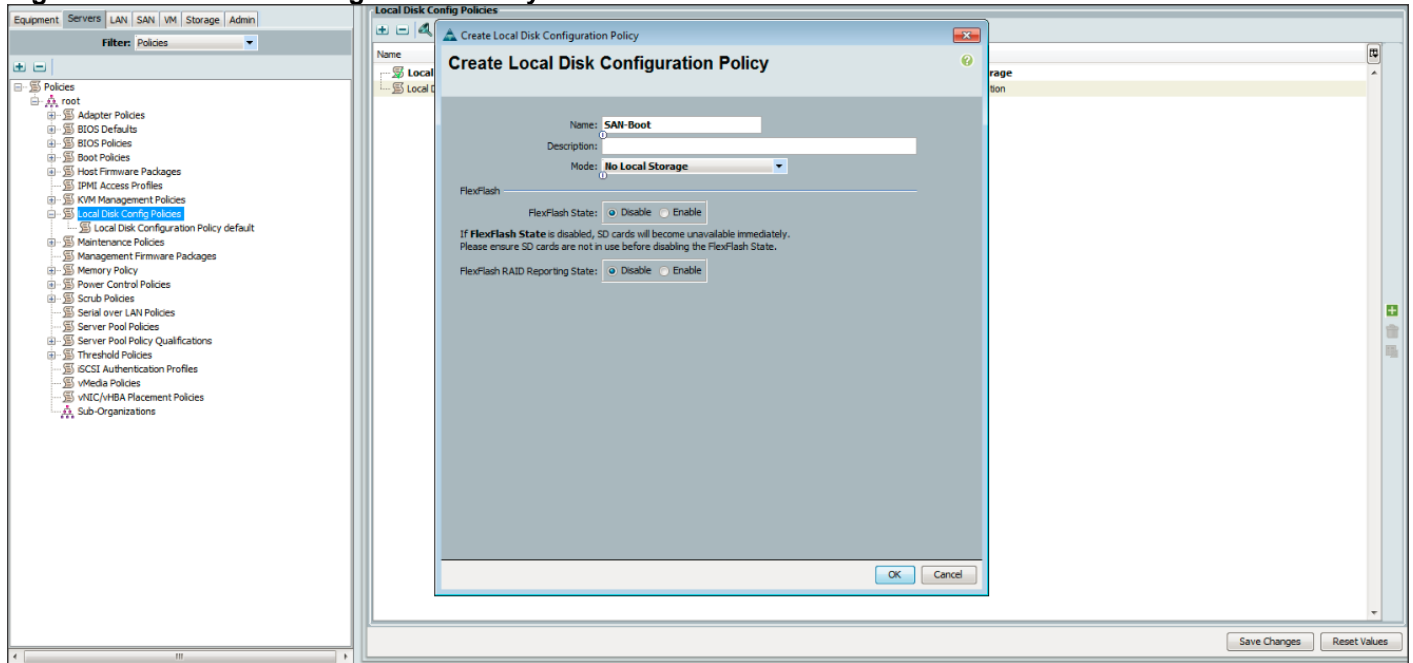
A Local disk configuration for the Cisco UCS is necessary if the servers in the environments have a local disk.

To configure Local disk policy, complete the following steps:

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

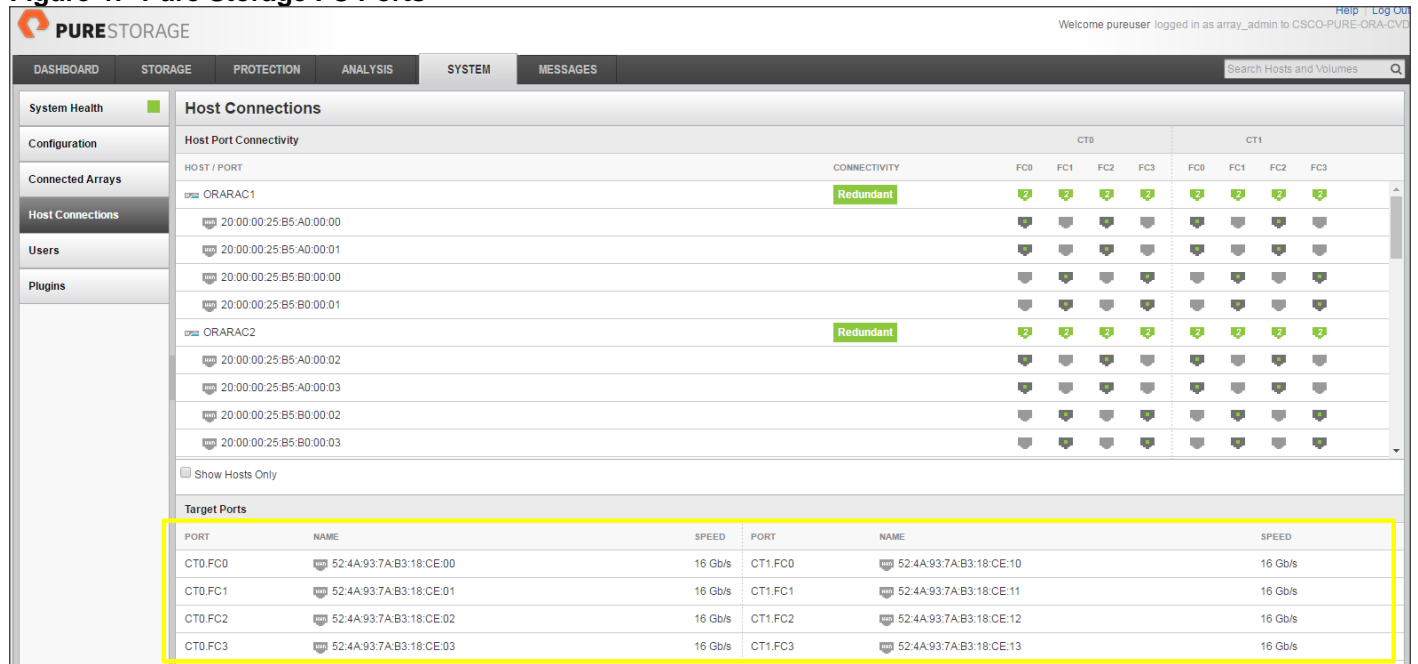
2. Click OK to create the policy as shown in the below figure.

**Figure 46 Local Disk Configuration Policy**



The SAN Ports CT0.FC0, CT0.FC2 of Pure Storage Controller 0 are connected to Cisco MDS 9148S Switch A and CT0.FC1, CT0.FC3 are connected to Cisco MDS 9148S Switch B. Similarly, the SAN Ports CT1.FC0, CT1.FC2 of Pure Storage Controller 1 are connected to Cisco MDS 9148S Switch A and CT1.FC1, CT1.FC3 are connected to Cisco MDS 9148S Switch B.

Figure 47 Pure Storage FC Ports



### Create SAN Boot Policy A

The SAN boot (SAN-Boot-A) configures the SAN primary's primary-target to be port CT0.FC0 on Pure storage cluster and SAN primary's secondary-target to be port CT1.FC0 on Pure storage cluster. Similarly, the SAN secondary's primary-target should be port CT1.FC1 on Pure storage cluster and SAN secondary's secondary-target should be port CT0.FC1 on Pure storage cluster. Log into the storage controller and verify all the port information is correct.

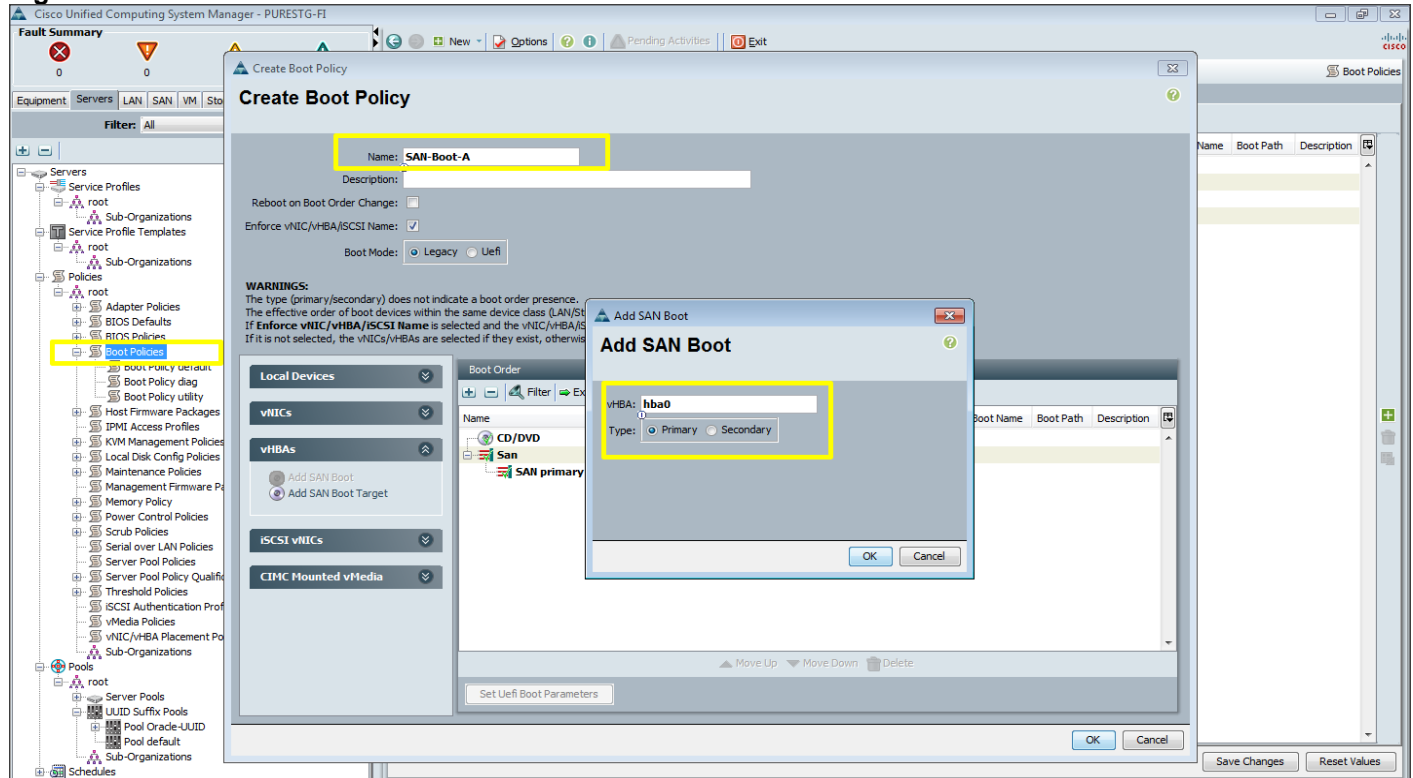
Now create SAN Boot primary (hba0) and SAN Boot secondary (hba1) in create boot policy by entering WWPN of Pure Storage FC Ports.

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

1. Go to tab Servers > Policies > root > Boot Policies. Right-click and create SAN-Boot-A as the name of the boot policy as shown in below figure.
2. Expand the Local Devices drop-down menu and Choose Add CD-ROM. Expand the vHBA drop-down menu and Choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba0" in the vHBA field and make sure type is selected as "Primary".

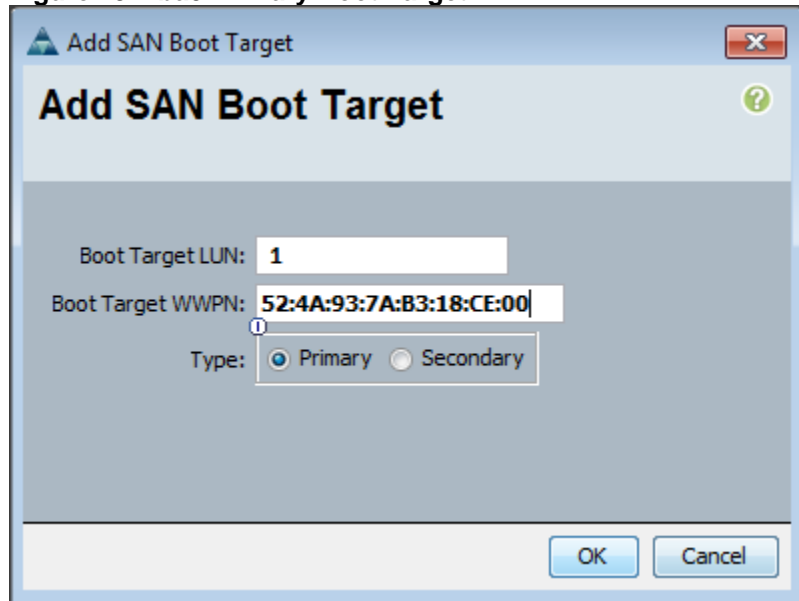


**Figure 48 SAN Boot A hba0**

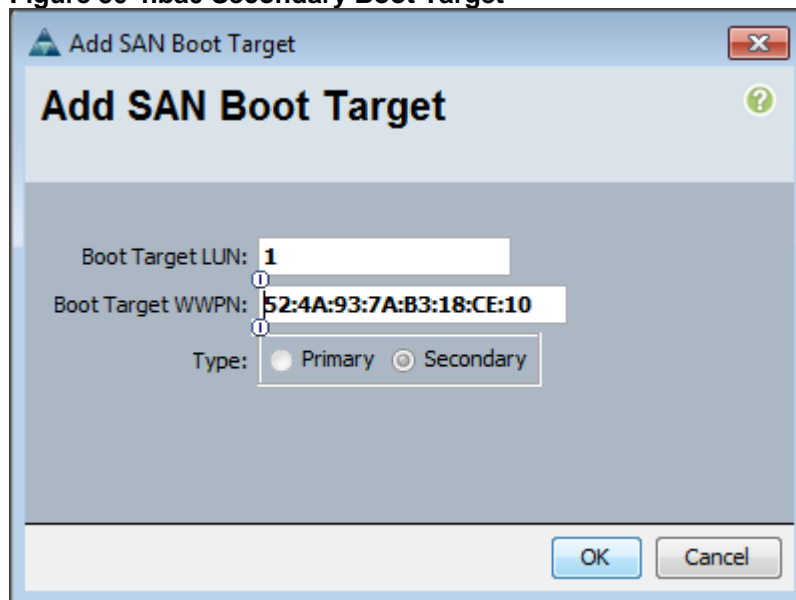


3. Click OK to add SAN Boot. Then choose add SAN Boot Target.
4. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC0 of Pure Storage and add SAN Boot Primary Target.

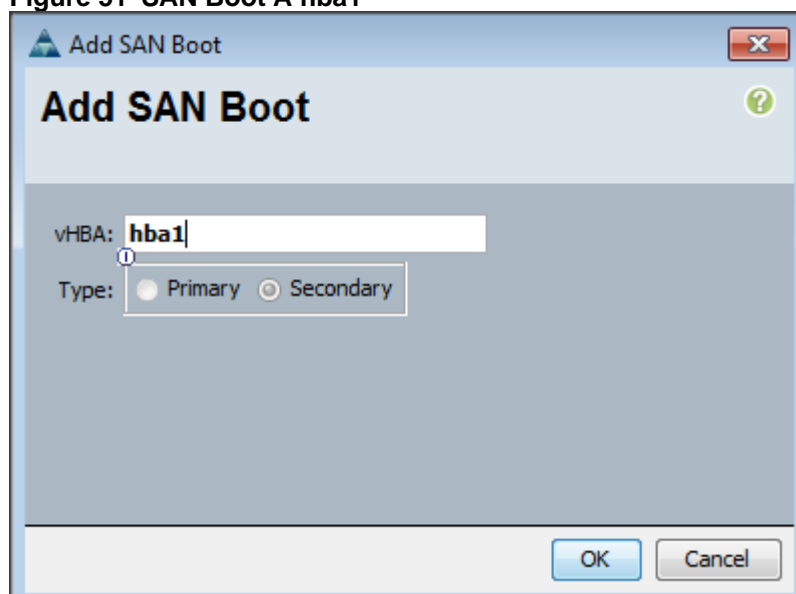
**Figure 49 hba0 Primary Boot Target**



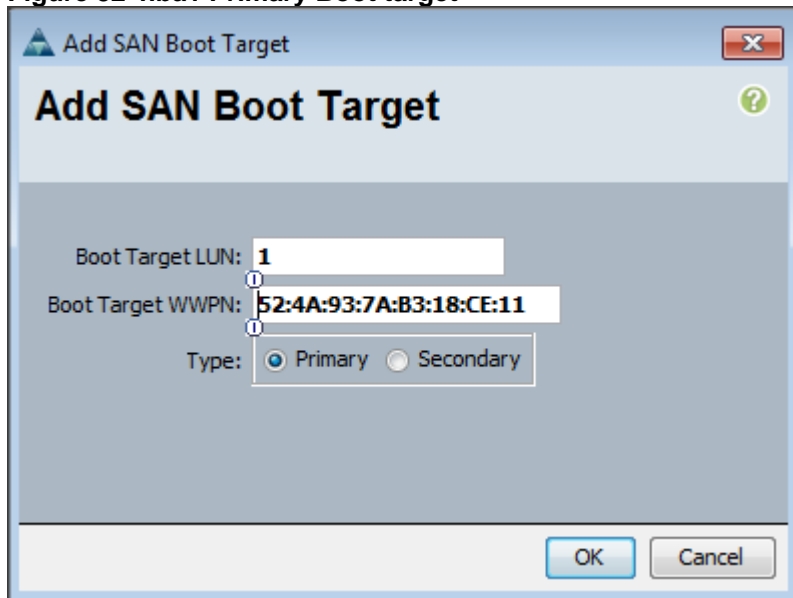
5. Add secondary SAN Boot target into same hba0 and enter boot target LUN as 1 and WWPN for FC port CT1.FC0 of Pure Storage and add SAN Boot Secondary Target.

**Figure 50 hba0 Secondary Boot Target**

6. Again, from the vHBA drop-down menu and Choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba1" in the vHBA field.

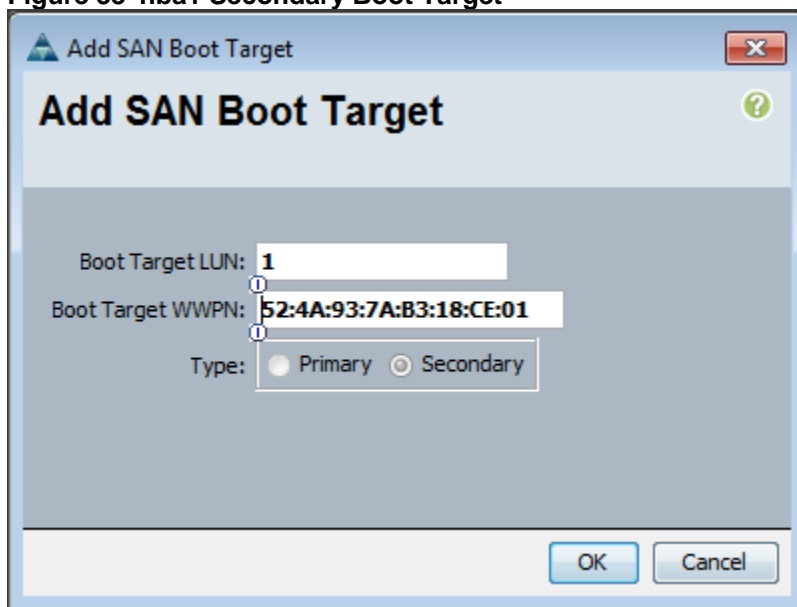
**Figure 51 SAN Boot A hba1**

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

**Figure 52 hba1 Primary Boot target**

The screenshot shows a dialog box titled "Add SAN Boot Target" with a close button (X) in the top right corner. The main title "Add SAN Boot Target" is displayed in a large font. Below the title, there are three input fields: "Boot Target LUN:" with the value "1", "Boot Target WWPN:" with the value "52:4A:93:7A:B3:18:CE:11", and "Type:" with radio buttons for "Primary" (selected) and "Secondary". At the bottom of the dialog, there are "OK" and "Cancel" buttons.

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

**Figure 53 hba1 Secondary Boot Target**

The screenshot shows a dialog box titled "Add SAN Boot Target" with a close button (X) in the top right corner. The main title "Add SAN Boot Target" is displayed in a large font. Below the title, there are three input fields: "Boot Target LUN:" with the value "1", "Boot Target WWPN:" with the value "52:4A:93:7A:B3:18:CE:01", and "Type:" with radio buttons for "Primary" and "Secondary" (selected). At the bottom of the dialog, there are "OK" and "Cancel" buttons.

10. 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 UCS Manager as shown below

Figure 54 SAN Boot A

The screenshot displays the Cisco UCS Manager interface for configuring a boot policy. The left-hand navigation pane shows the hierarchy: Servers > Policies > root > Boot Policies > SAN-Boot-A. The main content area is titled 'Boot Policy SAN-Boot-A' and includes the following sections:

- Actions:** Delete, Show Policy Usage, Use Global.
- Properties:** Name: SAN-Boot-A, Description: (empty), Owner: Local, Reboot on Boot Order Change: (unchecked), Enforce vNIC/vHBA/iSCSI Name: (checked), Boot Mode: Legacy (selected) / Uefi.
- Warning:** The type (primary/secondary) does not indicate a boot order preference. The effective order of boot devices within the same device class (LAN/storage/iSCSI) is determined by PCIe bus scan order. If Enforce vNIC/vHBA/iSCSI Name is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported. If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.
- Boot Order Table:**

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN	Slot Number	Boot Name	Boot Path	Description
CD/DVD	1								
San	2								
SAN primary		hba0	Primary						
SAN Target primary			Primary	1	52:4A:93:7A:B3:18:CE:00				
SAN Target secondary			Secondary	1	52:4A:93:7A:B3:18:CE:10				
SAN secondary		hba1	Secondary						
SAN Target primary			Primary	1	52:4A:93:7A:B3:18:CE:11				
SAN Target secondary			Secondary	1	52:4A:93:7A:B3:18:CE:01				

### Create SAN Boot Policy B

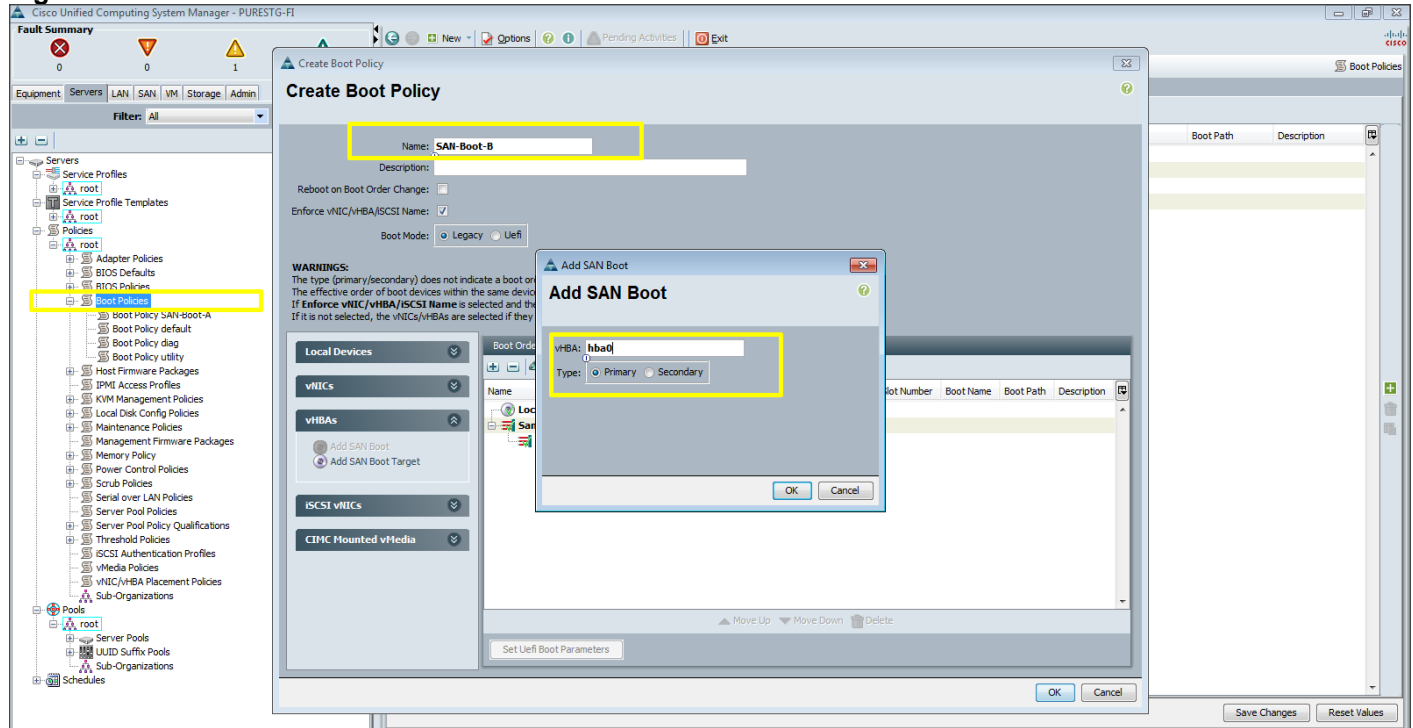
The SAN boot (SAN-Boot-B) configures the SAN primary's primary-target to be port CT0.FC2 on Pure storage cluster and SAN primary's secondary-target to be port CT1.FC2 on Pure storage cluster. Similarly, the SAN secondary's primary-target should be port CT1.FC3 on Pure storage cluster and SAN secondary's secondary-target should be port CT0.FC3 on Pure storage cluster. Log into the storage controller and verify all the port information is correct.

Now create SAN Boot primary (hba0) and SAN Boot secondary (hba1) in create boot policy by entering WWPN of Pure Storage FC Ports.

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

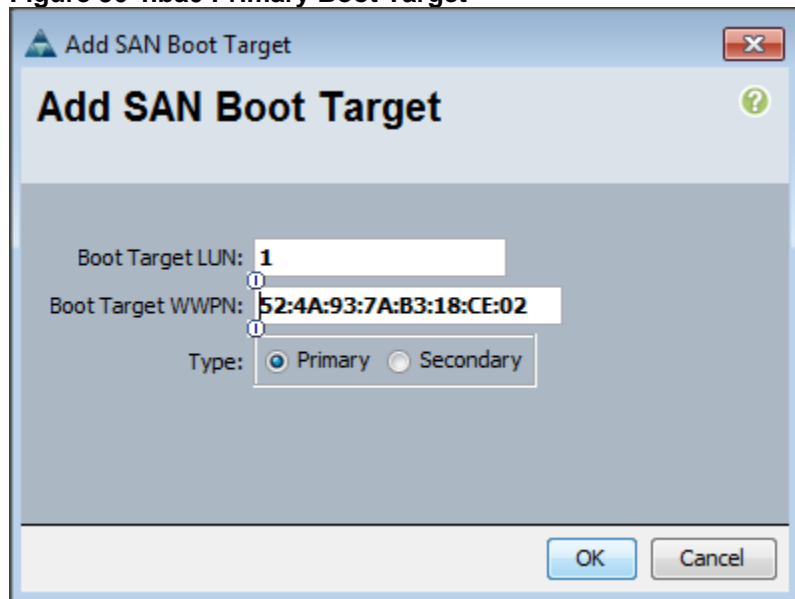
1. Go to the Servers tab > Policies > root > Boot Policies. Right-click and create SAN-Boot-B as the name of the boot policy as shown in below figure.
2. Expand the Local Devices drop-down menu and Choose Add CD-ROM. Expand the vHBA drop-down menu and Choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba0" in the vHBA field and make sure type is selected as "Primary".

**Figure 55 SAN Boot B hba0**



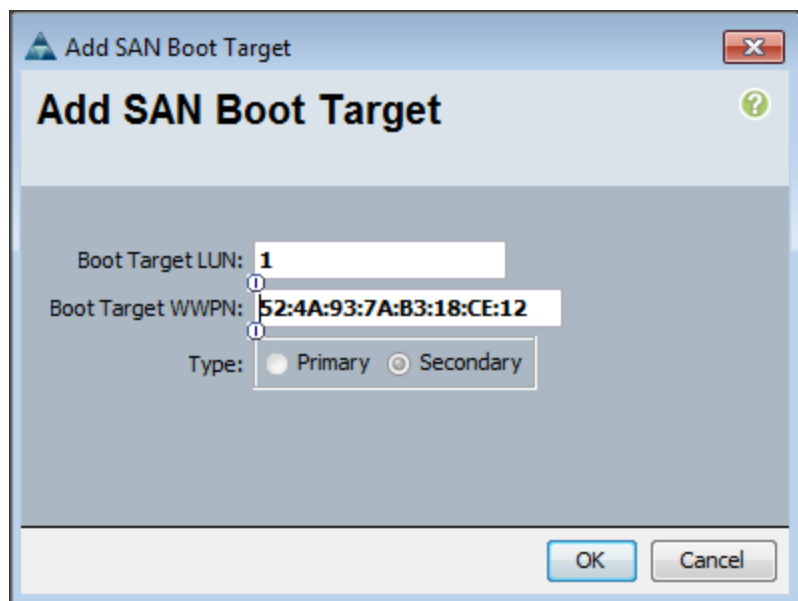
3. Click OK to add SAN Boot. Then choose add SAN Boot Target.
4. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC2 of Pure Storage and add SAN Boot Primary Target.

**Figure 56 hba0 Primary Boot Target**



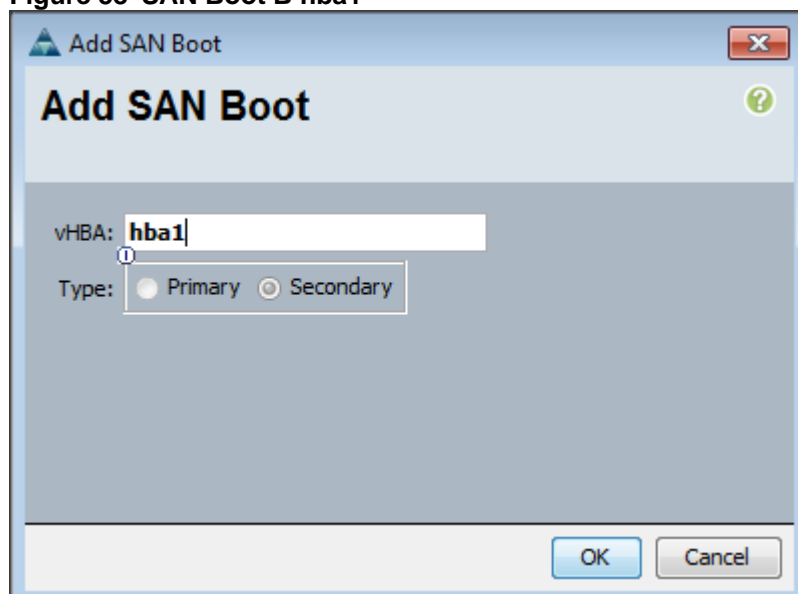
5. Add a secondary SAN Boot target into same hba0 and enter boot target LUN as 1 and WWPN for FC port CT1.FC2 of Pure Storage and add SAN Boot Secondary Target.

**Figure 57 hba0 Secondary Boot Target**



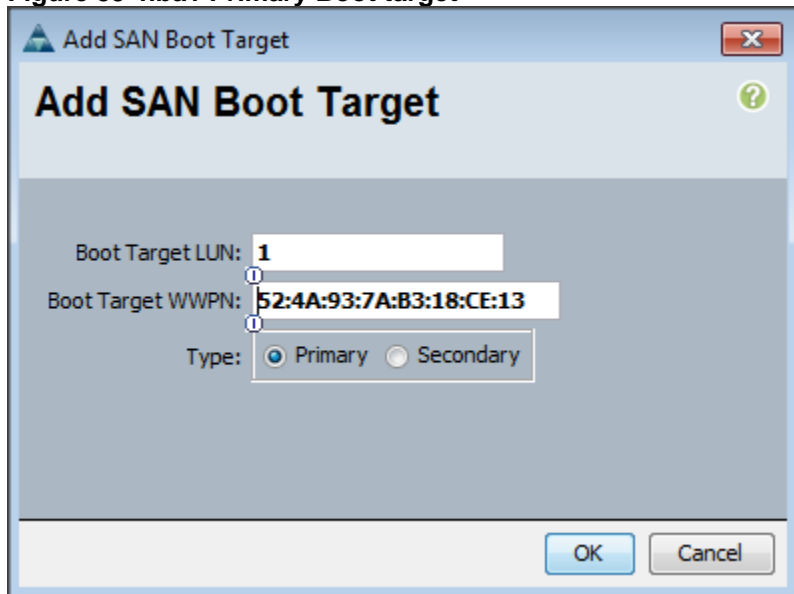
6. From the vHBA drop-down menu and Choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba1" in the vHBA field.

**Figure 58 SAN Boot B hba1**



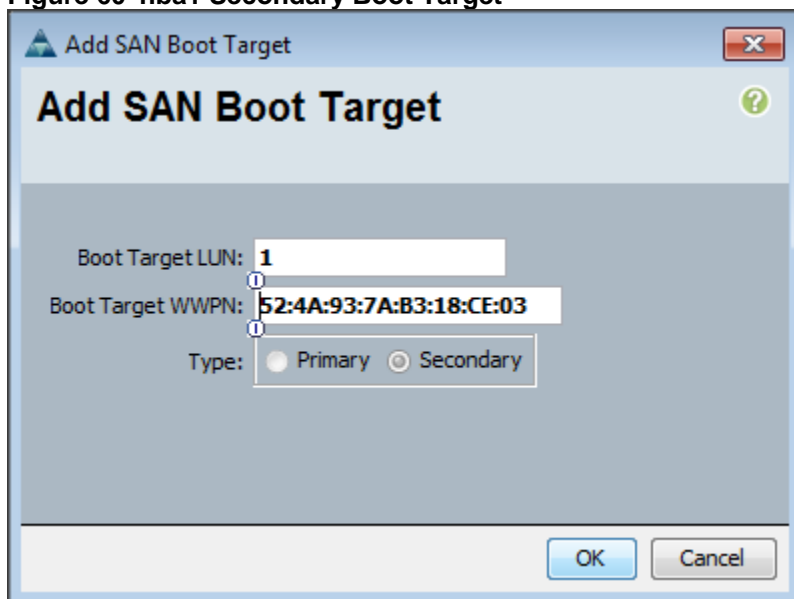
7. Click OK to SAN Boot. Then choose add SAN Boot Target.
8. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT1.FC3 of Pure Storage and add SAN Boot Primary Target.



**Figure 59 hba1 Primary Boot target**

The screenshot shows a dialog box titled "Add SAN Boot Target" with a close button (X) in the top right corner. The main title "Add SAN Boot Target" is displayed in large bold text. Below the title, there are three input fields: "Boot Target LUN:" with the value "1", "Boot Target WWPN:" with the value "52:4A:93:7A:B3:18:CE:13", and "Type:" with radio buttons for "Primary" (selected) and "Secondary". At the bottom of the dialog, there are "OK" and "Cancel" buttons.

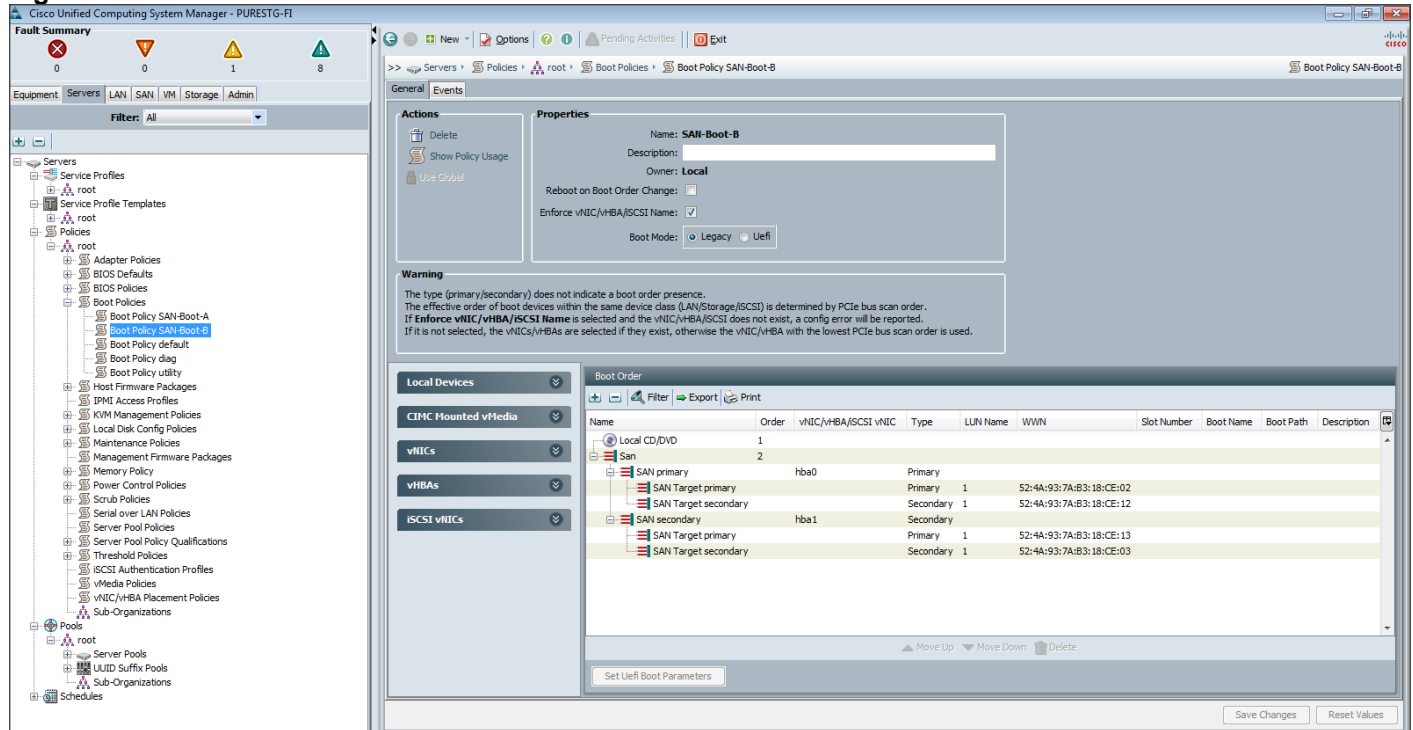
9. Add secondary SAN Boot target into same hba1 and enter boot target LUN as 1 and WWPN for FC port CT0.FC3 of Pure Storage and add SAN Boot Secondary Target.

**Figure 60 hba1 Secondary Boot Target**

The screenshot shows a dialog box titled "Add SAN Boot Target" with a close button (X) in the top right corner. The main title "Add SAN Boot Target" is displayed in large bold text. Below the title, there are three input fields: "Boot Target LUN:" with the value "1", "Boot Target WWPN:" with the value "52:4A:93:7A:B3:18:CE:03", and "Type:" with radio buttons for "Primary" and "Secondary" (selected). At the bottom of the dialog, there are "OK" and "Cancel" buttons.

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

**Figure 61 SAN Boot B**



For this solution, we created two Boot Policy as “SAN Boot A” and “SAN Boot B”. For 8 Oracle RAC Nodes, we will assign 4 Service Profiles with SAN Boot A to first 4 RAC nodes (orac1, orac2, orac3 and orac4) and remaining 4 Service Profiles with SAN Boot B to remaining 4 RAC nodes (orac5, orac6, orac7 and orac8) as explained in a subsequent section.

## Configure and Create a Service Profile Template.

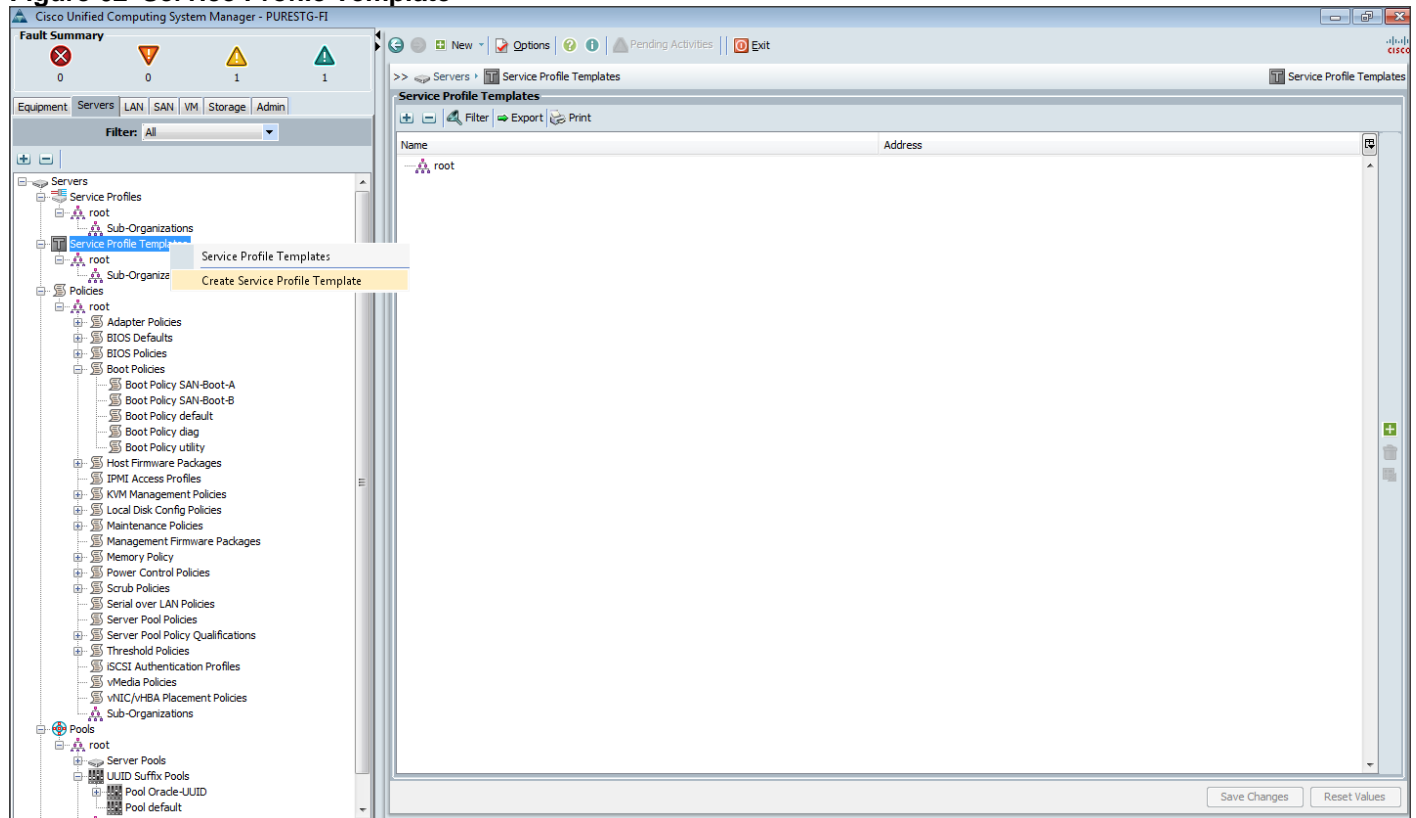
Service profile templates enable policy based server management that helps ensure consistent server resource provisioning suitable to meet predefined workload needs.

Create two Service Profile Template. First Service profile template “ORASAN-A” using “SAN Boot A” and second Service profile template “ORASAN-B” using “SAN Boot B”. We will create first ORASAN-A as explained below.

### Create Service Profile Template

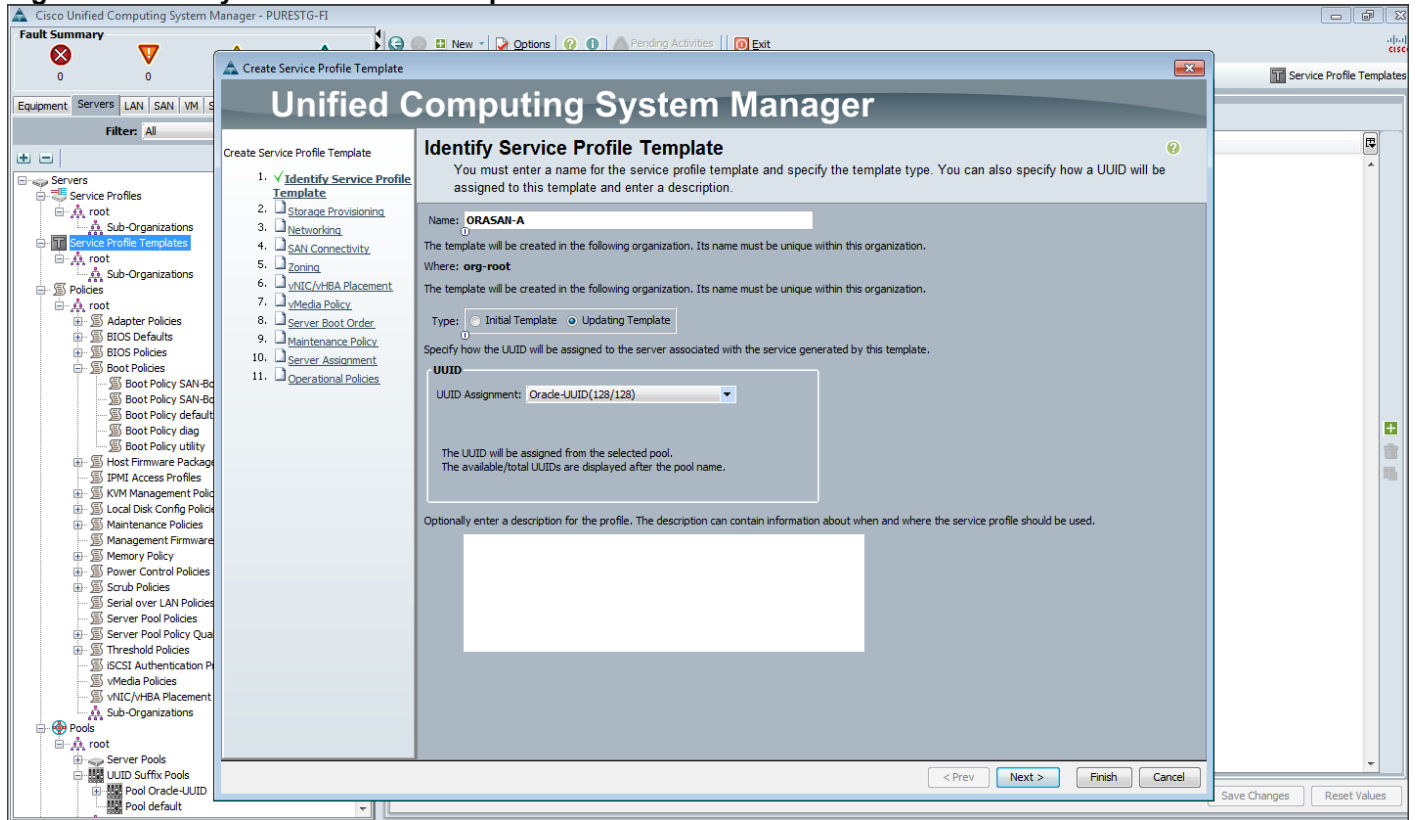
To create a service profile template, complete the following steps:

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

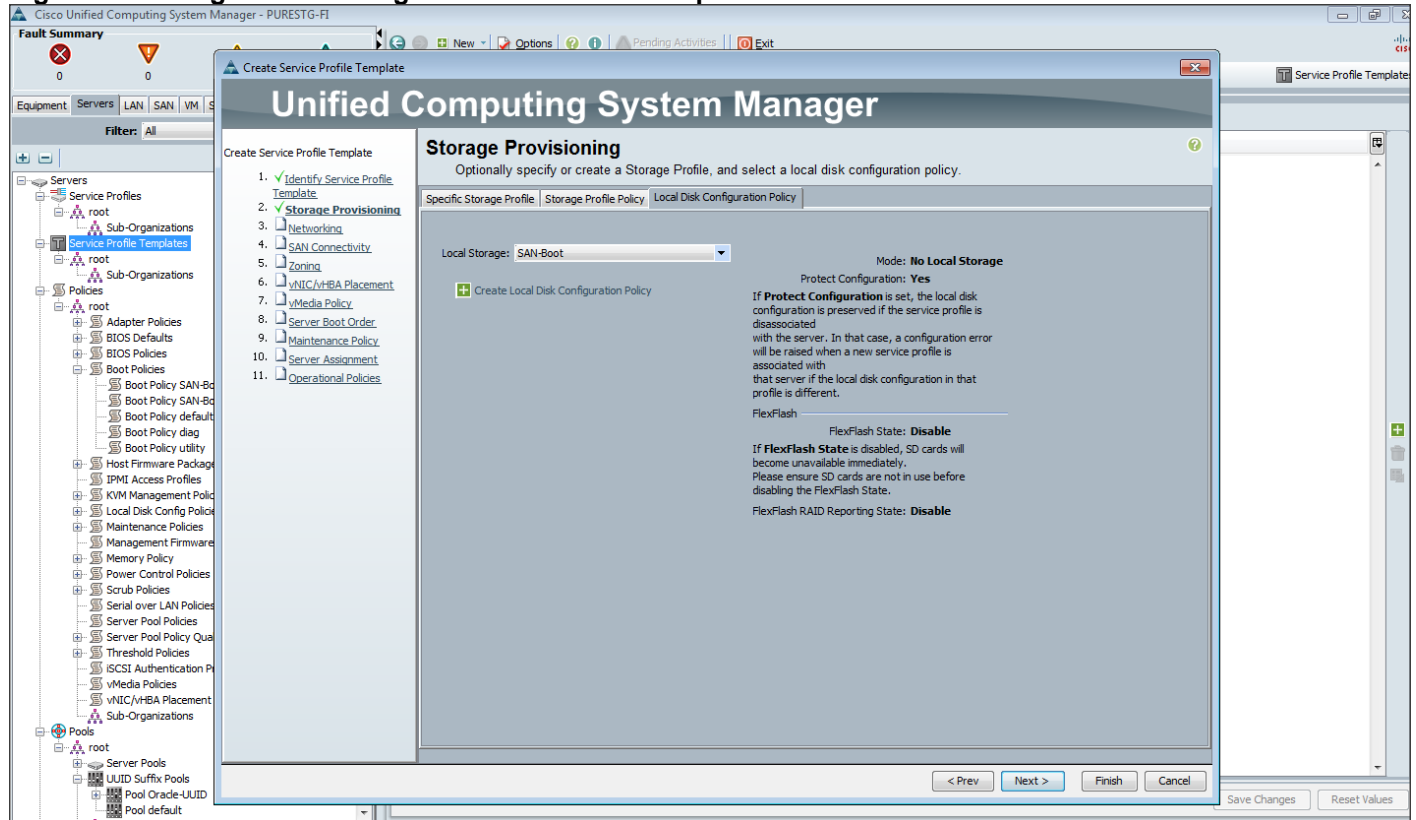
**Figure 62 Service Profile Template**

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

**Figure 63 Identify Service Profile Template**

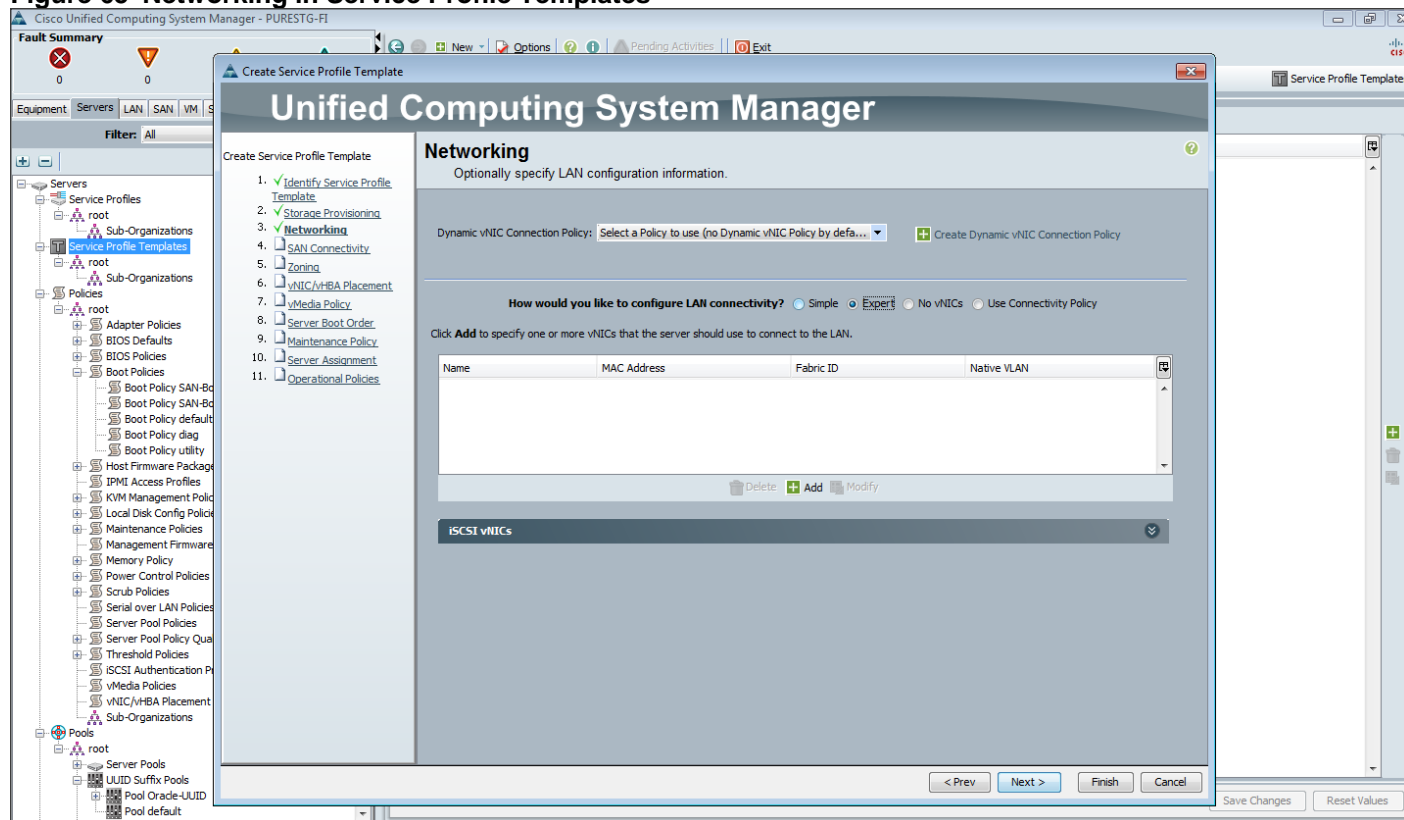


3. Select Local Disk Configuration Policy to SAN-Boot as no Local Storage.

**Figure 64 Storage Provisioning in Service Profile Templates**

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.

Figure 65 Networking in Service Profile Templates



5. Add vNIC as shown as below:

- a. In the Create vNIC menu, name the first vNIC as “eth0” and second vNIC as “eth1”.
- b. For “eth0” Select vNIC Template as Oracle-vNIC-A created earlier.
- c. For “eth1” Select vNIC Template as Oracle-vNIC-B created earlier.



Figure 66 Create vNIC-A “eth0”

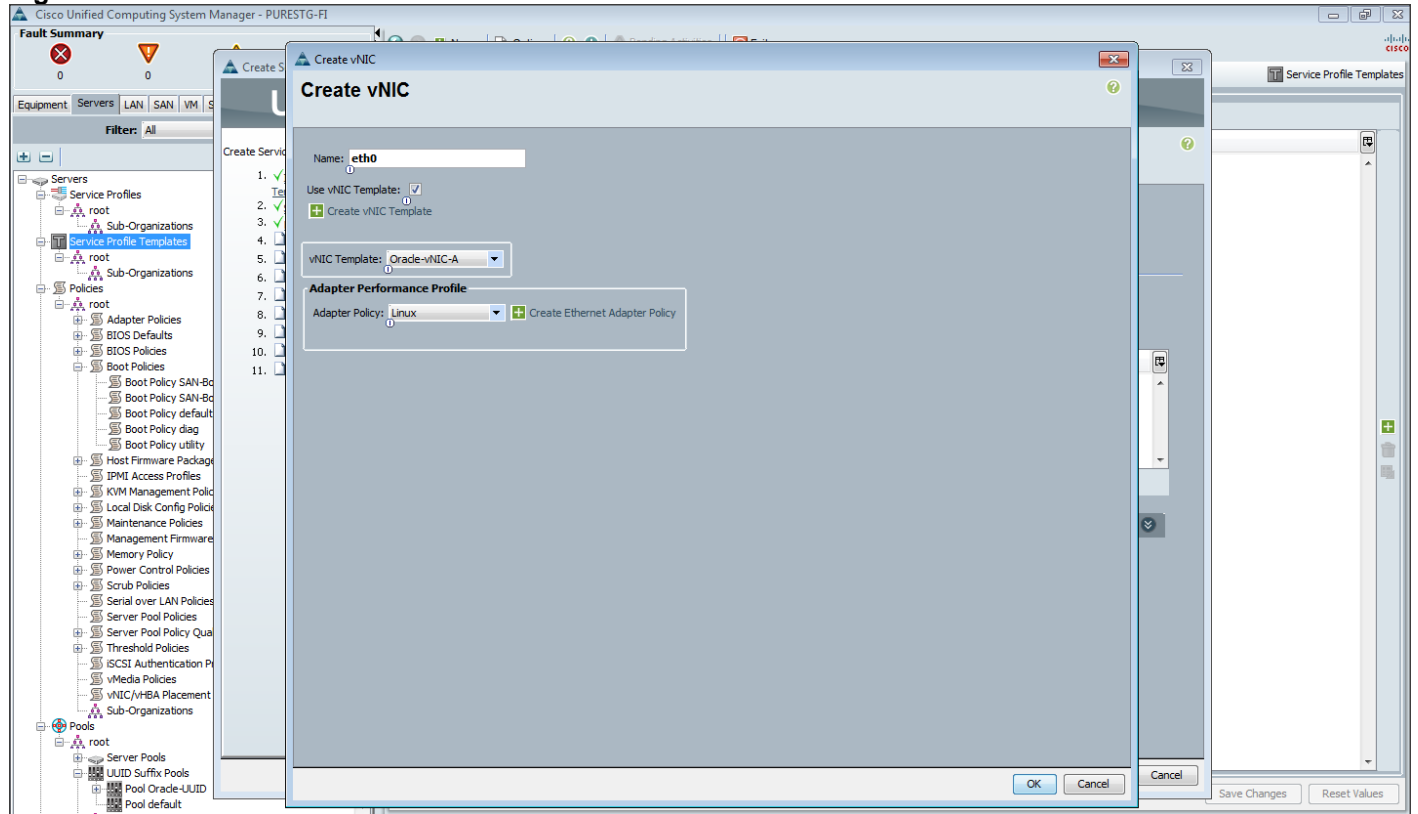


Figure 67 Create vNIC-B eth1

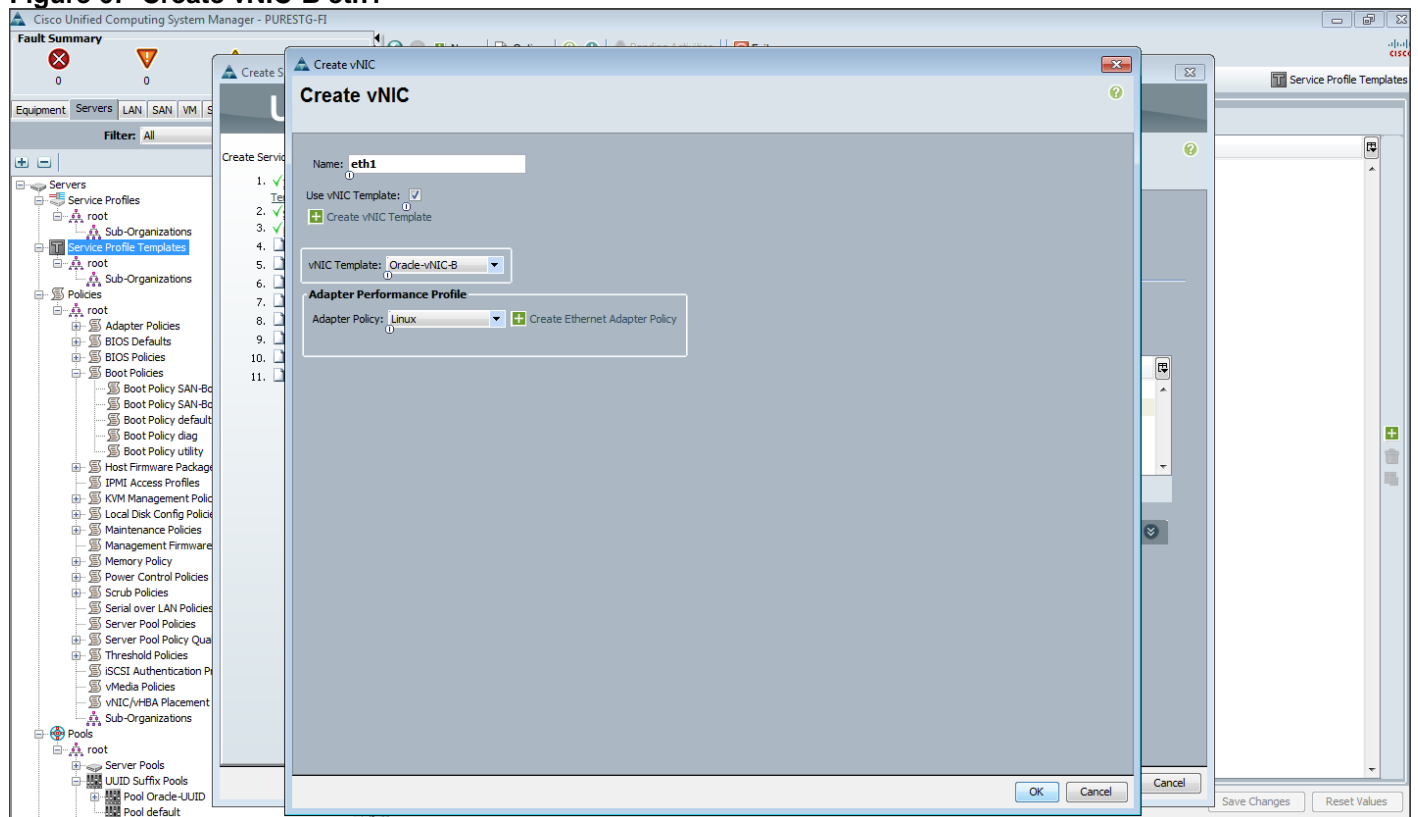
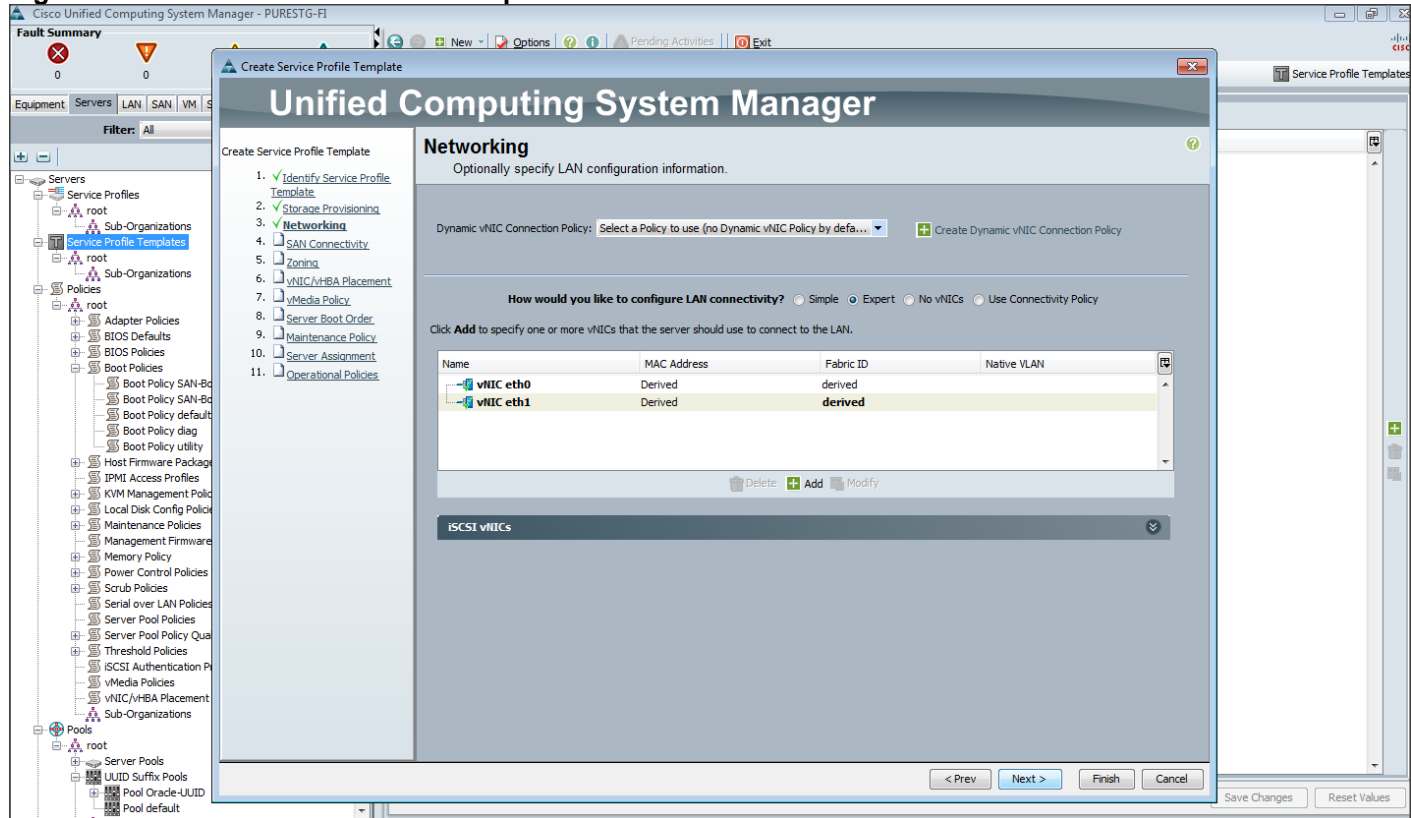


Figure 68 vNICs in Service Profile Template



6. eth0 and eth1 vNICs are created for the Servers to connect to the LAN.
7. Once vNICs are created, create vHBAs. In the SAN Connectivity menu, select Expert to configure as SAN connectivity. Select WWNN (World Wide Node Name) pool which we created earlier. Click Add to add vHBAs as shown below.
8. The following four HBA have 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

Figure 69 vHBA in Service Profile Template

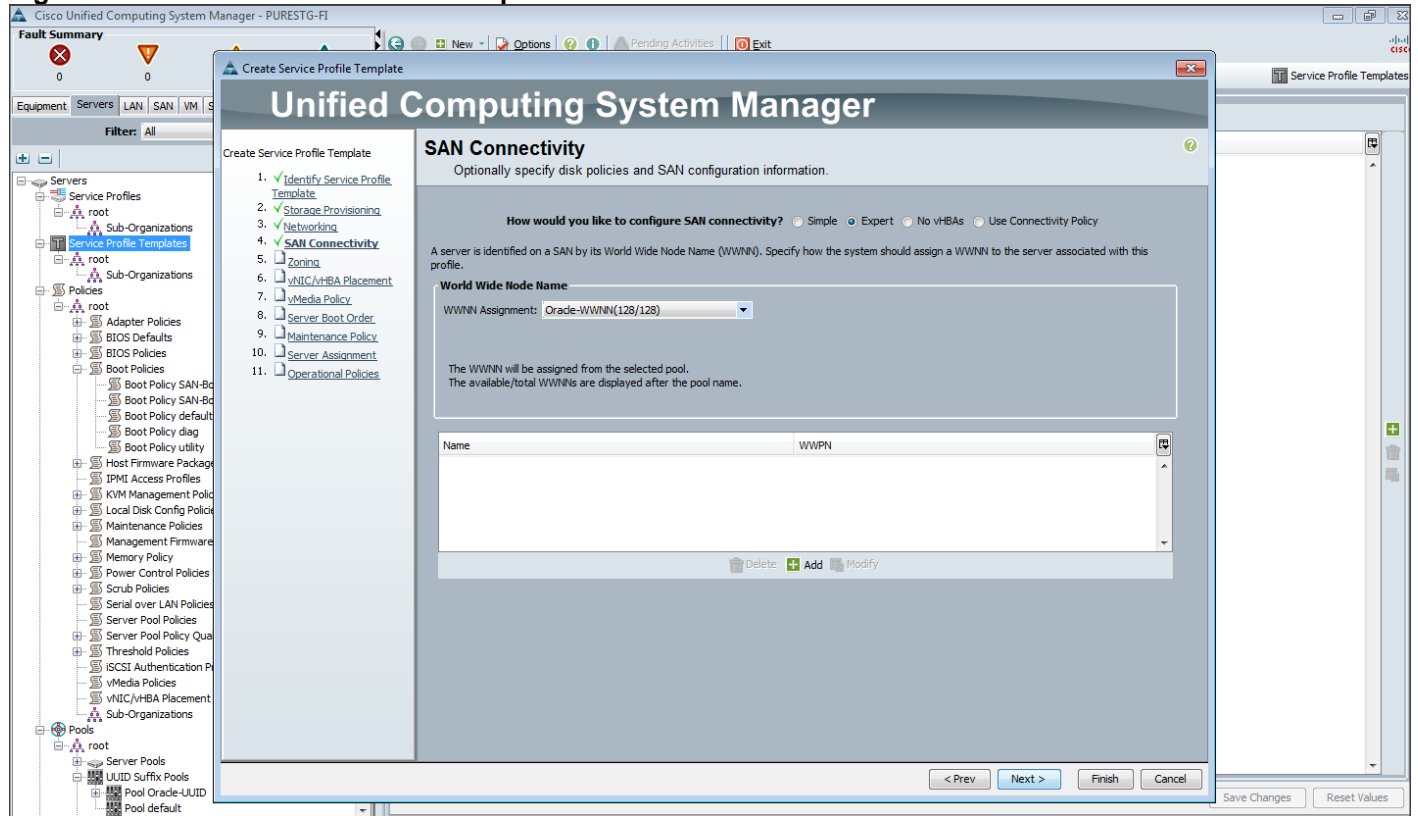


Figure 70 vHBA0

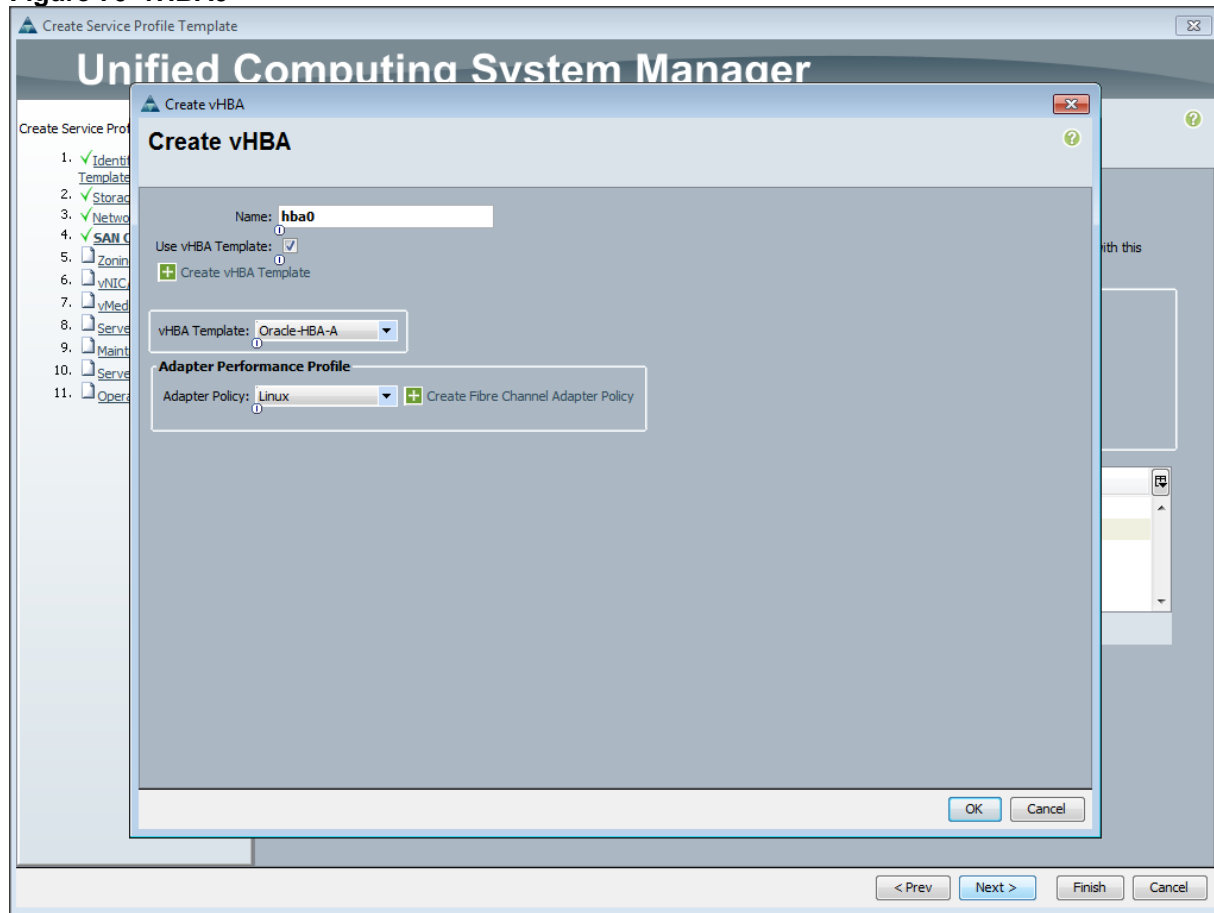


Figure 71 vHBA1

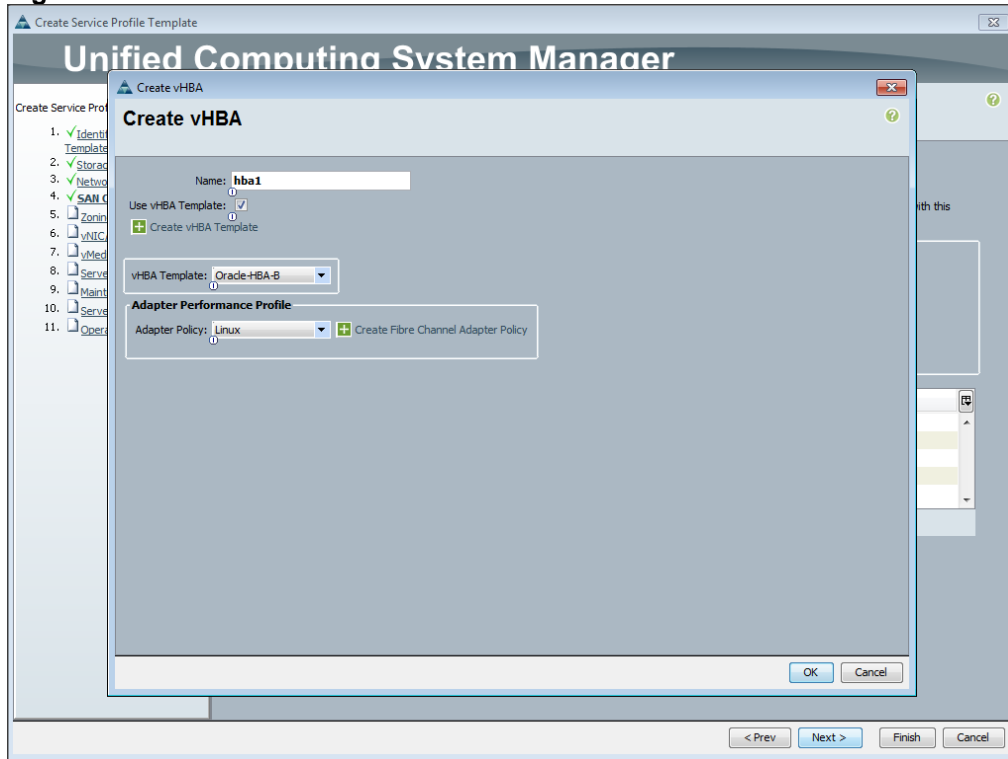


Figure 72 vHBA 2

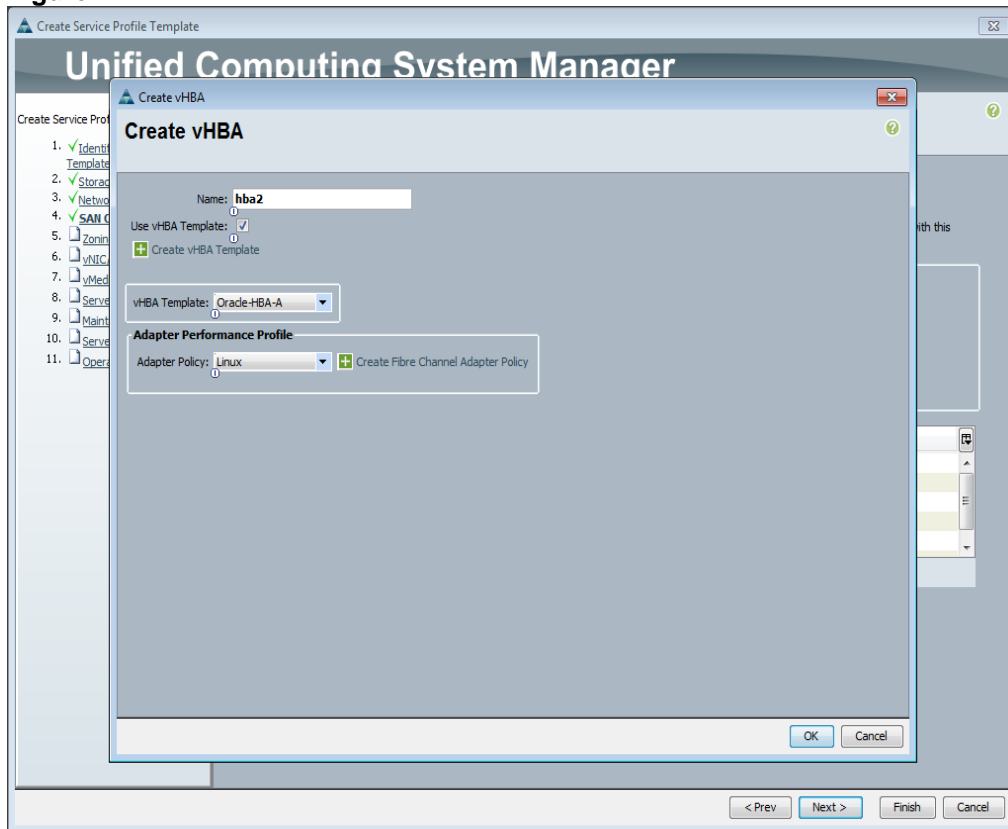


Figure 73 vHBA 3

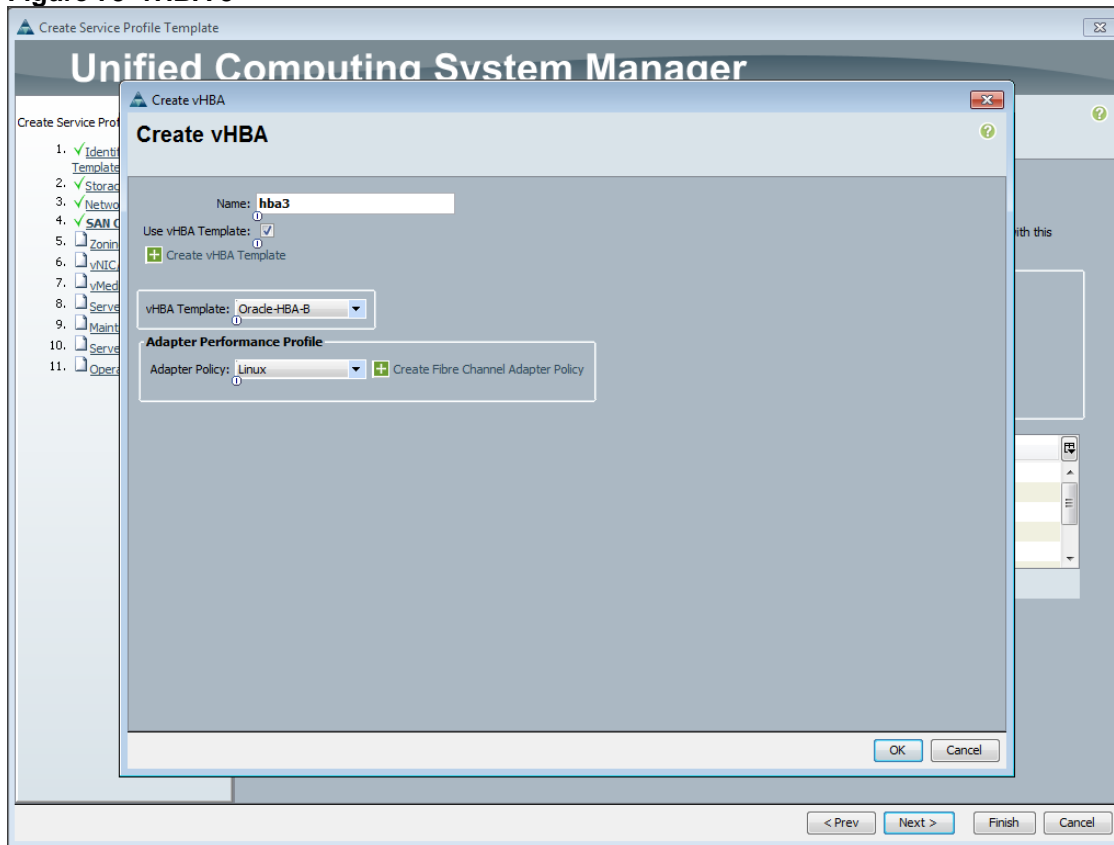
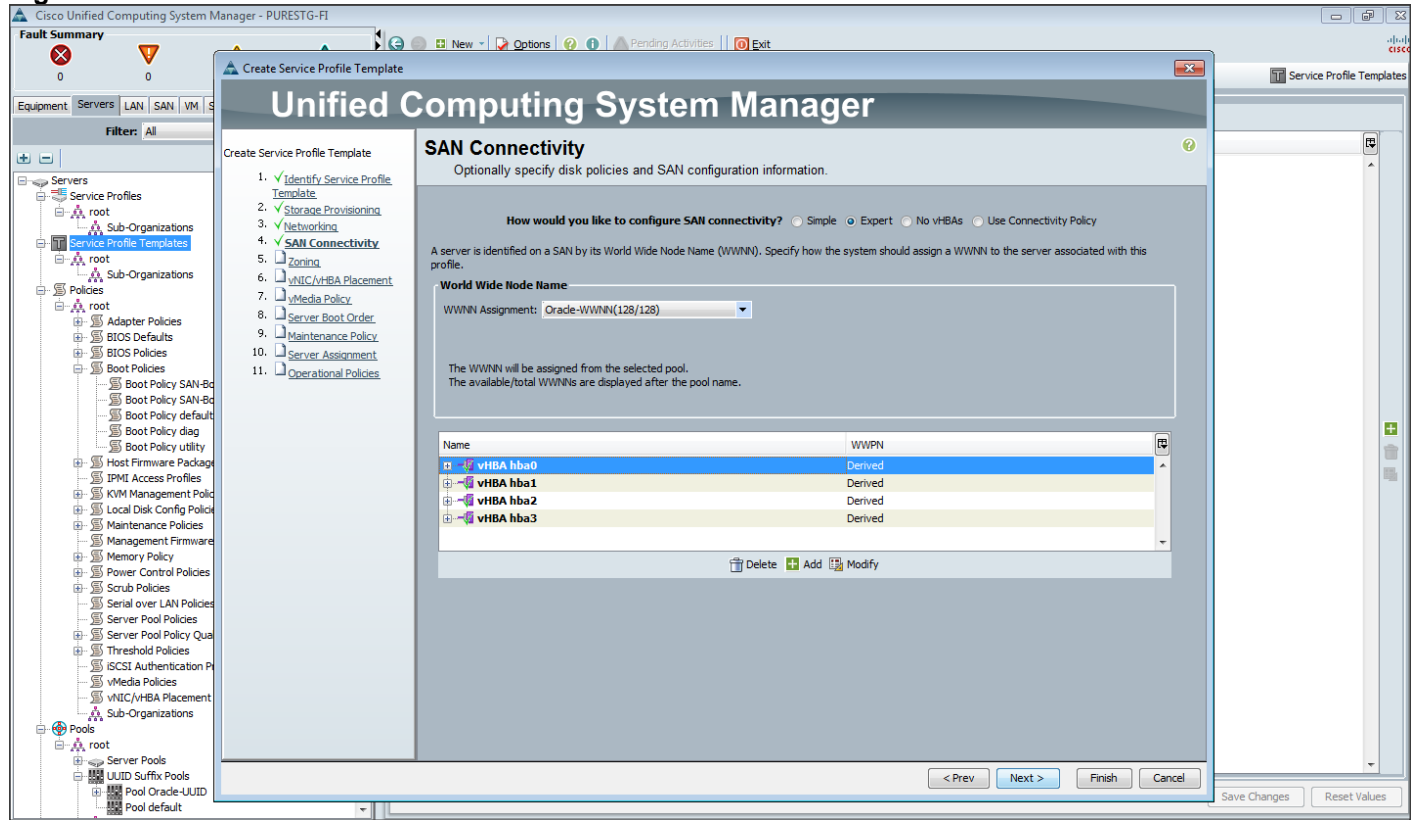


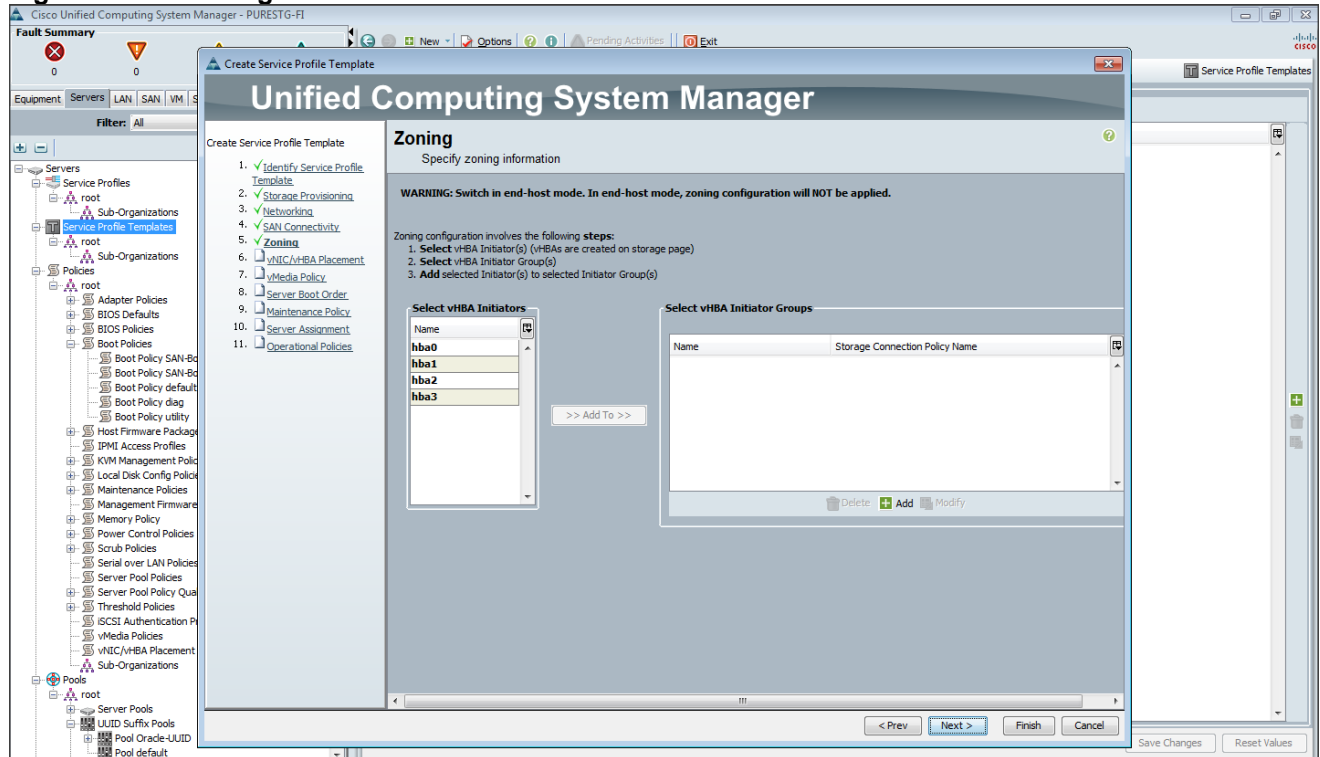
Figure 74 vHBAs



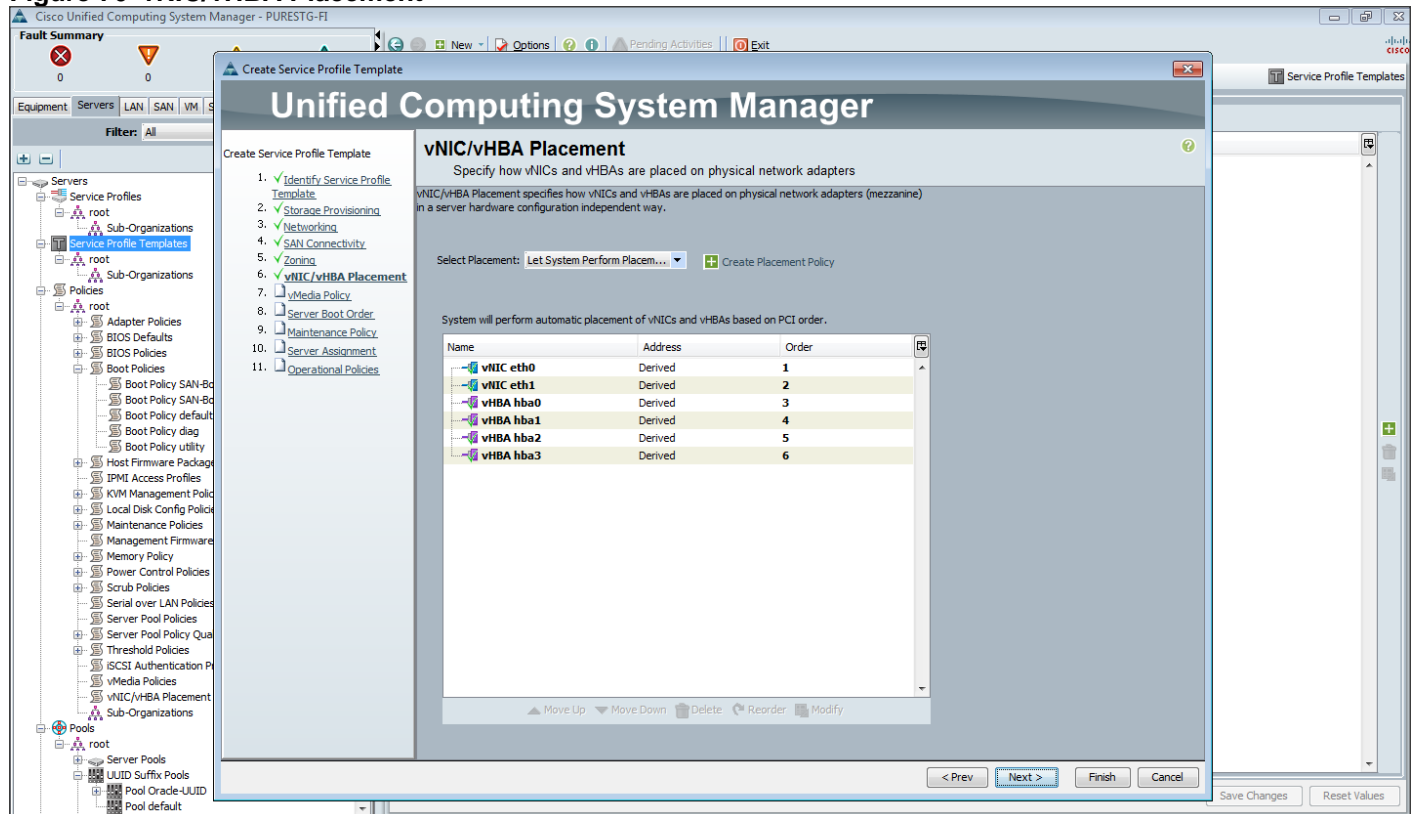
For this Oracle RAC Configuration, the Cisco MDS 9148S is used for zoning. Skip the zoning step.



**Figure 75 No Zoning**

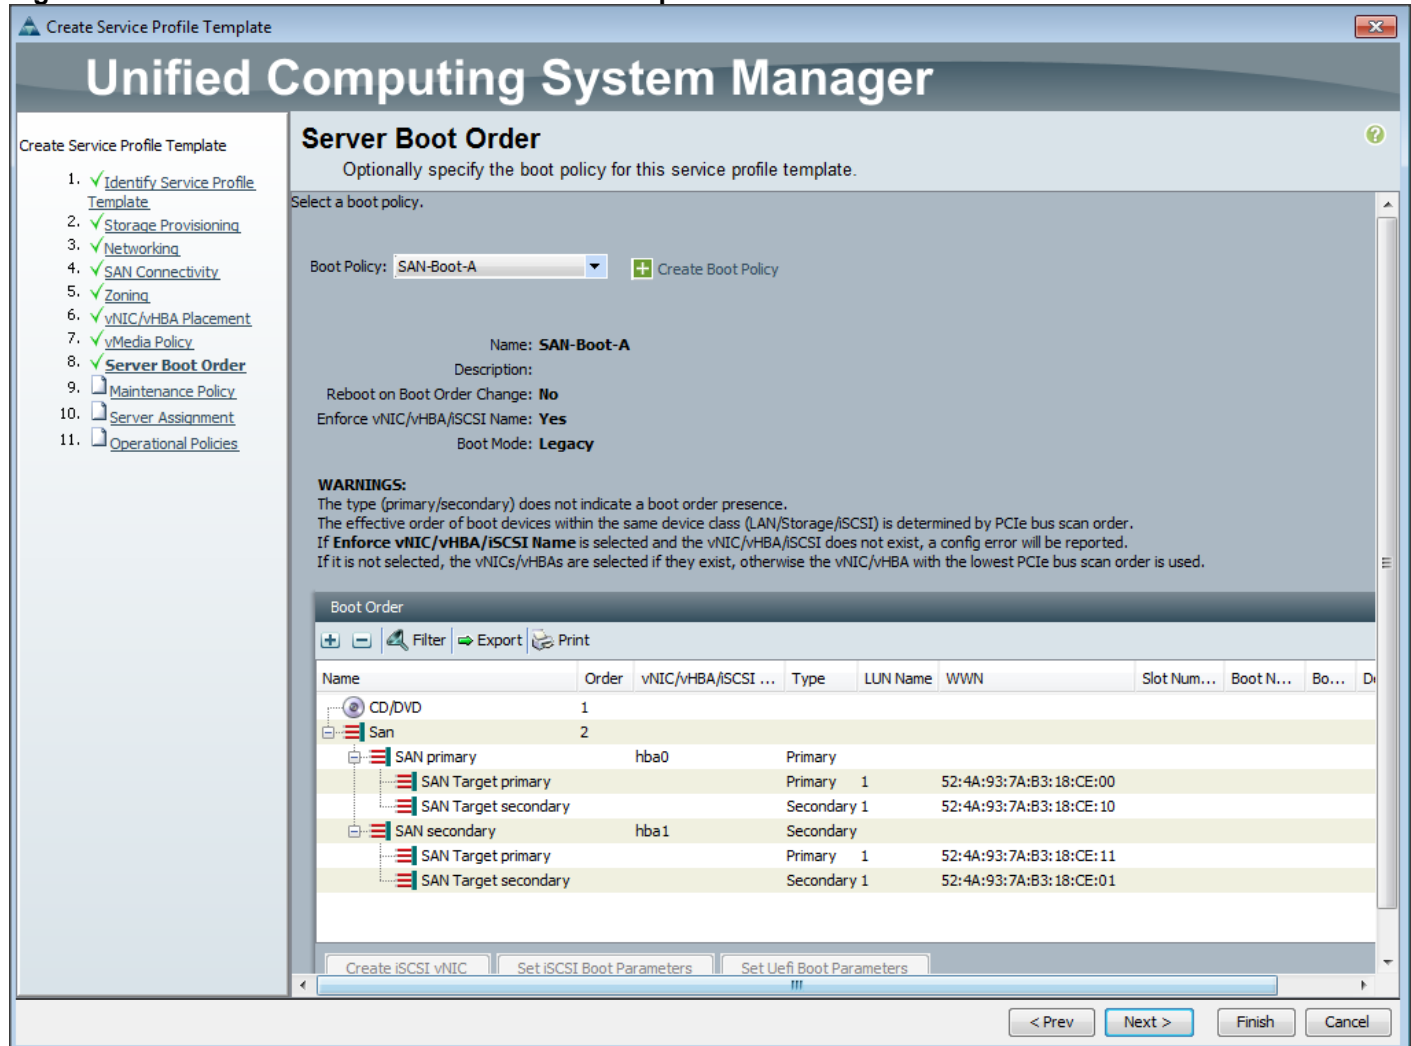


**Figure 76 vNIC/vHBA Placement**



9. For the Server Boot Order, select SAN-Boot-A as Boot Policy created earlier.

Figure 77 Server Boot Order in Service Profile Template



10. The rest maintenance and assignment policies are left as default in the configuration. However, they may vary from site-to-site depending on workloads, best practices, and policies.
11. Click Next then Finish to create Service profile template as ORASAN-A.
12. Create the Service profile template as ORASAN-B with Boot Order as SAN Boot B.
13. Repeat these steps to create the Service Profile Template ORASAN-B and during last step, select Server Boot Order as SAN-Boot-B.
14. Each Service profile template is ORASAN-A and ORASAN-B with each having four vHBAs and two vNICs.

## Create Service Profiles from Template and Associate to Servers

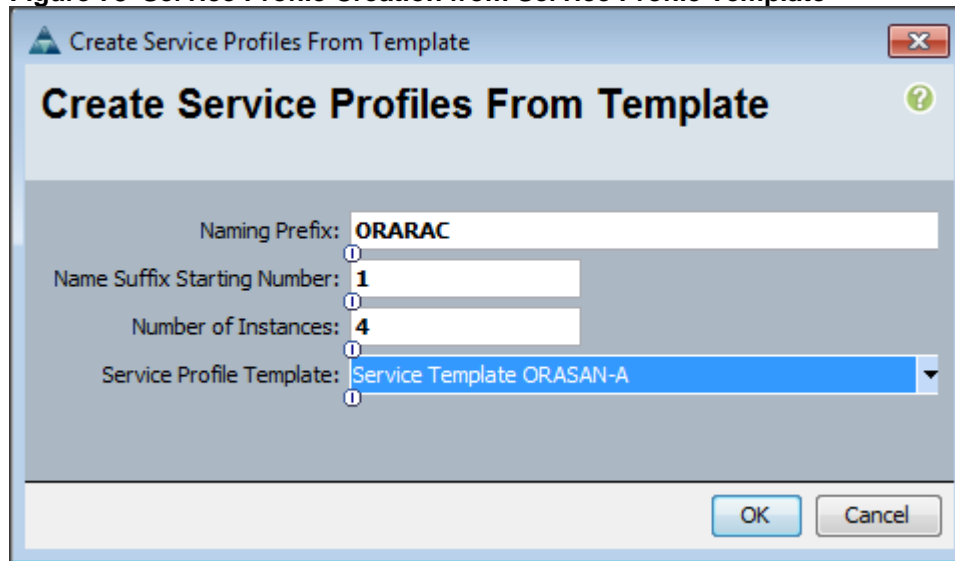
### Create Service Profiles from Template

Eight Service profiles are create for eight Oracle RAC nodes as explained below. The first four Oracle RAC Nodes (orarac1, orarac2, orarac3 and orarac4), need to have four Service Profile from Template "ORASAN-A". The remaining four Oracle RAC Nodes (orarac5, orarac6, orarac7 and orarac8), need to have four Service Profile from Template "ORASAN-B".

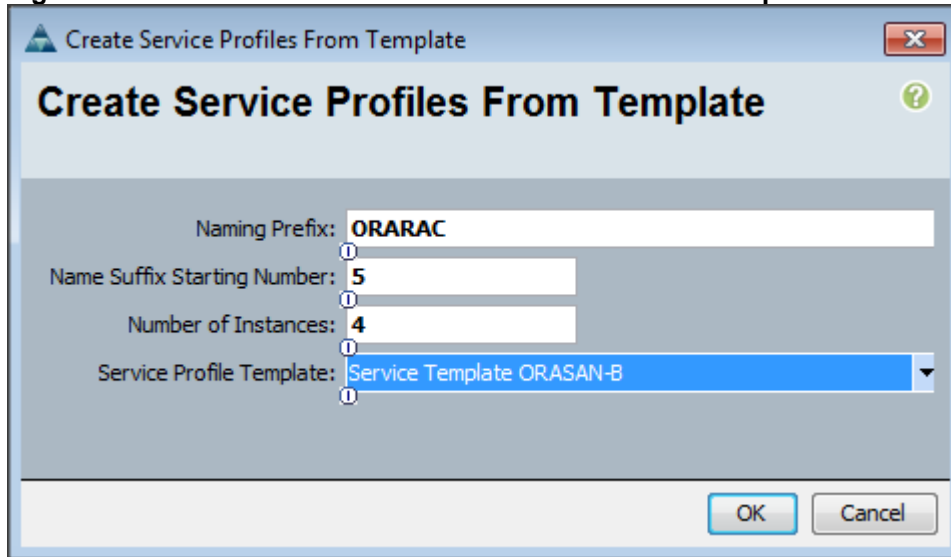
To create the first four Service Profiles from Template, complete the following steps:

1. Go to the Servers tab > Service Profiles > root > and right-click Create Service Profiles from Template.
2. Select the Service profile template ORASAN-A created earlier and name the service profile ORARAC. To create the four service profiles, enter "Number of Instances" as 4 as shown below.

**Figure 78 Service Profile Creation from Service Profile Template**



3. Create another four Service Profiles from Template "ORASAN-B" as shown below.

**Figure 79 Service Profile Creation from Service Profile Template**

4. Once the service profiles are created, associate them to the servers.
5. Assign Service Profile "ORARAC1" to Chassis 1 Server 1, Service Profile "ORARAC2" to Chassis 2 Server 1, Service Profile "ORARAC3" to Chassis 1 Server 2 and, Service Profile "ORARAC4" to Chassis 2 Server 2.
6. Assign Service Profile "ORARAC5" to Chassis 1 Server 3, Service Profile "ORARAC6" to Chassis 2 Server 3, Service Profile "ORARAC7" to Chassis 1 Server 4 and Service Profile "ORARAC8" to Chassis 2 Server 4.

#### Associate Service Profile to the Servers

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

1. Under the servers tab, select the desired service profile, and select change service profile association.
2. Right-click the name of service profile you want to associate with the server and select the option "Change Service Profile Association".
3. In the Change Service Profile Association page, from the Server Assignment drop-down, select existing server that you would like to assign, and click OK.
4. Repeat the same steps to associate remaining seven service profiles for the blade servers.
5. Make sure all the service profiles are associated as shown below.

**Figure 80 Service Profiles Association**

The screenshot displays the Cisco Unified Computing System Manager interface. On the left, a navigation tree shows 'Servers' expanded to 'Service Profiles', with ORARAC1 through ORARAC8 highlighted. The main window shows a table of Service Profiles:

Name	User Label	Overall Status	Assoc State	Server
Service Profile ORARAC1		OK	Associated	sys/chassis-1/blade-1
Service Profile ORARAC2		OK	Associated	sys/chassis-2/blade-1
Service Profile ORARAC3		OK	Associated	sys/chassis-1/blade-2
Service Profile ORARAC4		OK	Associated	sys/chassis-2/blade-2
Service Profile ORARAC5		OK	Associated	sys/chassis-1/blade-3
Service Profile ORARAC6		OK	Associated	sys/chassis-2/blade-3
Service Profile ORARAC7		OK	Associated	sys/chassis-1/blade-4
Service Profile ORARAC8		OK	Associated	sys/chassis-2/blade-4
Service Profile Windows-Client2		OK	Associated	sys/chassis-1/blade-5

Below the table, a green pie chart indicates the 'Associated' state, with a legend showing a green circle next to the label 'Associated'.



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

This completes the configuration required for Cisco UCS Manager Setup.

## Configure Cisco Nexus 9372PX-E Switches

This section details the steps for the Cisco Nexus 9372PX-E switch configuration. The full details of “show run” output is listed in the Appendix.

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

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

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

```

conf terminal

spanning-tree port type network default
spanning-tree port type edge bpduguard default
port-channel load-balance ethernet source-dest-port
policy-map type network-qos jumbo
class type network-qos class-default
mtu 9216

exit

```

```

class type network-qos class-fcoe
pause no-drop
mtu 2158
exit
exit
system qos
service-policy type network-qos jumbo
exit
copy run start

```

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

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

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

1. Login as admin user into Nexus Switch A.
2. Create VLAN 134 for Public traffic

```

PURESTG-NEXUS-A# config terminal
PURESTG-NEXUS-A(config)# VLAN 134
PURESTG-NEXUS-A(config-VLAN)# name Oracle_Public_Traffic
PURESTG-NEXUS-A(config-VLAN)# no shutdown
PURESTG-NEXUS-A(config-VLAN)# exit
PURESTG-NEXUS-A(config)# copy running-config startup-config
PURESTG-NEXUS-A(config)# exit

```

3. Create VLAN 10 for Private traffic

```

PURESTG-NEXUS-A# config terminal
PURESTG-NEXUS-A(config)# VLAN 10
PURESTG-NEXUS-A(config-VLAN)# name Oracle_Private_Traffic
PURESTG-NEXUS-A(config-VLAN)# no shutdown
PURESTG-NEXUS-A(config-VLAN)# exit
PURESTG-NEXUS-A(config)# copy running-config startup-config
PURESTG-NEXUS-A(config)# exit

```

4. Log in as admin user into Nexus Switch B and create VLAN 134 for Public traffic and VLAN 10 for Private Traffic sameway.

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

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

**Table 8 vPC Summary**


vPC Domain	vPC Name	vPC ID
1	Peer-Link	1
1	vPC Public	19
1	vPC Private	20

As listed in the table above, a single vPC domain with Domain ID 1 is created across two Cisco Nexus 9372PX-E member switches to define vPC members to carry specific VLAN network traffic. In this topology, we defined a total of 3 vPCs. vPC ID 1 is defined as Peer link communication between two Nexus switches in Fabric A and B. vPC IDs 19 and 20 are defined for public and private traffic from Cisco UCS fabric interconnects.

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

To create vPC peer-link between two Nexus switches, complete the following steps:

**Figure 81 Cisco Nexus Switch Peer-Link**



1. Log in to the Cisco Nexus 9372PX-E switch as “admin” user.



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

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

```
PURESTG-NEXUS-A# config terminal
PURESTG-NEXUS-A(config)#feature vpc
PURESTG-NEXUS-A(config)#feature lacp
PURESTG-NEXUS-A(config)#vpc domain 1
PURESTG-NEXUS-A(config-vpc-domain)# peer-keepalive destination 10.29.134.154
source 10.29.134.153
PURESTG-NEXUS-A(config-vpc-domain)# exit
```



```
PURESTG-NEXUS-A(config)# interface port-channel 1
PURESTG-NEXUS-A(config-if)# description VPC peer-link
PURESTG-NEXUS-A(config-if)# switchport mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type network
PURESTG-NEXUS-A(config-if)# vpc peer-link
PURESTG-NEXUS-A(config-if)# exit

PURESTG-NEXUS-A(config)# interface Ethernet1/1
PURESTG-NEXUS-A(config-if)# description Nexus5k-B-Cluster-Interconnect
PURESTG-NEXUS-A(config-if)# switchport mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-A(config-if)# channel-group 1 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit

PURESTG-NEXUS-A(config)# interface Ethernet1/2
PURESTG-NEXUS-A(config-if)# description Nexus5k-B-Cluster-Interconnect
PURESTG-NEXUS-A(config-if)# switchport mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-A(config-if)# channel-group 1 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit

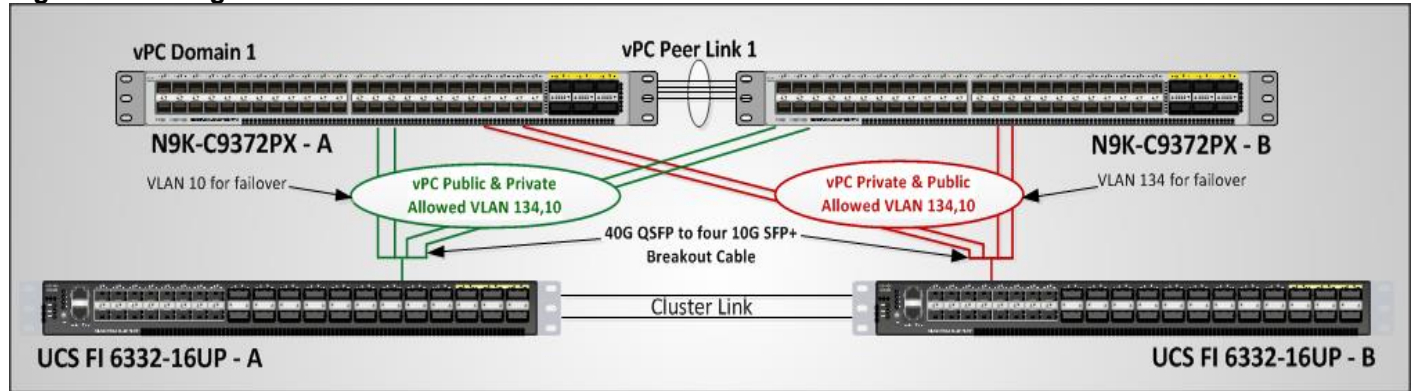
PURESTG-NEXUS-A(config)# interface Ethernet1/15
PURESTG-NEXUS-A(config-if)# description connect to uplink switch
PURESTG-NEXUS-A(config-if)# switchport access vlan 134
PURESTG-NEXUS-A(config-if)# speed 1000
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)#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

2. Log in to the Cisco Nexus 9372PX-E B switch as “admin” user and repeat steps to complete second switch configuration.

### Create vPC Configuration between Nexus 9372PX-E and Fabric Interconnects

Create and configure vPC 19 and 20 for Data network between Cisco Nexus 9372PX-E switches and Fabric Interconnects.

**Figure 82 Configuration between Nexus Switch and Fabric Interconnects**



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

**Table 9 VLAN IDs**

vPC Description	vPC ID Nexus 9372PX-E	Fabric Interconnect uplink ports	Cisco Nexus 9372PX-E ports	Allowed VLANs
<b>Port Channel FI-A</b>	19	FI-A P 1/19/1	N9KA P 17	134,10, Note: VLAN 10 needed for failover
		FI-A P 1/19/2	N9KA P 18	
		FI-A P 1/19/3	N9KB P 17	
		FI-A P 1/19/4	N9KB P 18	
<b>Port-Channel FI-B</b>	20	FI-A P 1/19/1	N9KA P 19	10,134 Note: VLAN 134 needed for failover
		FI-A P 1/19/2	N9KA P 20	
		FI-A P 1/19/3	N9KB P 19	
		FI-A P 1/19/4	N9KB P 20	

The following are the configuration details for Cisco Nexus 9372PX-E:

1. Log into the Cisco Nexus 9372PX-E A switch as “admin” user and complete the following steps.
2. To create the necessary port channels between devices, follow these steps on both Cisco Nexus Switches:

```
PURESTG-NEXUS-A# config Terminal
```

```
PURESTG-NEXUS-A(config)# interface port-channel19
```

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

```
PURESTG-NEXUS-A(config-if)# switchport mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# vpc 19
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface port-channel20
PURESTG-NEXUS-A(config-if)# description connect to Fabric Interconnect B
PURESTG-NEXUS-A(config-if)# switchport mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed VLAN 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# vpc 20
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit

PURESTG-NEXUS-A(config)# interface Ethernet1/17
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-A:1/19
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 19 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface Ethernet1/18
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-A:1/19
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 19 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
```

```
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface Ethernet1/19
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-B:1/19
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 20 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# interface Ethernet1/20
PURESTG-NEXUS-A(config-if)# description Fabric-Interconnect-B:1/19
PURESTG-NEXUS-A(config-if)# switch mode trunk
PURESTG-NEXUS-A(config-if)# switchport trunk allowed vlan 1,10,134
PURESTG-NEXUS-A(config-if)# spanning-tree port type edge trunk
PURESTG-NEXUS-A(config-if)# mtu 9216
PURESTG-NEXUS-A(config-if)# channel-group 20 mode active
PURESTG-NEXUS-A(config-if)# no shutdown
PURESTG-NEXUS-A(config-if)# exit
PURESTG-NEXUS-A(config)# copy running-config startup-config
```

3. Log into the Cisco Nexus 9372PX-E B switch as “admin” user and repeat steps to complete second switch configuration.

## Verify vPC Status on both Cisco Nexus 9372PX-E Switches

Figure 83 Cisco Nexus Switch A Status

```

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

```

Figure 84 Cisco Nexus Switch B Status

```

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

```

**Figure 85 vPC description for Cisco Nexus Switch A**

```

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

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

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

vPC status
-----
id   Port   Status Consistency Reason      Active vlans
--   --
19   Po19   up     success    success    1,10,134
20   Po20   up     success    success    1,10,134

PURESTG-NEXUS-A# █

```

**Figure 86 vPC description for Cisco Nexus Switch B**

```

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

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

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

vPC status
-----
id  Port  Status Consistency Reason           Active vlans
--  --
19  Po19  up    success  success  1,10,134
20  Po20  up    success  success  1,10,134

PURESTG-NEXUS-B# █

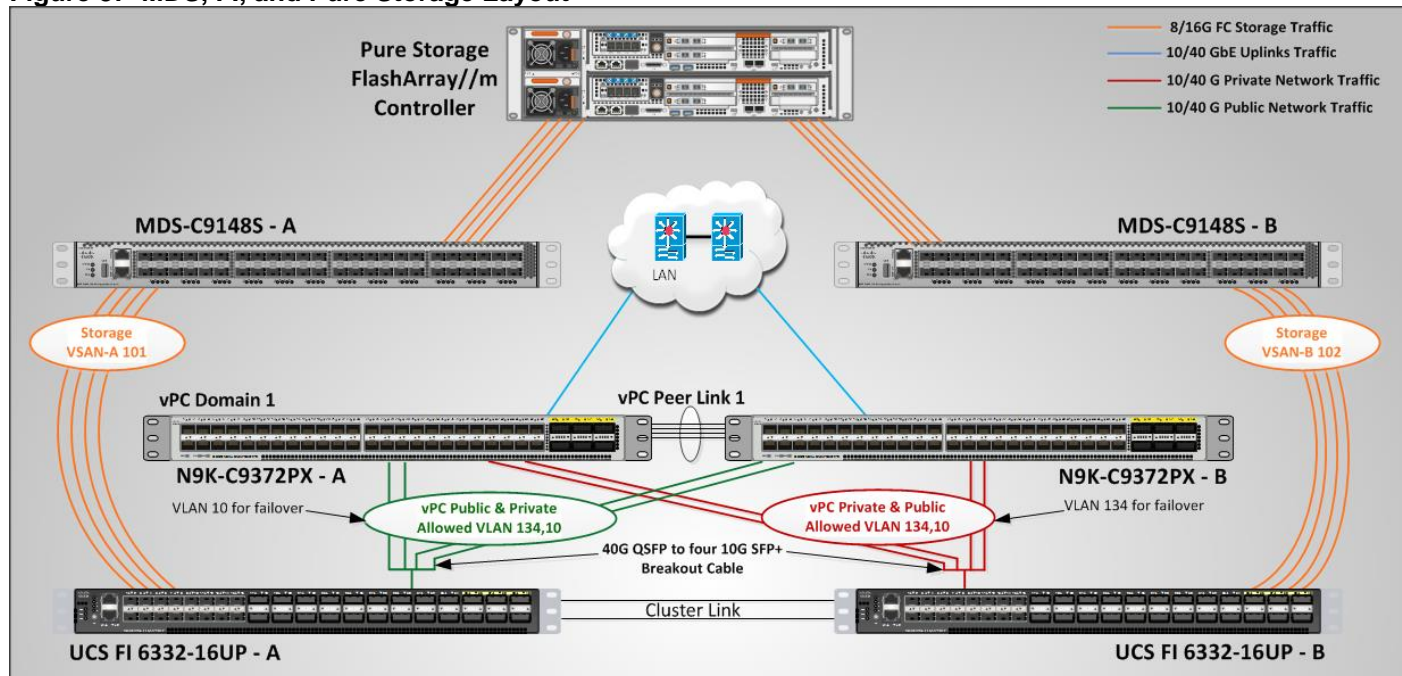
```

## Configure Cisco MDS 9148S Switches

The MDS Switches are connected to the Fabric Interconnects and Pure Storage System as shown in the figure below.



Figure 87 MDS, FI, and Pure Storage Layout



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

Table 10 MDS 9148S Port Connection to Fabric Interconnects

MDS Switch	MDS Switch Port	Fabric Interconnect	FI Port
MDS Switch A	fc1/1	Fabric Interconnect A (FI-A)	FI-A fc port 1/1
	fc1/2		FI-A fc port 1/2
	fc1/3		FI-A fc port 1/3
	fc1/4		FI-A fc port 1/4
MDS Switch B	fc1/1	Fabric Interconnect B (FI-B)	FI-B fc port 1/1
	fc1/2		FI-B fc port 1/2
	fc1/3		FI-B fc port 1/3
	fc1/4		FI-B fc port 1/4



For this solution, we connected four ports (ports 9-12) of MDS Switch A to Pure Storage System. Similarly, we connected four ports (ports 9-12) of MDS Switch B to Pure Storage System as shown in the table below. All ports carry 16 GB FC Traffic.

**Table 11 MDS 9148S Port Connection to Pure Storage System**

MDS Switch	MDS Switch Port	Pure Storage	Storage Port
<b>MDS Switch A</b>	fc1/9	Storage Controller-0	CT0-FC0
	fc1/10	Storage Controller-1	CT1-FC0
	fc1/11	Storage Controller-0	CT0-FC2
	fc1/12	Storage Controller-1	CT1-FC2
<b>MDS Switch B</b>	fc1/9	Storage Controller-0	CT0-FC1
	fc1/10	Storage Controller-1	CT1-FC1
	fc1/11	Storage Controller-0	CT0-FC3
	fc1/12	Storage Controller-1	CT1-FC3

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

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

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

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

2. Log in as admin user into MDS Switch B.

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

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

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

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

2. Create VSAN 101 for Storage Traffic

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

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

3. Login as admin user into MDS Switch B.

4. Create VSAN 102 for Storage Traffic.

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

PURESTG-MDS-B(config)# interface fc 1/1-12
PURESTG-MDS-B(config-if)# switchport trunk allowed vsan 102
PURESTG-MDS-B(config-if)# switchport trunk mode off
PURESTG-MDS-B(config-if)# port-license acquire
PURESTG-MDS-B(config-if)# no shutdown
PURESTG-MDS-B(config-if)# exit
PURESTG-MDS-B(config)# copy running-config startup-config
```

## Create and Configure Fiber Channel Zoning

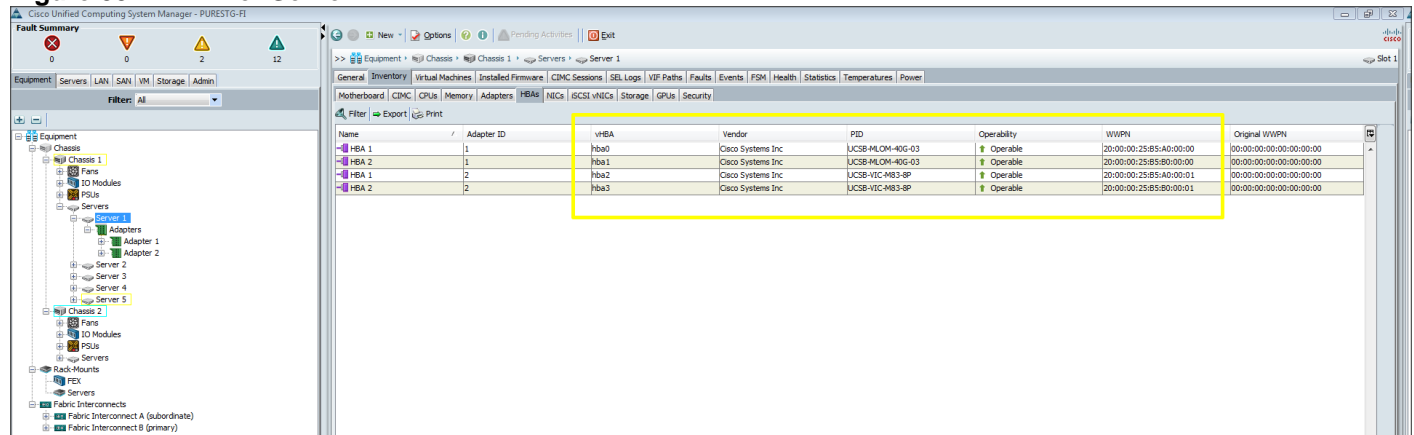
This procedure sets up the Fibre Channel connections between the Cisco MDS 9148S switches, the Cisco UCS Fabric Interconnects, and the Pure Storage FlashArray systems.

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

To create and configure the Fiber Channel Zoning, complete the following steps:

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

**Figure 88 WWPN of Server**



3. Connect to the Pure Storage System and extract the WWPN of FC Ports connected to the Cisco MDS Switches. We have connected 8 FC ports from Pure Storage System to Cisco MDS Switches. FC ports CT0.FC0, CT1.FC0, CT0.FC2, CT1.FC2 are connected to MDS Switch-A and similarly FC ports CT0.FC1, CT1.FC1, CT0.FC3, CT1.FC3 are connected to MDS Switch-B.

Figure 89 WWPN of Pure Storage

The screenshot shows the Pure Storage management interface. The 'Host Connections' section displays two controllers, ORARAC1 and ORARAC2, each with a 'Redundant' connectivity status. Below this, the 'Target Ports' table lists the specific FC ports and their WWPNs for both controllers.

Host Port Connectivity		CT0				CT1			
HOST / PORT	CONNECTIVITY	FC0	FC1	FC2	FC3	FC0	FC1	FC2	FC3
ORARAC1	Redundant	✓	✓	✓	✓	✓	✓	✓	✓
20:00:00:25:B5:A0:00:00		✗	✗	✗	✗	✗	✗	✗	✗
20:00:00:25:B5:A0:00:01		✗	✗	✗	✗	✗	✗	✗	✗
20:00:00:25:B5:B0:00:00		✗	✗	✗	✗	✗	✗	✗	✗
20:00:00:25:B5:B0:00:01		✗	✗	✗	✗	✗	✗	✗	✗
ORARAC2	Redundant	✓	✓	✓	✓	✓	✓	✓	✓
20:00:00:25:B5:A0:00:02		✗	✗	✗	✗	✗	✗	✗	✗
20:00:00:25:B5:A0:00:03		✗	✗	✗	✗	✗	✗	✗	✗
20:00:00:25:B5:B0:00:02		✗	✗	✗	✗	✗	✗	✗	✗
20:00:00:25:B5:B0:00:03		✗	✗	✗	✗	✗	✗	✗	✗

Target Ports		Target Ports	
PORT	NAME	PORT	NAME
CT0.FC0	52:4A:93:7A:B3:18:CE:00	CT1.FC0	52:4A:93:7A:B3:18:CE:10
CT0.FC1	52:4A:93:7A:B3:18:CE:01	CT1.FC1	52:4A:93:7A:B3:18:CE:11
CT0.FC2	52:4A:93:7A:B3:18:CE:02	CT1.FC2	52:4A:93:7A:B3:18:CE:12
CT0.FC3	52:4A:93:7A:B3:18:CE:03	CT1.FC3	52:4A:93:7A:B3:18:CE:13

## Create Device Aliases for Fiber Channel Zoning

### Cisco MDS Switch A

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

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

```

conf t
device-alias database
device-alias name ORARAC-A1-hba0 pwnn 20:00:00:25:b5:a0:00:00
device-alias name ORARAC-A1-hba2 pwnn 20:00:00:25:b5:a0:00:01
device-alias name ORARAC-A2-hba0 pwnn 20:00:00:25:b5:a0:00:02
device-alias name ORARAC-A2-hba2 pwnn 20:00:00:25:b5:a0:00:03
device-alias name ORARAC-A3-hba0 pwnn 20:00:00:25:b5:a0:00:04
device-alias name ORARAC-A3-hba2 pwnn 20:00:00:25:b5:a0:00:05
device-alias name ORARAC-A4-hba0 pwnn 20:00:00:25:b5:a0:00:06
device-alias name ORARAC-A4-hba2 pwnn 20:00:00:25:b5:a0:00:07
device-alias name ORARAC-A5-hba0 pwnn 20:00:00:25:b5:a0:00:08
device-alias name ORARAC-A5-hba2 pwnn 20:00:00:25:b5:a0:00:09
device-alias name ORARAC-A6-hba0 pwnn 20:00:00:25:b5:a0:00:0a
device-alias name ORARAC-A6-hba2 pwnn 20:00:00:25:b5:a0:00:0b
device-alias name ORARAC-A7-hba0 pwnn 20:00:00:25:b5:a0:00:0c
device-alias name ORARAC-A7-hba2 pwnn 20:00:00:25:b5:a0:00:0d
device-alias name ORARAC-A8-hba0 pwnn 20:00:00:25:b5:a0:00:0e
device-alias name ORARAC-A8-hba2 pwnn 20:00:00:25:b5:a0:00:0f
device-alias name Pure-STG-CT0-FC0 pwnn 52:4a:93:7a:b3:18:ce:00
device-alias name Pure-STG-CT0-FC2 pwnn 52:4a:93:7a:b3:18:ce:02
device-alias name Pure-STG-CT1-FC0 pwnn 52:4a:93:7a:b3:18:ce:10
device-alias name Pure-STG-CT1-FC2 pwnn 52:4a:93:7a:b3:18:ce:12
exit
device-alias commit

```

## Cisco MDS Switch B

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

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

```
conf t
device-alias database
device-alias name ORARAC-A1-hba1 pwn 20:00:00:25:b5:b0:00:00
device-alias name ORARAC-A1-hba3 pwn 20:00:00:25:b5:b0:00:01
device-alias name ORARAC-A2-hba1 pwn 20:00:00:25:b5:b0:00:02
device-alias name ORARAC-A2-hba3 pwn 20:00:00:25:b5:b0:00:03
device-alias name ORARAC-A3-hba1 pwn 20:00:00:25:b5:b0:00:04
device-alias name ORARAC-A3-hba3 pwn 20:00:00:25:b5:b0:00:05
device-alias name ORARAC-A4-hba1 pwn 20:00:00:25:b5:b0:00:06
device-alias name ORARAC-A4-hba3 pwn 20:00:00:25:b5:b0:00:07
device-alias name ORARAC-A5-hba1 pwn 20:00:00:25:b5:b0:00:08
device-alias name ORARAC-A5-hba3 pwn 20:00:00:25:b5:b0:00:09
device-alias name ORARAC-A6-hba1 pwn 20:00:00:25:b5:b0:00:0a
device-alias name ORARAC-A6-hba3 pwn 20:00:00:25:b5:b0:00:0b
device-alias name ORARAC-A7-hba1 pwn 20:00:00:25:b5:b0:00:0c
device-alias name ORARAC-A7-hba3 pwn 20:00:00:25:b5:b0:00:0d
device-alias name ORARAC-A8-hba1 pwn 20:00:00:25:b5:b0:00:0e
device-alias name ORARAC-A8-hba3 pwn 20:00:00:25:b5:b0:00:0f
device-alias name Pure-STG-CT0-FC1 pwn 52:4a:93:7a:b3:18:ce:01
device-alias name Pure-STG-CT0-FC3 pwn 52:4a:93:7a:b3:18:ce:03
device-alias name Pure-STG-CT1-FC1 pwn 52:4a:93:7a:b3:18:ce:11
device-alias name Pure-STG-CT1-FC3 pwn 52:4a:93:7a:b3:18:ce:13
exit
device-alias commit
```

## Create Zoning

### Cisco MDS Switch A

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

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

```
conf t
zone name chas1-server1-boot-hba vsan 101
  member pwn Pure-STG-CT0-FC0
  member pwn ORARAC-A1-hba0
  member pwn Pure-STG-CT1-FC0
  member pwn ORARAC-A1-hba2
  member pwn Pure-STG-CT0-FC2
  member pwn Pure-STG-CT1-FC2
exit
zone name chas1-server2-boot-hba vsan 101
  member pwn Pure-STG-CT0-FC0
  member pwn ORARAC-A2-hba0
  member pwn Pure-STG-CT1-FC0
  member pwn ORARAC-A2-hba2
  member pwn Pure-STG-CT0-FC2
  member pwn Pure-STG-CT1-FC2
exit
```

```

zone name chas1-server3-boot-hba vsan 101
  member pwnn Pure-STG-CT0-FC0
  member pwnn ORARAC-A3-hba0
  member pwnn Pure-STG-CT1-FC0
  member pwnn ORARAC-A3-hba2
  member pwnn Pure-STG-CT0-FC2
  member pwnn Pure-STG-CT1-FC2
zone name chas1-server4-boot-hba vsan 101
  member pwnn Pure-STG-CT0-FC0
  member pwnn ORARAC-A4-hba0
  member pwnn Pure-STG-CT1-FC0
  member pwnn ORARAC-A4-hba2
  member pwnn Pure-STG-CT0-FC2
  member pwnn Pure-STG-CT1-FC2
zone name chas2-server1-boot-hba vsan 101
  member pwnn Pure-STG-CT0-FC0
  member pwnn ORARAC-A5-hba0
  member pwnn Pure-STG-CT1-FC0
  member pwnn ORARAC-A5-hba2
  member pwnn Pure-STG-CT0-FC2
  member pwnn Pure-STG-CT1-FC2
zone name chas2-server2-boot-hba vsan 101
  member pwnn Pure-STG-CT0-FC0
  member pwnn ORARAC-A6-hba0
  member pwnn Pure-STG-CT1-FC0
  member pwnn ORARAC-A6-hba2
  member pwnn Pure-STG-CT0-FC2
  member pwnn Pure-STG-CT1-FC2
zone name chas2-server3-boot-hba vsan 101
  member pwnn Pure-STG-CT0-FC0
  member pwnn ORARAC-A7-hba0
  member pwnn Pure-STG-CT1-FC0
  member pwnn ORARAC-A7-hba2
  member pwnn Pure-STG-CT0-FC2
  member pwnn Pure-STG-CT1-FC2
zone name chas2-server4-boot-hba vsan 101
  member pwnn Pure-STG-CT0-FC0
  member pwnn ORARAC-A8-hba0
  member pwnn Pure-STG-CT1-FC0
  member pwnn ORARAC-A8-hba2
  member pwnn Pure-STG-CT0-FC2
  member pwnn Pure-STG-CT1-FC2

```

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

```

zoneset name Oracle-RAC-A vsan 101
  member chas1-server1-boot-hba
  member chas1-server2-boot-hba
  member chas1-server3-boot-hba
  member chas1-server4-boot-hba
  member chas2-server1-boot-hba
  member chas2-server2-boot-hba
  member chas2-server3-boot-hba
  member chas2-server4-boot-hba
  member chas2-server4-boot-hba
exit

```

4. Activate the zone set by running following commands.

```

zoneset activate name Oracle-RAC-A vsan 101
exit
copy run start

```

## Cisco MDS Switch B

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

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

```

conf t
zone name chas1-server1-boot1-hba vsan 102
  member pwn Pure-STG-CT1-FC1
  member pwn ORARAC-A1-hba1
  member pwn Pure-STG-CT0-FC1
  member pwn ORARAC-A1-hba3
  member pwn Pure-STG-CT1-FC3
  member pwn Pure-STG-CT0-FC3

zone name chas1-server2-boot1-hba vsan 102
  member pwn Pure-STG-CT1-FC1
  member pwn ORARAC-A2-hba1
  member pwn ORARAC-A2-hba3
  member pwn Pure-STG-CT0-FC1
  member pwn Pure-STG-CT0-FC3
  member pwn Pure-STG-CT1-FC3

zone name chas1-server3-boot1-hba vsan 102
  member pwn Pure-STG-CT1-FC1
  member pwn ORARAC-A3-hba1
  member pwn Pure-STG-CT0-FC1
  member pwn ORARAC-A3-hba3
  member pwn Pure-STG-CT1-FC3
  member pwn Pure-STG-CT0-FC3

zone name chas1-server4-boot1-hba vsan 102
  member pwn Pure-STG-CT1-FC1
  member pwn ORARAC-A4-hba1
  member pwn ORARAC-A4-hba3
  member pwn Pure-STG-CT0-FC1
  member pwn Pure-STG-CT0-FC3
  member pwn Pure-STG-CT1-FC3

zone name chas2-server1-boot1-hba vsan 102
  member pwn Pure-STG-CT1-FC1
  member pwn ORARAC-A5-hba1
  member pwn ORARAC-A5-hba3
  member pwn Pure-STG-CT0-FC1
  member pwn Pure-STG-CT0-FC3
  member pwn Pure-STG-CT1-FC3

zone name chas2-server2-boot1-hba vsan 102
  member pwn Pure-STG-CT1-FC1
  member pwn ORARAC-A6-hba1
  member pwn Pure-STG-CT0-FC1
  member pwn ORARAC-A6-hba3
  member pwn Pure-STG-CT1-FC3

```



```

member pwnn Pure-STG-CT0-FC3

zone name chas2-server3-boot1-hba vsan 102
member pwnn Pure-STG-CT1-FC1
member pwnn ORARAC-A7-hba1
member pwnn ORARAC-A7-hba3
member pwnn Pure-STG-CT0-FC1
member pwnn Pure-STG-CT0-FC3
member pwnn Pure-STG-CT1-FC3

zone name chas2-server4-boot1-hba vsan 102
member pwnn Pure-STG-CT1-FC1
member pwnn ORARAC-A8-hba1
member pwnn Pure-STG-CT0-FC1
member pwnn ORARAC-A8-hba3
member pwnn Pure-STG-CT1-FC3
member pwnn Pure-STG-CT0-FC3

```

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

```

zoneset name Oracle-RAC-B vsan 102
member chas1-server1-boot1-hba
member chas1-server2-boot1-hba
member chas1-server3-boot1-hba
member chas1-server4-boot1-hba
member chas2-server1-boot1-hba
member chas2-server2-boot1-hba
member chas2-server3-boot1-hba
member chas2-server4-boot1-hba

```

4. Activate the zone set by running following commands.
 

```

zoneset activate name Oracle-RAC-A vsan 101
exit
copy run start

```

## Configure Pure Storage

The design goal of the reference architecture was to best represent a real-world environment as closely as possible. The approach included features of Cisco UCS to rapidly deploy stateless servers and use Pure Storage FlashArray's Snapshot feature to clone the boot LUNs to provision the O.S on top of the stateless servers.

A Service Profile was created within Cisco UCS Manager to deploy the 8 servers quickly with a standard configuration. SAN boot volume for these servers were hosted on the same Pure Storage FlashArray. Once the stateless servers were provisioned, following process was performed to enable Rapid deployment of 7 additional RAC nodes.

Single LUN was provisioned for the very first node and the first stateless server was booted off SAN. Zoning was performed on the Cisco MDS 9148S switches to enable the initiators discover the targets during boot process. Installed Oracle Linux 7.2 and all pre-requisite packages for Oracle on this LUN.

Using Pure Storage FlashRecover snapshot feature, the initial boot LUN was snapshotted and cloned to create 7 additional LUNs that will be used to boot the 7 remaining nodes of Oracle RAC. Each cloned LUN was zoned to the specific UCS blade and the node was booted off of SAN.

The service profiles represent all the attributes of a logical server in Cisco UCS model. Using logical servers that are disassociated from the physical hardware removes many limiting constraints around how servers are provisioned.

In addition to the service profiles, the use of Pure Storage's FlashRecover snapshots with SAN boot policy brings the following benefits

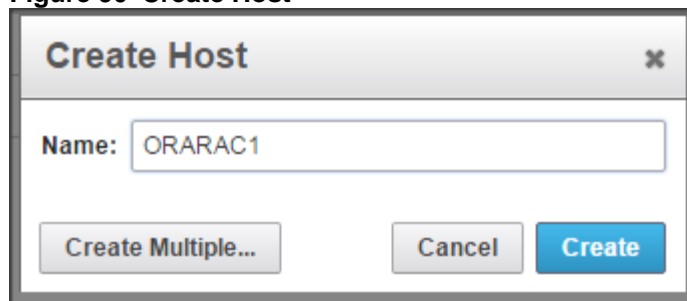
- Scalability - Rapid deployment of new servers to the environment in a very steps
- Manageability - Enables seamless hardware maintenance and upgrades without any restrictions. This is a huge benefit in comparison to other appliance model like Exadata
- Flexibility - Easy to repurpose physical servers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

### Configure Host

Before using a volume (LUN) on a host, the host has to be defined on Pure FlashArray. A host can be setup using the following steps in the Pure FlashArray GUI. Log into Pure Storage dashboard.

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

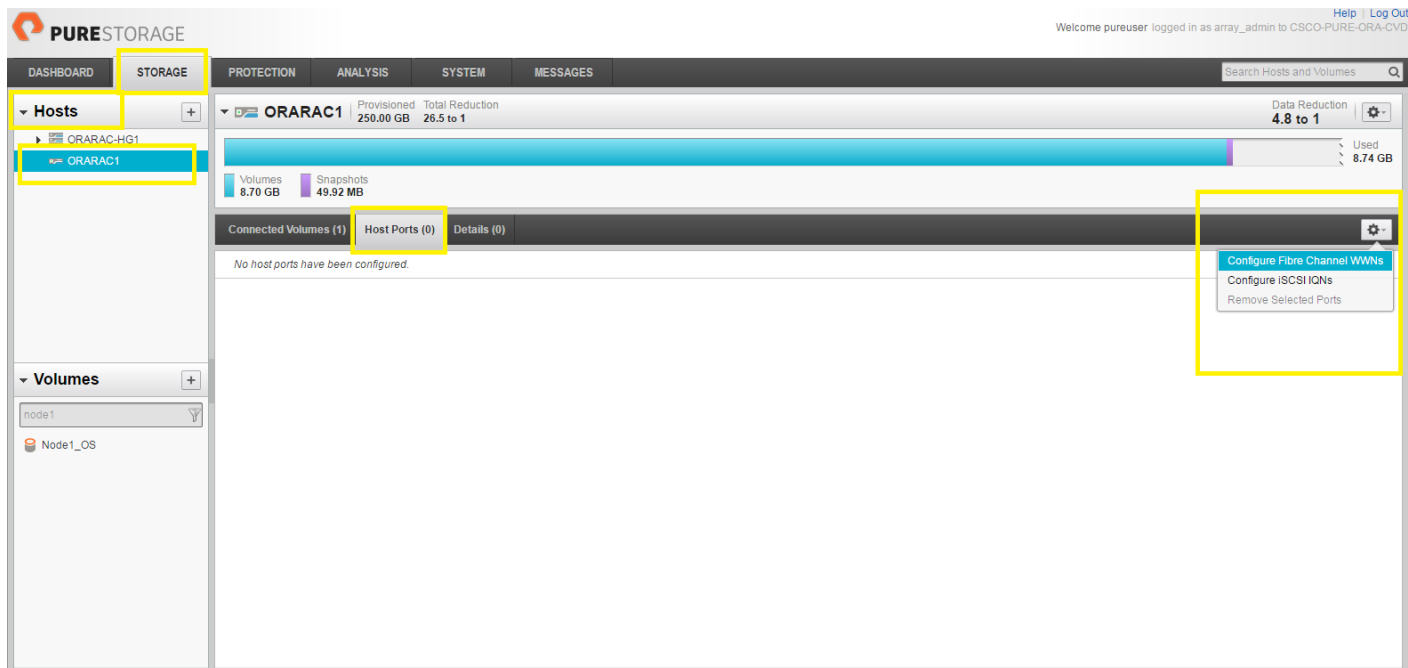
**Figure 90 Create Host**



The screenshot shows a 'Create Host' dialog box. The title bar contains the text 'Create Host' and a close button (X). Below the title bar is a text input field labeled 'Name:' with the text 'ORARAC1' entered. At the bottom of the dialog are three buttons: 'Create Multiple...', 'Cancel', and 'Create'.

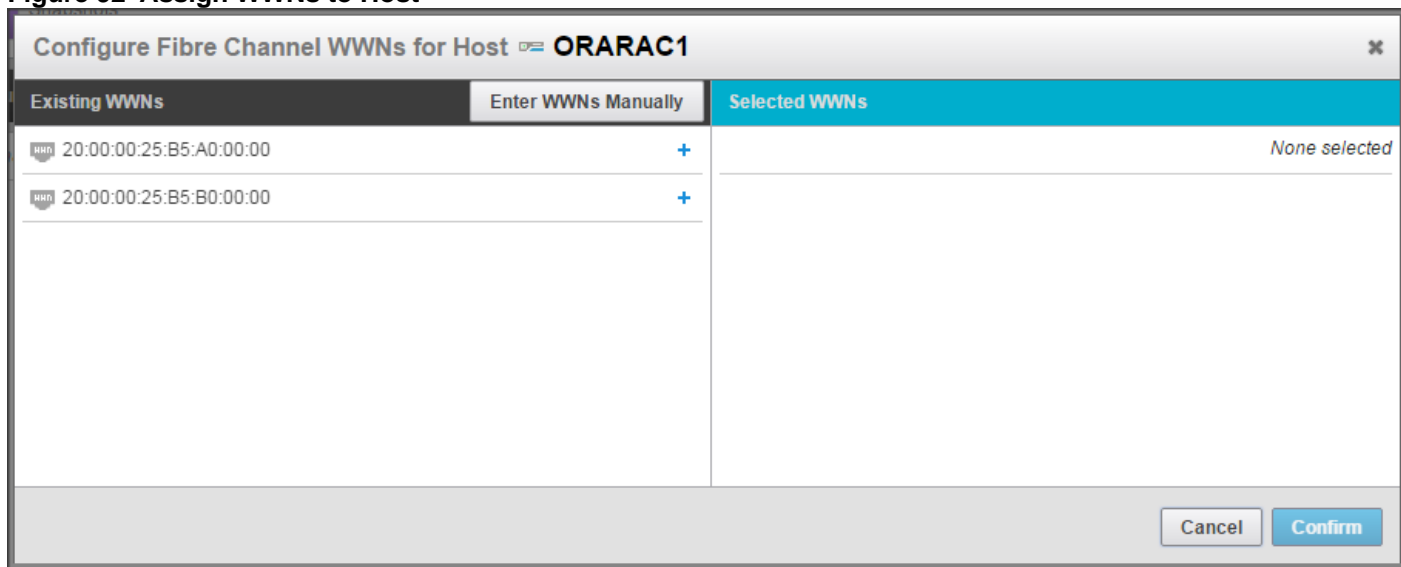
4. To update the host with the connectivity information by providing the Fibre Channel WWNs or iSCSI IQNs, click the Host that was created.

**Figure 91 Configure Host Port WWNs**



5. In the host context, click the Host Ports tab and click the settings button and select “Configure Fibre Channel WWNs” which should pop up a window with the available WWNs in the left side.

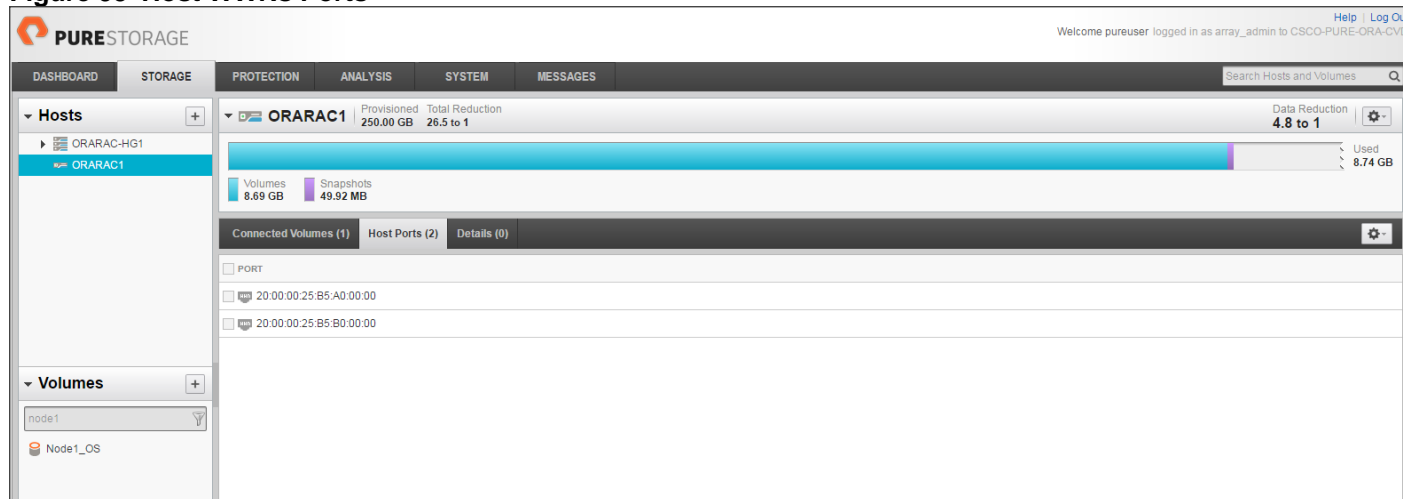
**Figure 92 Assign WWNs to Host**



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

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

**Figure 93 Host WWNs Ports**



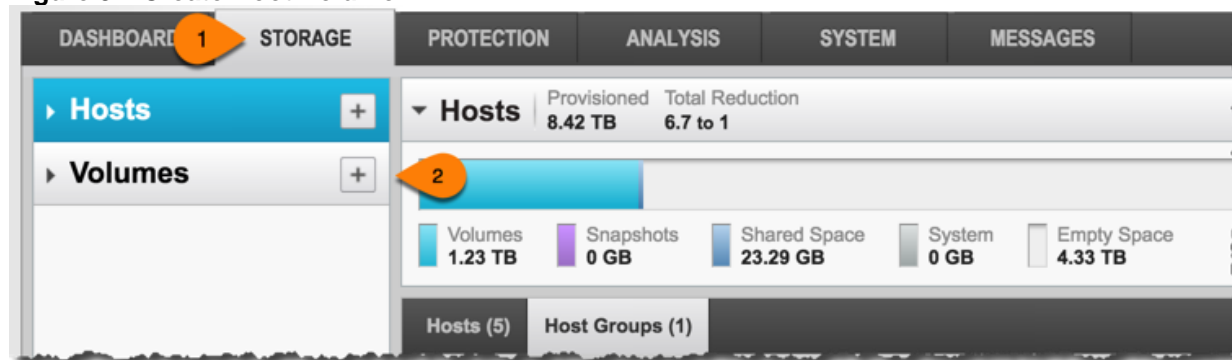
Make sure the zoning has been setup to include the WWNs details of the initiators along with the target, without which the SAN boot will not work.

**Configure Volume**

To configure a volume, complete the following steps:

1. Go to tab Storage > Volumes > and click the + sign to Create Volume.

**Figure 94 Create Boot Volume**



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

**Figure 95 Boot Volume Description**

The screenshot shows a 'Create Volume' dialog box with the following fields and buttons:

- Name:** Node1\_OS
- Provisioned Size:** 250 GB
- Buttons:** Create Multiple..., Cancel, Create

3. Attach the volume to a host by going to the “Connected Hosts and Host Groups” tab under the volume context menu, click the Settings icon and select Connect Hosts. Select the host where the volume should be attached and click Confirm.

**Figure 96 Connect Hosts to Volume**

The screenshot shows the 'Connect Hosts to Volume' dialog for 'Node1\_OS'. It features two main panels:

- Existing Hosts:** Contains a 'Filter Hosts' search bar and a 'Create New Host' button.
- Selected Hosts:** Shows a list with one entry: 'ORARAC1', which is selected. A status indicator shows '1 selected'.

Buttons at the bottom right include 'Cancel' and 'Confirm'.

This completes the connectivity of Storage LUN to the Server Node. We created one boot LUN (Node1\_OS) of 250GB and assigned this LUN to the first Oracle RAC node ORARAC1. Install the OS and perform the prerequisites for Oracle RAC Database to this LUN.

# OS and Database Deployment

## Operating System Configuration



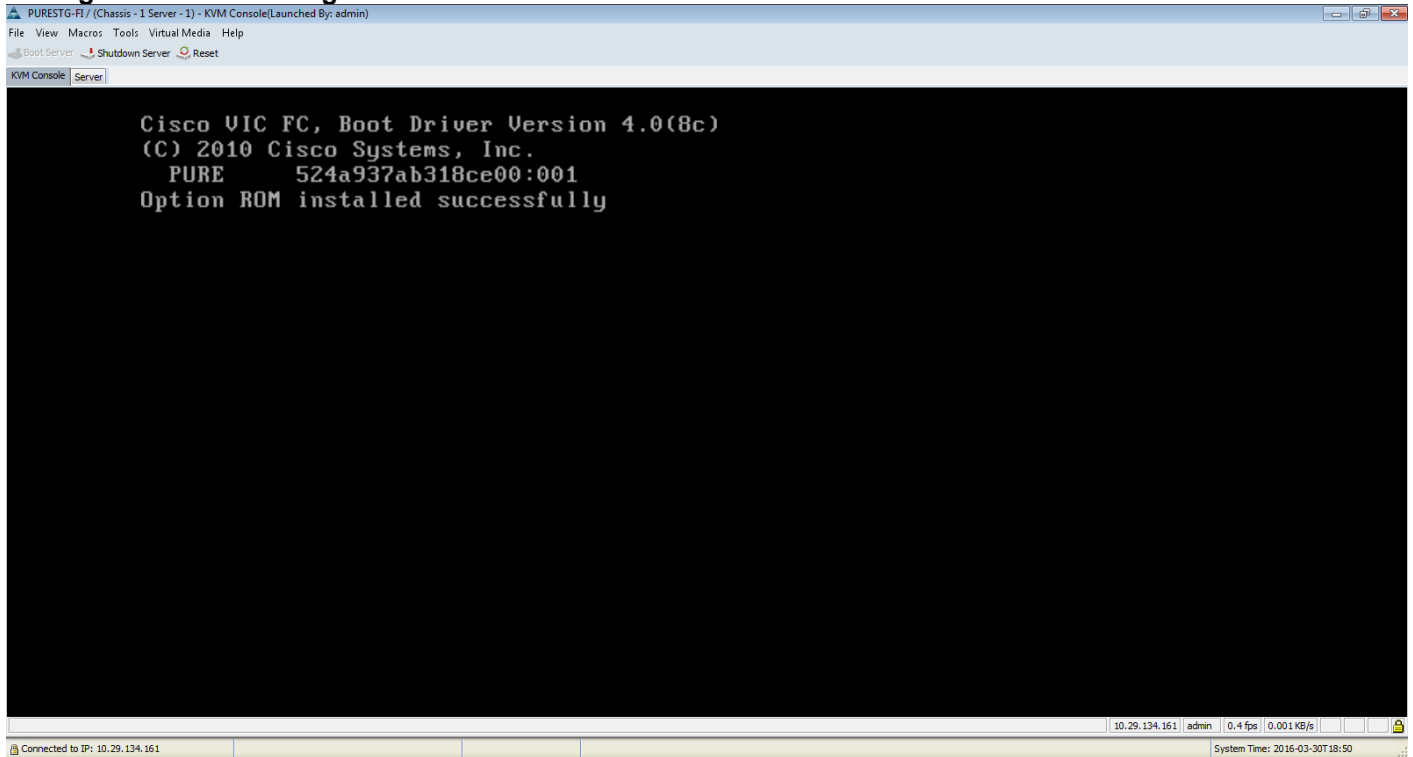
This document does not detail the step-by-step OS installation, but only provides an overview of the process.

1. Download Oracle Linux 7.2 OS image from <https://edelivery.oracle.com/linux>.
2. Launch KVM console on desired server, enable virtual media, map the Oracle Linux ISO image and reset the server.

**Figure 97 KVM Console Selection**

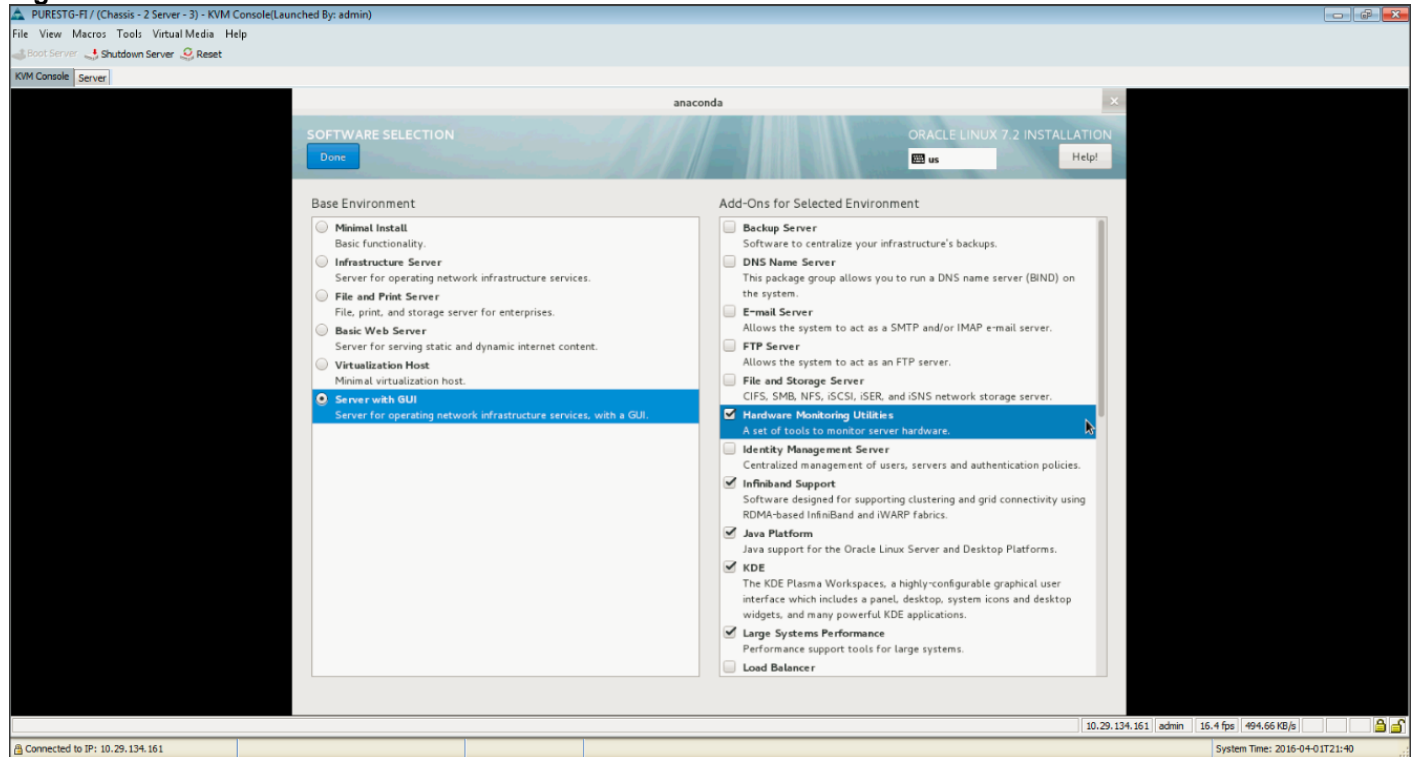
The screenshot displays the Cisco Unified Computing System Manager (CUCS) interface. On the left, a tree view shows the hierarchy: Equipment > Chassis > Chassis 1 > Servers > Server 1. The 'Server 1' item is highlighted in yellow. In the center, the 'Actions' menu is open, and the 'KVM Console' option is highlighted in yellow. The right side of the interface shows the 'Physical Display' and 'Properties' sections for the selected server. The 'Properties' section includes details such as Slot ID: 1, Product Name: Cisco UCS B200 M4, Vendor: Cisco Systems Inc, Revision: 0, Manufacturing Date: 2016-01-28, and various hardware specifications like Number of Processors: 2, Cores Enabled: 28, Effective Memory (MB): 262144, and Operating Memory Voltage: Regular Voltage.

3. When the Server start booting, it will detect the Pure Storage LUN 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.

**Figure 98 Pure Storage LUN**

4. During server boot order, it will detect the virtual media connected as Oracle Linux cd. It should launch the Oracle Linux installer. Select language and assign the Installation destination as Pure Storage FlashArray LUN. Apply hostname and click on “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.

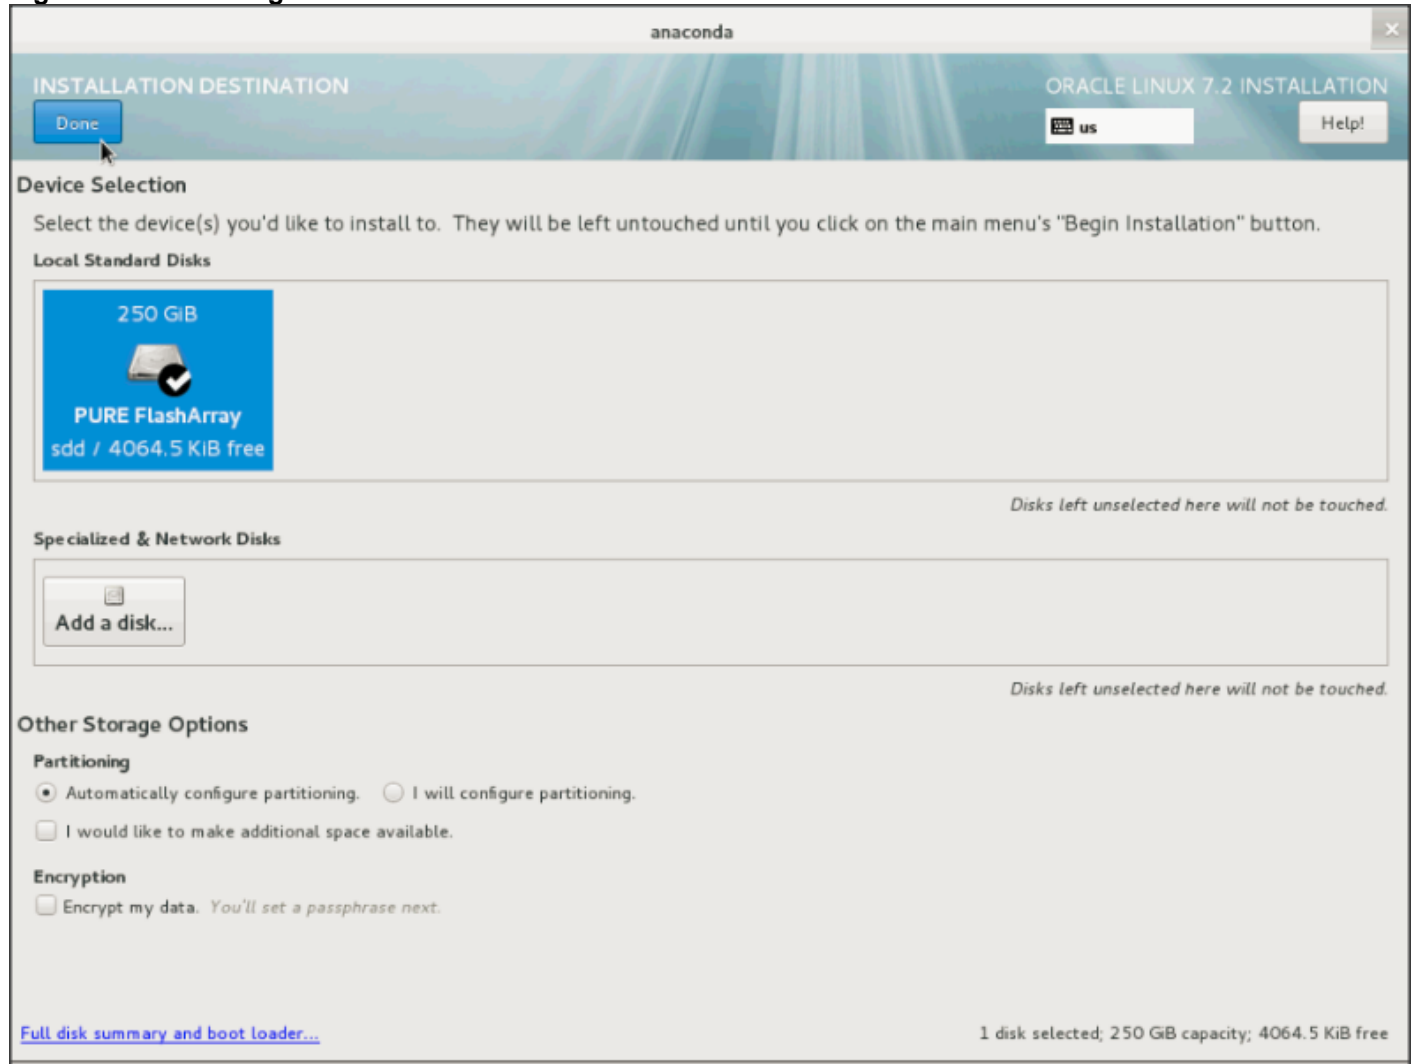
Figure 99 OS Software Selection



5. As a part of additional RPM package, it's recommended to select Customize Now and configure UEK kernel Repo.



Figure 100 OS Configuration



6. After 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.

## Miscellaneous Post-Installation Steps



All of them may have to be required on your setup. Please validate and change as needed. The following changes were made on the test bed where Oracle RAC install was done.

Once Oracle Linux 7.2 was installed on the initial server orarac1, Oracle 12c release 1 pre-requisite packages were installed as follows:

### Disable SELinux and Firewall

As most of organizations might already run hardware-based firewalls to protect their corporate networks, disable Security Enhanced Linux (SELinux) and the firewalls at the server level for this reference architecture.

Edit `/etc/selinux/config` and change to  
`SELINUX=disabled`

```
#SELINUXTYPE=targeted
```

### Disable firewall by running following commands

```
systemctl status firewalld
systemctl disable firewalld.service
systemctl status firewalld
```

### Configure sysctl.conf and limits.conf

Add or amend the following lines to the "/etc/sysctl.conf" file.

```
fs.file-max = 6815744
kernel.sem = 250 32000 100 128
kernel.shmmni = 4096
kernel.shmall = 1073741824
kernel.shmmax = 4398046511104
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
fs.aio-max-nr = 1048576
net.ipv4.ip_local_port_range = 9000 65500
```

Add the following lines to the "/etc/security/limits.conf" file

```
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft stack 10240
oracle hard stack 32768

grid soft nofile 1024
grid hard nofile 65536
grid soft nproc 2047
grid hard nproc 16384
grid soft stack 10240
grid hard stack 32768
```

### Add Groups and Users

Create and verify the Oracle users and groups in each cluster nodes. User accounts (oracle, grid) and groups (dba, oinstall, asmdba, oper) were created in the server. The "grid" user was setup as the owner of the Grid Infrastructure and "oracle" user was setup as the owner of the RAC database software. Both grid and oracle was configured to have oinstall as the common primary group for the Oracle central Inventory to work.

```
groupadd -g 501 oinstall
groupadd -g 502 dba
groupadd -g 504 asmdba
groupadd -g 505 asmadmin

useradd -u 1000 -g oinstall -G dba,asmdba,asmadmin oracle
passwd oracle

useradd -u 1001 -g oinstall -G dba,asmdba,asmadmin grid
passwd grid
```

## Configure Shell Limits

Shell limits were configured for the grid and oracle user which were saved in the `/etc/security/limits.d/20-nproc.conf` (Appendix D).

```
# Default limit for number of user's processes to prevent
root      soft      nproc     unlimited
```

## Configure Multipath

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

Modify `/etc/multipath.conf` file accordingly to give the alias name of each LUN id presented from Pure Storage as given below. Run `"multipath -ll"` command to view all the LUN id.

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

defaults {
    polling_interval 1
}
devices {
    device {
        vendor                "PURE"
        path_grouping_policy  multibus
        path_checker          tur
        path_selector         "queue-length 0"
        fast_io_fail_tmo     10
        dev_loss_tmo         30
    }
}
```

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

[https://support.purestorage.com/Solutions/Operating\\_Systems/Linux/Linux\\_Recommended\\_Settings](https://support.purestorage.com/Solutions/Operating_Systems/Linux/Linux_Recommended_Settings)

## Configure UDEV rules

As per Pure Storage FlashArray's best practice setup the queue settings with udev rules. Updated Linux best practices for Pure Storage FlashArray is available on Pure's support site.

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

```
# Recommended settings for PURE Storage FlashArray
# Use noop scheduler for high-performance SSD
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/scheduler}="noop"

# Reduce CPU overhead due to entropy collection
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/add_random}="0"

# Schedule I/O on the core that initiated the process
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/rq_affinity}="2"
```

These steps complete the prerequisites for the Oracle Database Installation at OS level. At this point, the initial node was shutdown to take the snapshot of the BOOT LUN and created 7 more clones of the same LUN for the 7 additional RAC nodes as explained in the next section.

## Cloning a Linux Host with SAN Boot on Pure Storage

After completion of OS configuration, the initial node “ORARAC1” was shutdown to take the snapshot of the BOOT LUN (Node1\_OS) and created 7 more clones of the source LUN for the 7 additional RAC nodes. This eliminated reinstalling most of the pre-requisite O.S level packages and kernel level settings as well as helped to provision additional nodes in minutes.

The nodes were reconfigured with appropriate network interfaces and WWNs. Also zoning information was updated with the WWN information of the host and the targets so the Oracle database LUNs can be presented to the hosts. This rapid deployment is one of the key features with Cisco UCS platform on Pure Storage that can improve IT productivity significantly.

The high-level steps for this process is as follows:

1. Pre-clone steps on the source Linux host
2. Clone the boot LUN on the Pure Storage
3. Post-clone steps on the source Linux host
4. Gather vHBA details for setup
5. Volume setup on Pure Storage for the clone
6. Target host setup

### Pre-clone Steps on the Source Linux Host

1. Unconfigure the networking setup on the source host. This is to avoid the cloned host to come up with the same IP addresses and causing IP conflicts.
  - a. Following set of commands identifies the device and disables them. If you have more than 2 interfaces, please disable them as well in the same format

```
[root]# nmcli device
DEVICE    TYPE        STATE        CONNECTION
enp13s0   ethernet    connected    enp13s0
enp6s0    ethernet    connected    enp6s0
lo        loopback    unmanaged    --

[root]# nmcli device disconnect enp13s0
[root]# nmcli device disconnect enp6s0

[root@donald ~]# nmcli device
DEVICE    TYPE        STATE        CONNECTION
enp6s0    ethernet    disconnected  --
enp13s0   ethernet    disconnected  --
lo        loopback    unmanaged    --
```

2. Create the following file under “/etc/sysconfig/network-scripts” directory and name it as “setup-network.sh”.

```
[root:/etc/sysconfig/network-scripts]$cat setup-network.sh
# User Input section.  Please provide the details to the following
variables.

hname=new-hostname
igw=<IP of your gateway>

# Network Interface 1
iface1=enp6s0 #or any name you wanted
iface1-ip=<new IP address> # in the cidr format like
192.168.10.10/24
iface1-mac=<new MAC address of the first interface>

# Network Interface 2
iface2=enp13s0 #If you have second interface
iface2-ip=<IP address for 2nd interface>
iface2-mac=<MAC address for the 2nd interface>

# If you have more than 2 interfaces, add similar entries below.
#
# Network interface setup section.
# Delete the old interfaces.  If you have more than 2 interfaces,
add them as well

nmcli conn del $iface1
nmcli conn del $iface2

# Add the new interfaces for the cloned server
# If you have more than 2 interfaces, repeat the following lines for
those interfaces

nmcli conn add type ethernet con-name $iface1 ifname $iface1 mac
$iface1-mac ip4 $iface1-ip gw4 igw
nmcli conn add type ethernet con-name $iface2 ifname $iface2 mac
$iface2-mac ip4 $iface2-ip

# Setup the new hostname
hostname $hname
echo $hname > /etc/hostname
```

3. Make sure the following entries are setup in the “/etc/sysconfig/network” file:

```
NETWORKING=yes
HOSTNAME=<source-hostname.domain>
GATEWAY=<gateway IP>
```

4. Check if multipathing service is setup to startup across reboots:

```
[root:/etc]$systemctl list-unit-files |grep multipath
multipathd.service disabled
```

```
[root:/etc]$chkconfig multipathd on
```

Note: Forwarding request to 'systemctl enable multipathd.service'.

```
[root:/etc]$systemctl list-unit-files |grep multipath
multipathd.service          enabled
```

5. Make sure the Linux best practices were followed in configuring the multipath.conf file



Preference is to use alias option instead of user\_friendly\_names under multipath.conf file which serves the same purpose but the naming is user driven and the user can name it to reflect the purpose like orahome or dg\_oradata. The user\_friendly\_names option is not enabled by default.

---

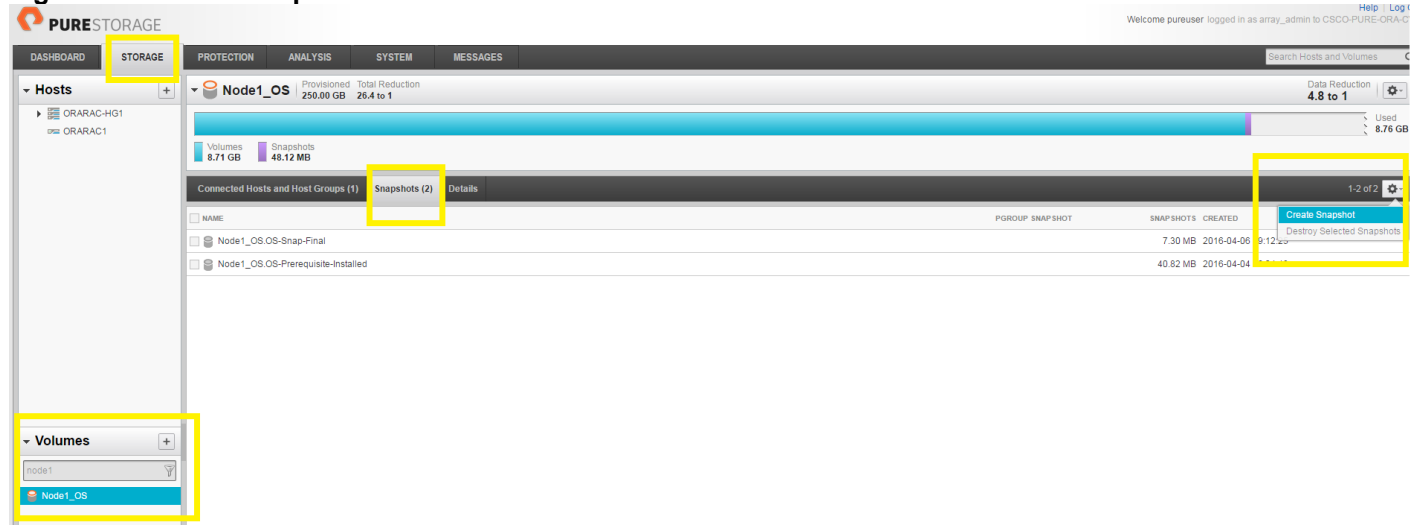
6. Shutdown the server. It is not mandatory to shut down the server to take a snapshot but to avoid the snapshot having any residual entries it is preferable to shut down the server.

## Clone the boot LUN on the Pure Storage

Use the Pure Storage FlashRecover snapshot functionality to clone the boot LUN, which delivers superior space efficiency, high availability and simplicity of volume snapshot management.

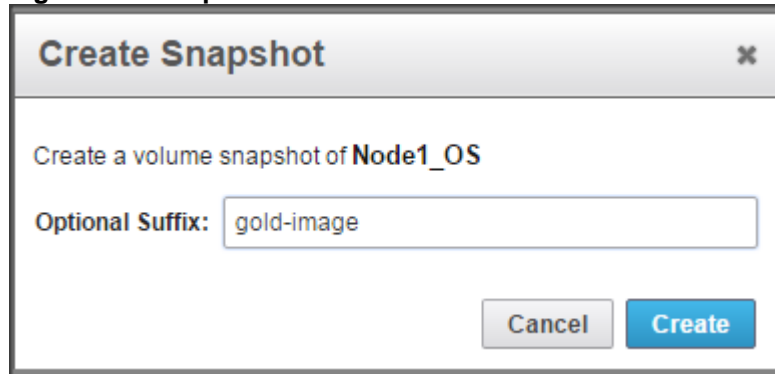
1. Take the snapshot of the source LUN either through the Pure Storage GUI or through the CLI
2. Log into Pure Storage and go to tab Storage > Volumes > Snapshots > Settings and "Create Snapshot" as shown below.


**Figure 101 Create Snapshot**



3. Name the snapshot and click Create.

**Figure 102 Snapshot of OS LUN**



 You will use this snapshot and clone this boot LUN to create additional boot LUNs for other Oracle nodes.

### Post-clone Steps on the Source Linux Host

1. Power on the source host
2. Enable the network interface that were disabled as part of the pre-clone steps on the source host.

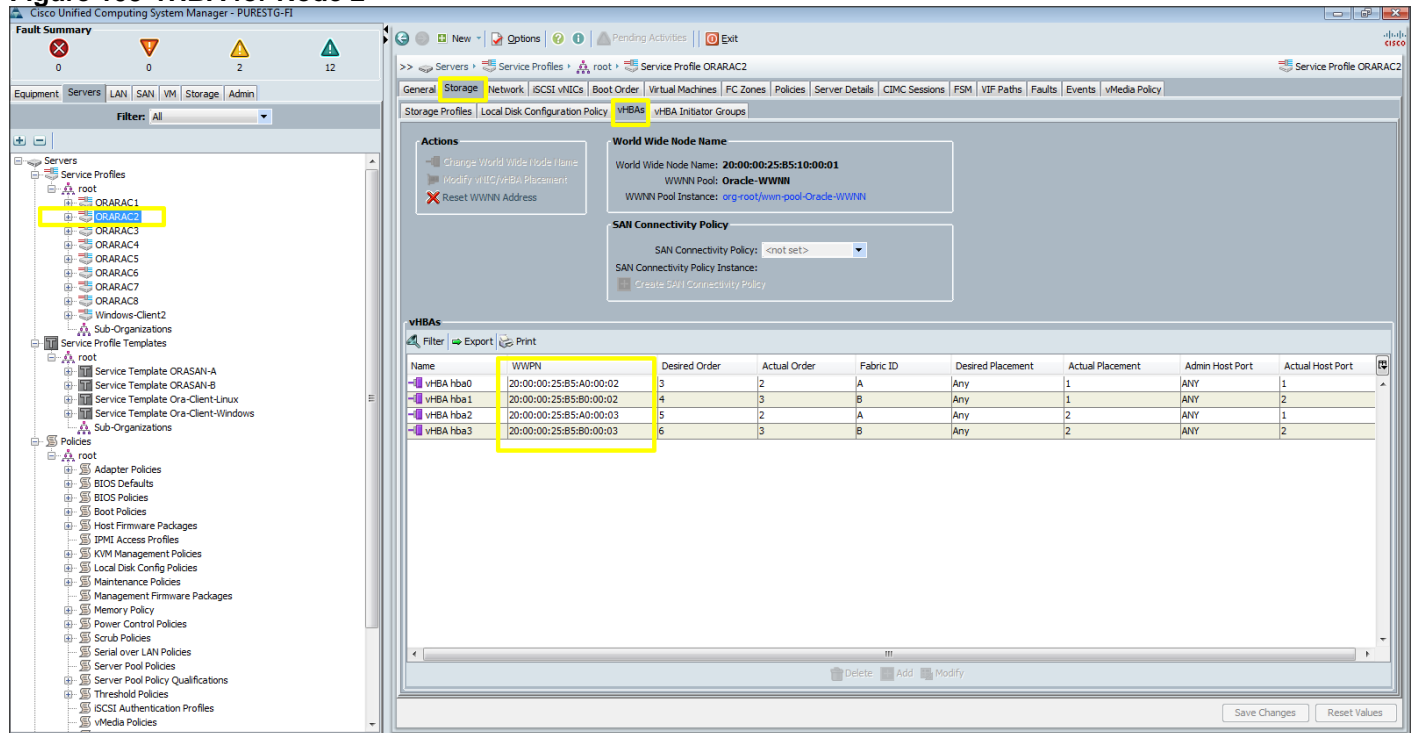
```
[root]# nmcli device connect enp13s0
[root]# nmcli device connect enp6s0
[root]# nmcli device
DEVICE    TYPE        STATE        CONNECTION
enp13s0   ethernet    connected    enp13s0
enp6s0    ethernet    connected    enp6s0
lo        loopback    unmanaged    --
```

### Gather vHBA Details for Setup

1. Click the Service Profile of the new server. Go to tab Servers > Service Profiles > root in the UCS Manager.

2. Click the Storage tab in the right side of the window to get the vHBA details.

**Figure 103 vHBA for Node 2**




Write down the WWPN information as seen in the screen for configuring the SAN zoning so the initiators can see the target for the SAN boot to work.

### Volume Setup on Pure Storage for the Clone

1. Setup the new Host in the Pure GUI. Go to tab Storage > Hosts > click the + sign to “Create Host” as shown below.

**Figure 104 Create Host ORARAC2**



2. To update the host with the connectivity information by providing the Fibre Channel WWNs or iSCSI IQNs, click the Host that was created.
3. In the host context, click the Host Ports tab and click the settings button and select “Configure Fibre Channel WWNs” which should pop up a window with the available WWNs in the left side.
4. Select the relevant WWNs for the host if it is discovered. If not, you can manually enter the WWNs based on the vHBA details gathered from the previous step and add as shown below. We have 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.



Figure 105 Host Ports for ORARAC2

The screenshot displays the Pure Storage management console. The 'STORAGE' tab is active. In the left-hand 'Hosts' sidebar, 'ORARAC2' is highlighted. The main content area shows details for 'ORARAC2', including 'Provisioned 1.24 TB' and 'Total Reduction 35.5 to 1'. Below this, there are sections for 'Volumes' (32.61 GB) and 'Snapshots' (0 GB). A 'Host Ports (4)' section is expanded, showing a table of four host ports. The 'Volumes' sidebar on the left shows 'Node1\_OS' under 'node1'.

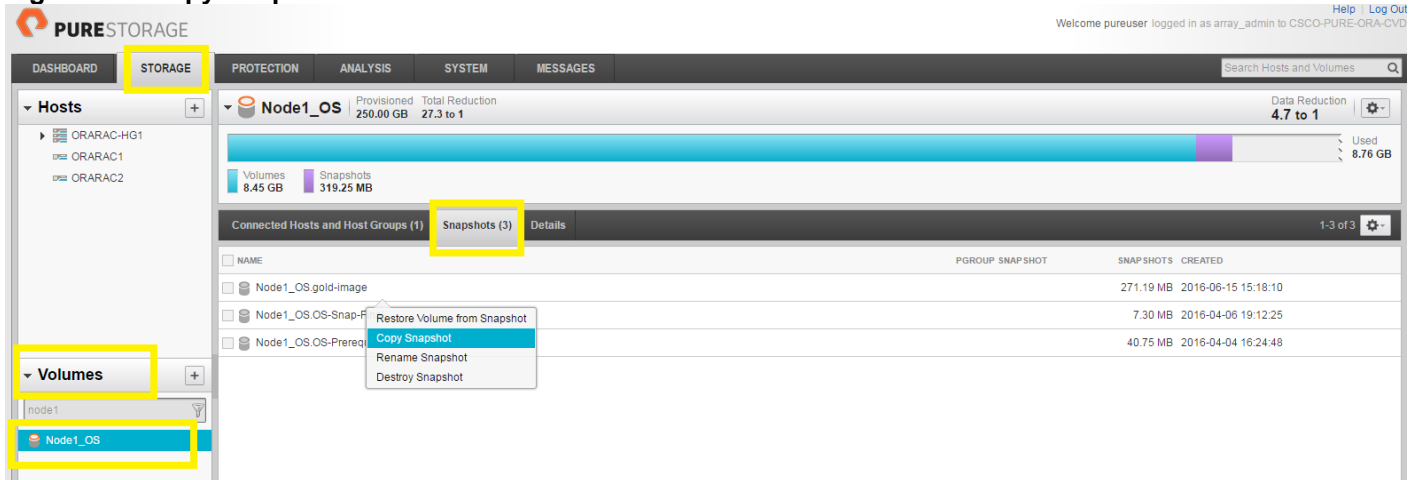
PORT
20:00:00:25:B5:A0:00:02
20:00:00:25:B5:A0:00:03
20:00:00:25:B5:B0:00:02
20:00:00:25:B5:B0:00:03



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

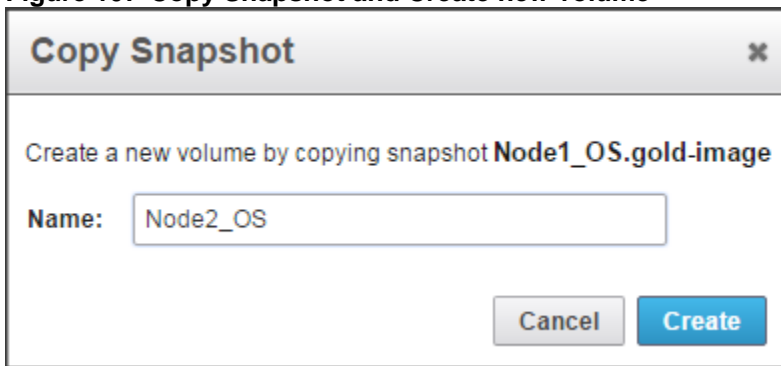
- Copy the snapshot to a new volume. This is equivalent to instantiating the snapshot to be a read/write volume. FlashRecover Snapshots are thin provisioned with no dedicated space allocated. They also inherit the data reduction characteristics of their parent volume and as such it will not consume the same amount of space as the source volume rather very minimal initially to reflect the metadata.
- To copy snapshot go to tab Storage > Volumes > Node1\_OS > Snapshots > Node1\_os.gold\_image snapshot which we created earlier. Select the snapshot and click on settings to "Copy Snapshot" as shown below.

Figure 106 Copy Snapshot for ORARAC2



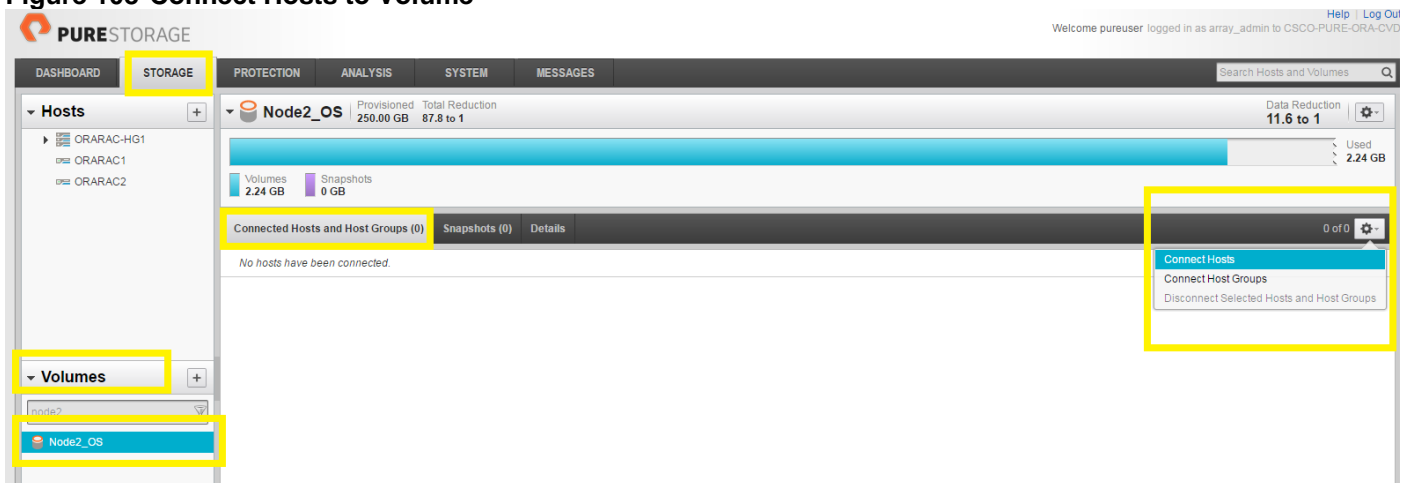
7. Provide a name for snapshot and click on Create to create volume as "Node2\_OS".

Figure 107 Copy Snapshot and Create new volume



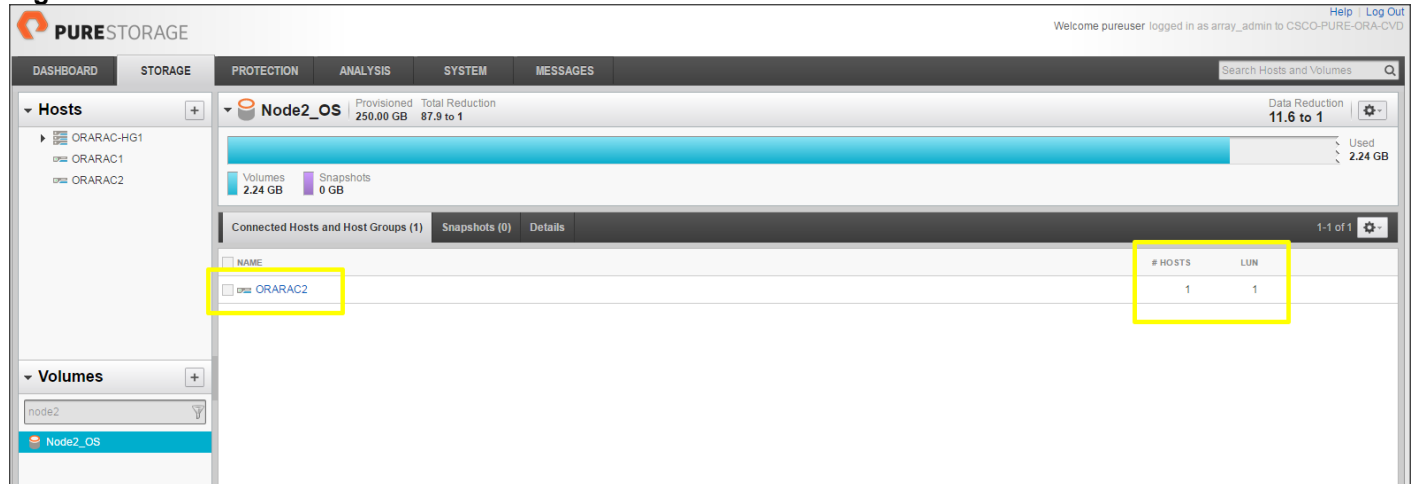
8. Attach the volume to a host by going to the "Connected Hosts and Host Groups" tab under the volume context menu, click the Settings icon and select "Connect Hosts". Select the host where the volume should be attached and click Confirm.

Figure 108 Connect Hosts to Volume




9. Make sure the LUN number shows up as 1. It will be 1 unless you have added other LUNS ahead of this. This LUN has to be the first lun, as the boot policy within Cisco UCS service profile template will be referring the SAN Boot Lun# to be 1.

Figure 109 Connect Hosts to Volume



10. Boot the new server through the KVM console on the UCSM.
11. 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.

Figure 110 Connect Hosts to Volume



## Target Host Setup

1. Update the script `/etc/sysconfig/network-scripts/setup-network.sh` with the IP address, MAC addresses of the interfaces, hostname for the target host.
2. Setup the network interfaces on the new host by running the script

```
/etc/sysconfig/network-scripts/setup-network.sh
```

```
[root:/etc/sysconfig/network-scripts]$ ./setup-network.sh
Connection 'enp6s0' (a7d8caaa-ad03-4f99-96dc-b27cb1805061) success-
fully added.
Connection 'enp13s0' (958f3c26-485a-4d42-adb2-8a780f3538e5) success-
fully added.
```

3. Edit the `/etc/sysconfig/network` file and update the `HOSTNAME` entry to reflect the new hostname with domain.

```
NETWORKING=yes
HOSTNAME=<hostname.domain>
GATEWAY=<gateway IP>
```

4. Verify the `/etc/resolv.conf` file reflects the search order and includes your DNS servers. If not, update the same accordingly.

```
[root]$ more /etc/resolv.conf
# Generated by NetworkManager
search cisco.com
nameserver xx.xx.xx.xx
nameserver xx.xx.xx.xx
```

5. Check if the interfaces are up and running and you have network connectivity between the hosts.

```
[root]$ ifconfig -a
```

6. Check if the multipath service is up and running.

```
[root@orarac2 ~]# systemctl status multipathd.service
multipathd.service - Device-Mapper Multipath Device Controller
Loaded: loaded (/usr/lib/systemd/system/multipathd.service; ena-
bled; vendor preset: enabled)
Active: active (running) since Fri 2016-06-03 10:16:36 PDT; 1
weeks 5 days ago
Process: 1590 ExecStart=/sbin/multipathd (code=exited, sta-
tus=0/SUCCESS)
Process: 1587 ExecStartPre=/sbin/multipath -A (code=exited, sta-
tus=0/SUCCESS)
Process: 1562 ExecStartPre=/sbin/modprobe dm-multipath
(code=exited, status=0/SUCCESS)
Main PID: 1592 (multipathd)
CGroup: /system.slice/multipathd.service
└─1592 /sbin/multipathd
```

7. If multipath service is not running, bring it up using the following command.

```
[root@orarac2 ~]# systemctl start multipathd.service
```

8. Check the hostname reflects the new hostname.
9. Perform the necessary steps to register the new host with Oracle Linux subscription and reboot the Server.

Now the OS setup for Oracle RAC node 2 (ORARAC2).



For this FlashStack Solution, we used 8 identical Cisco UCS B-series B200 M4 blade servers for hosting the 8-node Oracle RAC database. Repeat the same steps for cloning a Linux host with SAN Boot on all remaining 6 nodes and configure all the remaining Oracle RAC nodes. This process will complete all 8 Oracle RAC Node with OS and all prerequisites to install Oracle Database Software.

## Configure Host Groups

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

The Host Group can be setup using the following steps in the Pure FlashArray GUI. Log into Pure Storage dashboard.

To configure the Host groups, complete the following steps:

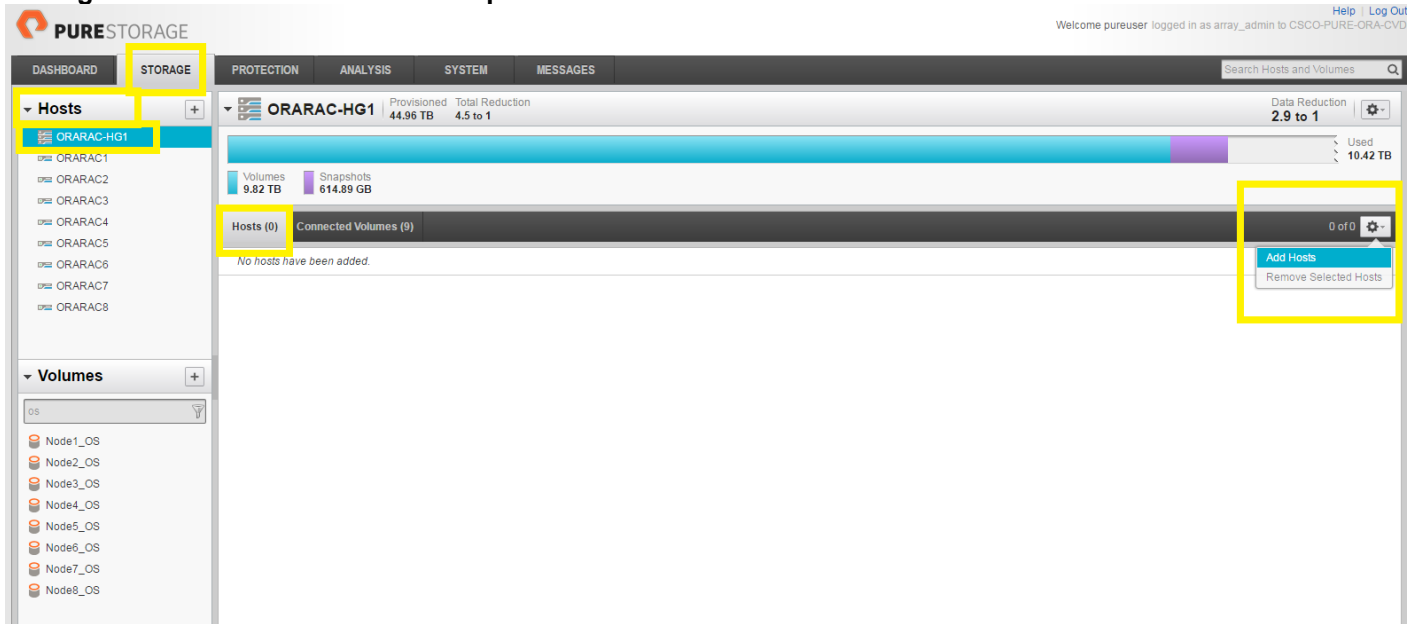
1. In the PURE GUI, go to tab Storage > Hosts > click the + sign to “Create Host Group” as shown below.

**Figure 111 Create Host Group**

The screenshot shows a dialog box titled "Create Host Group" with a close button (X) in the top right corner. Below the title bar, there is a text input field labeled "Name:" containing the text "ORARAC-HG1". At the bottom of the dialog, there are three buttons: "Create Multiple...", "Cancel", and "Create".

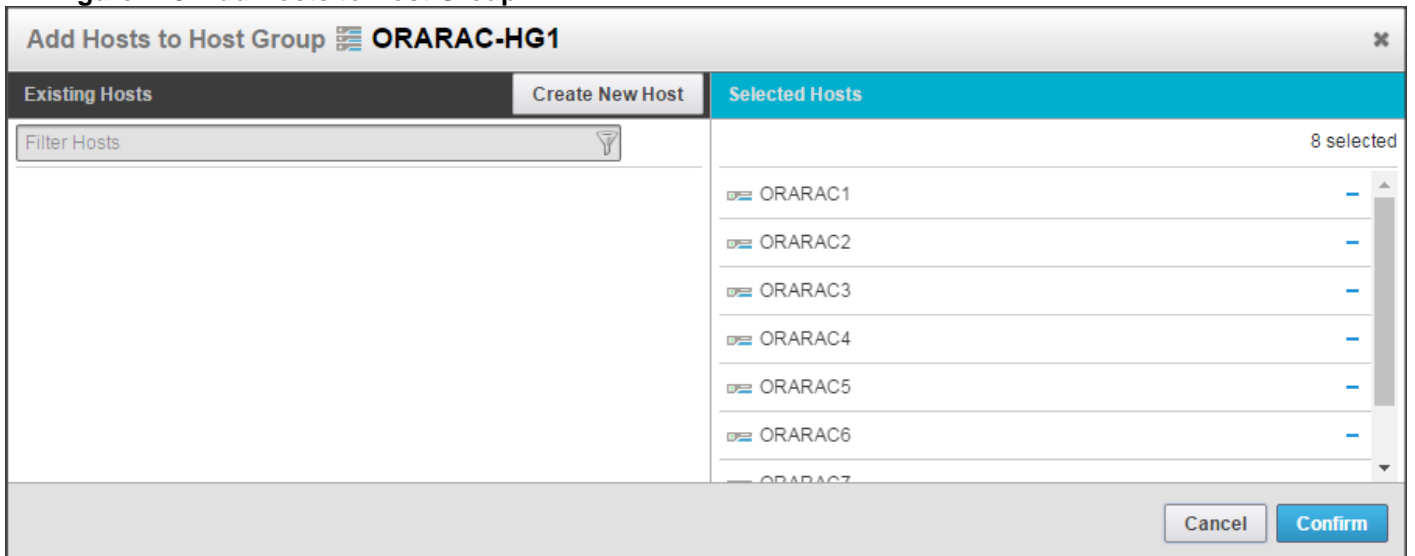
2. Select Host Group “ORARAC-HG1” > Hosts > click Settings and Add Hosts.

Figure 112 Add Host to Host Group



3. Add all nodes from "Existing Hosts" to "Selected Hosts" and click Confirm.

Figure 113 Add Hosts to Host Group



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

**Figure 114 Host Group ORARAC-HG1**

The screenshot shows the Pure Storage console interface for Host Group ORARAC-HG1. The top navigation bar includes 'DASHBOARD', 'STORAGE', 'PROTECTION', 'ANALYSIS', 'SYSTEM', and 'MESSAGES'. The main content area displays the host group details, including a bar chart for Volumes (9.82 TB) and Snapshots (614.89 GB). Below this is a table listing the hosts and their connected volumes.

NAME	INTERFACE	PROVISIONED	VOLUMES	REDUCTION
ORARAC1	FC	46.21 TB	9.85 TB	2.9 to 1
ORARAC2	FC	46.21 TB	9.85 TB	2.9 to 1
ORARAC3	FC	46.21 TB	9.85 TB	2.9 to 1
ORARAC4	FC	46.21 TB	9.85 TB	2.9 to 1
ORARAC5	FC	46.21 TB	9.84 TB	2.9 to 1
ORARAC6	FC	46.21 TB	9.84 TB	2.9 to 1
ORARAC7	FC	46.21 TB	9.84 TB	2.9 to 1
ORARAC8	FC	46.21 TB	9.84 TB	2.9 to 1

### Configure CRS, Data and Redo Log Volumes

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

For this FlashStack solution, create two OLTP Database (OLTPDB1 and OLTPDB2) and one DSS Database (DSSDB1). For each database, create Data Volume of 7TB size and Redo log Volume of 500GB size.

1. Create CRS Volume of 10GB as shown below.

**Figure 115 CRS Volume**

The screenshot shows the 'Create Volume' dialog box. The 'Name' field contains 'dg\_orarac\_crs' and the 'Provisioned Size' field contains '10 GB'. There are three buttons at the bottom: 'Create Multiple...', 'Cancel', and 'Create'.

2. Create Data Volume of 7TB as shown below.

**Figure 116 Data Volume**

3. Create Redo Volume of 500GB as shown below.

**Figure 117 Redolog Volume**

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

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

**Figure 118 Connected Volumes to Host Group**

NAME	LUN	PROVISIONED	VOLUMES	SNAPSHOTS	REDUCTION
dg_oradata_dss1	12	20 TB	4.70 TB	194.02 GB	2.4 to 1
dg_oradata_dss2	17	10 TB	2.57 TB	0 GB	2.8 to 1
dg_oradata_oltp1	10	7 TB	1.32 TB	33.84 GB	3.9 to 1
dg_oradata_oltp2	14	6 TB	1.20 TB	358.17 GB	4.2 to 1
dg_orarac_crs	13	10 GB	1.02 GB	420.12 MB	4.6 to 1
dg_redolog_dss1	15	500 GB	5.42 GB	9.58 GB	4.7 to 1
dg_redolog_dss2	18	500 GB	179.53 MB	0 GB	4.4 to 1
dg_redolog_oltp1	11	500 GB	6.88 GB	7.57 GB	2.4 to 1
dg_redolog_oltp2	16	500 GB	7.60 GB	11.31 GB	4.2 to 1

The following table shows the LUNs are created and the description of the LUNs.



**Table 12 LUN Description**

LUN Name	Size	Description
dg_orarac_crs	10 GB	For OCR and Voting Disks
dg_oradata_oltp1	7 TB	For Database Files OLTPDB1
dg_oradata_oltp2	7 TB	For Database Files OLTPDB2
dg_oradata_dss1	10 TB	For Database Files DSSDB1
dg_redolog_oltp1	500 GB	For Redolog Files OLTPDB1
dg_redolog_oltp2	500 GB	For Redolog Files OLTPDB2
dg_redolog_dss1	500 GB	For Redolog Files DSSDB1

## Configure UDEV Rule

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



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

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

```
#All volumes which starts with dg_orarac_* #
ENV{DM_NAME}=="dg_orarac_*", OWNER=="grid", GROUP=="oinstall", MODE=="660"

#All volumes which starts with dg_oradata_* #
ENV{DM_NAME}=="dg_oradata_*", OWNER=="oracle", GROUP=="oinstall", MODE=="660"

#All volumes which starts with dg_oraredo_* #
ENV{DM_NAME}=="dg_oraredo_*", OWNER=="oracle", GROUP=="oinstall", MODE=="660"
```



Once all the O.S level prerequisites are completed, you are ready to install Oracle Grid Infrastructure as grid user. Download Oracle Database 12c Release 1 (12.1.0.2.0) for Linux x86-64 and Oracle Database 12c Release 1 Grid Infrastructure (12.1.0.2.0) for Linux x86-64 software from Oracle Software Delivery Cloud site. Copy these software binaries to Oracle RAC Node 1 and Unzip all files into appropriate directories.

## Oracle Database 12c GRID Infrastructure Deployment

For this FlashStack Solution, you will install Oracle Grid and Database software on the first four node (orarac1, orarac2, orarac3 and orarac4). Then configure the remaining four Oracle Nodes (orarac5, orarac6, orarac7 and orarac8) to scale up the system to add more nodes. The installation guides through gathering all nodes information and configuring ASM devices and all the pre-requisite 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 partial summary of details that might be relevant.

This section describes high level steps for Oracle Database 12c R1 RAC install. Prior to GRID and database install, verify all the prerequisites are completed. As an alternate, you can install Oracle validated RPM that will ensure all prerequisites are met before Oracle grid install.

Use the following Oracle document for pre-installation tasks, such as setting up the kernel parameters, RPM packages, user creation, etc. [http://docs.oracle.com/database/121/LADBI/pre\\_install.htm#LADBI7495](http://docs.oracle.com/database/121/LADBI/pre_install.htm#LADBI7495)



Prior to GRID install, we recommend completing the following steps.

---

## Create the Local Directory Structure and Ownership on each RAC Node

```
mkdir -p /u01/oracle/app
mkdir -p /u01/grid/app
chown -R oracle:oinstall /u01/oracle
chmod -R 775 /u01/oracle
chown -R grid:oinstall /u01/grid
chmod -R 775 /u01/grid
```

## Configure Public and Private NICs on each RAC Node

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

To configure public and private NICs, complete the following steps:

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

## Configure “/etc/hosts” on each RAC Node

To configure /etc/hosts on each RAC node, complete the following steps:

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

```
[root@oraracl network-scripts]# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain

#Public IP
10.10.10.171 orarac1    orarac1.cisco.com
10.10.10.172 orarac2    orarac2.cisco.com
10.10.10.173 orarac3    orarac3.cisco.com
10.10.10.174 orarac4    orarac4.cisco.com
10.10.10.175 orarac5    orarac5.cisco.com
10.10.10.176 orarac6    orarac6.cisco.com
10.10.10.177 orarac7    orarac7.cisco.com
10.10.10.178 orarac8    orarac8.cisco.com
```

```

#Virtual IP
10.10.10.179 orarac1-vip orarac1-vip.cisco.com
10.10.10.180 orarac2-vip orarac2-vip.cisco.com
10.10.10.181 orarac3-vip orarac3-vip.cisco.com
10.10.10.182 orarac4-vip orarac4-vip.cisco.com
10.10.10.183 orarac5-vip orarac5-vip.cisco.com
10.10.10.184 orarac6-vip orarac6-vip.cisco.com
10.10.10.185 orarac7-vip orarac7-vip.cisco.com
10.10.10.186 orarac8-vip orarac8-vip.cisco.com

#Private IP
192.168.76.171 orarac1-priv orarac1-priv.cisco.com
192.168.76.172 orarac2-priv orarac2-priv.cisco.com
192.168.76.173 orarac3-priv orarac3-priv.cisco.com
192.168.76.174 orarac4-priv orarac4-priv.cisco.com
192.168.76.175 orarac5-priv orarac5-priv.cisco.com
192.168.76.176 orarac6-priv orarac6-priv.cisco.com
192.168.76.177 orarac7-priv orarac7-priv.cisco.com
192.168.76.178 orarac8-priv orarac8-priv.cisco.com

#SCAN IP
10.10.10.187 orarac-cluster orarac-cluster.cisco.com
10.10.10.188 orarac-cluster orarac-cluster.cisco.com
10.10.10.189 orarac-cluster orarac-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 each cluster

## Establish SSH Connectivity

You can configure the ssh option (with no password) for the Oracle and Grid user during pre-installation steps. You can also configure this step during grid software installation process.

To establish SSH connectivity, complete the following steps:

1. Log in as Grid User in Oracle RAC Node 1.
2. Go to the directory where the oracle grid software binaries are located.
3. Go to sshsetup directory. To enable and check ssh connectivity in all eight nodes, run the script file named as "sshUserSetup.sh" by running following command:

```
[grid@orarac1 sshsetup]$ ./sshUserSetup.sh -user oracle -hosts "orarac1 orarac2
orarac3 orarac4 orarac5 orarac6 orarac7 orarac8" -noPromptPassphrase -confirm -
advanced
```

4. The command will ping all nodes and check remote host reachability. If all hosts are reachable then it will prompt you password for oracle user for each node. Enter the oracle user password for all nodes and complete the SSH connectivity.

Make sure you get the following script output:

```
The following hosts are reachable: orarac1 orarac2 orarac3 orarac4 orarac5 orarac6
orarac7 orarac8.
```

```

Verifying SSH connectivity has been setup from orarac1 to orarac1
Verifying SSH connectivity has been setup from orarac1 to orarac2
Verifying SSH connectivity has been setup from orarac1 to orarac3
Verifying SSH connectivity has been setup from orarac1 to orarac4
Verifying SSH connectivity has been setup from orarac1 to orarac5
Verifying SSH connectivity has been setup from orarac1 to orarac6
Verifying SSH connectivity has been setup from orarac1 to orarac7
Verifying SSH connectivity has been setup from orarac1 to orarac8
-Verification from complete-
SSH verification complete.

```

## Run Cluster Verification Utility

The steps in this section verify that all prerequisites are met to install Oracle Grid Infrastructure Software. Oracle Grid Infrastructure ships with the Cluster Verification Utility (CVU) that can be 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 directory where oracle grid software binaries are located. Run script named as "runcluvfy.sh" as follows:

```

[grid@orarac1 grid]$ ./runcluvfy.sh stage -pre crsinst -n
orarac1,orarac2,orarac3,orarac4,orarac5,orarac6,orarac7,orarac8 -verbose

```

## Configure HugePages

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

### Advantages of HugePages

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

Please refer to Oracle support (formerly metalink) document 361323.1 for HugePages configuration details.

## Install and Configure Oracle Database Grid Infrastructure Software

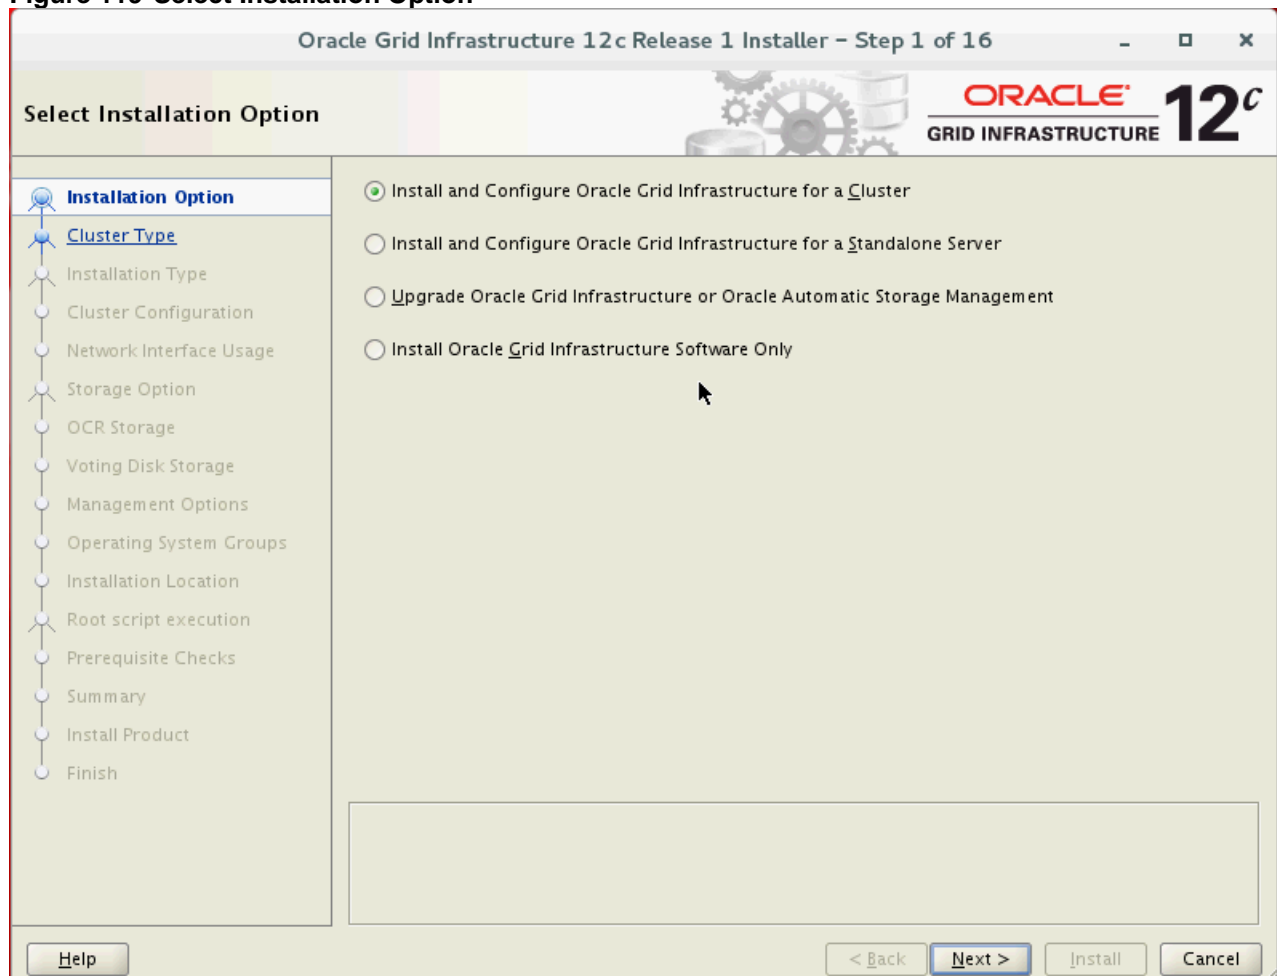
It is not within the scope of this document to include the specifics of an Oracle RAC installation. However, we will provide partial summary of details that might be relevant. Please refer to the Oracle installation documentation for specific installation instructions for your environment.

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

1. Launch the installer as the "grid" user from the staging area where the Oracle 12c R1 Grid Infrastructure software is located

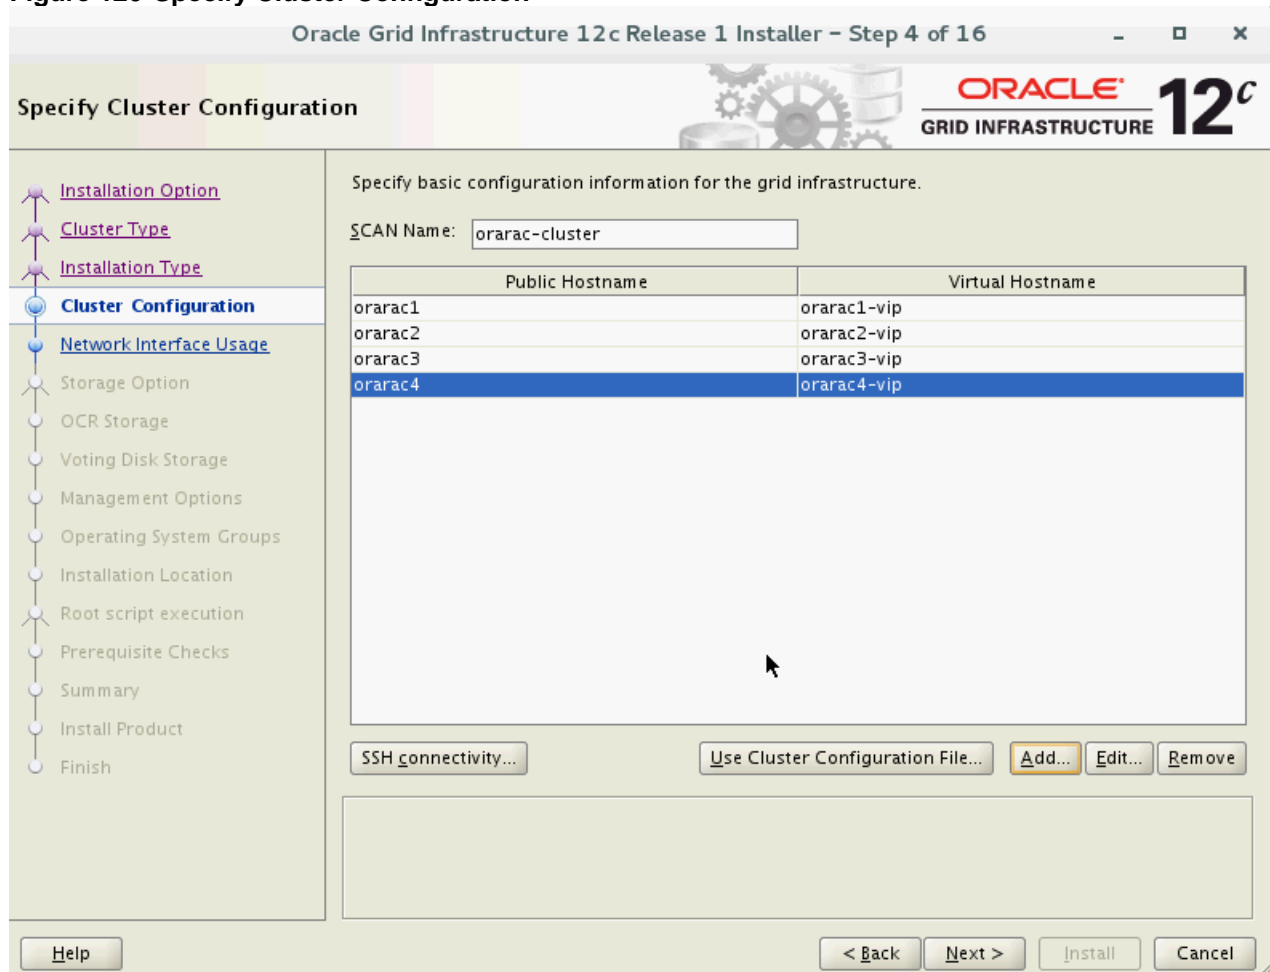
2. Select option "Install and Configure Oracle Grid Infrastructure for a Cluster" as shown below.

Figure 119 Select Installation Option



3. Select option type of cluster as "Configure a Standard Cluster".
4. Select Installation type as "Advanced Installation" for this setup.
5. On the "Specify Cluster Configuration" screen, enter the correct SCAN Name and click the "Add" button. Enter the details of the first four node in the cluster as shown below.

Figure 120 Specify Cluster Configuration



6. However, you are only adding first four node (orarac1, orarac2, orarac3 and orarac4) to the cluster for now to check the scalability and performance of 4 node RAC Cluster. Later in the solution, you will add remaining four nodes (orarac5, orarac6, orarac7 and orarac8) and upgrade the Storage while database workload is running to show the scalability of FlashStack solution as demand grow.
7. Click the SSH Connectivity... button and enter the password for the "grid" user. Click Setup to configure SSH connectivity, and click Test to test it once it is complete. Once the test is complete, click Next.

Figure 121 Specify Cluster Configuration SSH Connectivity

Oracle Grid Infrastructure 12c Release 1 Installer – Step 4 of 16

Specify Cluster Configuration

ORACLE 12c  
GRID INFRASTRUCTURE

Specify basic configuration information for the grid infrastructure.

SCAN Name:

Public Hostname	Virtual Hostname
orac1	orac1-vip
orac2	orac2-vip
orac3	orac3-vip
orac4	orac4-vip

SSH connectivity...

OS Username:  OS Password:

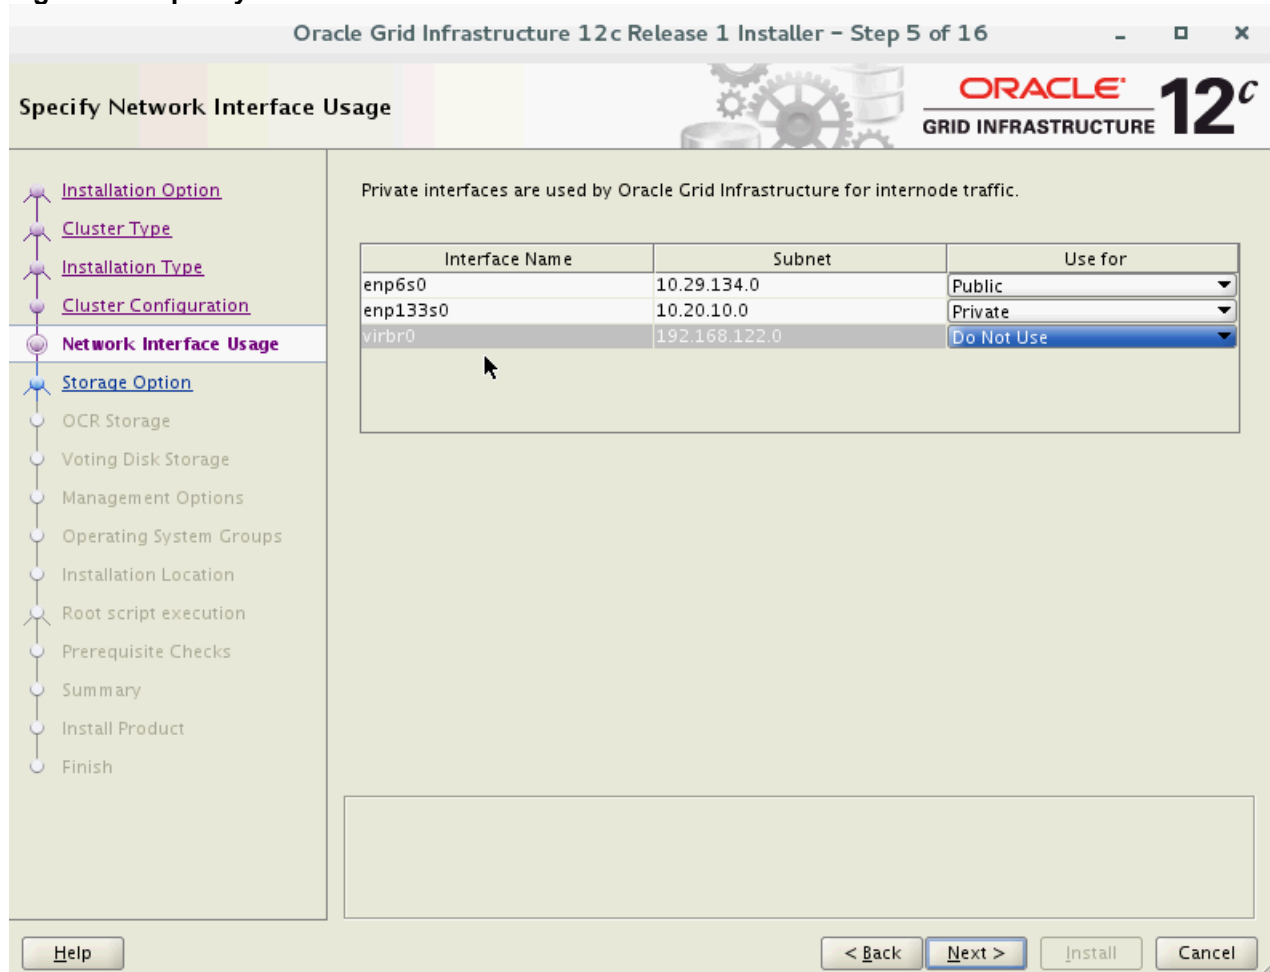
User home is shared by the selected nodes

Reuse private and public keys existing in the user home

8. Check and verify the public and private networks are specified correctly. If the NAT interface is displayed, remember to mark it as "Do Not Use". Click Next.



Figure 122 Specify Network Interface



- Enter Software Location as "/u01/grid/app/12.1.0/grid" and select "Oracle Automatic Storage Manager" as the cluster registry storage type. Enter the ASM password, select "oinstall" as the OSASM group and click Next.

Figure 123 Specify Install Locations

Oracle Grid Infrastructure 12c Release 1 Installer – Step 6 of 10

**Specify Install Locations**

Specify locations for Oracle base, where to install the software, where to place the Oracle Cluster Registry (OCR), and which operating system group should be given the administrative privileges (SYSASM) for Oracle Automatic Storage Management.

Oracle Base:

Software Location:

Cluster Registry Storage Type:

Cluster Registry Location:

SYSASM Password:

Confirm Password:

OSASM group:

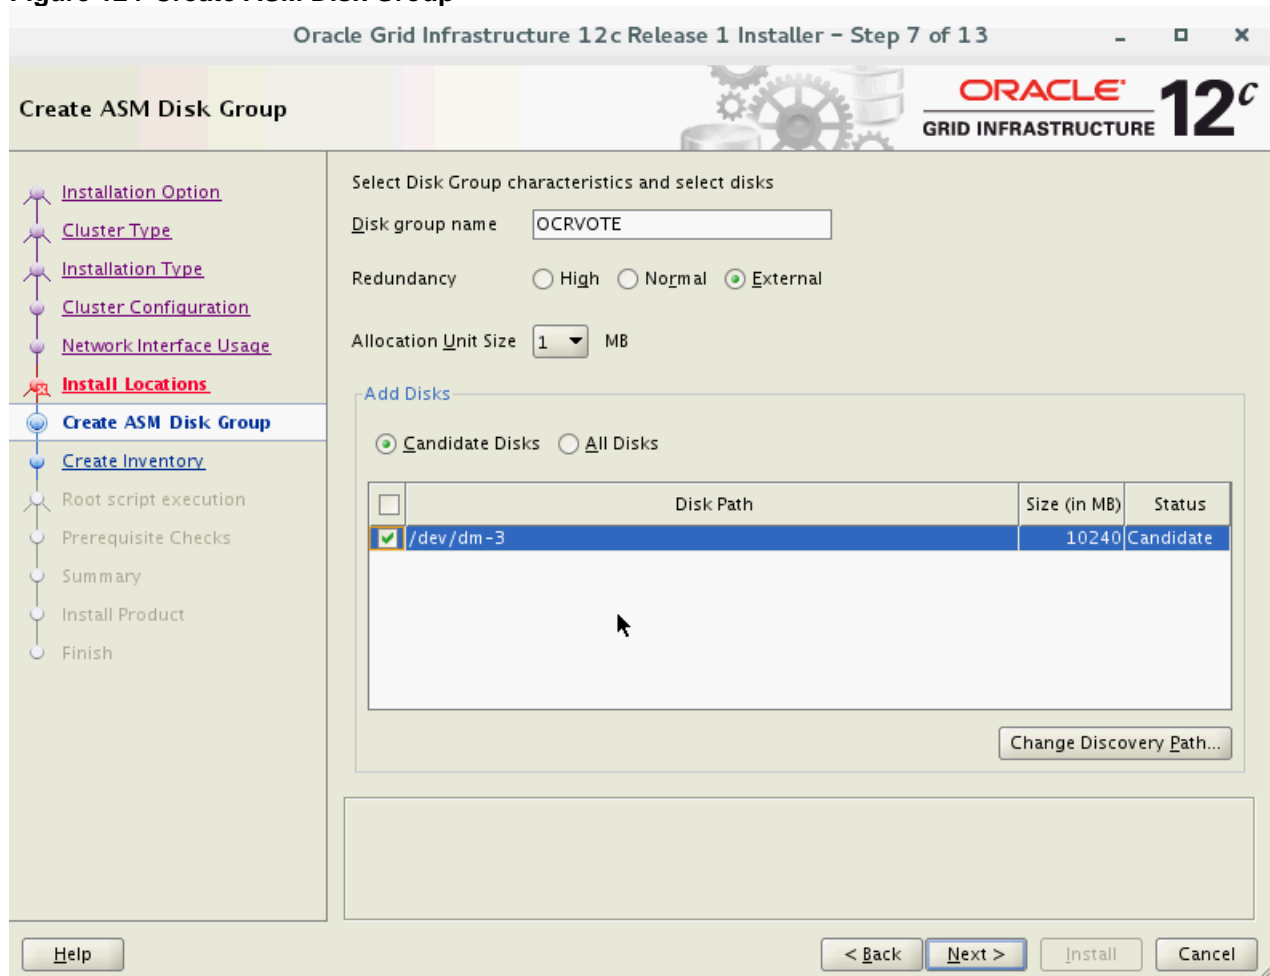
Messages:

**SYSASM Password:[INS-30011] The SYS password entered does not conform to the Oracle recommended standards.**

Help

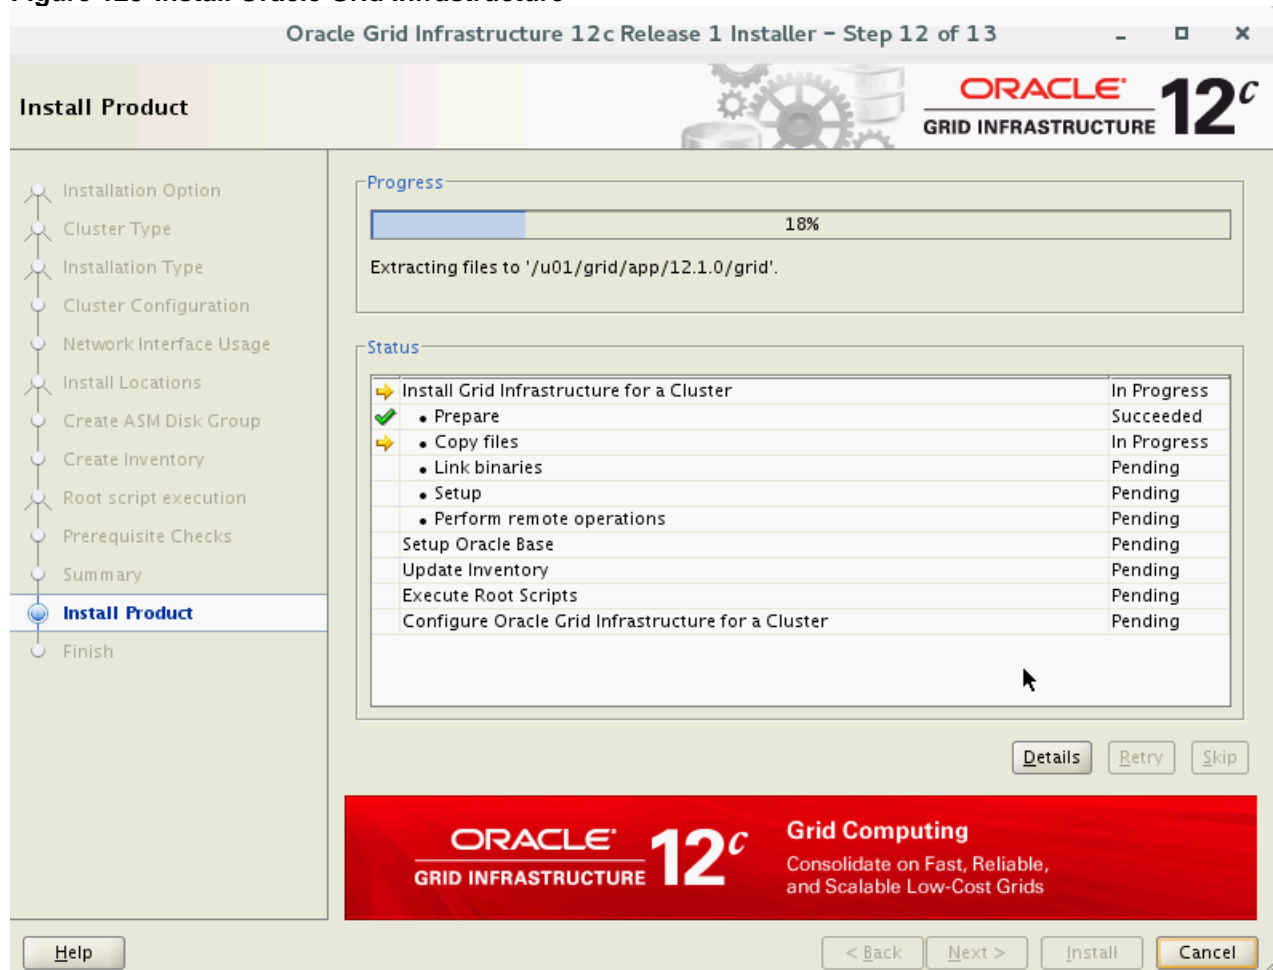
10. Enter Disk group name and set the redundancy. For this solution we have setup redundancy as "External". Click Change Discovery Path and set the path to "/dev/dm-3". Return to the main screen and select the OCR LUN assigned from Pure Storage to store OCR and Voting disks. Click Next.

Figure 124 Create ASM Disk Group



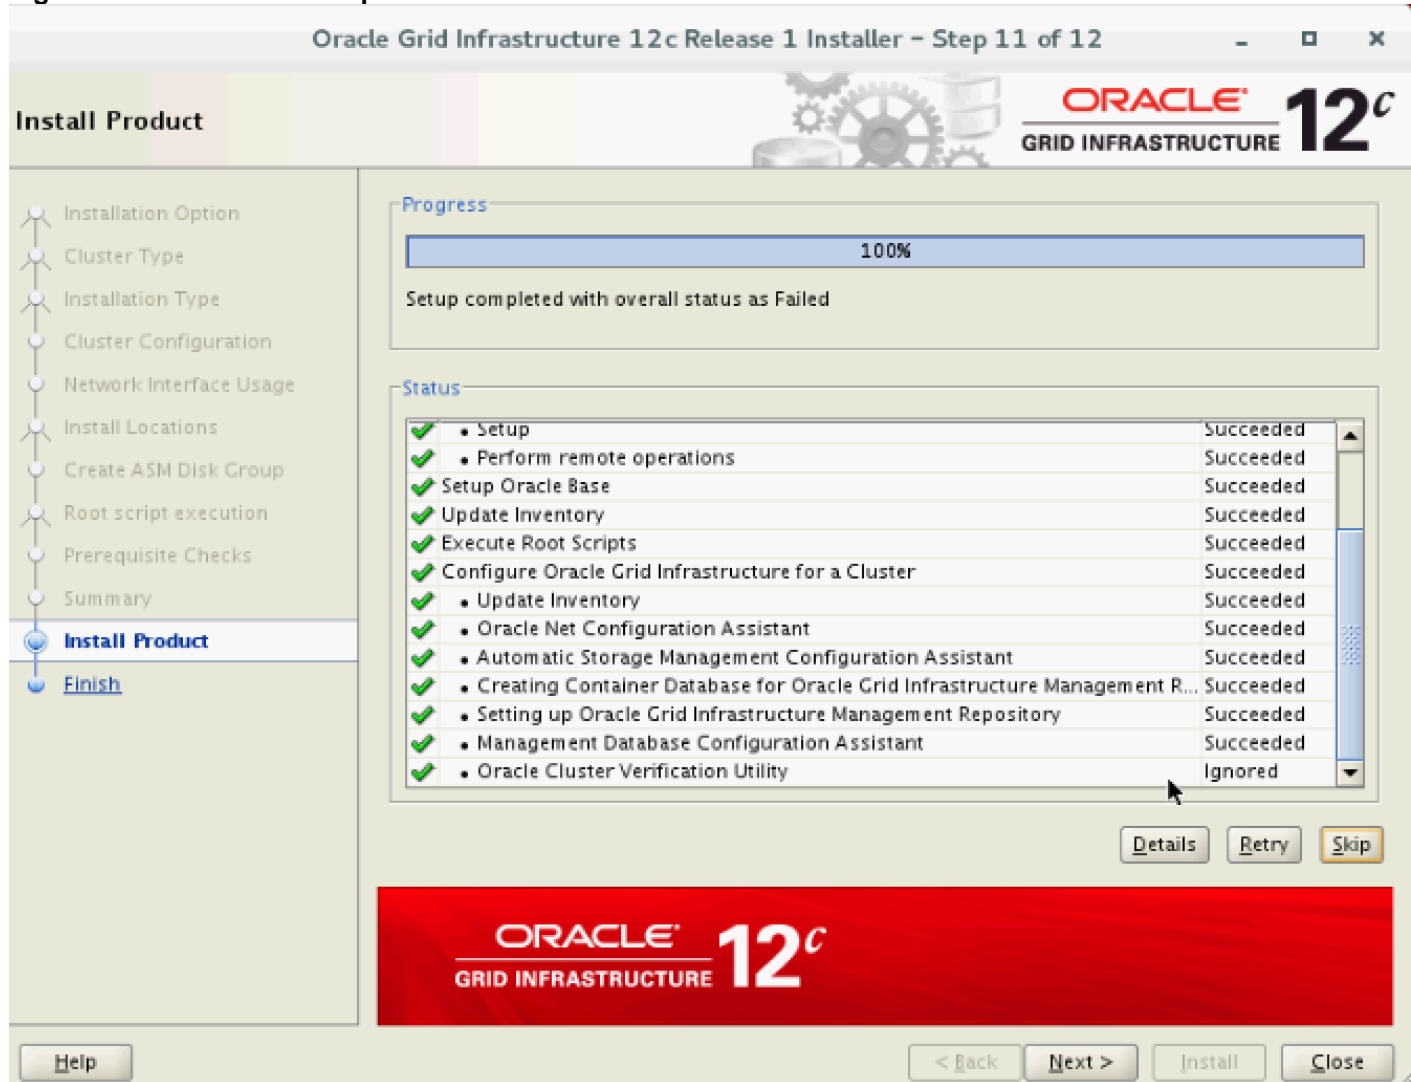
11. In next screen, accept the default inventory directory and click Next.
12. Click Automatically run configuration scripts to run scripts automatically and enter the relevant root user credentials. Click Next.
13. Wait while the prerequisite checks complete. If there are any issues, click the Fix & Check Again button.
14. Check and verify all the summary information and click Install to start the grid installation.

Figure 125 Install Oracle Grid Infrastructure



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

Figure 126 Installation Completed for Oracle Grid Infrastructure



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



Once GRID install is successful, log into each of the nodes and perform minimum health checks to make sure that Cluster state is healthy.

## Install Oracle Database Software

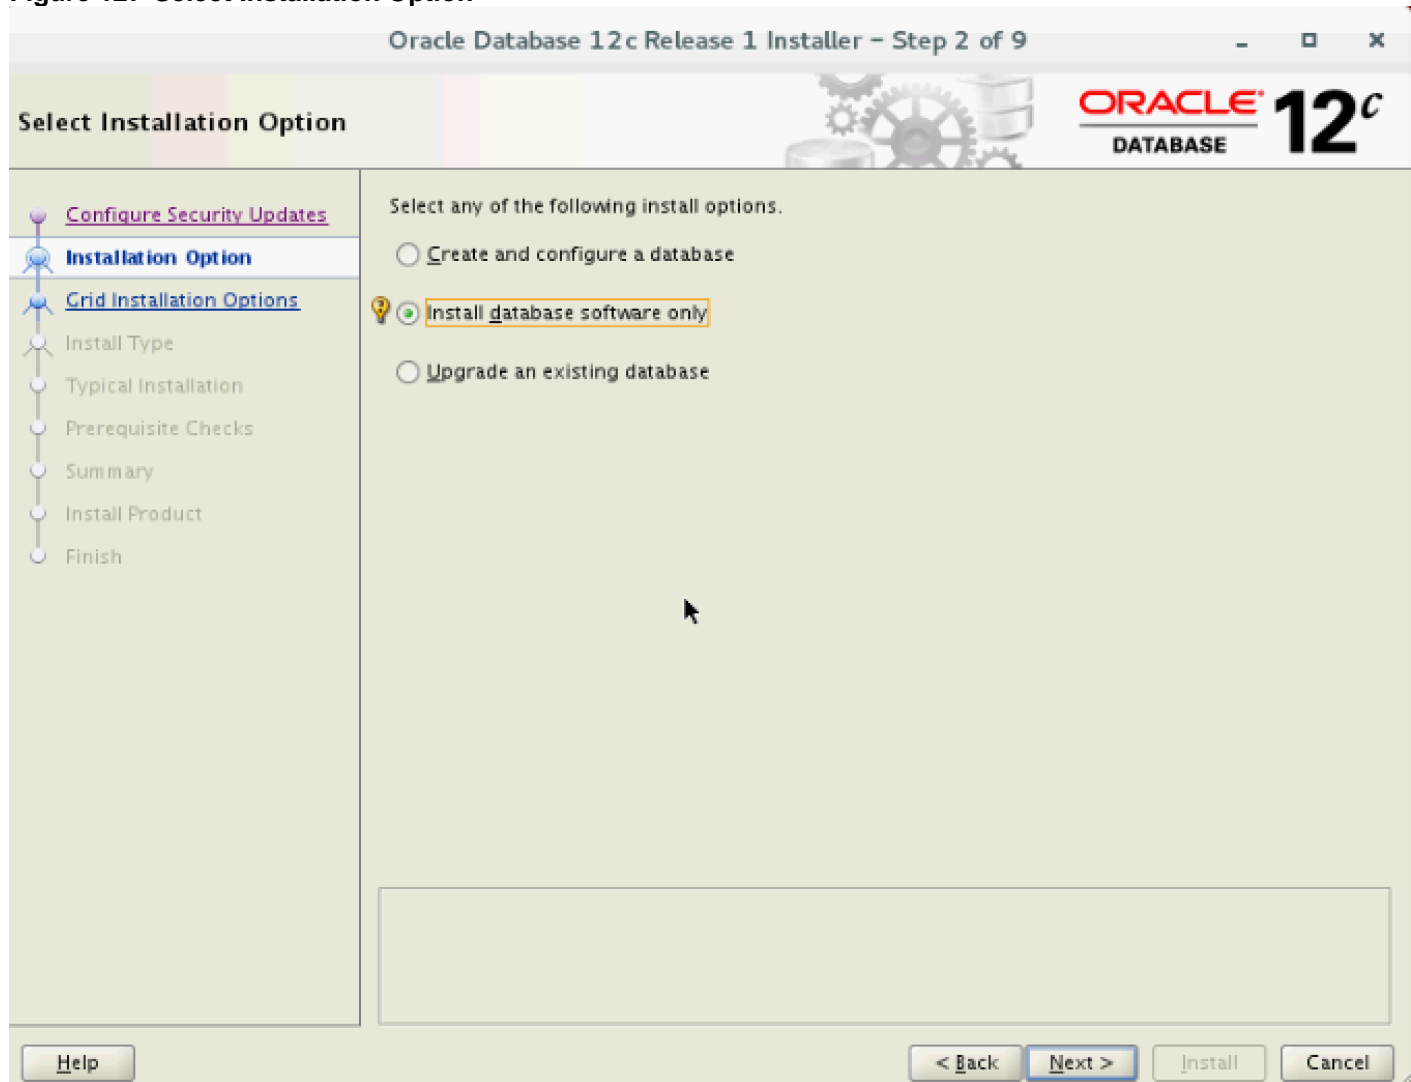
After successful GRID install, we recommend to install Oracle Database 12c software only. You can create databases using DBCA or database creation scripts at later stage.

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

To install Oracle Database Software, complete the following steps:

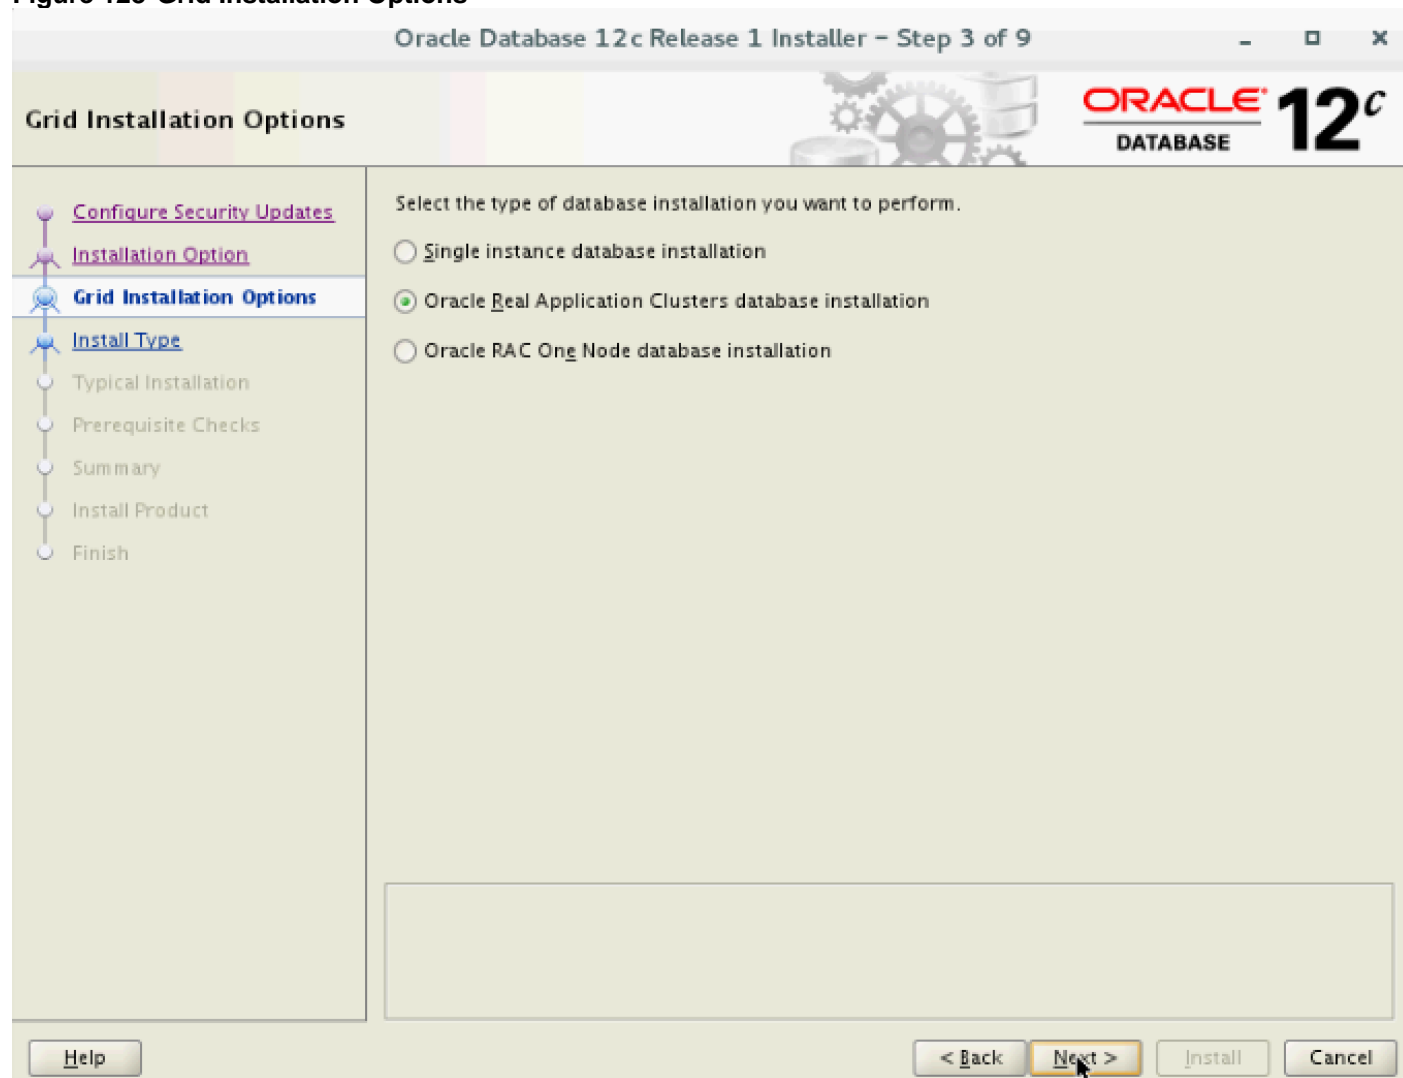
1. Launch the installer as the "oracle" user from the staging area where the Oracle database software is located.
2. Select option "Install database software only" into Select Installation Option.

Figure 127 Select Installation Option



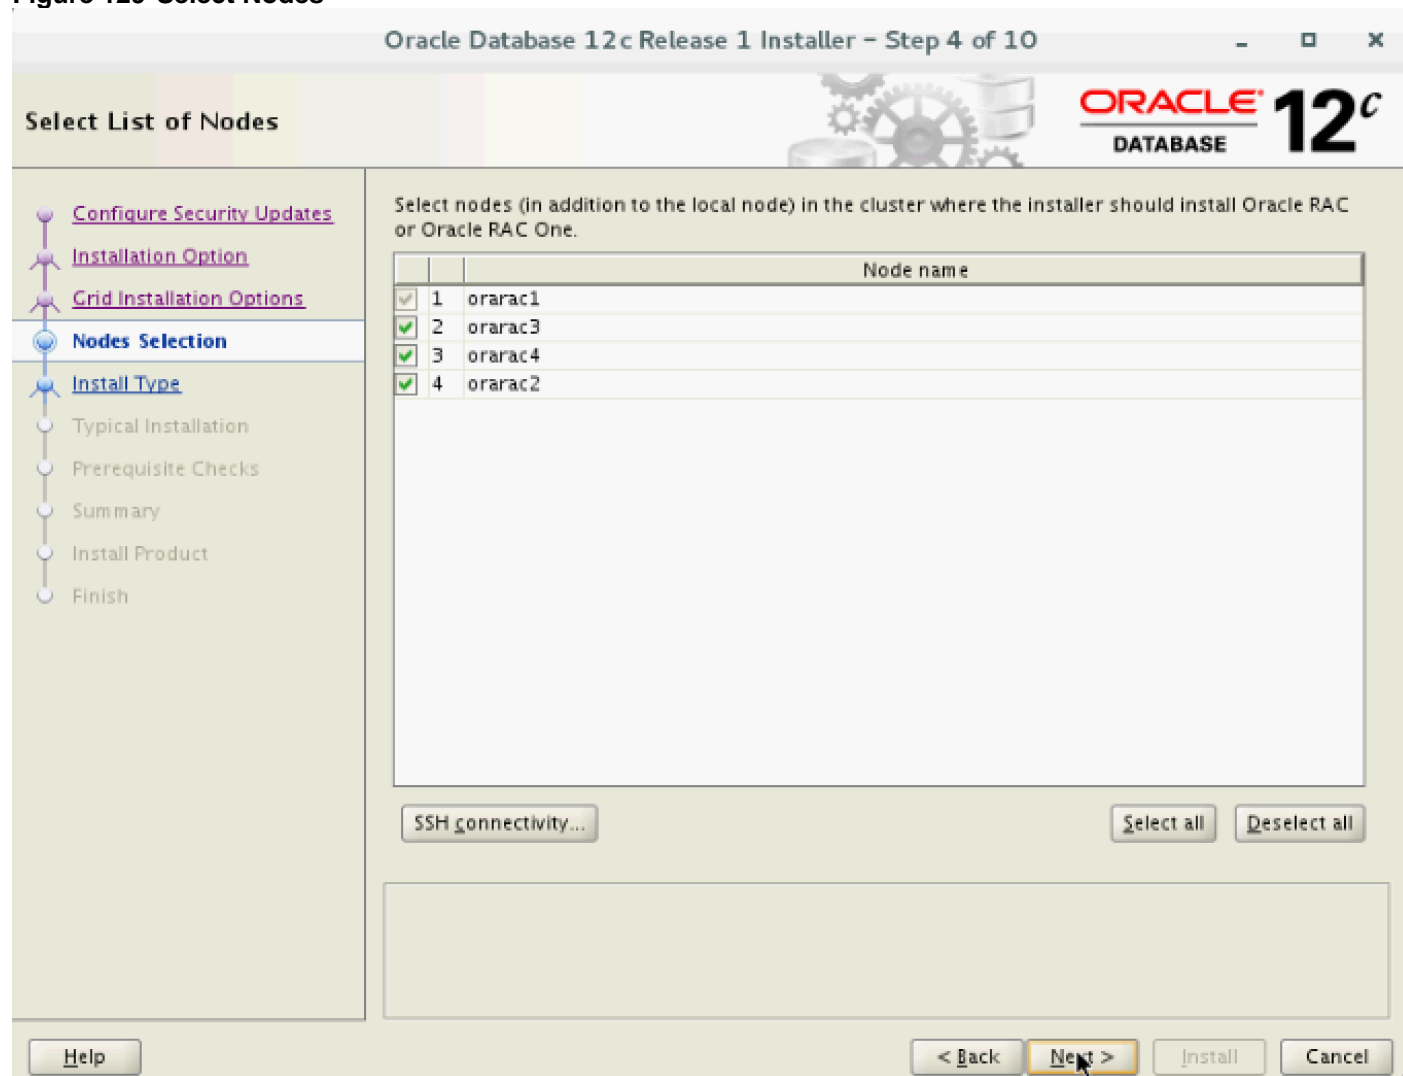
3. Select option "Oracle Real Application Clusters database installation" and click Next.

Figure 128 Grid Installation Options



4. Select nodes in the cluster where installer should install Oracle RAC. For this setup, install four Oracle nodes as shown below. Also scale up the system by adding additional nodes in later section.

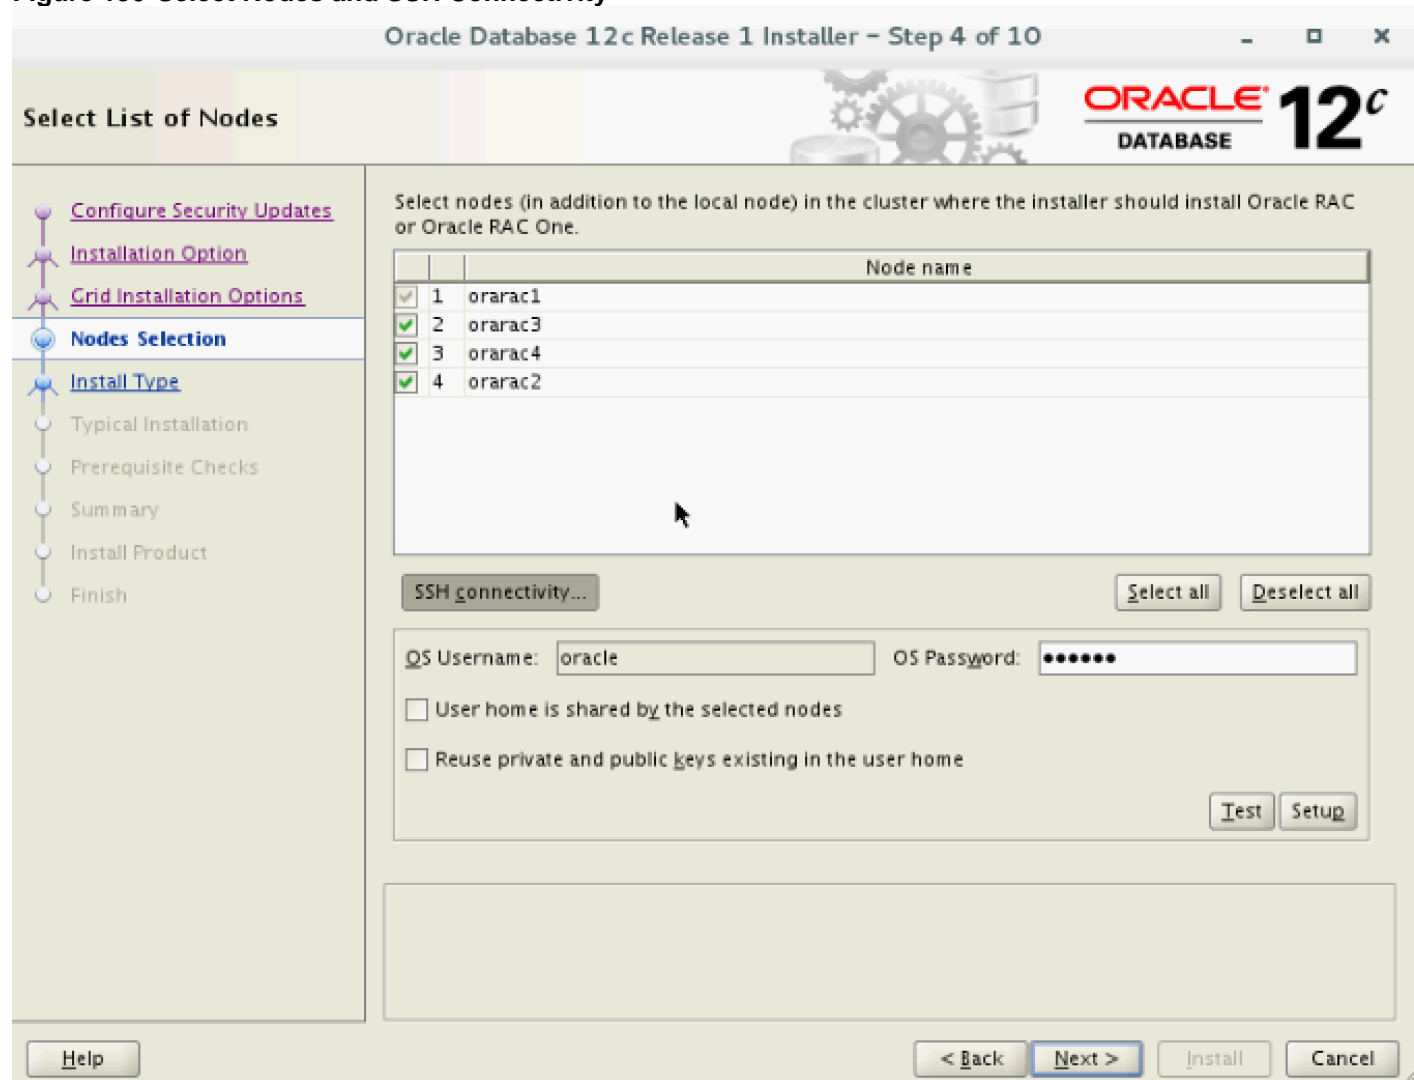
Figure 129 Select Nodes



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



Figure 130 Select Nodes and SSH Connectivity



6. Select the installation type as "Typical install" and click Next.

Figure 131 Select Install Type



7. Enter Oracle Base as "/u01/oracle/app" and "/u01/oracle/app/product/12.1.0.2/db\_1" as the software location, then click Next.
8. Select the desired operating system groups, 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 have been run on each node, click OK.
12. Click Close to exit the installer.



We have only added first four nodes (orac1, orac2, orac3 and orac4) to the cluster for now to check the scalability and performance of 4 node RAC Cluster with Pure Storage FlashArray //m20. Later in the solution, we will add remaining four nodes (orac5, orac6, orac7 and orac8) to join cluster and upgrade the Storage to Pure Storage FlashArray //m70 while database workload is running to show the scalability of FlashStack solution as demand grow.

## Scalability Test and Results

Before configuring any database for workload tests, it is extremely important to validate that this is indeed a balanced configuration that is capable of delivering the expected performance.

In this FlashStack solution, we will first test and validate node and user scalability for Pure Storage FlashArray //m20 on 4 Oracle RAC Nodes. Then we will add the remaining 4 nodes to join the existing RAC cluster and upgrade the storage to FlashArray //m70 while database and database workload is running. We will repeat all the tests for scalability and check the performance for the new storage. This will show the scalability of FlashStack Solution as demand grows and upgrade of the system without any downtime.


In this solution, we use widely adopted I/O generation tools such as Linux FIO, SLOB (Silly Little Oracle Benchmark) and Swingbench to test and validate node and user scalability on both FlashArray //m System as explained below.

### Hardware Calibration using FIO on FlashArray //m20

Flexible IO (FIO), 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 storage device over a given period of time.

We run three FIO tests by changing Read/Write ratio and record the results over a given period of time. We run 100% Write, 80/20% Read/Write and 100% read and scale the IOPS and throughput as shown below.

**Figure 132 FIO results for FlashArray //m20**



### Database Creation with DBCA

We used Oracle Database Configuration Assistant (DBCA) to create two OLTP (OLTPDB1 and OLTPDB2) and one DSS (DSSDB1) databases. Alternatively, you can use Database creation scripts to create the databases as well. Please ensure to place the data files, redo logs and control files in appropriate directory paths discussed in the storage layout section.

## Database Workload Configuration

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

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

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

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

For this solution, we created two OLTP (Order Entry) and one DSS (Sales History) database to demonstrate database consolidation, multi-tenancy capability, performance and sustainability. The OLTP1 and OLTP2 databases are approximately 6 TB each in size while DSS database is approximately 10 TB in size.

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

- OLTP database user scalability and OLTP database node scalability representing small and random transactions
- DSS database workload representing larger transactions
- Mixed workload featuring OLTP and DSS database workloads running simultaneously for 24 hours

## SLOB Performance on FlashArray //m20

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

SLOB issues single block reads for the read workload that are generally 8K (as the database block size was 8K). Following tests were performed and various metrics like IOPS and latency were captured along with Oracle AWR reports for each test.

## User Scalability Performance on FlashArray //m20


SLOB was configured to run against all the 4 RAC nodes and the concurrent users were equally spread across all the nodes. For Pure Storage FlashArray //m20, we scale users from 16 to 128 for Oracle RAC 4 nodes and identify the maximum IOPS and latency as explained.

- User Scalability testing with 16, 32, 64 and 128 users for Oracle RAC 4 nodes
- Varying workloads
  - 90% read, 10% update
  - 70% read, 30% update

— 50% read, 50% update

The following graphs illustrate user scalability in terms of the total IOPS (both read and write) when run at 16, 32, 64 and 128 concurrent users with 90% read, 70% read and 50% read. SLOB was configured to run against all the 4 RAC nodes and the concurrent users were equally spread across all the nodes.

**Figure 133 User Scalability IOPS for FlashArray //m20**



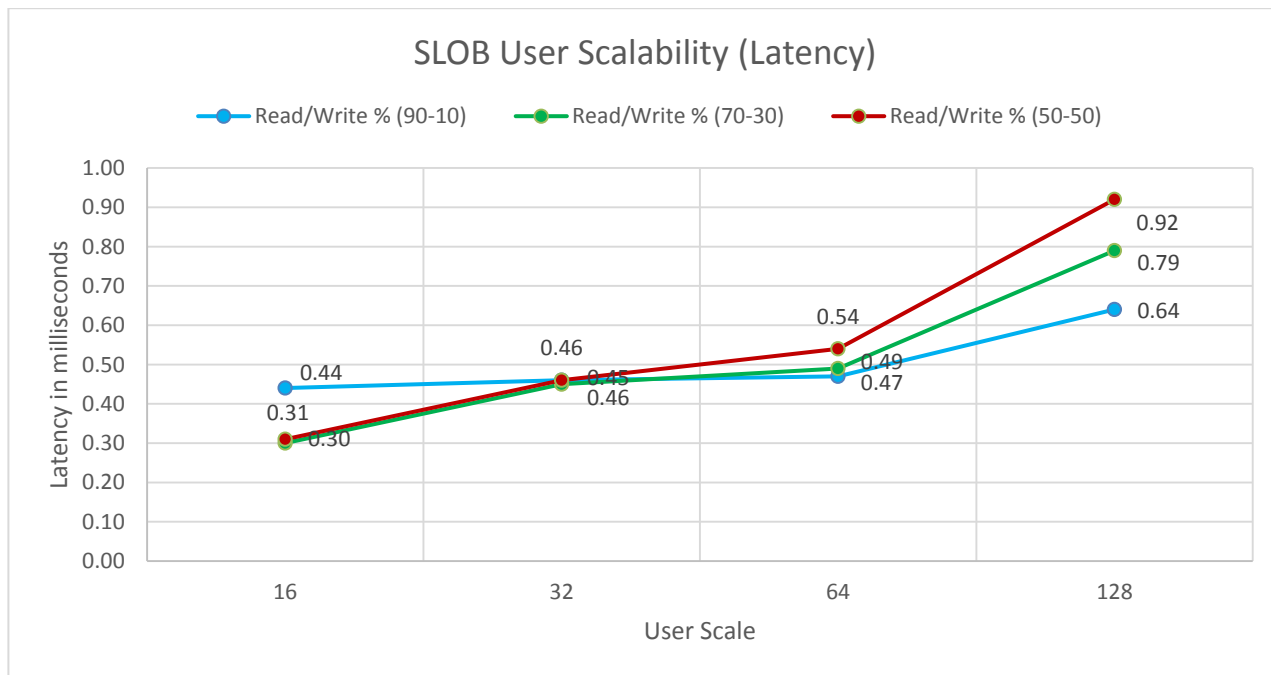
As expected, the graph illustrates the linear scalability with increased users and similar IOPS (around 205K) until 128 users across all workloads. Beyond 170k IOPS, the additional users yield higher IOPS but not at the same IOPS/user rate when the write percentages are higher.

The 50% update (50% read) resulted in 200k IOPS at 128 users, whereas 70% read (30% update) resulted in 205k IOPS at 128 users. The 90% read (10% update) resulted in 208k IOPS at 128 users which are excellent considering the tests were performed on the low-range FlashArray //m20.

Even though the FlashArray //m20 can scale up to 3GBps of reads we were limited by the total number of IOPS and not on the bandwidth. The maximum bandwidth is validated with the DSS query.

The following graph illustrates the latency exhibited by the //m20 FlashArray across different workloads. All the workloads experienced less than 1ms and it varies based on the workload and the 50% read (50% update) exhibited higher latencies with increased user counts.

**Figure 134 User Scalability Latency for FlashArray //m20**

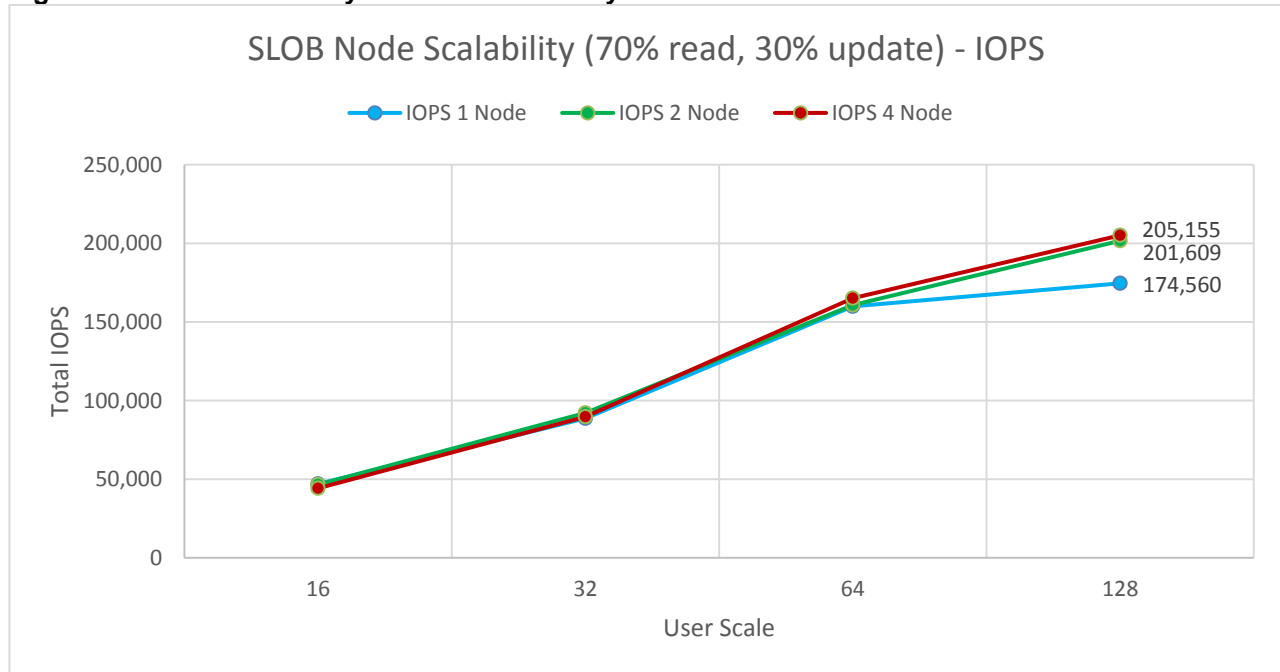


### Node Scalability Performance on FlashArray //m20

To validate the node scalability, we ran the same SLOB tests for the following scenarios.

- Node Scalability testing with 1, 2 and 4 nodes
  - 70% read, 30% update workload
  - 16, 32, 64 and 128 users

The following graph illustrates the behavior of the node scalability with the IOPS metrics. Running the 70% read and 30% update operations workload on 4 nodes had generated 205k IOPS at 128 users.

**Figure 135 Node Scalability IOPS for FlashArray //m20**

## SwingBench Performance on FlashArray //m20

### Scalability Performance on FlashArray //m20

The first step after the databases creation is calibration; about the number of concurrent users, nodes, OS and database optimization. For Pure Storage FlashArray //m20, scale the system for 1 to 4 Oracle RAC Nodes. Also, for this FlashStack solution, test system performance with different databases running at a time and capture the results as explained below.

#### OLTPDB1 Database Performance


For OLTP database workload featuring Order Entry schema, we used one OLTPDB1 database. For the OLTPDB1 database (6 TB), we used 64GB size of System Global Area (SGA). We also ensured that HugePages were in use. The OLTP Database scalability test was run for at least 12 hours and ensured that results are consistent for the duration of the full run.

Run the SwingBench scripts on each node to start OLTP1DB1 database and generate AWR reports for each scenario as shown below.

- User Scalability

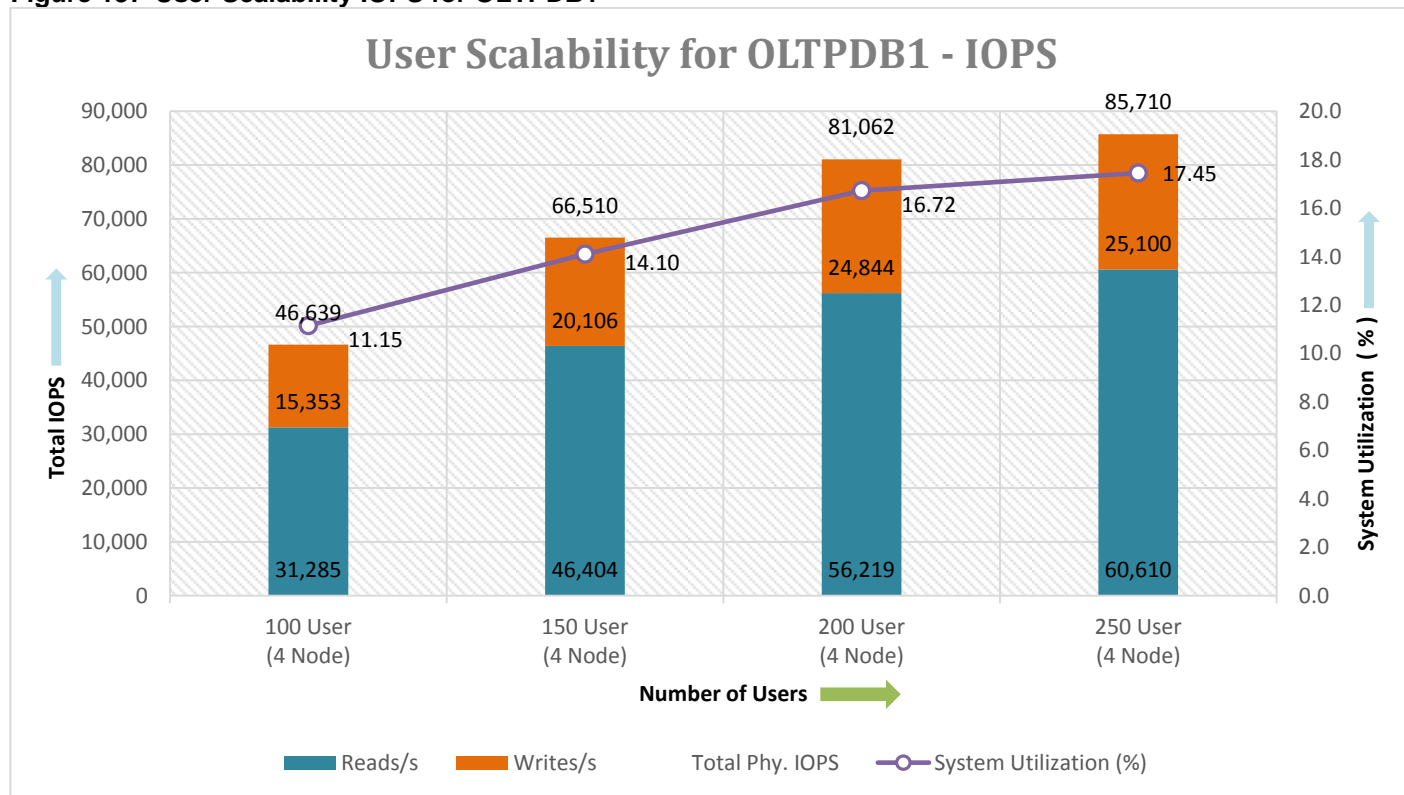
The graph illustrates the TPM for OLTPDB1 database user scale on 4 node. TPM for 100, 150, 200 and 250 user are around 333k, 623k, 746k and 774k with latency under 3 milliseconds all the time.

Figure 136 User Scalability TPM for OLTPDB1



The graph illustrates the Total IOPS for OLTPDB1 database user scale on 4 node. Total IOPS for 100, 150, 200 and 250 user are around 46k, 66k, 81k and 85k with system utilization under 17% all the time.

Figure 137 User Scalability IOPS for OLTPDB1

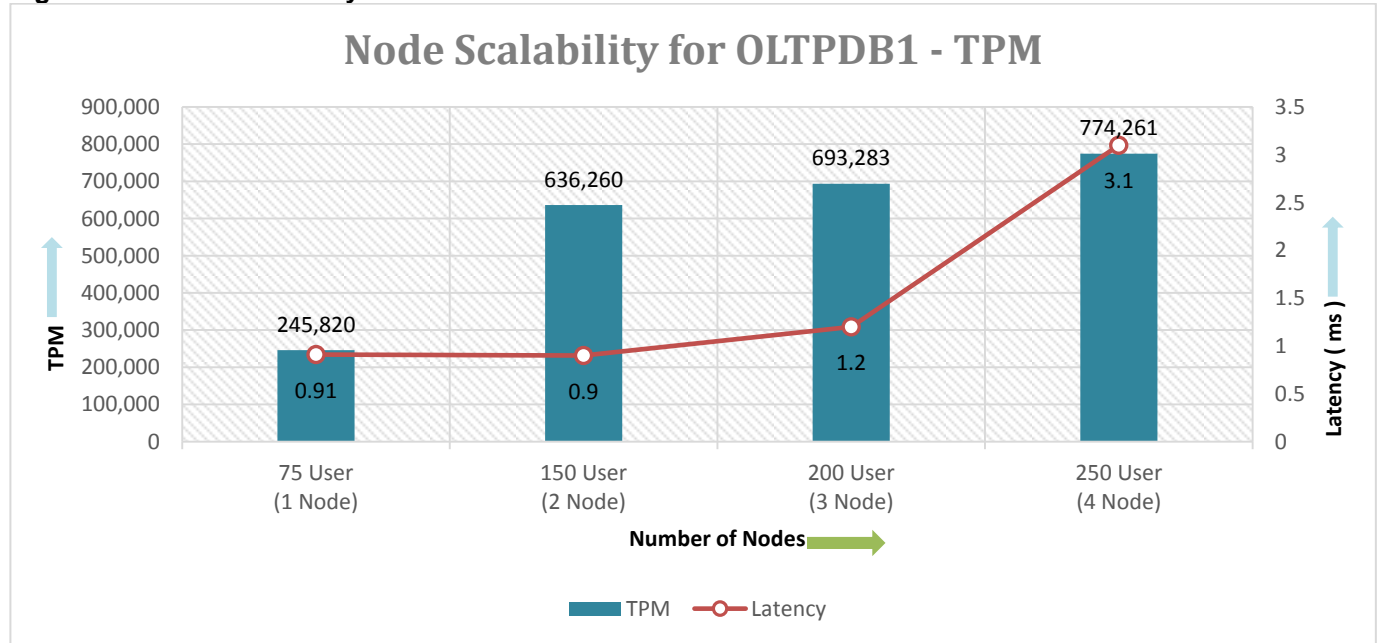




- Node Scalability


The graph illustrates the TPM for OLTPDB1 database node scale. TPM for 1, 2, 3 and 4 nodes are around 245k, 636k, 693k and 774k with latency under 3 milliseconds all the time.

**Figure 138 Node Scalability TPM for OLTPDB1**



The graph illustrates the Total IOPS for OLTPDB1 database node scale. Total IOPS for 1, 2, 3 and 4 nodes are around 62k, 72k, 86k and 90k with system utilization under 25% all the time.

**Figure 139 Node Scalability IOPS for OLTPDB1**



### OLTPDB1 + OLTPDB2 Database Performance

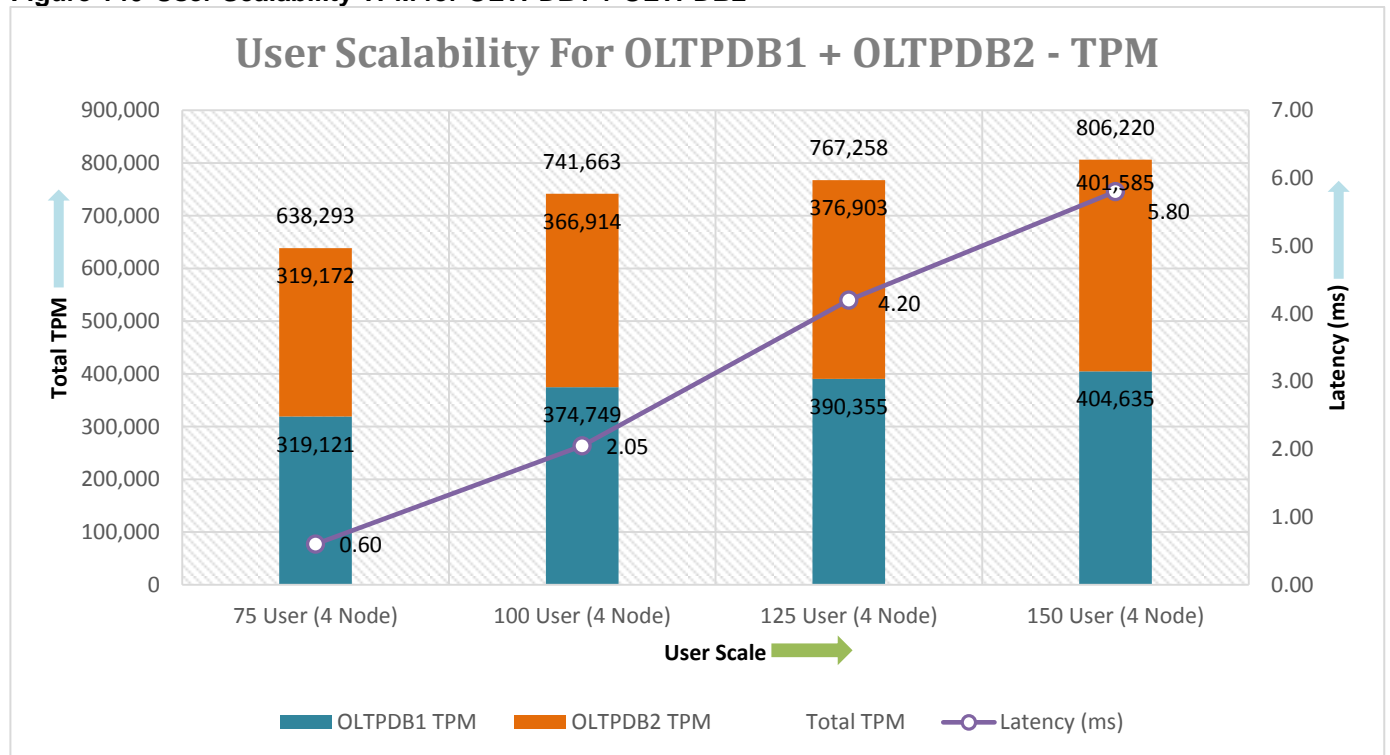
For OLTPDB1 + OLTPDB2 database workload featuring Order Entry schema, we used two OLTPDB1 and OLTPDB2 databases. For the OLTPDB1 database (6 TB), we used 64GB size of System Global Area (SGA) and for the OLTPDB2 database (6 TB), we used 64GB size of System Global Area (SGA). We also ensured that HugePages were in use. The OLTPDB1 + OLTPDB2 Database scalability test was run for at least 12 hours and ensured that results are consistent for the duration of the full run.

Run the SwingBench scripts on each node to start OLTP1DB1 and OLTPDB2 database and generate AWR reports for each scenario as shown below.

- User Scalability


The graph illustrates the TPM for OLTPDB1 + OLTPDB2 database user scale on 4 nodes. TPM for 75, 100, 125 and 150 user are around 638k, 741k, 767k and 806k with latency under 5 milliseconds all the time.

**Figure 140 User Scalability TPM for OLTPDB1 + OLTPDB2**



The graph illustrates the Total IOPS for OLTPDB1 + OLTPDB2 database user scale on 4 nodes. Total IOPS for 75, 100, 125 and 150 user are around 71k, 83k, 87k and 92k with system utilization under 22% all the time.


**Figure 141 User Scalability IOPS for OLTPDB1 + OLTPDB2**



- Node Scalability


The below graph illustrates the TPM for OLTPDB1 + OLTPDB2 database node scale. TPM for 1, 2, 3 and 4 nodes are around 363k, 671k, 719k and 806k with latency under 6 milliseconds all the time.

Figure 142 Node Scalability TPM for OLTPDB1 + OLTPDB2



The graph illustrates the Total IOPS for OLTPDB1 + OLTPDB2 database node scale. Total IOPS for 1, 2, 3 and 4 nodes are around 42k, 74k, 85k and 92k with system utilization under 22% all the time.

Figure 143 Node Scalability IOPS for OLTPDB1 + OLTPDB2




### DSSDB1 Database Performance

DSSDB1 database workloads are generally sequential in nature, read intensive and exercise large IO size. DSSDB1 database workload runs a small number of users that typically exercise extremely complex queries that run for hours. DSSDB1 Database activity is captured for four Oracle RAC Instances using Oracle Enterprise Manager for 24 hours workload test.

For this test, we ran Swingbench Sales history workload with 4 users on 4 nodes. The following charts show DSSDB1 database workload results and average throughput around 3.75 GB/s.


**Figure 144 DSSDB1 Database Performance**



### Mixed Workload


The next test is to run both OLTP and DSS database workloads simultaneously. This test will ensure that configuration in this test is able to sustain small random queries presented via OLTP database along with large and sequential transactions submitted via DSS database workload. DSSDB1, OLTPDB1 and OLTPDB2 Database activity is captured for four Oracle RAC Instances using Oracle Enterprise Manager for 24 hours mixed workload test. We ran the test for OLTPDB1 + OLTPDB2 and DSSDB1 database together as shown in the chart below.

Figure 145 DSSDB1 Database Throughput Performance




The charts above show DSSDB1 database average Throughput around 2.1 GB/s for database workload OLTPDB1, OLTPDB2 and DSSDB1 running together for 4 nodes.

Figure 146 OLTPDB1 Database IOPS Performance



The charts above show OLTPDB1 database average IOPS around 15,000 and average TPM around 104k for database workload OLTPDB1, OLTPDB2 and DSSDB1 running together for 4 nodes.

Figure 147 OLTPDB2 Database IOPS Performance



The charts above show OLTPDB2 database average IOPS around 14,000 and average TPM around 104k for database workload OLTPDB1, OLTPDB2 and DSSDB1 running together for 4 nodes.

## FlashStack Scalability

This section describes how you can add nodes into existing Oracle RAC Cluster and upgrade Pure Storage from FlashArray //m20 to FlashArray //m70 without any downtime.

In this FlashStack solution, we initially tested and validated node and user scalability with 4 node Oracle RAC on Pure Storage FlashArray //m20. We then added remaining 4 nodes to join the existing RAC cluster and upgraded the storage to FlashArray //m70 non disruptively. We then repeated all the tests for node and user scalability and checked the performance for new storage. This shows the scalability of FlashStack Solution when demand grows as well as the flexibility of storage upgrade without any downtime.

We worked with Pure support resource to perform the Pure Storage FlashArray upgrade from //m20 to //m70 non disruptively while the database and workloads were up and running. Pure support swapped the secondary //m20 controller with that of //m70 and once the secondary //m70 controller is active, failed over the primary controller from //m20 to //m70 and swapped the secondary //m20 controller with //m70. Also the support added two NVRAM modules as //m70 is equipped with 4 NVRAM modules. The NDU activity took less than 30 minutes to complete and then additional shelves were added to complete the upgrade process.



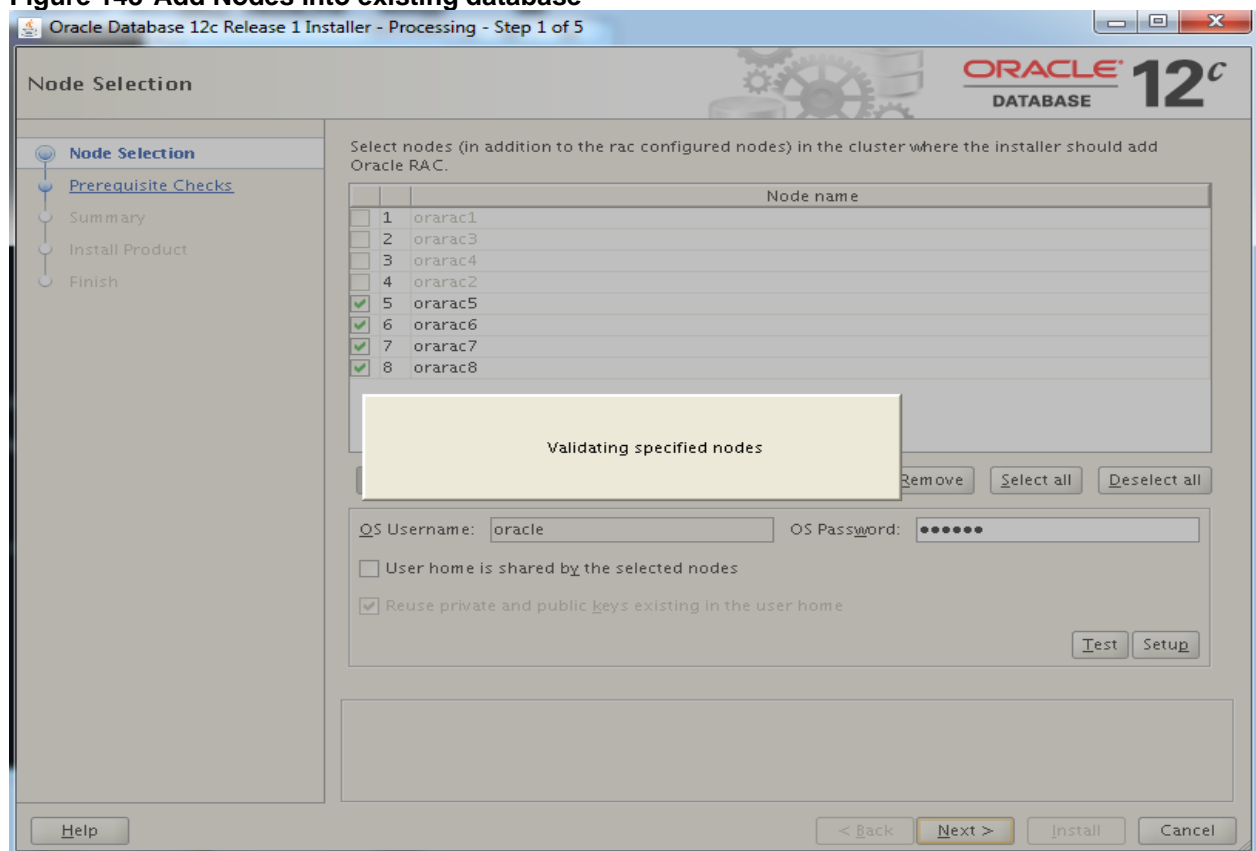
The NDU process is initiated and performed by the Pure Support team.

Now, extend an existing Oracle Real Application Clusters (Oracle RAC) home to other remaining four nodes (orarac5, orarac6, orarac7 and orarac8) and instances in the cluster to make the system 8 node Oracle RAC.

We used a local (non-shared) Oracle home for all Oracle RAC nodes, so we must extend the Oracle RAC database home that is on an existing node (orarac1) to a target node (orarac5, orarac6, orarac7 and orarac8).

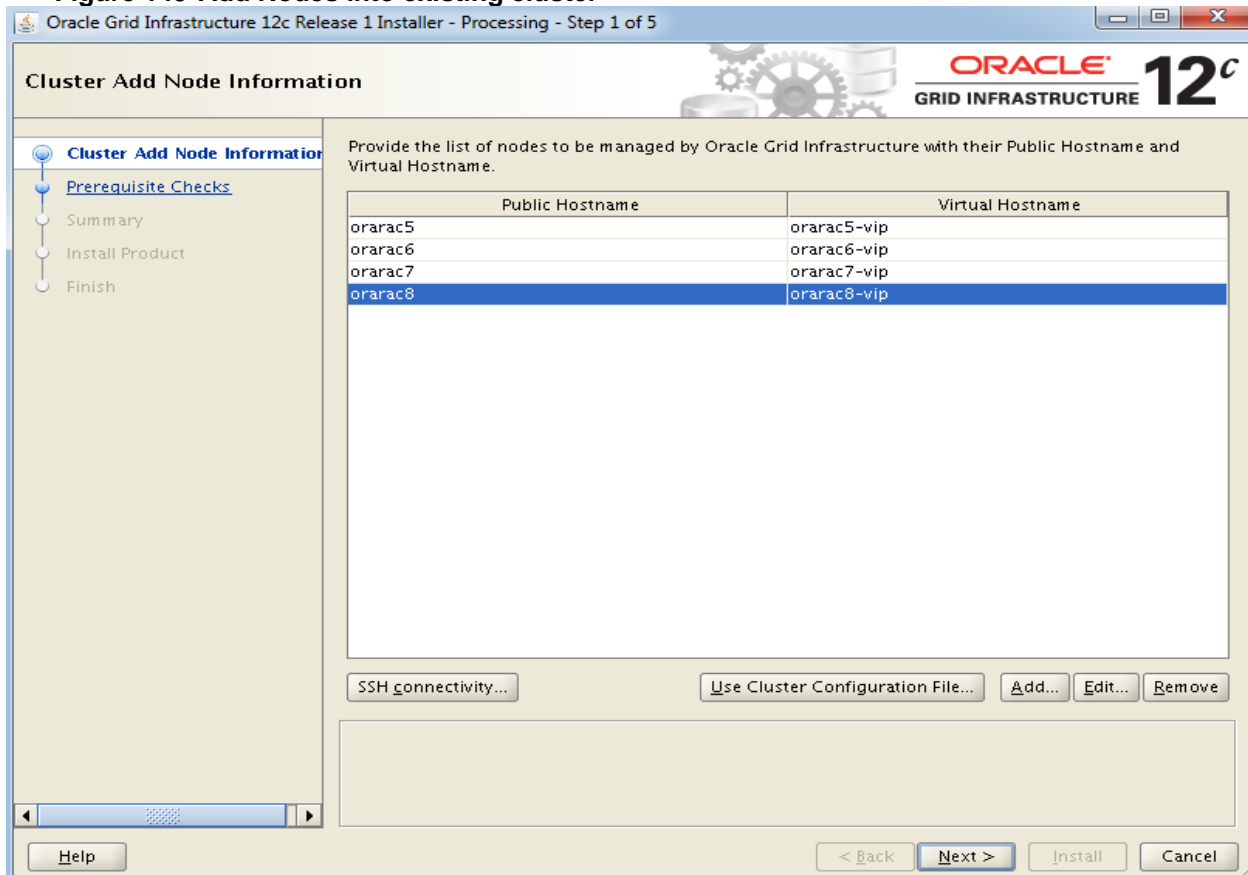
1. Go to orarac1 and run the script "addnode.sh" from "\$ORACLE\_HOME/addnode" directory.

2. Click Add to add remaining 4 nodes.
3. Click SSH Connectivity and enter the password for the "oracle" user.
4. Click Setup to configure SSH connectivity, then click Test to test it once it is complete. Once the test is complete, click Next. Review summary and then finish adding remaining four nodes as shown below.

**Figure 148 Add Nodes into existing database**



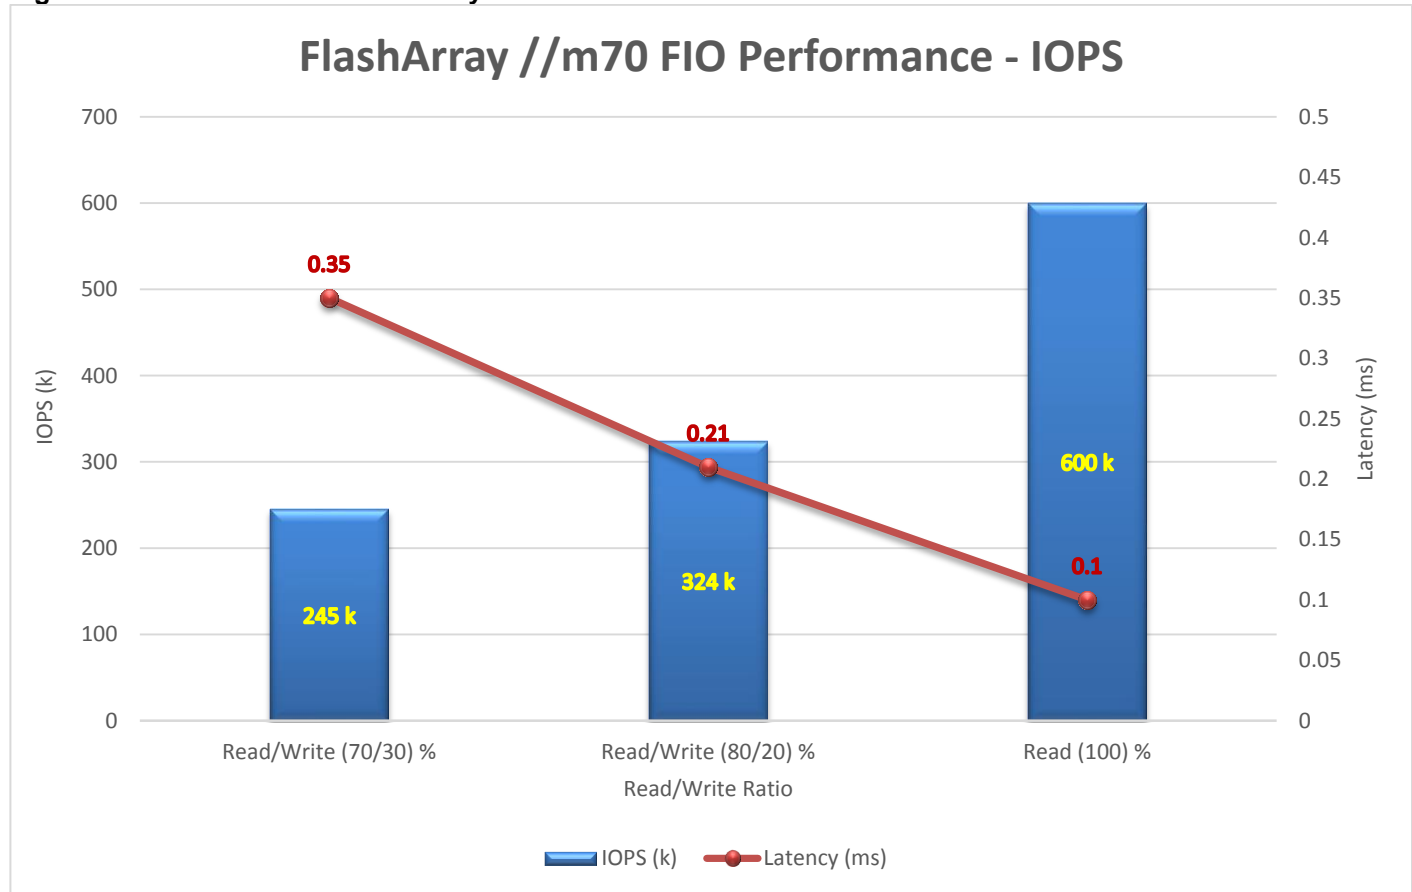
**Figure 149 Add Nodes into existing cluster**



## Hardware Calibration using FIO on FlashArray //m70

We ran two FIO test to record the IOPS and the Throughput and used 8k block size to record IOPS by changing Read/Write ratio. For IOPS, we run 70/30 % Read/Write, 80/20 % Read/Write and 100% read. We use 128k block size to record the throughput. The below graph illustrates the average IOPS around 600k for 8k block size. Also, we recorded the average Throughput 8.6 GB/s for 128k Block size.

Figure 150 FIO results for FlashArray //m70



## SLOB Performance on FlashArray //m70

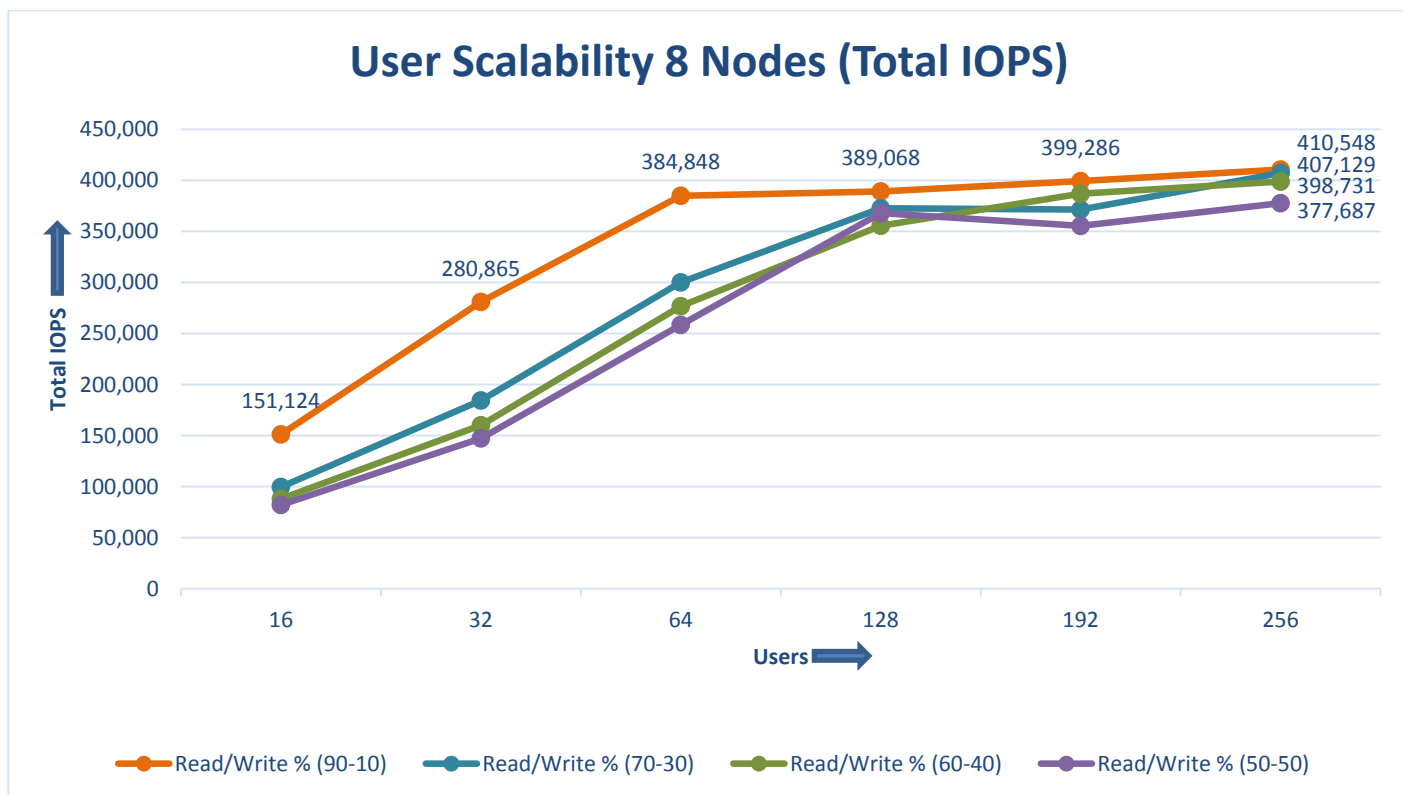
### User Scalability Performance on FlashArray //m20

SLOB was configured to run against all the 8 RAC nodes and the concurrent users were equally spread across all the nodes. For Pure Storage FlashArray //m70, we scale users from 16 to 256 for Oracle RAC 8 nodes and identify the maximum IOPS and latency as explained.

- User Scalability testing with 16, 32, 64, 128, 192 and 256 users for Oracle RAC 8 nodes
- Varying workloads
  - 90% read, 10% update
  - 70% read, 30% update
  - 60% read, 40% update
  - 50% read, 50% update

The following graphs illustrate user scalability in terms of the total IOPS (both read and write) when run at 16, 32, 64, 128, 192 and 256 concurrent users with 90% read, 70% read, 60% read and 50% read. SLOB was configured to run against all the 8 RAC nodes and the concurrent users were equally spread across all the nodes.

**Figure 151 User Scalability IOPS for FlashArray //m70**

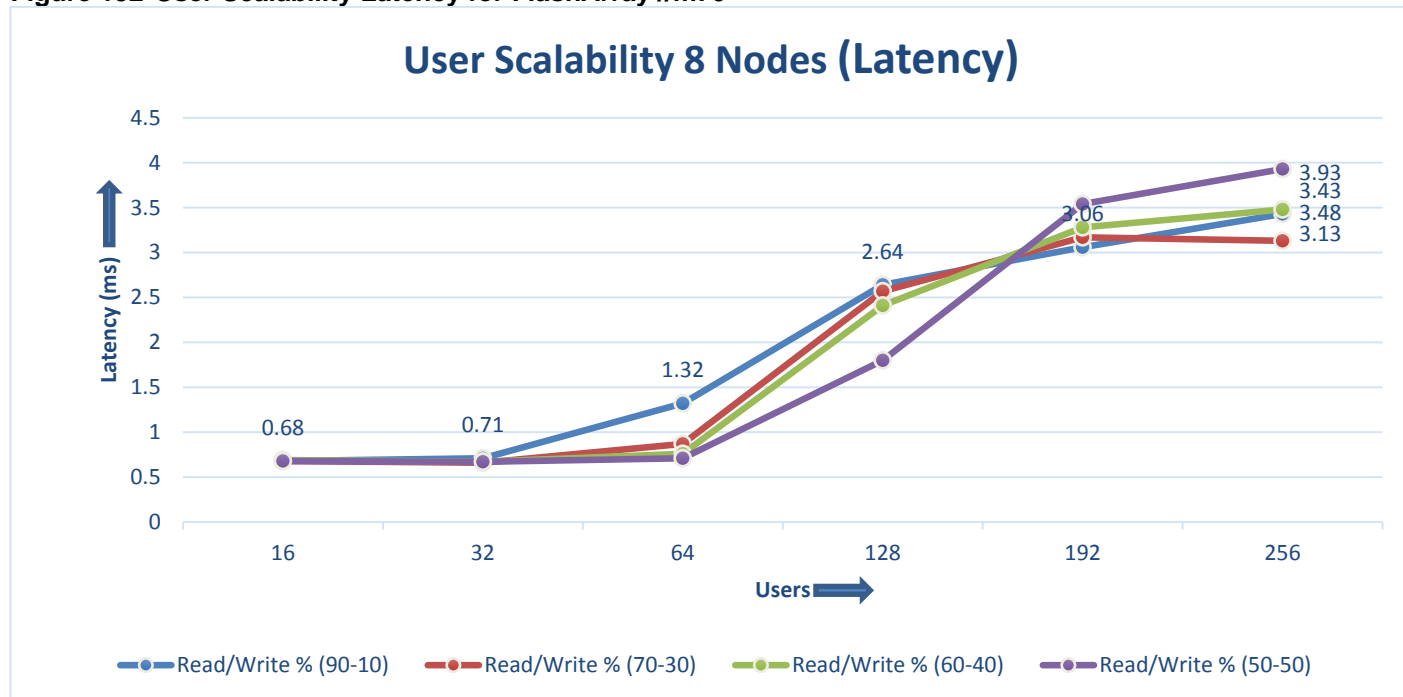


The graph illustrates the linear scalability with increased users and similar IOPS (around 389K) until 128 users across all workloads. Beyond 389k IOPS, the additional users yield higher IOPS but not at the same IOPS/user rate when the write percentages are higher.

The 50% update (50% read) resulted in 377k IOPS at 256 users, whereas 70% read (30% update) resulted in 407k IOPS at 256 users. The 90% read (10% update) resulted in 410k IOPS at 256 users which are excellent considering the tests were performed on the FlashArray //m70.

The following graph illustrates the latency exhibited by the //m70 FlashArray across different workloads. All the workloads experienced less than 4 millisecond and it varies based on the workload and the 50% read (50% update) exhibited higher latencies with increased user counts.

Figure 152 User Scalability Latency for FlashArray //m70



## SwingBench Performance on FlashArray //m70

### Scalability Performance on FlashArray //m70

For Pure Storage FlashArray //m70, we will scale the system for 1 to 8 Oracle RAC Nodes. Also, for this FlashStack solution, we will test system performance with different databases running at a time and capture the results as explained below.

#### OLTPDB2 Database Performance

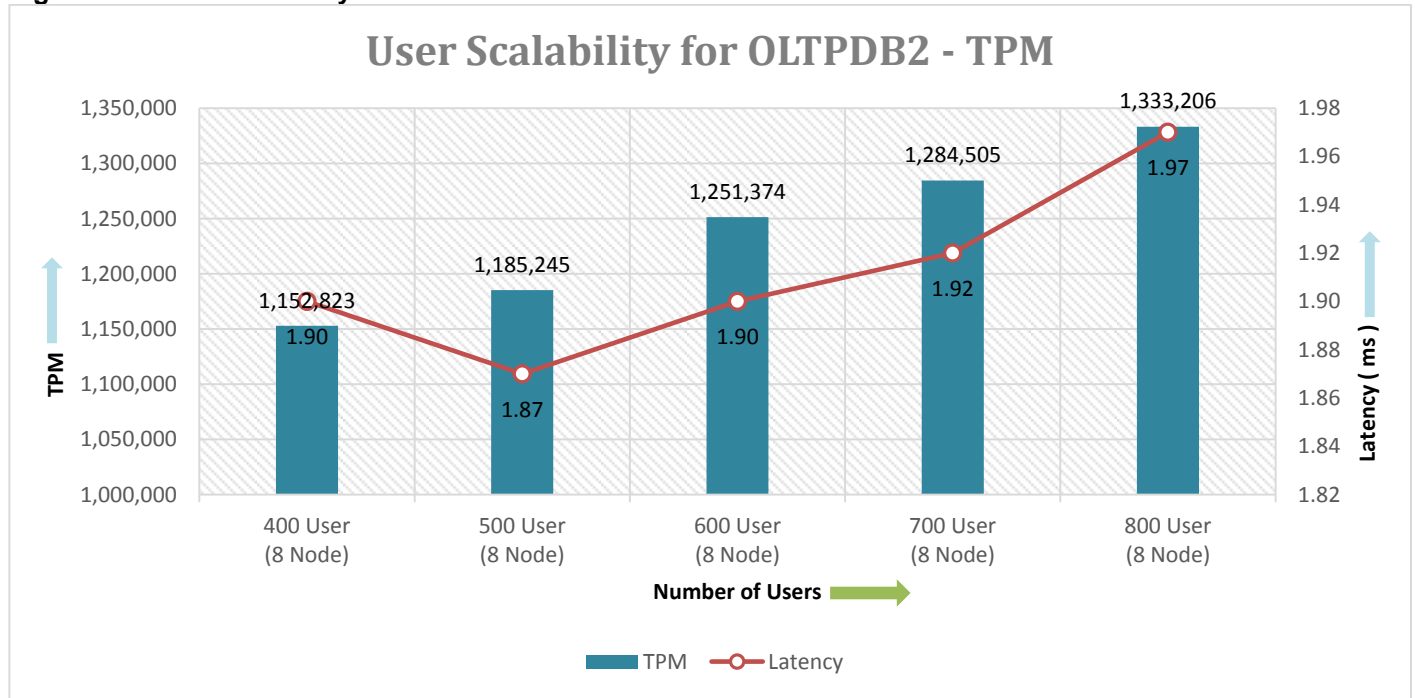
For OLTP database workload featuring Order Entry schema, we used one OLTPDB2 database. For the OLTPDB2 database (6 TB), we used 64GB size of System Global Area (SGA). We also ensured that HugePages were in use. The OLTP Database scalability test was run for at least 12 hours and ensured that results are consistent for the duration of the full run.

We run the SwingBench scripts on each node to start OLTPDB2 database and generate AWR reports for each scenario as shown below.

- User Scalability

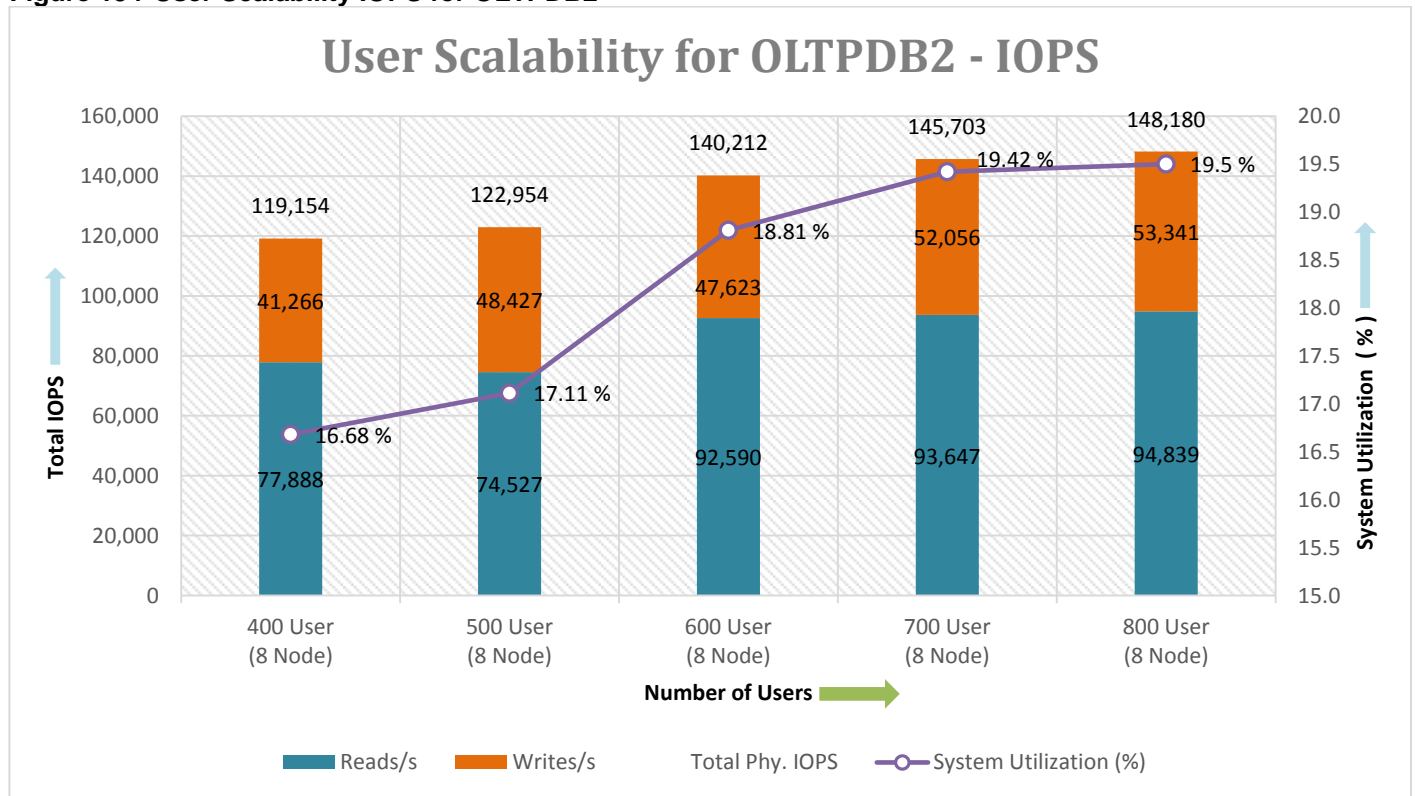
The following graph illustrates the TPM for OLTPDB2 database user scale on 8 node. TPM for 400, 500, 600, 700 and 800 user are around 1152k, 1185k, 1251k, 1284k and 1333k with latency under 2 milliseconds all the time.

Figure 153 User Scalability TPM for OLTPDB2



The following graph illustrates the Total IOPS for OLTPDB2 database user scale on 8 nodes. Total IOPS for 400, 500, 600, 700 and 800 user are around 119k, 122k, 140k, 145k and 148k with system utilization under 20% all the time.

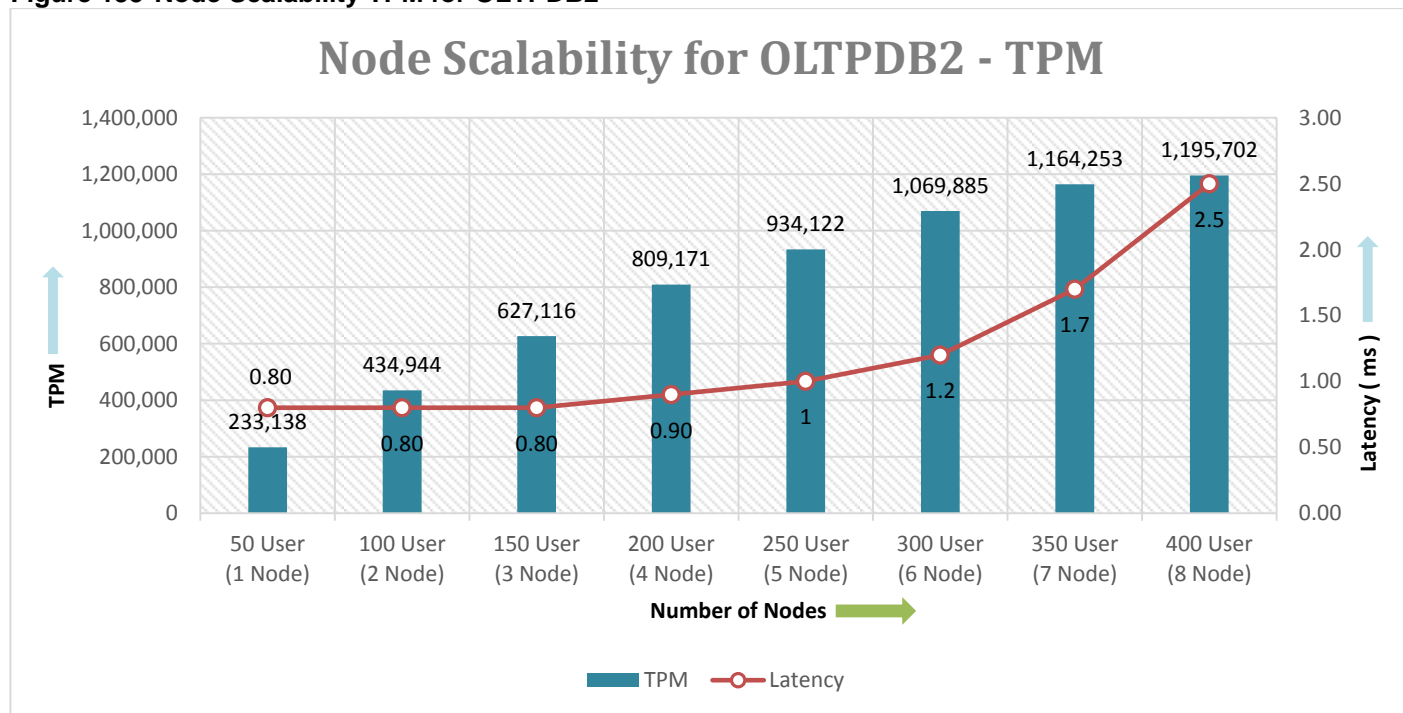
Figure 154 User Scalability IOPS for OLTPDB2



- Node Scalability

The following graph illustrates the TPM for OLTPDB2 database node scale. TPM for 1, 2, 3, 4, 5, 6, 7 and 8 nodes are around 233k, 434k, 627k, 809k, 934k, 1069k, 1164k and 1195k with latency under 3 milliseconds all the time.

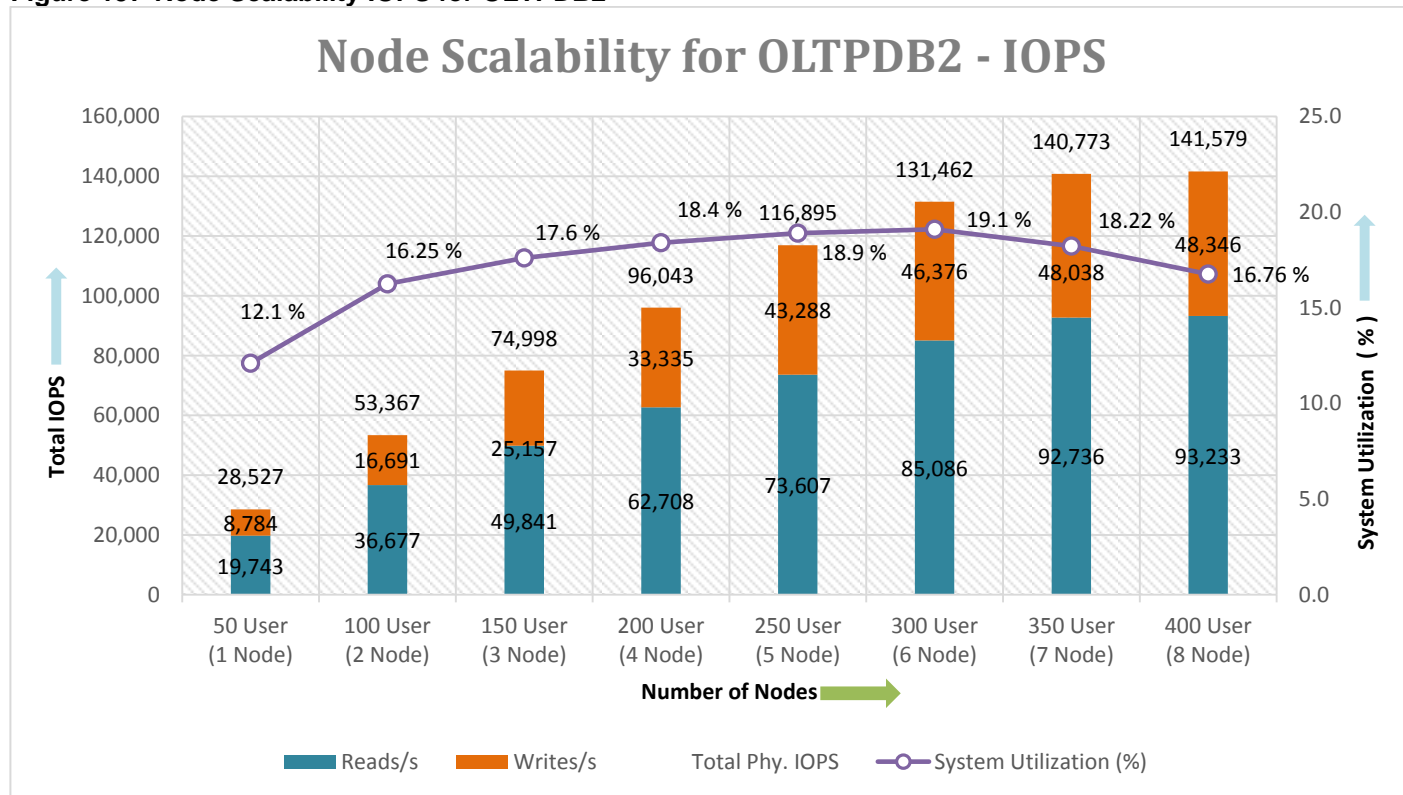
**Figure 155 Node Scalability TPM for OLTPDB2**



The following graph illustrates the Total IOPS for OLTPDB2 database node scale. Total IOPS for 1, 2, 3, 4, 5, 6, 7 and 8 nodes are around 28k, 53k, 74k, 96k, 116k, 131k, 140k and 141k with system utilization under 19% all the time.

**Figure 156**

Figure 157 Node Scalability IOPS for OLTPDB2



OLTPDB1 + OLTPDB2 Database Performance


For OLTPDB1 + OLTPDB2 database workload featuring Order Entry schema, we used two OLTPDB1 and OLTPDB2 databases. For the OLTPDB1 database (6 TB), we used 64GB size of System Global Area (SGA) and for the OLTPDB2 database (6 TB), we used 64GB size of System Global Area (SGA). We also ensured that HugePages were in use. The OLTPDB1 + OLTPDB2 Database scalability test was run for at least 12 hours and ensured that results are consistent for the duration of the full run.

We ran the SwingBench scripts on each node to start OLTP1DB1 and OLTPDB2 database and generate AWR reports for each scenario as shown below.

- User Scalability

The following graph illustrates the TPM for OLTPDB1 + OLTPDB2 database user scale on 8 node. TPM for 100, 200, 300 and 400 user are around 842k, 1340k, 1419k and 1511k with latency under 5 milliseconds all the time.

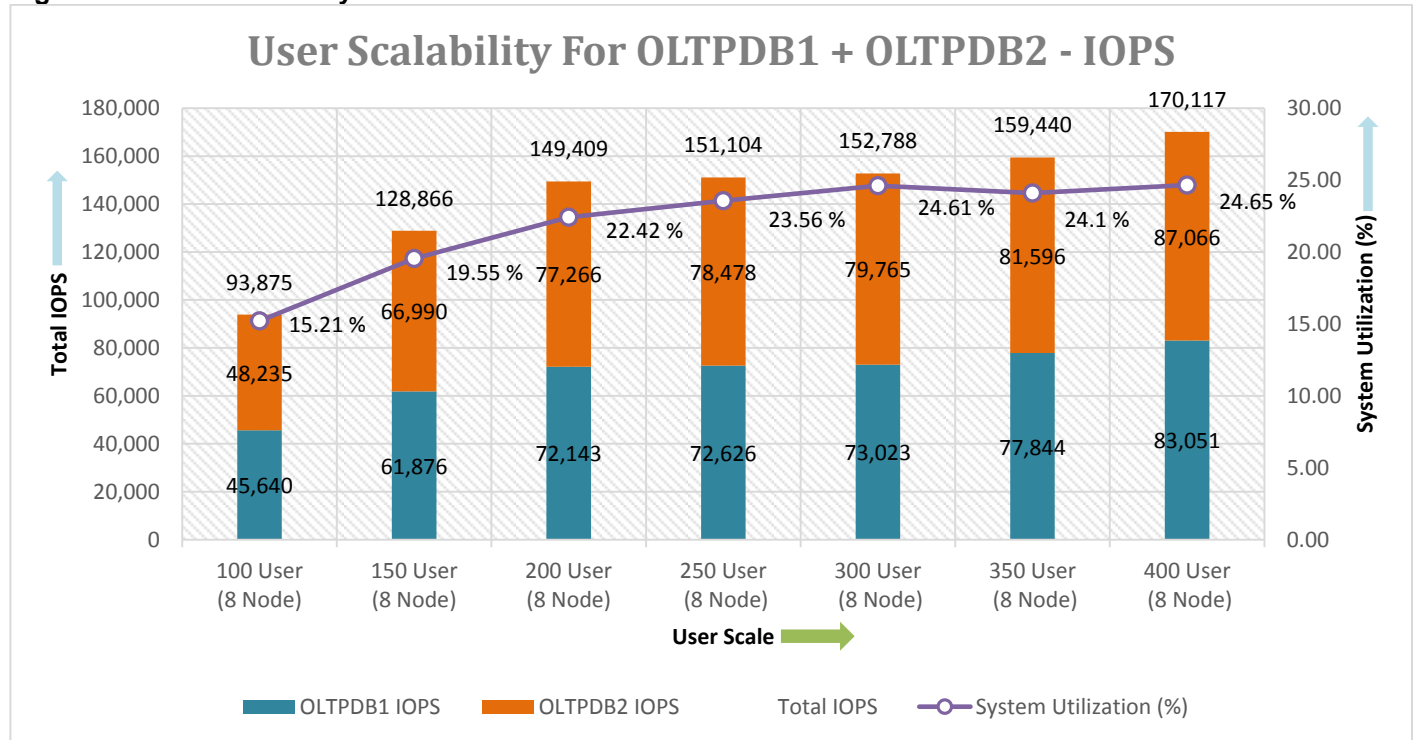
Figure 158 User Scalability TPM for OLTPDB1 + OLTPDB2



The following graph illustrates the Total IOPS for OLTPDB1 + OLTPDB2 database user scale on 8 nodes. Total IOPS for 100, 200, 300 and 400 user are around 93k, 149k, 152k and 170k with system utilization under 25% all the time.



Figure 159 User Scalability IOPS for OLTPDB1 + OLTPDB2




### DSSDB1 Database Performance

DSSDB1 database workloads are generally sequential in nature, read intensive and exercise large IO size. DSSDB1 database workload runs a small number of users that typically exercise extremely complex queries that run for hours. DSSDB1 Database activity is captured for eight Oracle RAC Instances using Oracle Enterprise Manager for 24 hours workload test.

For this test, we ran Swingbench Sales history workload with 8 users. The charts below show DSSDB1 database workload results and average Throughput around 6.4 GB/s

**Figure 160 DSSDB1 Database Performance**

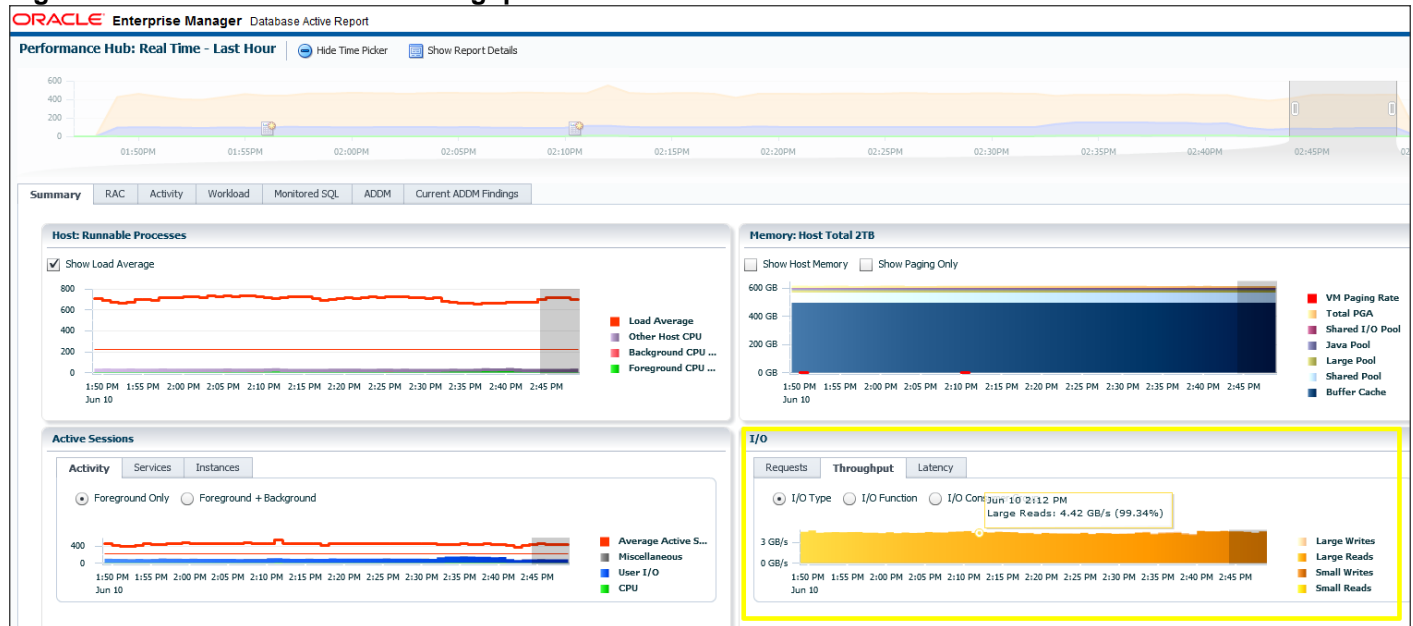


The chart above illustrates the DSSDB1 workload Throughput is around 6.4 GB/s.

**Mixed Workload**

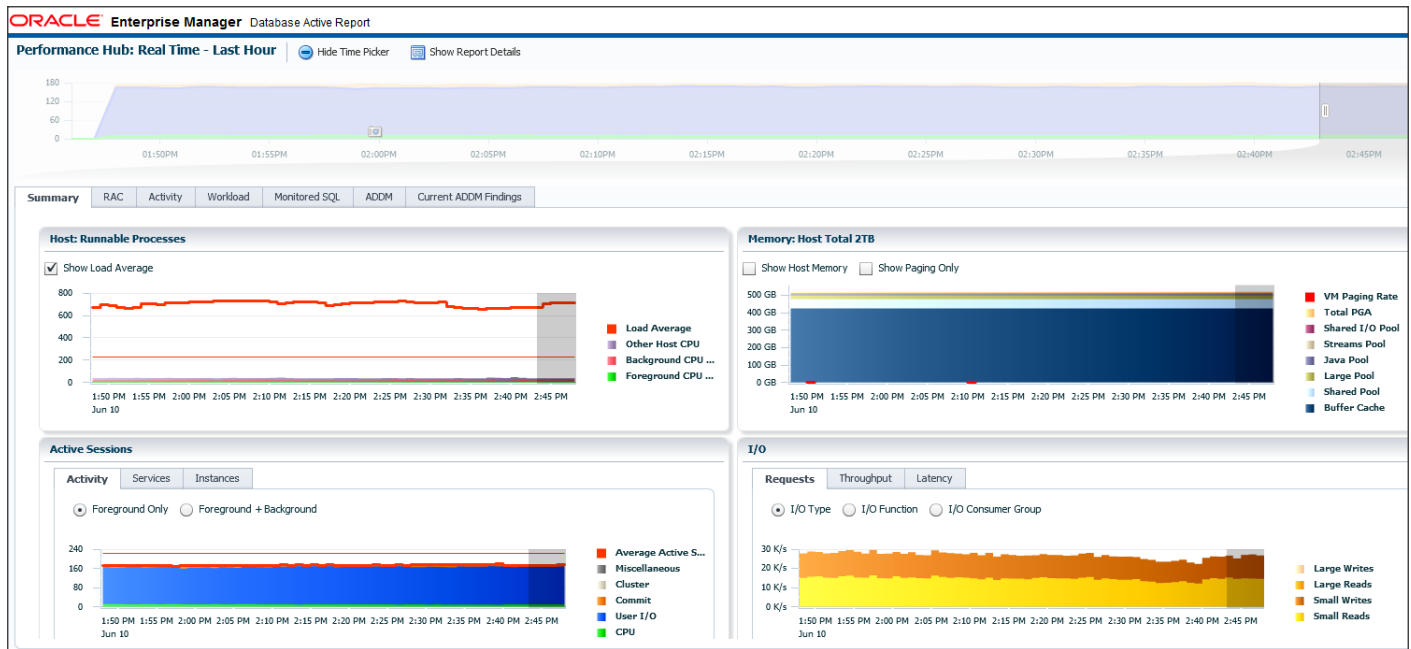
The next test is to run both OLTP and DSS database workloads simultaneously. This test will ensure that configuration in this test is able to sustain small random queries presented via OLTP database along with large and sequential transactions submitted via DSS database workload. DSSDB1, OLTPDB1 and OLTPDB2 Database activity is captured for eight Oracle RAC Instances using Oracle Enterprise Manager for 24 hours mixed workload test. We ran the test for OLTPDB1 + OLTPDB2 and DSSDB1 database together as shown in the chart below.

**Figure 161 DSSDB1 Database Throughput Performance**



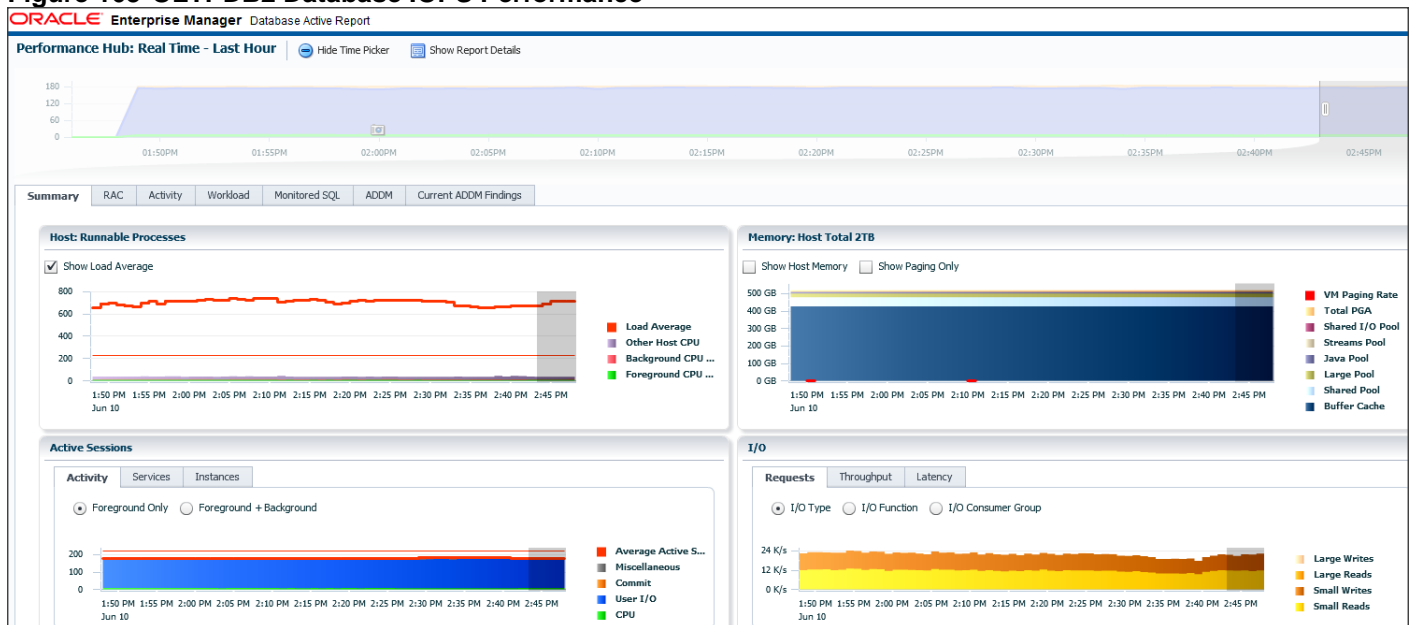
The charts above show DSSDB1 database average Throughput around 4.4 GB/s for database workload OLTPDB1, OLTPDB2 and DSSDB1 running together for 8 nodes.

**Figure 162 OLTPDB1 Database IOPS Performance**



The charts above show OLTPDB1 database average IOPS around 30k for database workload OLTPDB1, OLTPDB2 and DSSDB1 running together for 8 nodes.

**Figure 163 OLTPDB2 Database IOPS Performance**



The charts above show OLTPDB2 database average IOPS around 28k for database workload OLTPDB1, OLTPDB2 and DSSDB1 running together for 8 nodes.

## Resiliency and Failure Tests

The goal of these tests is to ensure that reference architecture withstands commonly occurring failures either due to unexpected crashes, hardware failures or human errors. We conduct many hardware, 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 VIC 1240 adapter. We have highlighted some of those test cases below.

We have highlighted some of the test cases below.

**Table 13 Hardware Failover Tests**

Scenario	Test	Status
<b>Test 1 – UCS Fabric Interconnect – A Failure</b>	Run the system on Full work Load. Power Off Fabric Interconnect – A and check network traffic on Fabric Interconnect – B.	Fabric Interconnect Failover did not cause any disruption to Private, Public and Storage Traffic
<b>Test 2 – UCS Fabric Interconnect – B Failure</b>	Run the system on Full work Load. Power Off Fabric Interconnect – B and check network traffic on Fabric Interconnect – A	Fabric Interconnect Failover did not cause any disruption to Private, Public and Storage Traffic
<b>Test 3 – UCS Nexus Switch – A Failure</b>	Run the system on Full work Load. Power Off Nexus Switch – A and check network traffic on Nexus Switch – B.	Nexus Switch Failover did not cause any disruption to Private and Public network Traffic
<b>Test 4 – UCS Nexus Switch – B Failure</b>	Run the system on Full work Load. Power Off Nexus Switch – B and check network traffic on Nexus Switch – A.	Nexus Switch Failover did not cause any disruption to Private and Public network Traffic
<b>Test 5 – UCS MDS Switch – A Failure</b>	Run the system on Full work Load. Power Off MDS Switch – A and check storage traffic on MDS Switch – B	MDS Switch Failover did not cause any disruption to Storage network Traffic
<b>Test 6 – UCS MDS Switch – B Failure</b>	Run the system on Full work Load. Power Off MDS Switch – B and check storage traffic on MDS Switch – A	MDS Switch Failover did not cause any disruption to Storage network Traffic
<b>Test 7 – UCS Chassis 1 and 2 IOM Link Failure</b>	Run the system on full work load. Disconnect two links from IOM 1 and IOM2 by pulling it out and reconnect it after 5 minutes.	No disruption in network traffic.

**Figure 164 FlashStack Infrastructure**

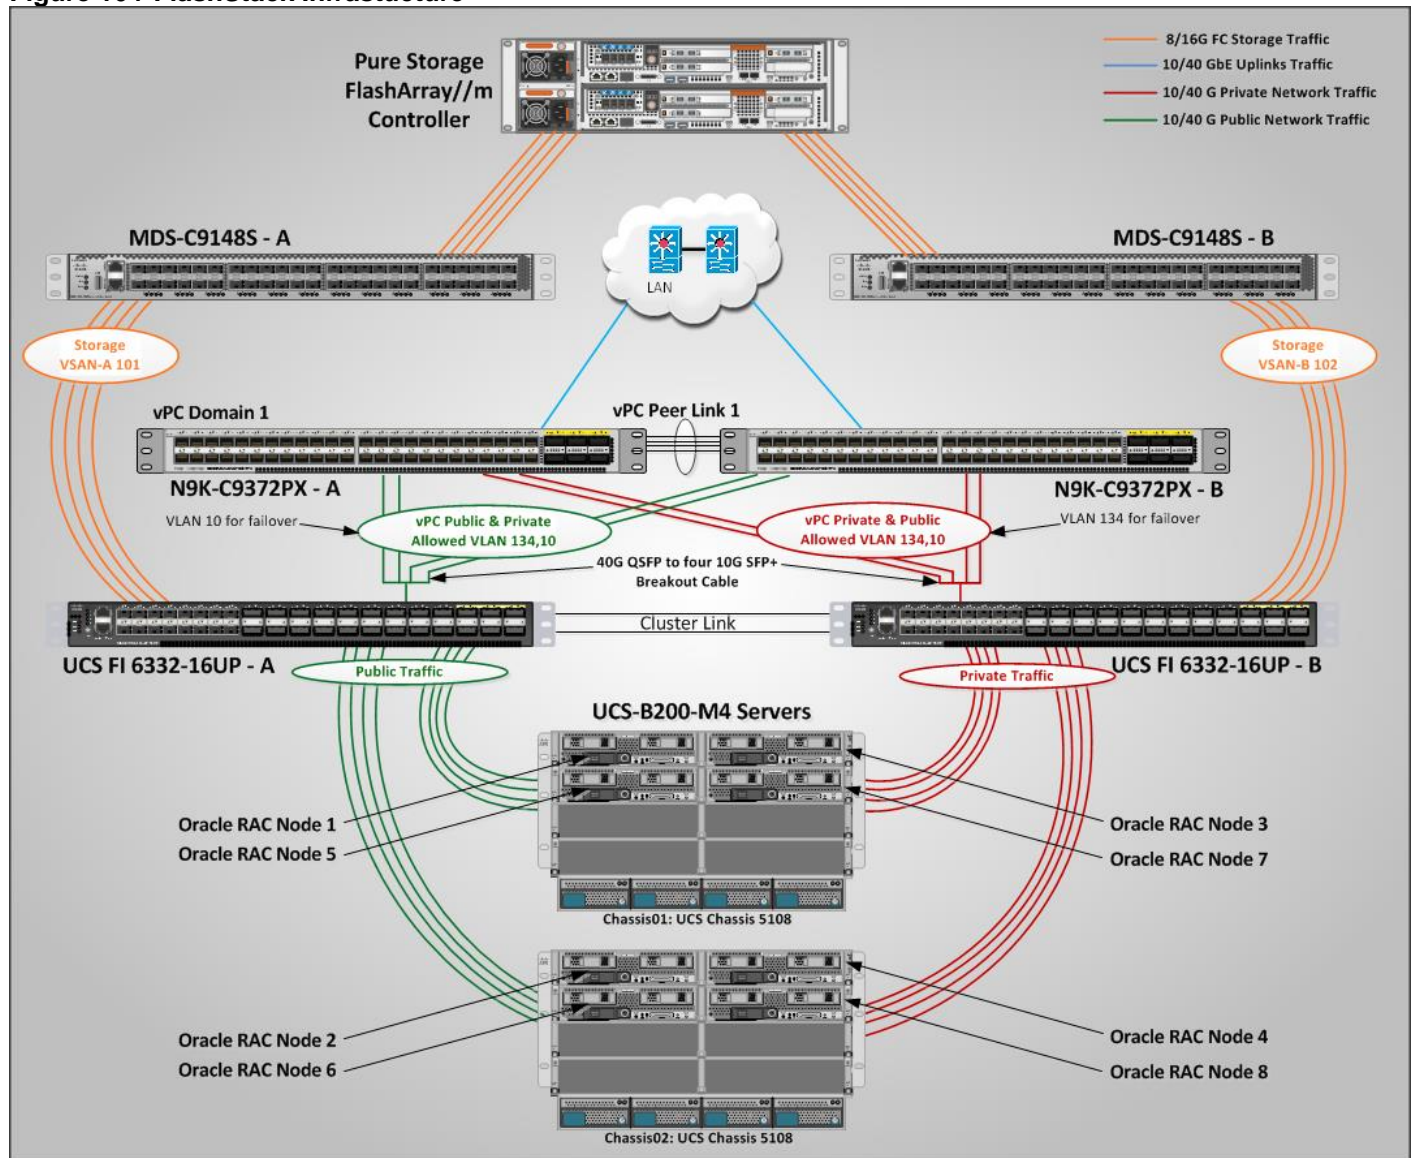


Figure 164 illustrates the FlashStack solution infrastructure diagram under normal conditions. Cisco UCS 6332-16UP Fabric Interconnect carries both storage and network traffic from the blades with the help of Cisco Nexus 9372PX-E and Cisco MDS 9148S switches. Two virtual Port-Channels (vPCs) are configured to provide public network and private network paths for the blades to northbound switches. Eight (four per chassis) links go to Fabric Interconnect – A. Similarly, eight links go to Fabric Interconnect – B. Fabric Interconnect – A links are used for Oracle Public network traffic shown as green lines. Fabric Interconnect – B links are used for Oracle private interconnect traffic shown as red lines. FC Storage access from Fabric Interconnect – A and Fabric Interconnect – B show as orange line.

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

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

**Figure 165 Fabric Interconnect – A Network Traffic**

```
PURESTG-FI-A(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 134      0025.b500.a000   static    0         F      F      Veth1225
* 134      0025.b500.a001   static    0         F      F      Veth1183
* 134      0025.b500.a002   static    0         F      F      Veth1189
* 134      0025.b500.a003   static    0         F      F      Veth1195
* 134      0025.b500.a004   static    0         F      F      Veth1201
* 134      0025.b500.a005   static    0         F      F      Veth1207
* 134      0025.b500.a006   static    0         F      F      Veth1213
* 134      0025.b500.a007   static    0         F      F      Veth1219
```

Fabric Interconnect – A carry Oracle Public Network traffic on VLAN 134 as shown above before failover test.

2. Log into 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 table below.

**Figure 166 Fabric Interconnect – B Network Traffic**

```
PURESTG-FI-B(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 10       0025.b500.b000   static    0         F      F      Veth1226
* 10       0025.b500.b001   static    0         F      F      Veth1184
* 10       0025.b500.b002   static    0         F      F      Veth1190
* 10       0025.b500.b003   static    0         F      F      Veth1196
* 10       0025.b500.b004   static    0         F      F      Veth1202
* 10       0025.b500.b005   static    0         F      F      Veth1208
* 10       0025.b500.b006   static    0         F      F      Veth1214
* 10       0025.b500.b007   static    0         F      F      Veth1220
```

Fabric Interconnect – B carry Oracle Private Network traffic on VLAN 10 as shown above before the failover test.

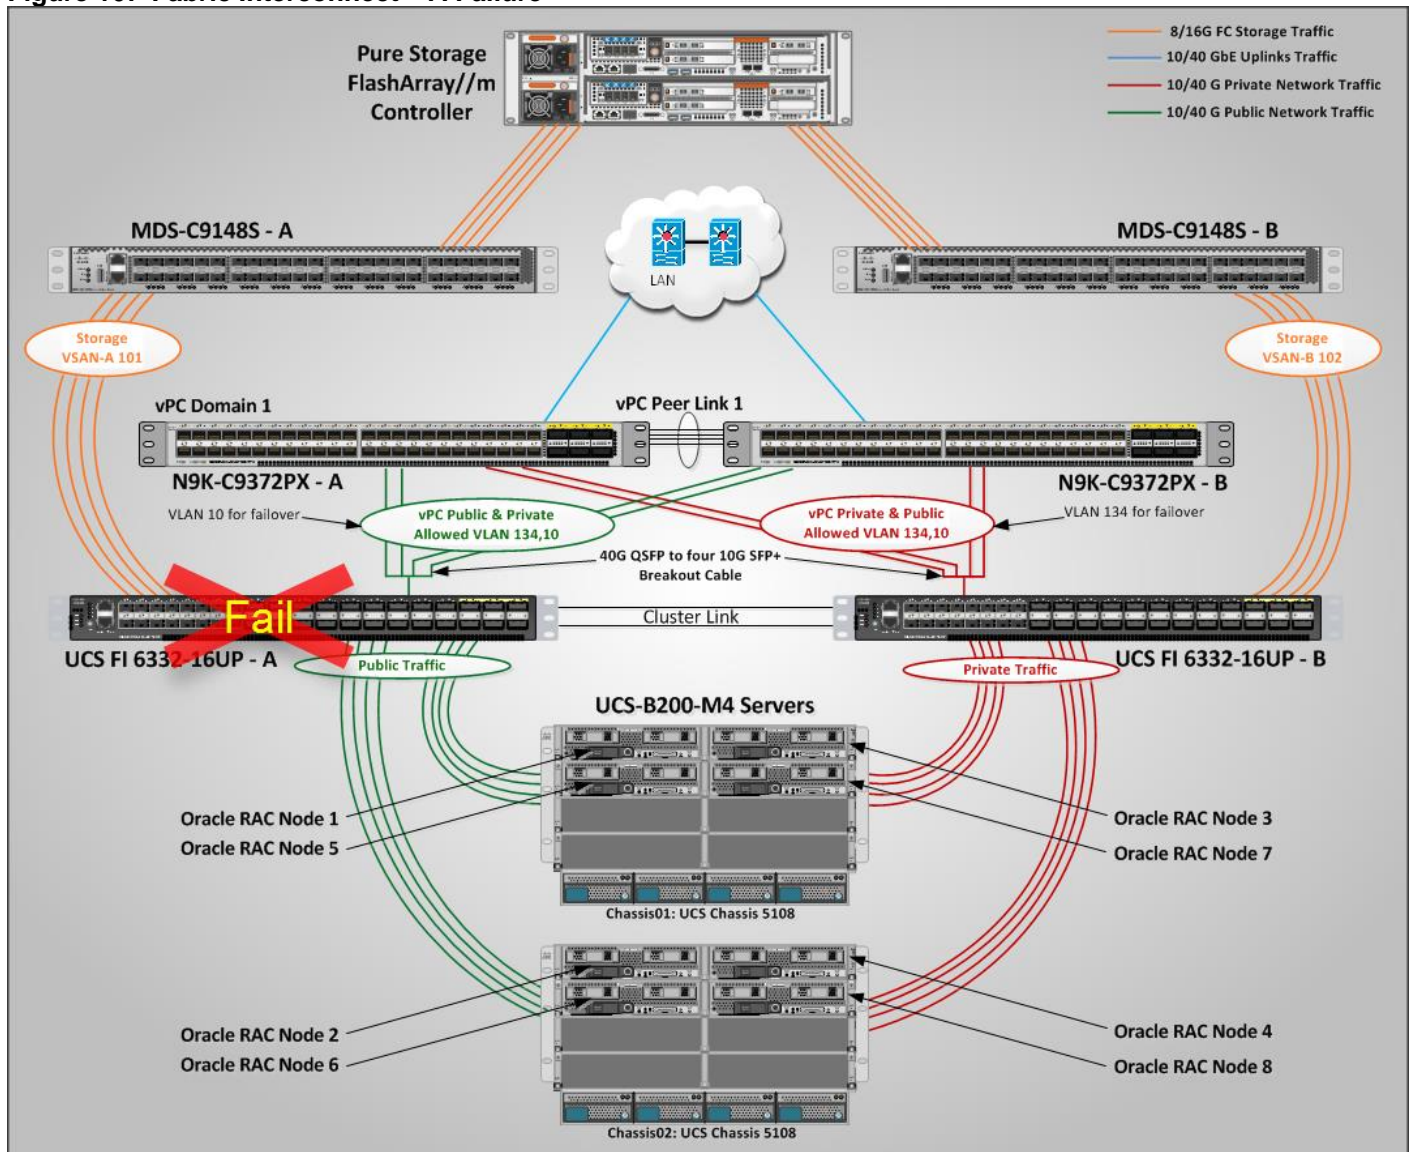
#### Test 1 – UCS 6332-16UP Fabric Interconnect – A Failure Test

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

The below figure illustrates how during Fabric Interconnect – A switch failure, the respective blades (ORARAC1, ORARAC3, ORARAC5 and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 and ORARAC8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to fabric interconnect – B.



Figure 167 Fabric Interconnect – A Failure



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

**Figure 168 Fabric Interconnect – B Network Traffic**

```
PURESTG-FI-B(nxos)#
PURESTG-FI-B(nxos)# show mac address-table | grep static
* 134      0025.b500.a000      static      0           F      F      Veth1286
* 134      0025.b500.a001      static      0           F      F      Veth1284
* 134      0025.b500.a002      static      0           F      F      Veth1282
* 134      0025.b500.a003      static      0           F      F      Veth1280
* 134      0025.b500.a004      static      0           F      F      Veth1278
* 134      0025.b500.a005      static      0           F      F      Veth1276
* 134      0025.b500.a006      static      0           F      F      Veth1274
* 134      0025.b500.a007      static      0           F      F      Veth1272
* 10       0025.b500.b000      static      0           F      F      Veth1301
* 10       0025.b500.b001      static      0           F      F      Veth1299
* 10       0025.b500.b002      static      0           F      F      Veth1297
* 10       0025.b500.b003      static      0           F      F      Veth1295
* 10       0025.b500.b004      static      0           F      F      Veth1293
* 10       0025.b500.b005      static      0           F      F      Veth1291
* 10       0025.b500.b006      static      0           F      F      Veth1289
* 10       0025.b500.b007      static      0           F      F      Veth1287
* 10       0025.b5dd.0001     static      0           F      F      Veth1270
```

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

Plug the back power cable to Fabric Interconnect – A Switch, the respective blades (ORARAC1, ORARAC3, ORARAC5 & ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 & ORARAC8) on chassis 2 will route back the MAC addresses and its VLAN traffic to Fabric Interconnect – A.

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

**Figure 169 Fabric Interconnect – A Network Traffic**

```
PURESTG-FI-A(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 134      0025.b500.a000      static      0           F      F      Veth1225
* 134      0025.b500.a001      static      0           F      F      Veth1183
* 134      0025.b500.a002      static      0           F      F      Veth1189
* 134      0025.b500.a003      static      0           F      F      Veth1195
* 134      0025.b500.a004      static      0           F      F      Veth1201
* 134      0025.b500.a005      static      0           F      F      Veth1207
* 134      0025.b500.a006      static      0           F      F      Veth1213
* 134      0025.b500.a007      static      0           F      F      Veth1219
```

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



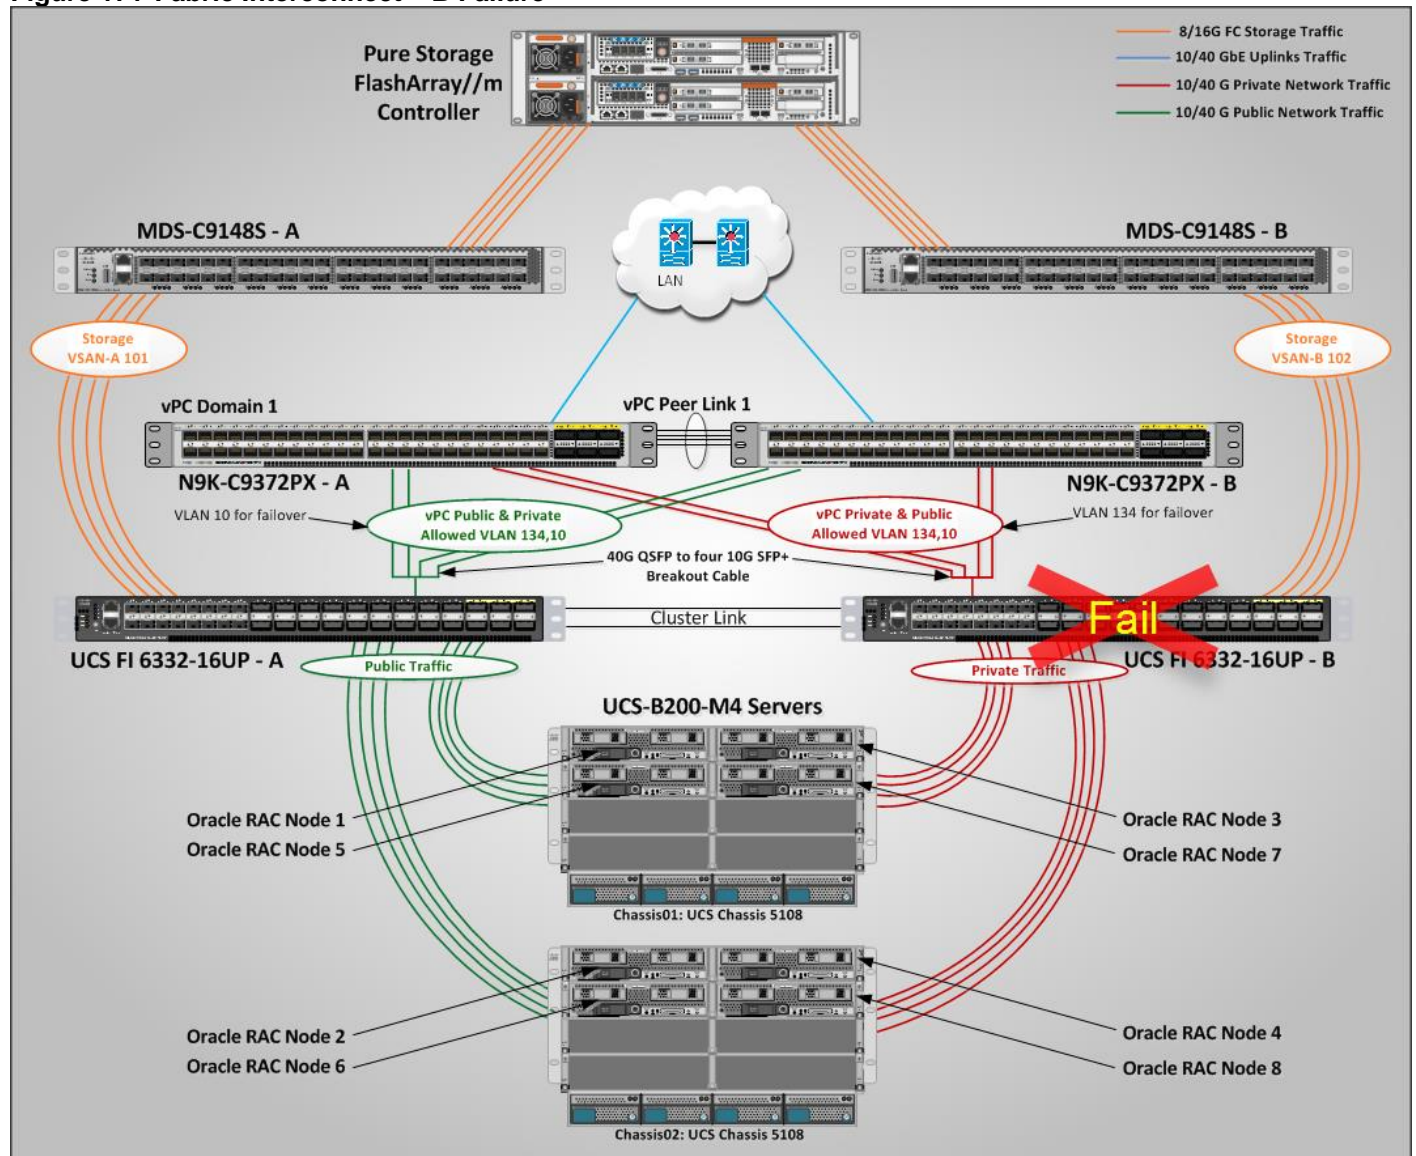
**Figure 170 Fabric Interconnect – B Network Traffic**

```
PURESTG-FI-B(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 10      0025.b500.b000   static    0        F      F      Veth1226
* 10      0025.b500.b001   static    0        F      F      Veth1184
* 10      0025.b500.b002   static    0        F      F      Veth1190
* 10      0025.b500.b003   static    0        F      F      Veth1196
* 10      0025.b500.b004   static    0        F      F      Veth1202
* 10      0025.b500.b005   static    0        F      F      Veth1208
* 10      0025.b500.b006   static    0        F      F      Veth1214
* 10      0025.b500.b007   static    0        F      F      Veth1220
```

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

The following figure illustrates how during Fabric Interconnect – B switch failure, the respective blades (ORARAC1, ORARAC3, ORARAC5 and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 and ORARAC8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to fabric interconnect – A.

Figure 171 Fabric Interconnect – B Failure



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

**Figure 172 Fabric Interconnect – A Network Traffic**

```

PURESTG-FI-A(nxos)#
PURESTG-FI-A(nxos)# show mac address-table static
Legend:
  * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
  age - seconds since first seen,+ - primary entry using vPC Peer-Link
  VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 134      0025.b500.a000      static    0         F      F      Veth1285
* 134      0025.b500.a001      static    0         F      F      Veth1283
* 134      0025.b500.a002      static    0         F      F      Veth1281
* 134      0025.b500.a003      static    0         F      F      Veth1279
* 134      0025.b500.a004      static    0         F      F      Veth1277
* 134      0025.b500.a005      static    0         F      F      Veth1275
* 134      0025.b500.a006      static    0         F      F      Veth1273
* 134      0025.b500.a007      static    0         F      F      Veth1271
* 134      0025.b5cc.0001      static    0         F      F      Veth1269
* 10       0025.b500.b000      static    0         F      F      Veth1302
* 10       0025.b500.b001      static    0         F      F      Veth1300
* 10       0025.b500.b002      static    0         F      F      Veth1298
* 10       0025.b500.b003      static    0         F      F      Veth1296
* 10       0025.b500.b004      static    0         F      F      Veth1294
* 10       0025.b500.b005      static    0         F      F      Veth1292
* 10       0025.b500.b006      static    0         F      F      Veth1290
* 10       0025.b500.b007      static    0         F      F      Veth1288

```

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

Plug the back power cable to Fabric Interconnect – B Switch, the respective blades (ORARAC1, ORARAC3, ORARAC5 & ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 & ORARAC8) on chassis 2 will route back the MAC addresses and its VLAN traffic to Fabric Interconnect – B.

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

**Figure 173 Fabric Interconnect – A Network Traffic**

```

PURESTG-FI-A(nxos)# show mac address-table
Legend:
  * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
  age - seconds since first seen,+ - primary entry using vPC Peer-Link
  VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----
* 134      0025.b500.a000      static    0         F      F      Veth1225
* 134      0025.b500.a001      static    0         F      F      Veth1183
* 134      0025.b500.a002      static    0         F      F      Veth1189
* 134      0025.b500.a003      static    0         F      F      Veth1195
* 134      0025.b500.a004      static    0         F      F      Veth1201
* 134      0025.b500.a005      static    0         F      F      Veth1207
* 134      0025.b500.a006      static    0         F      F      Veth1213
* 134      0025.b500.a007      static    0         F      F      Veth1219

```

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

**Figure 174 Fabric Interconnect – B Network Traffic**

```
PURESTG-FI-B(nxos)# show mac address-table
```

```
Legend:
```

```
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
```

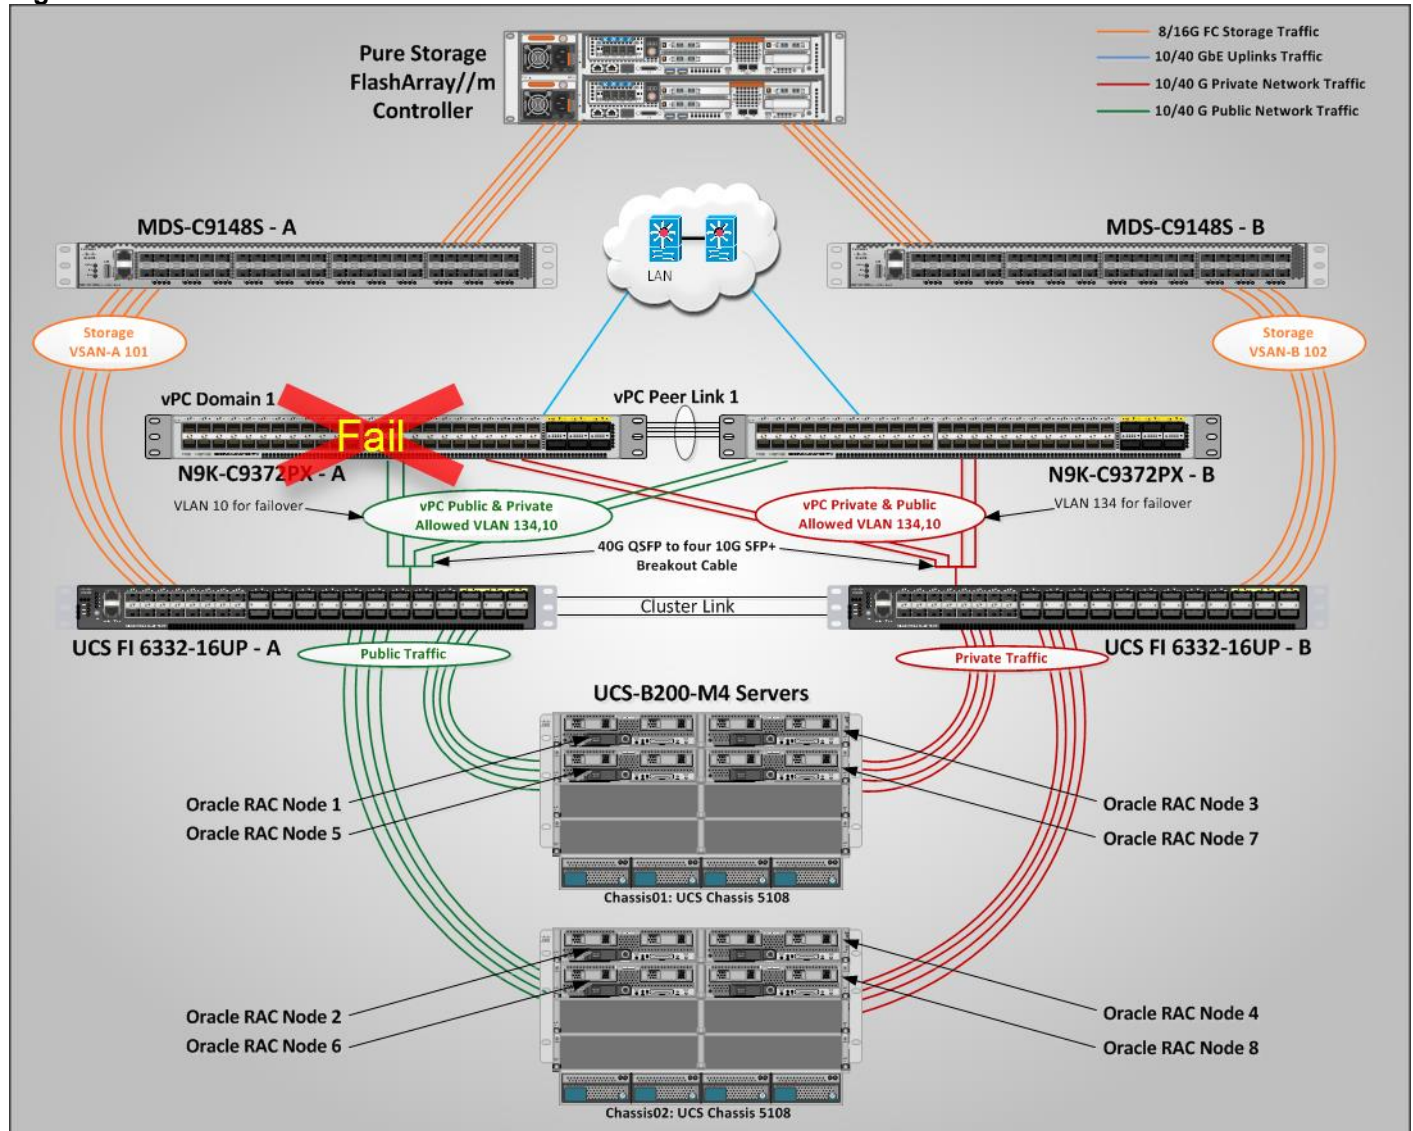
VLAN	MAC Address	Type	age	Secure	NTFY	Ports/SWID.SSID.LID
* 10	0025.b500.b000	static	0	F	F	Veth1226
* 10	0025.b500.b001	static	0	F	F	Veth1184
* 10	0025.b500.b002	static	0	F	F	Veth1190
* 10	0025.b500.b003	static	0	F	F	Veth1196
* 10	0025.b500.b004	static	0	F	F	Veth1202
* 10	0025.b500.b005	static	0	F	F	Veth1208
* 10	0025.b500.b006	static	0	F	F	Veth1214
* 10	0025.b500.b007	static	0	F	F	Veth1220

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

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

The below figure illustrates how during Nexus Switch – A failure, the respective blades (ORARAC1, ORARAC3, ORARAC5 and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 and ORARAC8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to Nexus Switch – B.

**Figure 175 Cisco Nexus Switch – A Failure**



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

Plug the back power cable to Nexus Switch – A Switch, the respective blades (ORARAC1, ORARAC3, ORARAC5 & ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 & ORARAC8) on chassis 2 will route back the MAC addresses and its VLAN traffic to Nexus Switch – A.

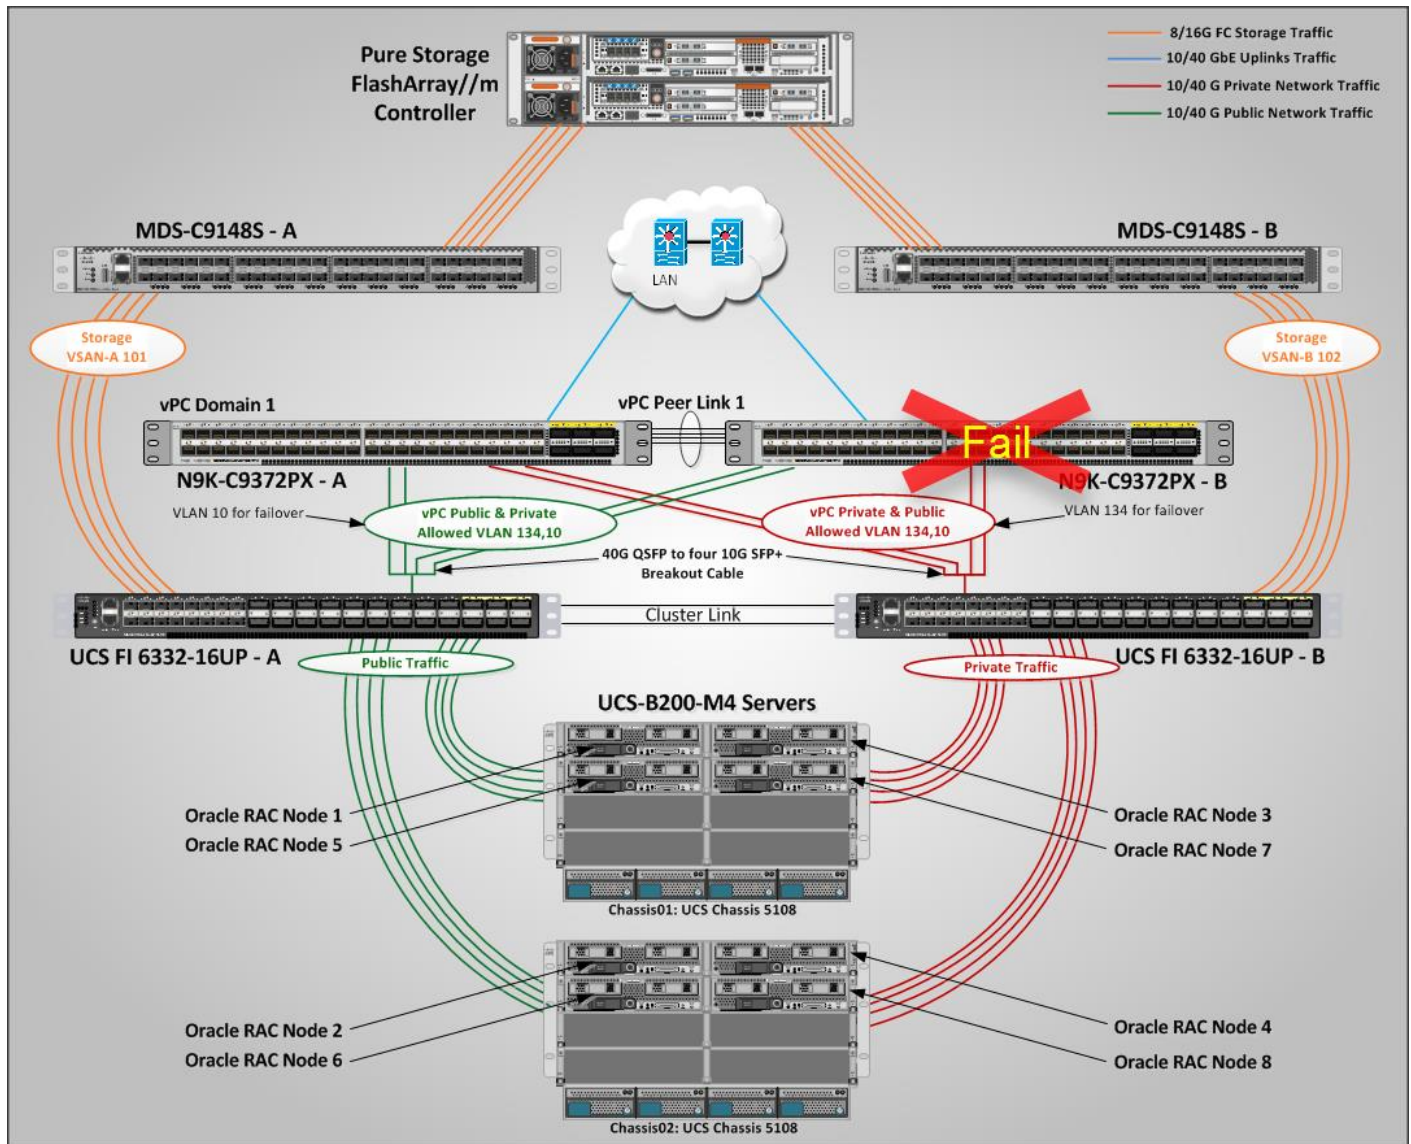
**Test 4 – Nexus 9372PX-E Switch – B Failure Test**

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

The following figure illustrates how during Nexus Switch – B failure, the respective blades (ORARAC1, ORARAC3, ORARAC5 and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 and ORARAC8) on chassis 2 will fail over the MAC addresses and its VLAN network traffic to Nexus Switch – A.

**Figure 176 Cisco Nexus Switch – B Failure**





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

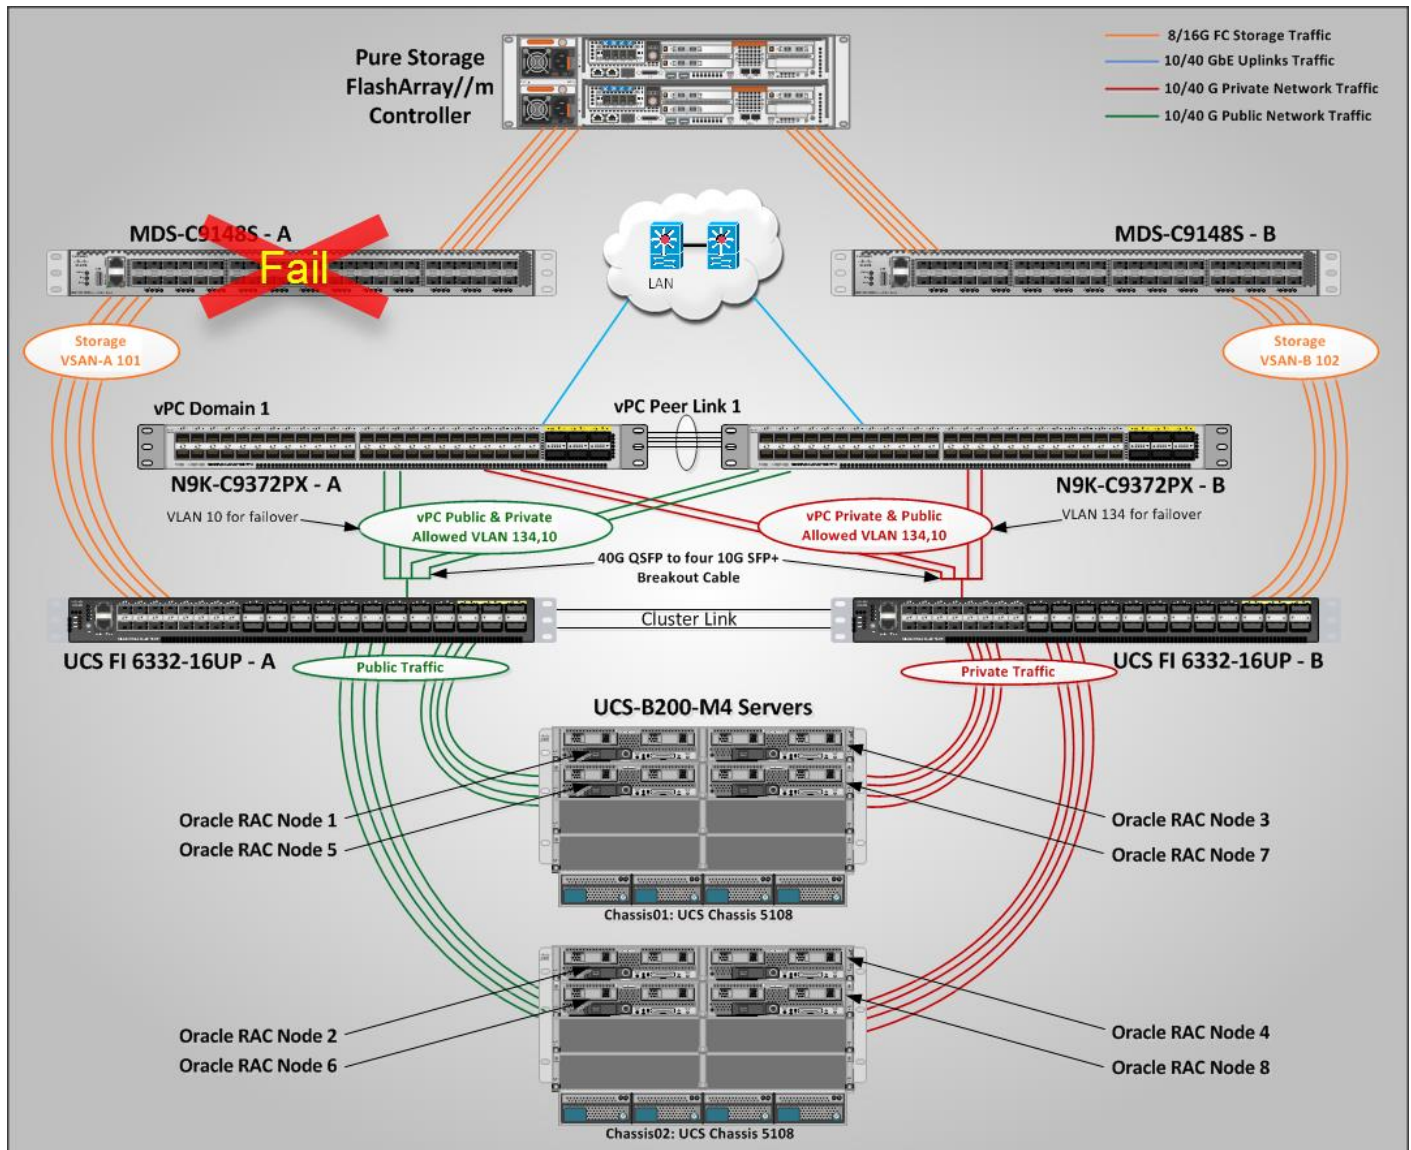
Plug the back power cable to Nexus Switch – B Switch, the respective blades (ORARAC1, ORARAC3, ORARAC5 & ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 & ORARAC8) on chassis 2 will route back the MAC addresses and its VLAN traffic to Nexus Switch – B.

### Test 5 – MDS 9148S Switch – A Failure Test

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

The below figure illustrates how during MDS Switch – A failure, the respective blades (ORARAC1, ORARAC3, ORARAC5, and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6, and ORARAC8) on chassis 2 will failover the MAC addresses and its storage traffic to MDS Switch B same way as Fabric Switch failure.

**Figure 177 MDS Switch A Failure**

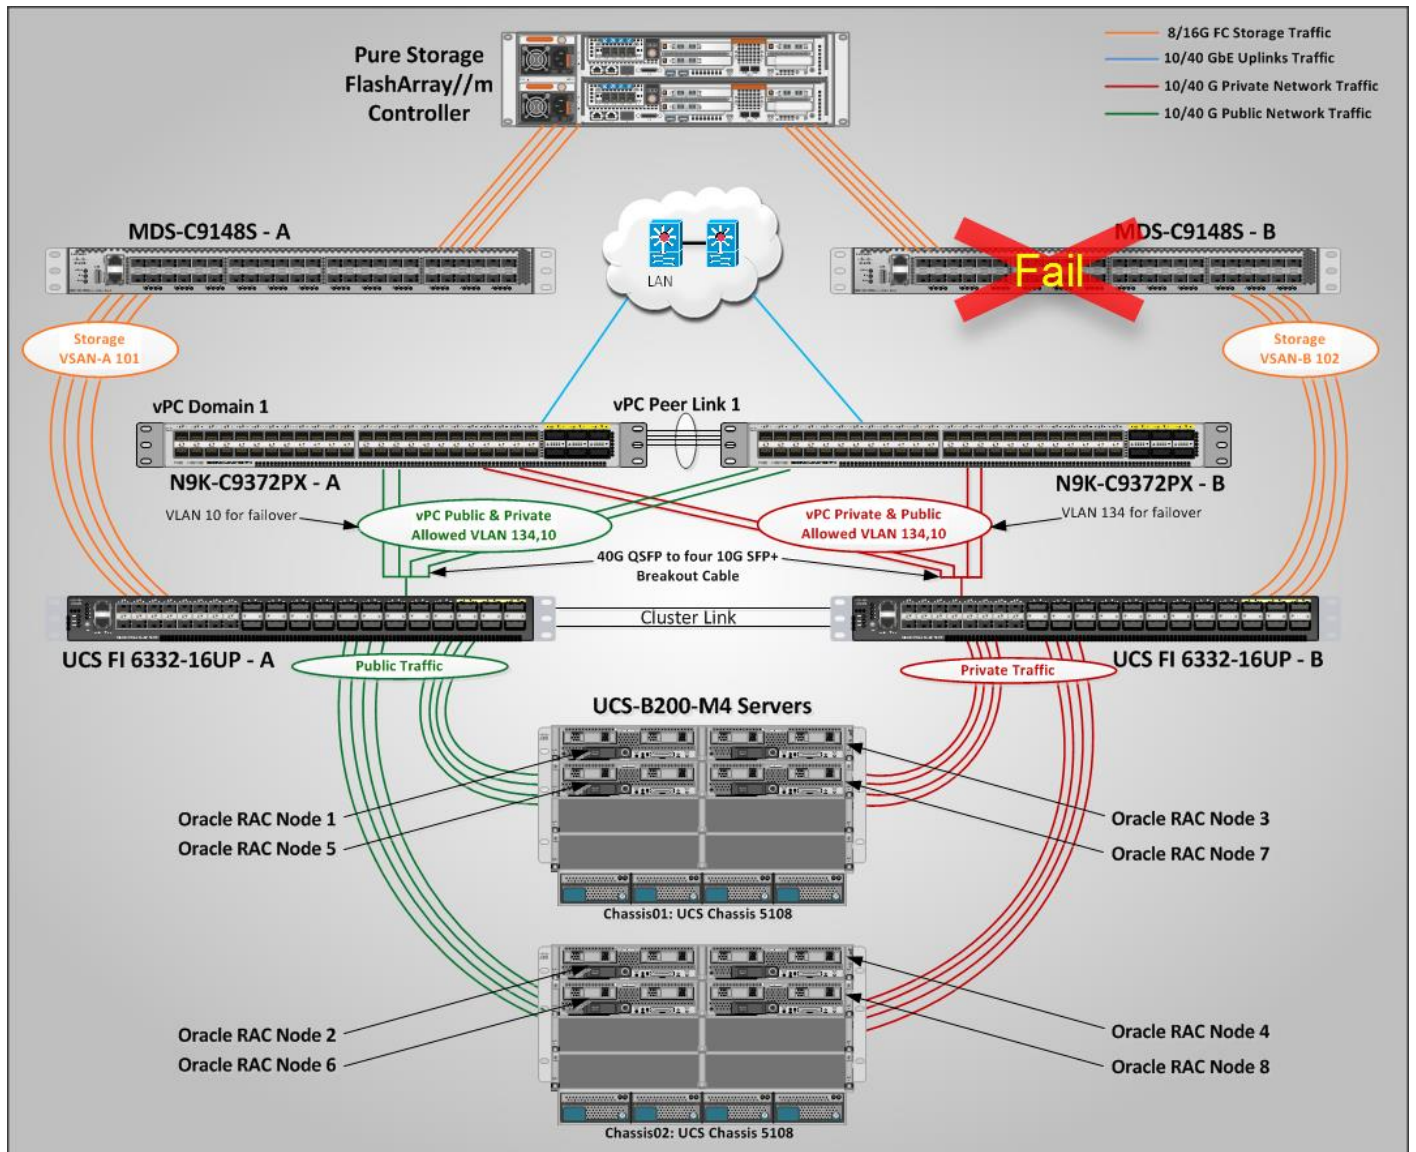


Test 6 – MDS 9148S Switch – B Failure Test

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

The below figure illustrates how during MDS Switch – B failure, the respective blades (ORARAC1, ORARAC3, ORARAC5, and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6, and ORARAC8) on chassis 2 will failover the MAC addresses and its storage traffic to MDS Switch B same way as Fabric Switch failure.

Figure 178 MDS Switch B Failure



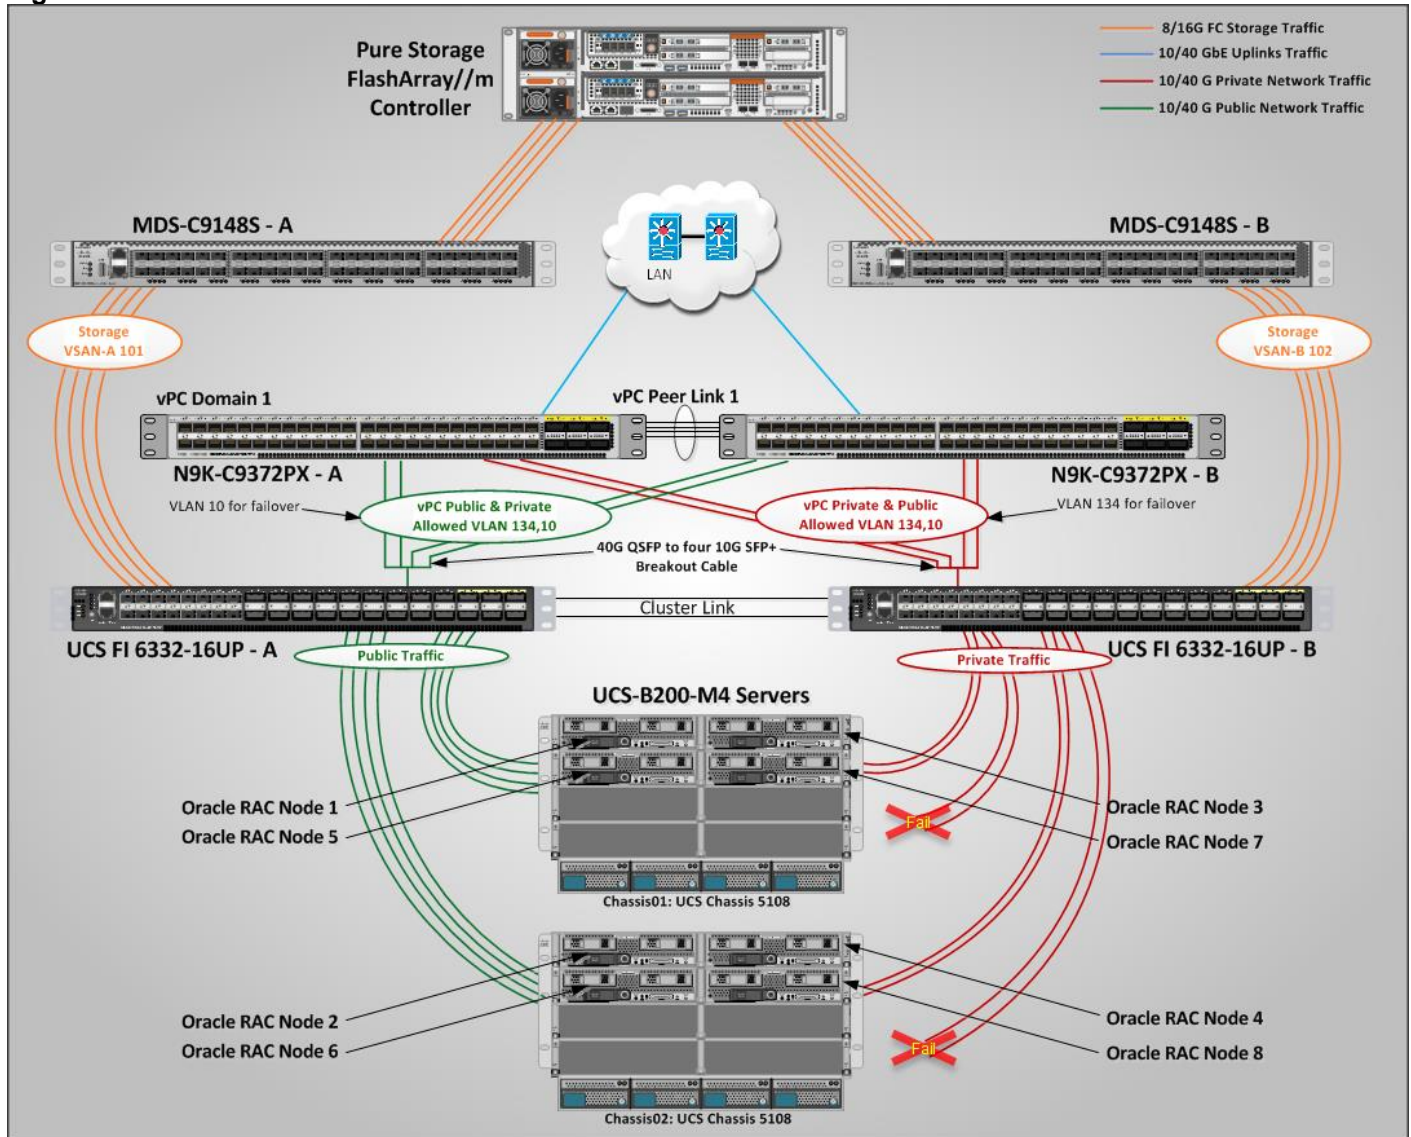
### Test 7 – UCS Chassis 1 and 2 IOM Link Failure

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

The following figure illustrates IOM Link failure for Chassis 1 and Chassis 2.



Figure 179 Chassis 1 and 2 IOM Link Failure



Now, unplug the two server port cables from each Chassis 1 and Chassis 2 and check the MAC address and VLAN traffic information on both Cisco UCS Fabric Interconnects.

**Figure 180 Fabric Interconnect – B Network Traffic**

```
PURESTG-FI-B(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
C
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN MAC Address Type age Secure NTFY Ports/SWID.SSID
.LID
-----+-----+-----+-----+-----+-----+-----
-
* 10 0025.b500.b000 static 0 F F Veth1323
* 10 0025.b500.b001 static 0 F F Veth1315
* 10 0025.b500.b002 static 0 F F Veth1339
* 10 0025.b500.b003 static 0 F F Veth1355
* 10 0025.b500.b004 static 0 F F Veth1331
* 10 0025.b500.b005 static 0 F F Veth1347
* 10 0025.b500.b006 static 0 F F Veth1363
* 10 0025.b500.b007 static 0 F F Veth1371
* 10 0025.b5dd.0001 static 0 F F Veth1378
```

As shown in the above figure, when two links from Chassis 1 and two links from Chassis 2 Failed, it will have no effect on Private Network traffic links of VLAN 10 to Fabric Interconnect – B.

**Figure 181 Fabric Interconnect – A Network Traffic**

```
PURESTG-FI-A(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
C
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN MAC Address Type age Secure NTFY Ports/SWID.SSID
.LID
-----+-----+-----+-----+-----+-----+-----
-
* 134 0025.b500.a000 static 0 F F Veth1321
* 134 0025.b500.a001 static 0 F F Veth1313
* 134 0025.b500.a002 static 0 F F Veth1337
* 134 0025.b500.a003 static 0 F F Veth1353
* 134 0025.b500.a004 static 0 F F Veth1329
* 134 0025.b500.a005 static 0 F F Veth1345
* 134 0025.b500.a006 static 0 F F Veth1361
* 134 0025.b500.a007 static 0 F F Veth1369
* 134 0025.b5cc.0001 static 0 F F Veth1377
```

As shown in the above figure, network traffic on Fabric Interconnect – A has no effect of IOM link failure. So Chassis 1 and 2 IOM Link Failure did not cause any disruption to Private Network Traffic.

Plug the back two server port cables from Chassis 1 and Chassis 2, the respective blades (ORARAC1, ORARAC3, ORARAC5 and ORARAC7) on chassis 1 and (ORARAC2, ORARAC4, ORARAC6 and ORARAC8) on chassis 2 will have same MAC addresses and its VLAN traffic to Fabric Interconnect – A and Fabric Interconnect – B.

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

**Figure 182 Fabric Interconnect – A Network Traffic**

```
PURESTG-FI-A(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----+-----
* 134      0025.b500.a000   static    0         F      F      Veth1225
* 134      0025.b500.a001   static    0         F      F      Veth1183
* 134      0025.b500.a002   static    0         F      F      Veth1189
* 134      0025.b500.a003   static    0         F      F      Veth1195
* 134      0025.b500.a004   static    0         F      F      Veth1201
* 134      0025.b500.a005   static    0         F      F      Veth1207
* 134      0025.b500.a006   static    0         F      F      Veth1213
* 134      0025.b500.a007   static    0         F      F      Veth1219
```

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

**Figure 183 Fabric Interconnect – B Network Traffic**

```
PURESTG-FI-B(nxos)# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since first seen,+ - primary entry using vPC Peer-Link
VLAN      MAC Address      Type      age      Secure NTFY      Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----+-----
* 10       0025.b500.b000   static    0         F      F      Veth1226
* 10       0025.b500.b001   static    0         F      F      Veth1184
* 10       0025.b500.b002   static    0         F      F      Veth1190
* 10       0025.b500.b003   static    0         F      F      Veth1196
* 10       0025.b500.b004   static    0         F      F      Veth1202
* 10       0025.b500.b005   static    0         F      F      Veth1208
* 10       0025.b500.b006   static    0         F      F      Veth1214
* 10       0025.b500.b007   static    0         F      F      Veth1220
```

From the test shown above, we can conclude that in case of partial links failure from any IOM would not affect any disruption to the Network Traffic.

## Summary

---

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

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

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

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

FlashStack differs from other solutions by providing:

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

## Appendix

---

### Cisco Nexus 9372PX-E Running Configuration

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

```
!Command: show running-config
```

```
!Time: Fri Jun 17 21:17:44 2016
```

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

cfs eth distribute
feature lacp
feature vpc

username admin password 5 $1$pBQrHGEG$Yt6qBcYWDQlWQt8qBujqp. role network-admin
no password strength-check
ip domain-lookup
system qos
  service-policy type network-qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0xa9aed036608b5fe64ec2cf3f9a27f3ab priv
0xa9ae
d036608b5fe64ec2cf3f9a27f3ab localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

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

spanning-tree port type edge bpduguard default
spanning-tree port type network default
vrf context management
  ip route 0.0.0.0/0 10.29.134.1
port-channel load-balance src-dst l4port
vpc domain 1
```

## Appendix

```
role priority 10
peer-keepalive destination 10.29.134.154 source 10.29.134.153
auto-recovery

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

interface port-channel19
description connect to Fabric Interconnect A
switchport mode trunk
switchport trunk allowed vlan 1,10,134
spanning-tree port type edge trunk
vpc 19

interface port-channel20
description connect to Fabric Interconnect B
switchport mode trunk
switchport trunk allowed vlan 1,10,134
spanning-tree port type edge trunk
vpc 20

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

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

interface Ethernet1/3
interface Ethernet1/4
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
description connect to uplink switch
switchport access vlan 134
speed 1000

interface Ethernet1/16
interface Ethernet1/17
description Fabric-Interconnect-A:1/19
switchport mode trunk
```

## Appendix

```
switchport trunk allowed vlan 1,10,134
channel-group 19 mode active

interface Ethernet1/18
description Fabric-Interconnect-A:1/19
switchport mode trunk
switchport trunk allowed vlan 1,10,134
channel-group 19 mode active

interface Ethernet1/19
description Fabric-Interconnect-B:1/19
switchport mode trunk
switchport trunk allowed vlan 1,10,134
channel-group 20 mode active

interface Ethernet1/20
description Fabric-Interconnect-B:1/19
switchport mode trunk
switchport trunk allowed vlan 1,10,134
channel-group 20 mode active

interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
interface Ethernet1/47
interface Ethernet1/48
interface Ethernet2/1
interface Ethernet2/2
interface Ethernet2/3
interface Ethernet2/4
interface Ethernet2/5
interface Ethernet2/6
interface Ethernet2/7
interface Ethernet2/8
interface Ethernet2/9
```

```

interface Ethernet2/10
interface Ethernet2/11
interface Ethernet2/12

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

```

## MDS 9148S FC Switch Running Configuration

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

```
!Command: show running-config
!Time: Wed Jun 15 23:44:51 2016
```

```

version 6.2(9)
power redundancy-mode redundant
feature npiv
feature telnet
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 $1$eHBzCc/e$N0yxbSDo5.z3ktgoI1m2k0 role network-admin
no password strength-check
ip domain-lookup
ip host PURESTG-MDS-A 10.29.134.155
aaa group server radius radius
snmp-server user admin network-admin auth md5 0x5bb23f0db04581372c1cc0e36b65a303 priv
0x5bb2
3f0db04581372c1cc0e36b65a303 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vsan database
  vsan 101
device-alias database
  device-alias name ORARAC-A1-hba0 pwwn 20:00:00:25:b5:a0:00:00
  device-alias name ORARAC-A1-hba2 pwwn 20:00:00:25:b5:a0:00:01
  device-alias name ORARAC-A2-hba0 pwwn 20:00:00:25:b5:a0:00:02
  device-alias name ORARAC-A2-hba2 pwwn 20:00:00:25:b5:a0:00:03
  device-alias name ORARAC-A3-hba0 pwwn 20:00:00:25:b5:a0:00:04
  device-alias name ORARAC-A3-hba2 pwwn 20:00:00:25:b5:a0:00:05
  device-alias name ORARAC-A4-hba0 pwwn 20:00:00:25:b5:a0:00:06
  device-alias name ORARAC-A4-hba2 pwwn 20:00:00:25:b5:a0:00:07
  device-alias name ORARAC-A5-hba0 pwwn 20:00:00:25:b5:a0:00:08
  device-alias name ORARAC-A5-hba2 pwwn 20:00:00:25:b5:a0:00:09
  device-alias name ORARAC-A6-hba0 pwwn 20:00:00:25:b5:a0:00:0a
  device-alias name ORARAC-A6-hba2 pwwn 20:00:00:25:b5:a0:00:0b

```



## Appendix

```
device-alias name ORARAC-A7-hba0 pwwn 20:00:00:25:b5:a0:00:0c
device-alias name ORARAC-A7-hba2 pwwn 20:00:00:25:b5:a0:00:0d
device-alias name ORARAC-A8-hba0 pwwn 20:00:00:25:b5:a0:00:0e
device-alias name ORARAC-A8-hba2 pwwn 20:00:00:25:b5:a0:00:0f
device-alias name Pure-STG-CT0-FC0 pwwn 52:4a:93:7a:b3:18:ce:00
device-alias name Pure-STG-CT0-FC2 pwwn 52:4a:93:7a:b3:18:ce:02
device-alias name Pure-STG-CT1-FC0 pwwn 52:4a:93:7a:b3:18:ce:10
device-alias name Pure-STG-CT1-FC2 pwwn 52:4a:93:7a:b3:18:ce:12
```

```
device-alias commit
```

```
fcdomain fcid database
```

```
vsan 1 wwn 52:4a:93:7a:b3:18:ce:02 fcid 0x3a0000 dynamic
!
[Pure-STG-CT0-FC2]
vsan 1 wwn 52:4a:93:7a:b3:18:ce:12 fcid 0x3a0100 dynamic
!
[Pure-STG-CT1-FC2]
vsan 1 wwn 20:01:8c:60:4f:bd:31:80 fcid 0x3a0200 dynamic
vsan 1 wwn 20:02:8c:60:4f:bd:31:80 fcid 0x3a0300 dynamic
vsan 101 wwn 20:01:8c:60:4f:bd:31:80 fcid 0x5d0000 dynamic
vsan 101 wwn 20:02:8c:60:4f:bd:31:80 fcid 0x5d0100 dynamic
vsan 101 wwn 20:00:00:25:b5:a0:00:06 fcid 0x5d0702 dynamic
!
[ORARAC-A4-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:00 fcid 0x5d0001 dynamic
!
[ORARAC-A1-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:04 fcid 0x5d0602 dynamic
!
[ORARAC-A3-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:02 fcid 0x5d0104 dynamic
!
[ORARAC-A2-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:0a fcid 0x5d0106 dynamic
!
[ORARAC-A6-hba0]
vsan 101 wwn 52:4a:93:7a:b3:18:ce:02 fcid 0x5d0200 dynamic
!
[Pure-STG-CT0-FC2]
vsan 101 wwn 52:4a:93:7a:b3:18:ce:12 fcid 0x5d0300 dynamic
!
[Pure-STG-CT1-FC2]
vsan 101 wwn 52:4a:93:7a:b3:18:ce:00 fcid 0x5d0400 dynamic
!
[Pure-STG-CT0-FC0]
vsan 101 wwn 52:4a:93:7a:b3:18:ce:10 fcid 0x5d0500 dynamic
!
[Pure-STG-CT1-FC0]
vsan 101 wwn 20:00:00:25:b5:a0:00:01 fcid 0x5d0105 dynamic
!
[ORARAC-A1-hba2]
vsan 101 wwn 20:00:00:25:b5:a0:00:03 fcid 0x5d0603 dynamic
!
[ORARAC-A2-hba2]
vsan 101 wwn 20:00:00:25:b5:a0:00:08 fcid 0x5d0006 dynamic
!
[ORARAC-A5-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:0e fcid 0x5d0701 dynamic
!
[ORARAC-A8-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:0c fcid 0x5d0601 dynamic
!
[ORARAC-A7-hba0]
vsan 101 wwn 20:00:00:25:b5:a0:00:07 fcid 0x5d0004 dynamic
!
[ORARAC-A4-hba2]
vsan 101 wwn 20:00:00:25:b5:a0:00:05 fcid 0x5d0703 dynamic
!
[ORARAC-A3-hba2]
vsan 101 wwn 20:03:8c:60:4f:bd:31:80 fcid 0x5d0600 dynamic
vsan 101 wwn 20:04:8c:60:4f:bd:31:80 fcid 0x5d0700 dynamic
vsan 101 wwn 20:00:00:25:b5:a0:00:09 fcid 0x5d0002 dynamic
!
[ORARAC-A5-hba2]
vsan 101 wwn 20:00:00:25:b5:a0:00:0d fcid 0x5d0604 dynamic
!
```

## Appendix

```
vsan 101 wwn 20:00:00:25:b5:a0:00:0b fcid 0x5d0101 dynamic
!      [ORARAC-A6-hba2]
vsan 101 wwn 20:00:00:25:b5:a0:00:0f fcid 0x5d0704 dynamic
!      [ORARAC-A8-hba2]
vsan 101 wwn 52:4a:93:7a:b3:18:ce:03 fcid 0x5d0800 dynamic

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 fc1/9
vsan 101 interface fc1/10
vsan 101 interface fc1/11
vsan 101 interface fc1/12
switchname PURESTG-MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9.bin
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
```

## Appendix

```
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
!Active Zone Database Section for vsan 101
zone name chas1-server1-boot-hba vsan 101
  member pwwn 52:4a:93:7a:b3:18:ce:00
!      [Pure-STG-CT0-FC0]
  member pwwn 20:00:00:25:b5:a0:00:00
!      [ORARAC-A1-hba0]
  member pwwn 52:4a:93:7a:b3:18:ce:10
!      [Pure-STG-CT1-FC0]
  member pwwn 20:00:00:25:b5:a0:00:01
!      [ORARAC-A1-hba2]
  member pwwn 52:4a:93:7a:b3:18:ce:02
!      [Pure-STG-CT0-FC2]
  member pwwn 52:4a:93:7a:b3:18:ce:12
!      [Pure-STG-CT1-FC2]

zone name chas1-server2-boot-hba vsan 101
  member pwwn 52:4a:93:7a:b3:18:ce:00
!      [Pure-STG-CT0-FC0]
  member pwwn 20:00:00:25:b5:a0:00:02
!      [ORARAC-A2-hba0]
  member pwwn 52:4a:93:7a:b3:18:ce:10
!      [Pure-STG-CT1-FC0]
  member pwwn 20:00:00:25:b5:a0:00:03
!      [ORARAC-A2-hba2]
  member pwwn 52:4a:93:7a:b3:18:ce:02
!      [Pure-STG-CT0-FC2]
  member pwwn 52:4a:93:7a:b3:18:ce:12
!      [Pure-STG-CT1-FC2]

zone name chas1-server3-boot-hba vsan 101
  member pwwn 52:4a:93:7a:b3:18:ce:00
!      [Pure-STG-CT0-FC0]
  member pwwn 20:00:00:25:b5:a0:00:04
!      [ORARAC-A3-hba0]
  member pwwn 52:4a:93:7a:b3:18:ce:10
!      [Pure-STG-CT1-FC0]
  member pwwn 20:00:00:25:b5:a0:00:05
!      [ORARAC-A3-hba2]
  member pwwn 52:4a:93:7a:b3:18:ce:02
!      [Pure-STG-CT0-FC2]
  member pwwn 52:4a:93:7a:b3:18:ce:12
!      [Pure-STG-CT1-FC2]

zone name chas1-server4-boot-hba vsan 101
  member pwwn 52:4a:93:7a:b3:18:ce:00
!      [Pure-STG-CT0-FC0]
  member pwwn 20:00:00:25:b5:a0:00:06
!      [ORARAC-A4-hba0]
```

## Appendix

```
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:07
!           [ORARAC-A4-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server1-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:08
!           [ORARAC-A5-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:09
!           [ORARAC-A5-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server2-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:0a
!           [ORARAC-A6-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:0b
!           [ORARAC-A6-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server3-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:0c
!           [ORARAC-A7-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:0d
!           [ORARAC-A7-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server4-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:0e
!           [ORARAC-A8-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
```

## Appendix

```
    member pwnn 20:00:00:25:b5:a0:00:0f
!           [ORARAC-A8-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zoneset name Oracle-RAC-A vsan 101
    member chas1-server1-boot-hba
    member chas1-server2-boot-hba
    member chas1-server3-boot-hba
    member chas1-server4-boot-hba
    member chas2-server1-boot-hba
    member chas2-server2-boot-hba
    member chas2-server3-boot-hba
    member chas2-server4-boot-hba

zoneset activate name Oracle-RAC-A vsan 101
do clear zone database vsan 101
!Full Zone Database Section for vsan 101
zone name chas1-server1-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:00
!           [ORARAC-A1-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:01
!           [ORARAC-A1-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas1-server2-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:02
!           [ORARAC-A2-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:03
!           [ORARAC-A2-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
    member pwnn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas1-server3-boot-hba vsan 101
    member pwnn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
    member pwnn 20:00:00:25:b5:a0:00:04
!           [ORARAC-A3-hba0]
    member pwnn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
    member pwnn 20:00:00:25:b5:a0:00:05
!           [ORARAC-A3-hba2]
    member pwnn 52:4a:93:7a:b3:18:ce:02
```

## Appendix

```
!           [Pure-STG-CT0-FC2]
  member pwn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas1-server4-boot-hba vsan 101
  member pwn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
  member pwn 20:00:00:25:b5:a0:00:06
!           [ORARAC-A4-hba0]
  member pwn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
  member pwn 20:00:00:25:b5:a0:00:07
!           [ORARAC-A4-hba2]
  member pwn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
  member pwn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server1-boot-hba vsan 101
  member pwn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
  member pwn 20:00:00:25:b5:a0:00:08
!           [ORARAC-A5-hba0]
  member pwn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
  member pwn 20:00:00:25:b5:a0:00:09
!           [ORARAC-A5-hba2]
  member pwn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
  member pwn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server2-boot-hba vsan 101
  member pwn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
  member pwn 20:00:00:25:b5:a0:00:0a
!           [ORARAC-A6-hba0]
  member pwn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
  member pwn 20:00:00:25:b5:a0:00:0b
!           [ORARAC-A6-hba2]
  member pwn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
  member pwn 52:4a:93:7a:b3:18:ce:12
!           [Pure-STG-CT1-FC2]

zone name chas2-server3-boot-hba vsan 101
  member pwn 52:4a:93:7a:b3:18:ce:00
!           [Pure-STG-CT0-FC0]
  member pwn 20:00:00:25:b5:a0:00:0c
!           [ORARAC-A7-hba0]
  member pwn 52:4a:93:7a:b3:18:ce:10
!           [Pure-STG-CT1-FC0]
  member pwn 20:00:00:25:b5:a0:00:0d
!           [ORARAC-A7-hba2]
  member pwn 52:4a:93:7a:b3:18:ce:02
!           [Pure-STG-CT0-FC2]
  member pwn 52:4a:93:7a:b3:18:ce:12
```

## Appendix

```
!                [Pure-STG-CT1-FC2]

zone name chas2-server4-boot-hba vsan 101
  member pwwn 52:4a:93:7a:b3:18:ce:00
!                [Pure-STG-CT0-FC0]
  member pwwn 20:00:00:25:b5:a0:00:0e
!                [ORARAC-A8-hba0]
  member pwwn 52:4a:93:7a:b3:18:ce:10
!                [Pure-STG-CT1-FC0]
  member pwwn 20:00:00:25:b5:a0:00:0f
!                [ORARAC-A8-hba2]
  member pwwn 52:4a:93:7a:b3:18:ce:02
!                [Pure-STG-CT0-FC2]
  member pwwn 52:4a:93:7a:b3:18:ce:12
!                [Pure-STG-CT1-FC2]

zoneset name Oracle-RAC-A vsan 101
  member chas1-server1-boot-hba
  member chas1-server2-boot-hba
  member chas1-server3-boot-hba
  member chas1-server4-boot-hba
  member chas2-server1-boot-hba
  member chas2-server2-boot-hba
  member chas2-server3-boot-hba
  member chas2-server4-boot-hba

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
  port-license acquire

interface fc1/6
  port-license acquire

interface fc1/7
  port-license acquire
```

```
interface fc1/8
  port-license acquire

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
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
```



```

interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface mgmt0
  ip address 10.29.134.155 255.255.255.0
no system default switchport shutdown
ip default-gateway 10.29.134.1

```

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

```

[oracle@oraracl etc]$ cat multipath.conf
# This is a basic configuration file with some examples, for device mapper
# multipath.
blacklist {
    devnode "^(ram|zram|raw|loop|fd|md|sr|scd|st) [0-9]*"
}
defaults {
    find_multipaths yes
    polling_interval 1
}
devices {
    device {
        vendor                "PURE"
        path_grouping_policy  multibus
        path_checker          tur
        path_selector         "queue-length 0"
        fast_io_fail_tmo     10
        dev_loss_tmo         60
        no_path_retry        0
    }
}
multipaths {
    multipath {
        wwid                  3624a9370ea38e05d8d3348d300011018
        alias                 osboot
    }
    multipath {
        wwid                  3624a9370ea38e05d8d3348d30001102c
        alias                 dg_orarac_crs
    }
    multipath {
        wwid                  3624a9370ea38e05d8d3348d300011025
        alias                 dg_oradata_oltp1
    }
    multipath {
        wwid                  3624a9370ea38e05d8d3348d300011026
        alias                 dg_oraredo_oltp1
    }
    multipath {
        wwid                  3624a9370ea38e05d8d3348d30001102d
        alias                 dg_oradata_oltp2
    }
    multipath {
        wwid                  3624a9370ea38e05d8d3348d30001102e
        alias                 dg_oraredo_oltp2
    }
}

```

```

    }
    multipath {
        wwid                3624a9370ea38e05d8d3348d30001102f
        alias                dg_oradata_dss1
    }
    multipath {
        wwid                3624a9370ea38e05d8d3348d300011030
        alias                dg_oraredo_dss1
    }
    multipath {
        wwid                3624a9370ea38e05d8d3348d300011031
        alias                dg_oradata_dss2
    }
    multipath {
        wwid                3624a9370ea38e05d8d3348d300011032
        alias                dg_oraredo_dss2
    }
}

```

### Configuration of "/etc/sysctl.conf"

```

### File located "/etc/sysctl.conf" directory
[grid@oraracl etc]$ cat sysctl.conf
# System default settings live in /usr/lib/sysctl.d/00-system.conf.
# To override those settings, enter new settings here, or in an
/etc/sysctl.d/<name>.conf file
# For more information, see sysctl.conf(5) and sysctl.d(5).

# prerequisites, Add the following lines to the "/etc/sysctl.conf" file.

fs.file-max = 6815744
kernel.sem = 250 32000 100 128
kernel.shmmni = 4096
kernel.shmall = 1073741824
kernel.shmmax = 4398046511104
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
fs.aio-max-nr = 1048576
net.ipv4.ip_local_port_range = 9000 65500

# Recommended value for kernel.panic_on_oops
kernel.panic_on_oops = 1

vm.nr_hugepages=102000

```

### Configuration of "/etc/security/limits.conf"

```

### File located "/etc/security/limits.conf" directory

[grid@oraracl security]$ cat limits.conf
# Prerequisites, Add the following lines to the "/etc/security/limits.conf"
oracle soft nofile 1024
oracle hard nofile 65536

```

## Appendix

```
oracle  soft  nproc  2047
oracle  hard  nproc  16384
oracle  soft  stack  10240
oracle  hard  stack  32768
```

```
grid    soft  nofile  1024
grid    hard  nofile  65536
grid    soft  nproc  2047
grid    hard  nproc  16384
grid    soft  stack  10240
grid    hard  stack  32768
```

```
*      soft  memlock 235929600
*      hard  memlock 235929600
```

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

```
### File located "/etc/udev/rules.d" directory
```

```
[grid@oraracl rules.d]$ cat 99-oracle-asmdevices.rules
#All volumes which starts with dg_orarac_* #
ENV{DM_NAME}=="dg_orarac_*", OWNER:="grid", GROUP:="oinstall", MODE:="660"

#All volumes which starts with dg_oradata_* #
ENV{DM_NAME}=="dg_oradata_*", OWNER:="oracle", GROUP:="oinstall", MODE:="660"

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

### Configuration of "/etc/udev/rules.d/99-pure-storage.rules"

```
### File located "/etc/udev/rules.d" directory
```

```
[grid@oraracl rules.d]$ cat 99-pure-storage.rules
# Recommended settings for PURE Storage FlashArray
# Use noop scheduler for high-performance SSD
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/scheduler}="noop"

# Reduce CPU overhead due to entropy collection
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/add_random}="0"

# Schedule I/O on the core that initiated the process
ACTION=="add|change", SUBSYSTEM=="block", ENV{ID_VENDOR}=="PURE",
ATTR{queue/rq_affinity}="2"
```

## About the Authors

---

### **Tushar Patel, Cisco Systems, Inc.**

Tushar Patel is a Principal Engineer in Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group and a specialist in Flash Storage technologies and Oracle RAC RDBMS. Tushar has over 20 years' 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.

### **Niranjan Mohapatra, Cisco Systems, Inc.**

Niranjan Mohapatra is a Technical Marketing Engineer in Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group, and specialist on Oracle RAC RDBMS. He has over 17 years of extensive experience on Oracle RAC Database and associated tools. Niranjan has worked as a TME and a DBA handling production systems in various organizations. Niranjan has worked on various Oracle RAC solutions on different platforms like FlexPod, vBlock and Hitachi Storage. He holds a Master of Science (MSc) degree in Computer Science and is also an Oracle Certified Professional (OCP -DBA). Niranjan also has strong background in Cisco UCS, Cisco Nexus and Cisco MDS, Pure Storage and Virtualization.

### **Hardik Vyas, Cisco Systems, Inc.**

Hardik Vyas is a Solution Engineer in Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group for developing and validating infrastructure best practices for Oracle database on Cisco UCS Servers, Cisco Nexus Products and Storage Technologies. Hardik Vyas holds a Master's degree in Electrical Engineering and over 3 years of experience in Oracle Database and applications. Hardik Vyas's main focus is developing Oracle RAC Database solutions on Cisco UCS Platform.

### **Somu Rajarathinam, Pure Storage**

Somu Rajarathinam is the Oracle Solutions Architect at Pure Storage responsible for defining database solution based on the company's products, performing benchmarks, preparing reference architecture and technical papers for Oracle databases on Pure. Somu has over 20 years of Oracle database experience, including as a member of Oracle Corporation's Systems Performance and Oracle Applications Performance Groups. His career also included assignments with Logitech, Inspirage, and Autodesk, ranging from providing database and performance solutions to managing infrastructure, to delivering database and application support, both in-house and in the cloud.

## Acknowledgements

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

- Radhakrishna Manga, Sr. Director, Pure Storage