NetApp Verified Architecture

FlexPod Select for High-Performance Oracle RAC

NVA Deployment

Rodrigo Nascimento, NetApp
December 2015 | NVA-0012-DEPLOY | Version 2.0

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Solution Overview

This NVA validates that the NetApp® FlexPod® Select solution can run Oracle databases in a highly resilient architecture while remaining competitive in both price and performance. This document addresses the scenarios that customers experience in their production Oracle databases.

The solution outlined in this NVA deployment guide was architected to deliver over 2 million input/output operations per second (IOPS) at microsecond-level average latency with a workload that is 100% random reads using an 8K request size. Additionally, NetApp validated that this solution delivers consistent performance using a variety of other mixed online transaction processing (OLTP) type workloads, such as:

- 90% reads, 10% writes, 100% random I/O using an 8K request size
- 80% reads, 20% writes, 100% random I/O using an 8K request size

1.1 Solution Technology

FlexPod Select uses technologies from Cisco and NetApp that are configured according to the companies' best practices. This section discusses the products and technologies leveraged in the solution.

NetApp EF560 Flash Array

The NetApp EF560 flash array is designed for performance-driven applications with microsecond-level latency requirements.

This solution uses the NetApp EF560 flash array as the underlying storage technology, which is built on storage architecture with more than 20 years of storage development experience and more than 750,000 systems in the field. Each EF560 flash array can deliver extreme consistent performance with microsecond-level response times, enabling business-critical applications to deliver faster results and improving the end-user experience. This combination of high IOPS and ultralow latency makes the EF560 flash array an ideal choice for database-driven applications that require extreme performance.

The EF560 flash array runs on the enterprise-proven NetApp SANtricity® platform, which is optimized for flash solutions and allows storage administrators to achieve maximum performance and capacity utilization. The extensive configuration flexibility, custom performance tuning, and complete control over data placement make it an ideal choice for mission-critical applications. Its GUI-based performance tools provide key information about storage input/output (I/O) from multiple viewpoints, allowing administrators to make informed decisions about configuration adjustments to further refine performance.

The NetApp EF560 flash array delivers extreme performance, reliability, and availability to drive greater speed and responsiveness from the applications controlling your key business operations.

The NetApp EF560 flash array can:

- Increase the speed of business with microsecond-level response times.
- Eliminate overprovisioning and improve IT efficiency.
- Achieve the transactional performance of 2,000 15K RPM drives in a two–rack unit (2RU) enclosure that requires just 5% of the available rack space, power, and cooling as compared to storage systems that run on spinning disks.
- Detect and resolve issues quickly with advanced monitoring and proactive repair.
- Protect against data loss and downtime with NetApp point-in-time images, remote replication, and other advanced data protection.
- Create copies of the database by using the NetApp Snapshot® volume feature.
- Replicate data to either an EF560 flash array or an E-Series system.

- Leverage the enterprise-proven SANtricity software platform.

By combining extreme IOPS, microsecond-level response times, scale-up capacity, and enterprise-grade reliability, the NetApp EF560 flash array helps you to increase productivity and achieve faster business results.

Cisco Unified Computing System

The Cisco Unified Computing System (Cisco UCS) is a next-generation solution for blade and rack server computing. The system integrates a low-latency, lossless 10 Gigabit Ethernet (10GbE) unified network fabric with enterprise-class, x86 architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain. Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems.

Cisco Nexus 5000

Cisco Nexus 5000 series switches are designed to deliver high-density, top-of-rack layer 2 and layer 3 10GbE with unified ports in compact 1RU and 2RU form factors. The Cisco Nexus 5000 series includes Cisco Nexus 5500 and 5600 platforms as part of the Cisco Unified Fabric portfolio.

The Cisco Nexus 5500 switches simplify convergence through broad connectivity support. This makes them ideal top-of-rack access switches for traditional and converged deployments. The Cisco 5500 switches are designed to meet the scalability demands of today's data centers. Key Cisco Nexus 5500 series features include:

- Up to 1,152 ports in a single management domain that uses Cisco Fabric Extender (FEX) architecture
- Up to 96 unified ports

Oracle Database

The Oracle Database 12c Enterprise Edition provides industry-leading performance, scalability, security, and reliability on clustered or single servers with a wide range of options to meet the business needs of critical enterprise applications. Oracle Real Application Cluster (RAC) brings an innovative approach to the challenges of rapidly increasing amounts of data and demand for high performance. Oracle RAC uses a scale-out model in which active-active clusters utilize multiple servers to deliver high performance, scalability, and availability.

Oracle Automatic Storage Management (Oracle ASM) provides an integrated cluster file system and volume-management features that remove the need for third-party volume management tools and reduce the complexity of the overall architecture.

Some of the key Oracle ASM features include:

- Automatic file and volume management
- Database file system with performance of raw I/O
- Automatic distribution and striping of data
- A choice of external (array-based) data protection, two-way, and three-way mirror protection
- Control over which copy of mirrored data should be used preferentially

With these capabilities, Oracle ASM provides an alternative to the third-party file system and volume-management solutions for database storage management tasks, such as creating or laying out databases and managing the use of disk space. Oracle ASM provides load balancing of I/O across all LUNs or files in an Oracle ASM disk group by distributing the contents of each data file evenly across the entire pool of storage in the disk group.

The NetApp SANtricity plug-in for Oracle Enterprise Manager (Oracle EM) provides Oracle database administrators (DBAs) with powerful capabilities designed to increase their productivity and simplify their jobs. The plug-in is designed to access E-Series and EF-Series storage arrays used in conjunction with Oracle EM database software. This allows Oracle DBAs to monitor and report on the storage subsystems, with the ultimate goal of confirming the performance and availability of the infrastructure they use. Performance views that come with the plug-in help DBAs easily identify bottlenecks in the system. The plug-in also gives a view of the end-to-end database mapping to the storage and allows DBAs to create a database-to-storage topology report without accessing the storage layers underneath. The plug-in is free and does not require a license.

Key features of the Oracle EM plug-in include:

- Integration with Oracle Enterprise Manager 12c
- Support for NetApp E-Series and EF-Series storage arrays
- End-to-end storage volume-to-database mapping
- Integrated business intelligence publisher reports
- Automatic metrics collection on key storage array components
- Integrated database performance homepage

Oracle Linux

Oracle Linux brings the latest Linux innovations to market, delivering extreme performance, advanced scalability, and reliability for enterprise applications and systems along with worldwide, enterprise-class, low-cost support. It is free to download and distribute, including patches and updates. It is certified for compliance with the Linux Standard Base (LSB) standard. Oracle Linux is completely free to download, deploy, and distribute. Oracle Linux Support delivers enterprise-class support for Linux with Ksplice zero-downtime updates, premier backports, comprehensive management, and indemnification at significantly lower cost. Only Oracle delivers the industry's most complete integrated apps-to-disk Linux solutions.

1.2 Use-Case Summary

The NetApp FlexPod Select for Oracle solution, which can run high-performance Oracle databases in a highly resilient architecture, is competitive in terms of both price and performance. As part of this solution, the following use cases were validated:

- Deliver an architecture and a prescriptive reference deployment that provides a high level of resiliency against component failure.
- Deliver over two million random read IOPS with microsecond-level latency using an 8K request size.
- Demonstrate consistent performance and response time utilizing a workload that consists of 90% random reads and 10% random writes using an 8K block size.
- Demonstrate consistent performance and response time utilizing a workload that consists of 80% random reads and 20% random writes using an 8K block size.

The FlexPod Select for Oracle high-performance solution provides extreme reliability when deploying tier 1 enterprise applications. This document describes deployment procedures for the FlexPod Select for Oracle high-performance solution. In the architecture described in the document, an Oracle RAC setup is used along with Cisco and NetApp components to demonstrate the performance and scalability of the solution. Section 2, “Solution Validation,” shows the performance characteristics of the solution in terms of IOPS and latency.

2 Solution Validation

2.1 Performance Testing Results

For all tests, NetApp used the Silly Little Oracle Benchmark (SLOB2) workload generator to simulate the I/O patterns that are likely to be encountered in actual Oracle production environments. SLOB2 drives different levels of simulated users, each generating the specific I/O patterns described previously in the use case section. After each test, NetApp recorded the physical database reads and average latency from the Oracle automatic workload repository reported by the Oracle database.

Figure 1 shows the IOPS and average latency observed by the database during testing with 100% random 8K reads. The load on the database was increased incrementally until the IOPS exceeded two million while simultaneously observing microsecond-level application latencies. The storage arrays are capable of delivering higher levels of IOPS provided the users have a tolerance for higher latency.

Figure 1) FlexPod Select for Oracle with eight RAC nodes showing 100% 8K random reads.

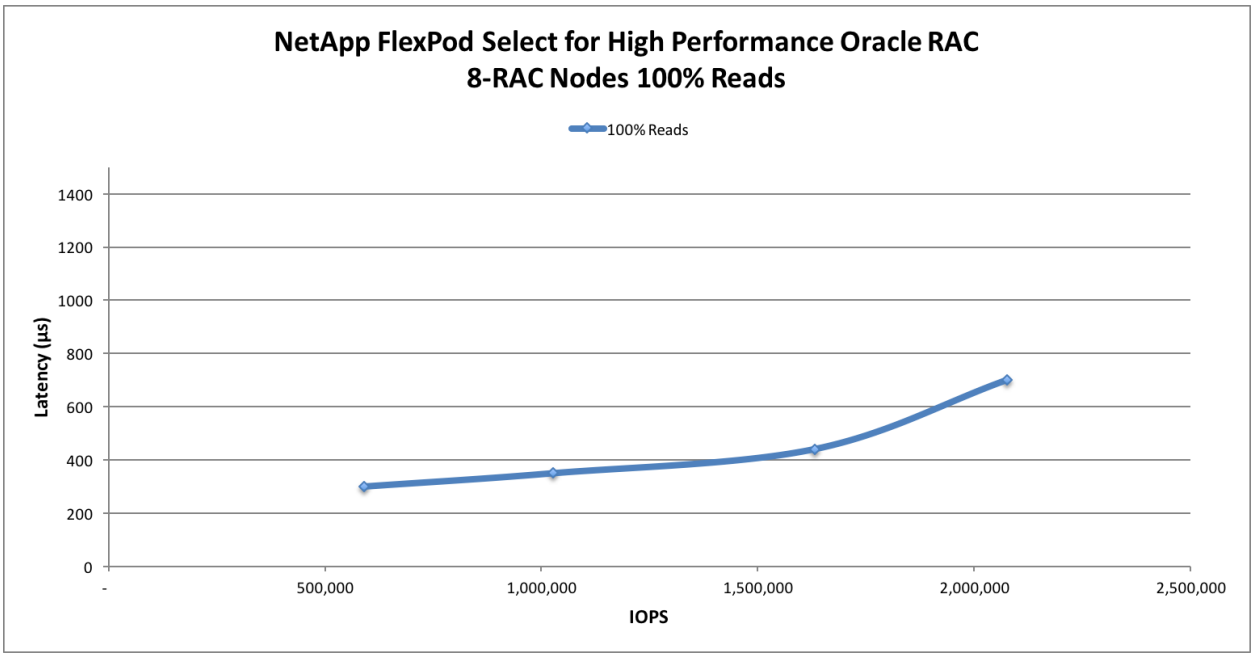

The results of the mixed workload use cases are shown in Table 1, which shows performance of the solution for all the use cases. Like 100% read workloads, we followed the same approach in running the workload to generate the IOPS with minimal latency.

Table 1) Solution validation results for all use cases.

Use Case	Max IOPS	Average Latency (µs)
100% random reads	2,076,637	700
90% random reads and 10% updates	1,400,395	600
80% random reads and 20% updates	1,300,909	1,040

3 Technology Requirements

This document is intended to provide an example of an environment that can be deployed using the best practices guidance for the FlexPod Select architecture. Additional resources can be added per the recommendations in the design guide for this solution (and companion to this document), NVA-0012-DESIGN: FlexPod Select for High-Performance Oracle RAC, and can be configured as necessary using the guidance in this document. Different models within the product families described as follows are acceptable as long as the models meet the physical cabling requirements specified in the design guide. Performance expectations for this environment are subject to change depending on the products used.

3.1 Hardware Requirements

Figure 2 shows the hardware components associated with this solution.

Figure 2) FlexPod Select for high-performance Oracle RAC solution architecture.

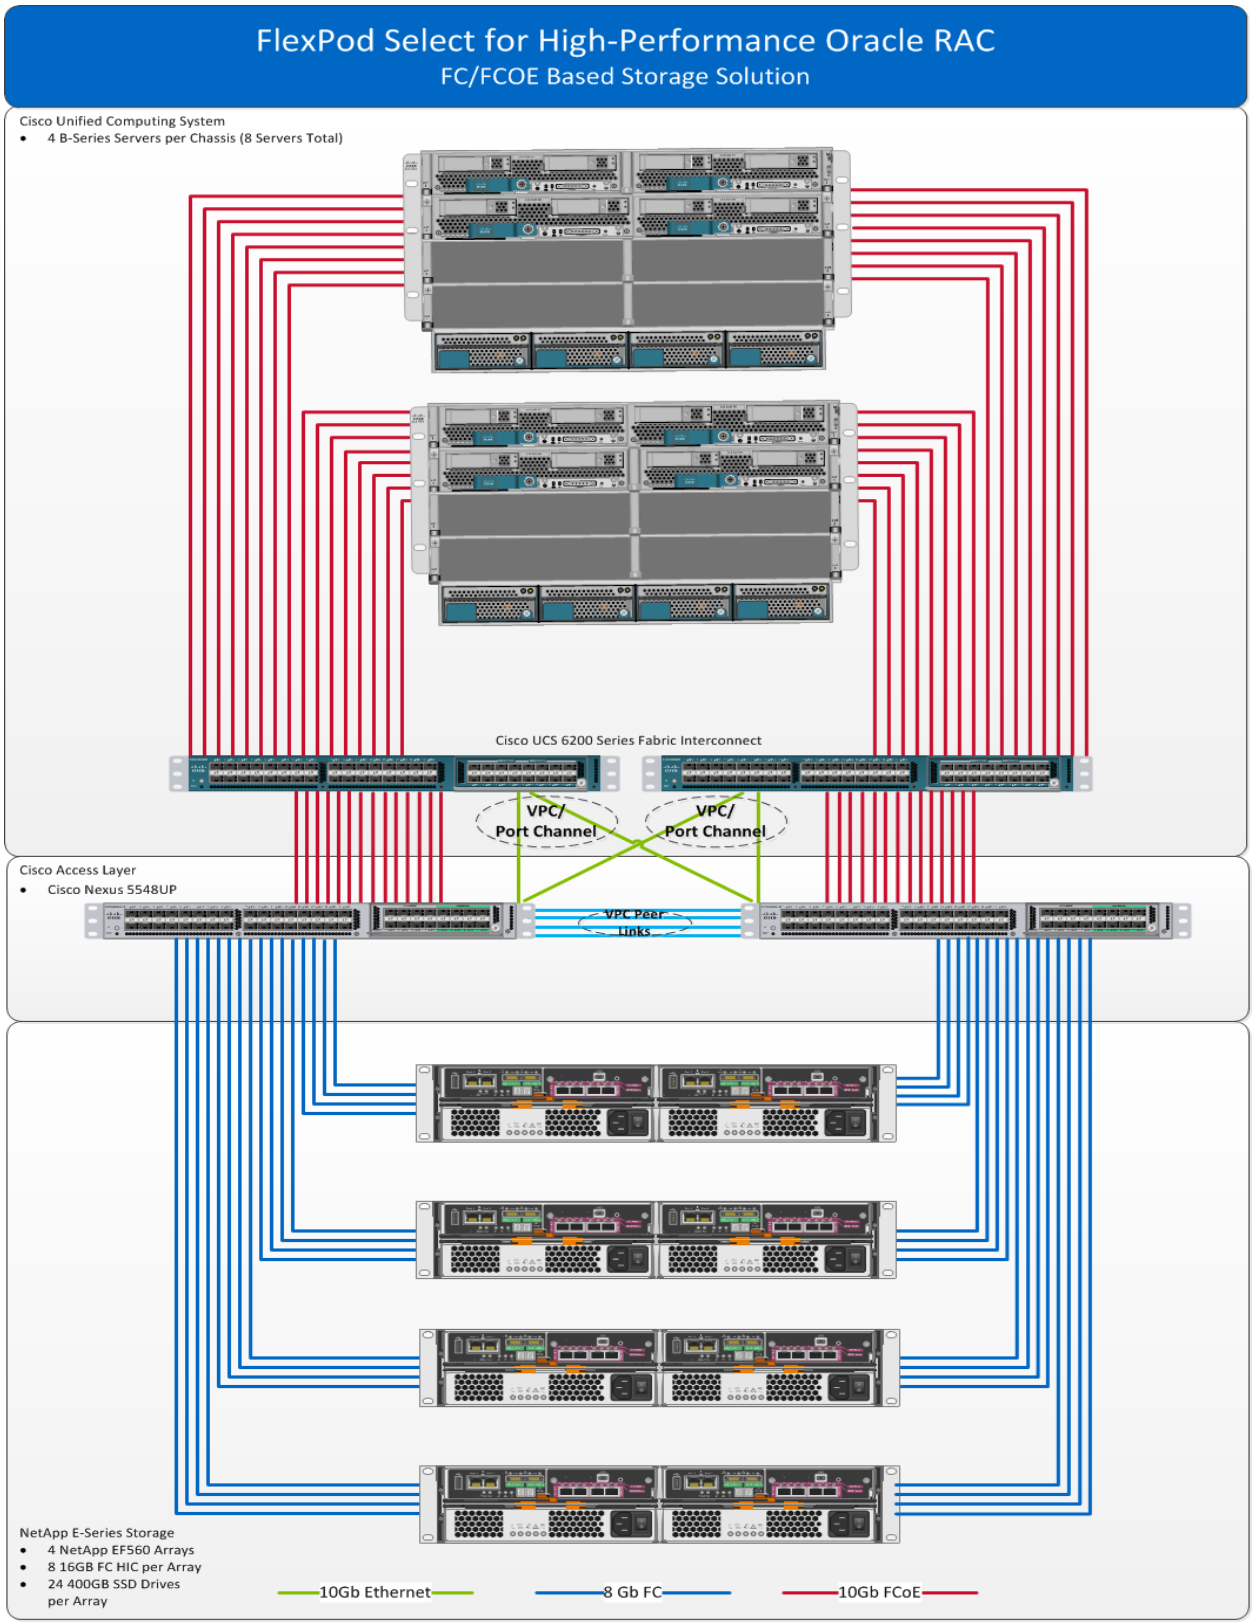

Table 2 lists the hardware components required to implement the solution and to achieve the previously defined performance objectives. The hardware components used in any particular implementation of this solution might vary based on customer requirements.

Table 2) Hardware requirements.

Hardware	Quantity
Storage	
NetApp EF560 dual-controller array	4
NetApp 16Gb Fibre Channel host interface card	8
NetApp 400GB SSD	96
Compute	
Cisco UCS fabric interconnect 6248	2
Cisco UCS 5108 chassis	2
Cisco B200M3 compute blade with 2 Intel Xeon E5-2620 v2 processors and 64GB RAM	4
Cisco B200M3 compute blade with 2 Intel Xeon E5-2650 v2 processors and 128GB RAM	4
Cisco UCS 1240 virtual interface card (VIC)	8
Cisco UCS 2208 FEX	4
Network	
Cisco Nexus 5548UP switch	2
Cisco 8Gb Fibre Channel SFP (DS-SFP-FC8G-SW)	32

3.2 Software Requirements

Table 3 lists the software components required to implement the solution. The software components used in any particular implementation of the solution may vary based on customer requirements.

Table 3) Software requirements.

Software	Version
Storage	
NetApp EF560 firmware version	8.20.08.00
NetApp SANtricity Storage Manager	11.20.0G00.0006
Compute	
Cisco UCS Manager	2.2(5b)
Cisco UCS firmware bundle	2.2(5b)
Cisco UCS 1240 VIC fnic driver	1.6.0.18
Cisco UCS 1240 VIC enic driver	2.1.1.67

Software	Version
Network	
Cisco Nexus NX-OS	6.0(2)N1(2a)
Operating System	
Oracle Enterprise Linux (OEL)	6.6
Application Software	
Oracle Real Application Clusters (RAC)	12.1.0.2.0
SLOB tool	2.3

3.3 Configuration Guidelines

This document describes the configuration of a fully redundant, highly available FlexPod unit with NetApp E-Series storage. We've clearly identified which component is being configured in each step or procedure, for example, Cisco Nexus 5548UP switch A and Cisco Nexus 5548UP switch B. The Cisco UCS fabric interconnects are similarly identified.

Additionally, this document provides steps for provisioning multiple Cisco UCS hosts.

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 4 describes the VLANs necessary for deployment as outlined in this guide. Table 5 lists the hosts necessary for deployment as outlined in this guide.

Table 4) Necessary VLANs.

VLAN Name	VLAN Purpose	Value Used in Validating This Document	Customer Value
Management	VLAN for management interfaces	3160	
Native	VLAN to which untagged frames are assigned	2	
RAC Cluster	VLAN for Oracle RAC intercluster node communication	3161	
FCoE-VLAN-A	VLAN for FCoE traffic for fabric A	101	
FCoE-VLAN-B	VLAN for FCoE traffic for fabric B	102	

Table 5) Host machines created.

Machine Description	Host Name
Qty 8 Oracle RAC servers	
SANtricity management system	

Table 6 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps.

Table 6) Configuration variables.

Description	Customer Value
EF560 controller 01a host name	

Description	Customer Value
EF560 controller 01a management IP address	
EF560 controller 01a channel 1 WWPN	
EF560 controller 01a channel 2 WWPN	
EF560 controller 01a channel 3 WWPN	
EF560 controller 01a channel 4 WWPN	
EF560 controller 01b host name	
EF560 controller 01b management IP address	
EF560 controller 01b channel 1 WWPN	
EF560 controller 01b channel 2 WWPN	
EF560 controller 01b channel 3 WWPN	
EF560 controller 01b channel 4 WWPN	
EF560 controller 02a host name	
EF560 controller 02a management IP address	
EF560 controller 02a channel 1 WWPN	
EF560 controller 02a channel 2 WWPN	
EF560 controller 02a channel 3 WWPN	
EF560 controller 02a channel 4 WWPN	
EF560 controller 02b host name	
EF560 controller 02b management IP address	
EF560 controller 02b channel 1 WWPN	
EF560 controller 02b channel 2 WWPN	
EF560 controller 02b channel 3 WWPN	
EF560 controller 02b channel 4 WWPN	
EF560 controller 03a host name	
EF560 controller 03a management IP address	
EF560 controller 03a channel 1 WWPN	
EF560 controller 03a channel 2 WWPN	
EF560 controller 03a channel 3 WWPN	
EF560 controller 03a channel 4 WWPN	
EF560 controller 03b host name	
EF560 controller 03b management IP address	
EF560 controller 03b channel 1 WWPN	
EF560 controller 03b channel 2 WWPN	
EF560 controller 03b channel 3 WWPN	

Description	Customer Value
EF560 controller 03b channel 4 WWPN	
EF560 controller 04a host name	
EF560 controller 04a management IP address	
EF560 controller 04a channel 1 WWPN	
EF560 controller 04a channel 2 WWPN	
EF560 controller 04a channel 3 WWPN	
EF560 controller 04a channel 4 WWPN	
EF560 controller 04b host name	
EF560 controller 04b management IP address	
EF560 controller 04b channel 1 WWPN	
EF560 controller 04b channel 2 WWPN	
EF560 controller 04b channel 3 WWPN	
EF560 controller 04b channel 4 WWPN	
Cisco Nexus A host name	
Cisco Nexus B host name	
Cisco Nexus VPC domain ID	
Global NTP server IP address	
Global DNS server IP address	
Cisco UCS manager cluster system name	
Cisco UCS manager password	
Cisco UCS fabric interconnect (FI) A management IP address	
Management network netmask	
Management network default gateway	
Cisco UCS manager cluster IP address	
Cisco UCS FI B management IP address	
RAC SCAN IP addresses 1	
RAC SCAN IP addresses 2	
RAC SCAN IP addresses 3	
RAC cluster subnet mask	
RAC private interconnect subnet mask	
Public virtual IP (VIP) subnet mask	
Management IP address of RAC-SERVER-01	
RAC cluster IP address of RAC-SERVER-01	

Description	Customer Value
RAC private interconnect address of RAC-SERVER-01	
Public virtual IP (VIP) addresses of RAC-SERVER-01	
WWPN of RAC-SERVER-01 A Port 1	
WWPN of RAC-SERVER-01 A Port 2	
WWPN of RAC-SERVER-01 B Port 1	
WWPN of RAC-SERVER-01 B Port 2	
Management IP address of RAC-SERVER-02	
RAC cluster IP address of RAC-SERVER-02	
RAC private interconnect address of RAC-SERVER-02	
Public VIP addresses of RAC-SERVER-02	
WWPN of RAC-SERVER-02 A Port 1	
WWPN of RAC-SERVER-02 A Port 2	
WWPN of RAC-SERVER-02 B Port 1	
WWPN of RAC-SERVER-02 B Port 2	
Management IP address of RAC-SERVER-03	
RAC cluster IP address of RAC-SERVER-03	
RAC private interconnect address of RAC-SERVER-03	
Public VIP addresses of RAC-SERVER-03	
WWPN of RAC-SERVER-03 A Port 1	
WWPN of RAC-SERVER-03 A Port 2	
WWPN of RAC-SERVER-03 B Port 1	
WWPN of RAC-SERVER-03 B Port 2	
Management IP address of RAC-SERVER-04	
RAC cluster IP address of RAC-SERVER-04	
RAC private interconnect address of RAC-SERVER-04	
Public VIP addresses of RAC-SERVER-04	
WWPN of RAC-SERVER-04 A Port 1	
WWPN of RAC-SERVER-04 A Port 2	
WWPN of RAC-SERVER-04 B Port 1	
WWPN of RAC-SERVER-04 B Port 2	
Management IP address of RAC-SERVER-05	
RAC cluster IP address of RAC-SERVER-05	

Description	Customer Value
RAC private interconnect address of RAC-SERVER-05	
Public VIP addresses of RAC-SERVER-05	
WWPN of RAC-SERVER-05 A Port 1	
WWPN of RAC-SERVER-05 A Port 2	
WWPN of RAC-SERVER-05 B Port 1	
WWPN of RAC-SERVER-05 B Port 2	
Management IP address of RAC-SERVER-06	
RAC cluster IP address of RAC-SERVER-06	
RAC private interconnect address of RAC-SERVER-06	
Public VIP addresses of RAC-SERVER-06	
WWPN of RAC-SERVER-06 A Port 1	
WWPN of RAC-SERVER-06 A Port 2	
WWPN of RAC-SERVER-06 B Port 1	
WWPN of RAC-SERVER-06 B Port 2	
Management IP address of RAC-SERVER-07	
RAC cluster IP address of RAC-SERVER-07	
RAC private interconnect address of RAC-SERVER-07	
Public VIP addresses of RAC-SERVER-07	
WWPN of RAC-SERVER-07 A Port 1	
WWPN of RAC-SERVER-07 A Port 2	
WWPN of RAC-SERVER-07 B Port 1	
WWPN of RAC-SERVER-07 B Port 2	
Management IP address of RAC-SERVER-08	
RAC cluster IP address of RAC-SERVER-08	
RAC private interconnect address of RAC-SERVER-08	
Public VIP addresses of RAC-SERVER-08	
WWPN of RAC-SERVER-08 A Port 1	
WWPN of RAC-SERVER-08 A Port 2	
WWPN of RAC-SERVER-08 B Port 1	
WWPN of RAC-SERVER-08 B Port 2	

4 Deployment Procedures

The following subsections describe the deployment of the overall solution and include specific steps to install and configure the technology components described in section 3, “Technology Requirements,” into a consumable solution.

Note: Note that some of the specific names used in these sections, for example, group names or file names, are examples and might not be the same for all environments.

4.1 Physical Infrastructure Layout

The physical layout of the integrated reference architecture is explained in this section. Included are graphical layouts, which provide helpful cabling diagrams for all equipment in the design. Correct cabling is instrumental in achieving correct and efficient operation of the infrastructure, both in the initial deployment and in the ongoing lifecycle.

NetApp EF-Series Storage Controllers

Follow the guidelines in Table 7 when unpacking and racking all of the storage controllers for this solution. Proper adherence to these guidelines is essential for the equipment to operate properly.

Table 7) Requirements for unpacking and racking of all storage controllers in solution.

Requirement	Reference	Comments
Physical site where storage system will be installed must be ready	NetApp EF560 Flash Array Site Preparation Guide	Refer to the section “Specifications of the EF560 Flash Array” of the NetApp EF560 Flash Array Site Preparation Guide.
Storage system connectivity requirements for out-of-band management	NetApp EF560 Flash Array Hardware Cabling Guide	Refer to the section “Cabling for Out-of-Band Management” of the NetApp EF560 Flash Array Hardware Cabling Guide.
Storage system power, cooling, air flow, temperature, and humidity requirements	NetApp EF560 Flash Array Site Preparation Guide	Refer to the section “Specifications of the EF560 Flash Array” of the NetApp EF560 Flash Array Site Preparation Guide.

Cisco Unified Computing System Fabric Interconnect and Chassis

Follow the guidelines in Table 8 when installing the Cisco UCS FI6248 fabric interconnects and the Cisco UCS 5108 chassis in the data center. Proper adherence to these guidelines is essential in order for the equipment to operate as expected.

Table 8) Requirements for installing Cisco UCS FI6248 fabric interconnects and Cisco UCS 5108 chassis in data center.

Requirement	Reference	Comments
Physical site where compute system will be installed must be ready	Cisco UCS Site Preparation Guide	
Compute system power, cooling, air flow, temperature, and humidity requirements	Cisco UCS Site Preparation Guide	

Cisco Nexus 5548UP Switches

Follow the guidelines in Table 9 when installing the Cisco Nexus 5548UP switches in the data center. Proper adherence to these guidelines is essential in order for the equipment to operate as expected.

Table 9) Requirements for installing Cisco Nexus 5548UP switches in data center.

Requirement	Reference	Comments
Physical site where switches need to be installed must be ready	Cisco Nexus 5000 Series Hardware Installation Guide	
Switch power, cooling, air flow, temperature, and humidity requirements	Cisco Nexus 5000 Series Hardware Installation Guide	Refer to section “Technical Specifications” in the Cisco Nexus 5000 Series Hardware Installation Guide.

Physical Cabling

Figure 3 and Figure 4 show the hardware components of the solution. Use them as a guide to visualize the physical connectivity between the solution components.

Note: Because of the size of the physical cabling diagram, the detail of the information contained within it, and the limitations of the PDF format for enlarging it, it has been broken up into two parts.

The prescribed connectivity was tested during the verification of this solution and provided the observed performance results. Following the diagram is a series of tables that correspond to each component of the solution.

Figure 3) Physical cabling diagram for Cisco UCS and Cisco Nexus switches.

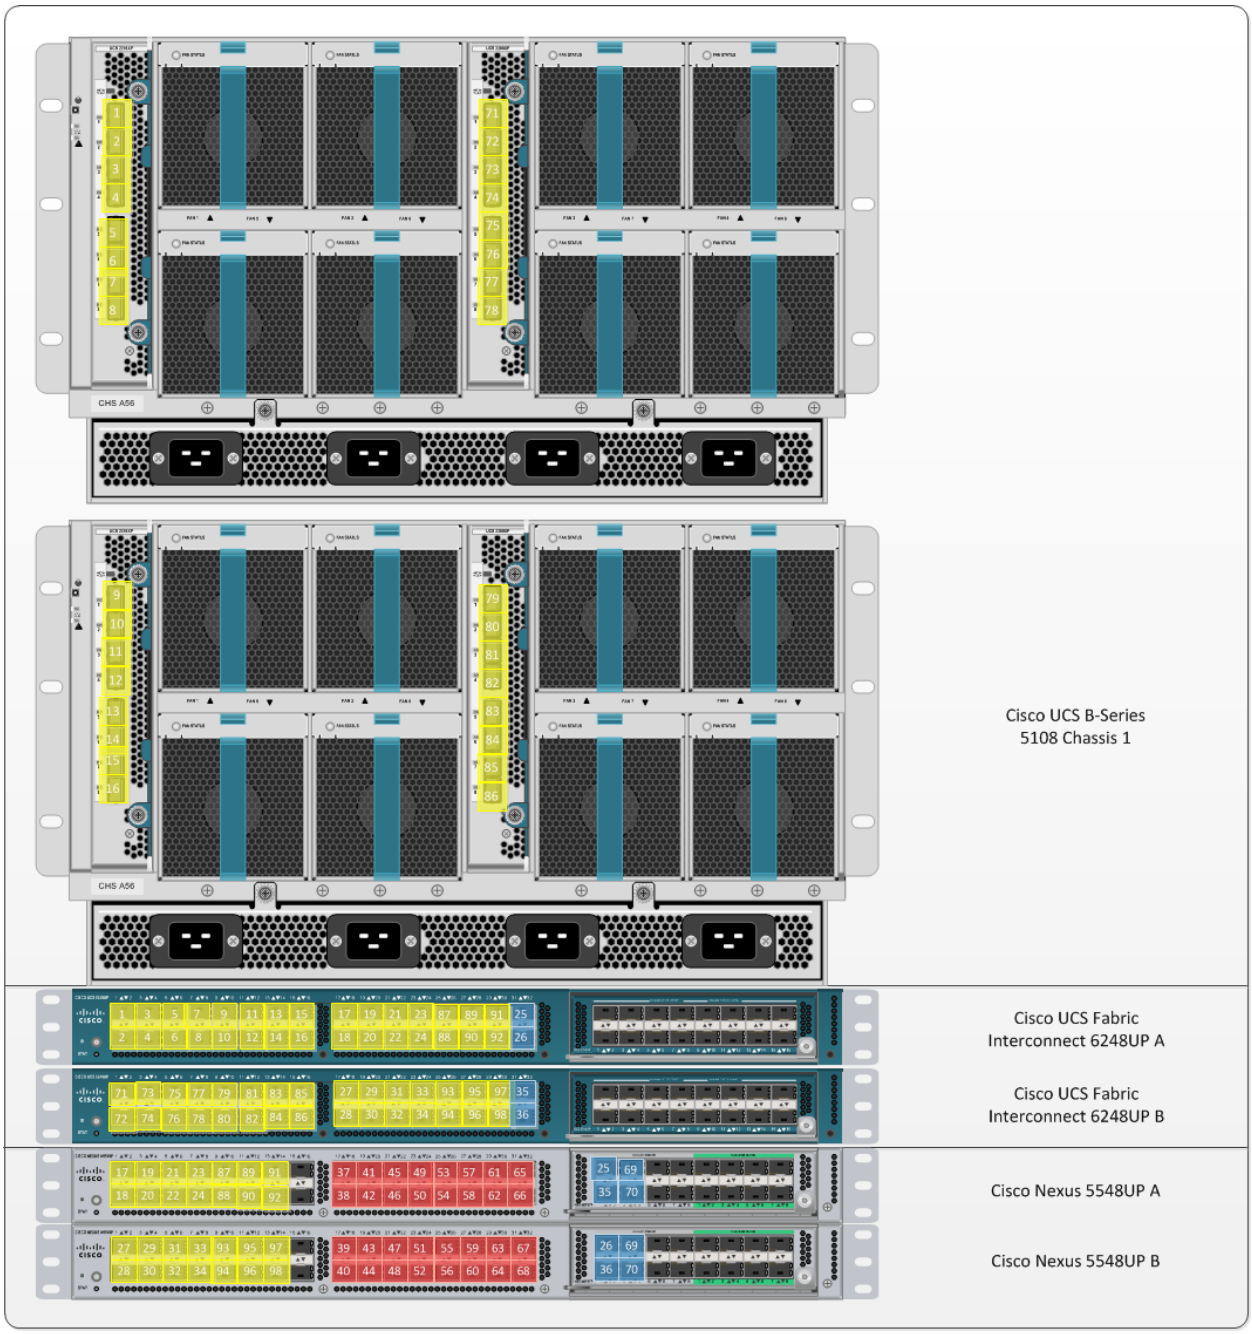

Figure 4) Physical cabling diagram for NetApp EF560 arrays.

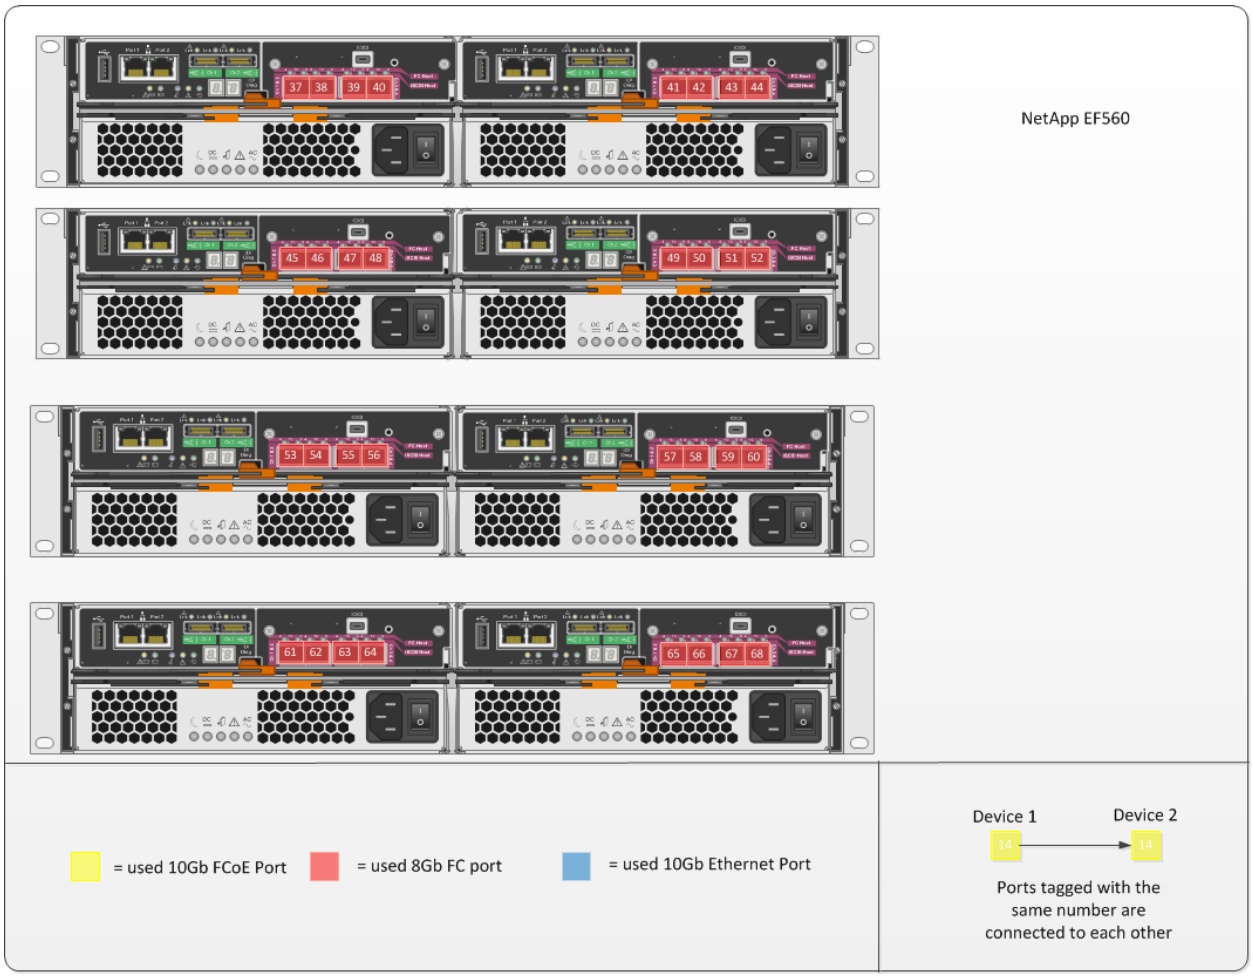

Table 10) Cisco UCS B-Series 5801 Chassis 1 cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS 5108 Chassis 1	FEX2208 A Port 1	10GbE	Cisco UCS FI6248 1	1/1	1
	FEX2208 A Port 2	10GbE	Cisco UCS FI6248 1	1/2	2
	FEX2208 A Port 3	10GbE	Cisco UCS FI6248 1	1/3	3
	FEX2208 A Port 4	10GbE	Cisco UCS FI6248 1	1/4	4
	FEX2208 A Port 5	10GbE	Cisco UCS FI6248 1	1/5	5
	FEX2208 A Port 6	10GbE	Cisco UCS FI6248 1	1/6	6
	FEX2208	10GbE	Cisco UCS FI6248 1	1/7	7

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	A Port 7				
	FEX2208 A Port 8	10GbE	Cisco UCS FI6248 1	1/8	8
	FEX2208 B Port 1	10GbE	Cisco UCS FI6248 2	1/1	71
	FEX2208 B Port 2	10GbE	Cisco UCS FI6248 2	1/2	72
	FEX2208 B Port 3	10GbE	Cisco UCS FI6248 2	1/3	73
	FEX2208 B Port 4	10GbE	Cisco UCS FI6248 2	1/4	74
	FEX2208 B Port 5	10GbE	Cisco UCS FI6248 2	1/5	75
	FEX2208 B Port 6	10GbE	Cisco UCS FI6248 2	1/6	76
	FEX2208 B Port 7	10GbE	Cisco UCS FI6248 2	1/7	77
	FEX2208 B Port 8	10GbE	Cisco UCS FI6248 2	1/8	78

Table 11) Cisco UCS B-Series 5108 Chassis 2 cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS 5108 Chassis 2	FEX2208 A Port 1	10GbE	Cisco UCS FI6248 1	1/9	9
	FEX2208 A Port 2	10GbE	Cisco UCS FI6248 1	1/10	10
	FEX2208 A Port 3	10GbE	Cisco UCS FI6248 1	1/11	11
	FEX2208 A Port 4	10GbE	Cisco UCS FI6248 1	1/12	12
	FEX2208 A Port 5	10GbE	Cisco UCS FI6248 2	1/13	13
	FEX2208 A Port 6	10GbE	Cisco UCS FI6248 2	1/14	14
	FEX2208 A Port 7	10GbE	Cisco UCS FI6248 2	1/15	15
	FEX2208 A Port 8	10GbE	Cisco UCS FI6248 2	1/16	16
	FEX2208	10GbE	Cisco UCS FI6248 1	1/9	79

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	B Port 1				
	FEX2208 B Port 2	10GbE	Cisco UCS FI6248 1	1/10	80
	FEX2208 B Port 3	10GbE	Cisco UCS FI6248 1	1/11	81
	FEX2208 B Port 4	10GbE	Cisco UCS FI6248 1	1/12	82
	FEX2208 B Port 5	10GbE	Cisco UCS FI6248 2	1/13	83
	FEX2208 B Port 6	10GbE	Cisco UCS FI6248 2	1/14	84
	FEX2208 B Port 7	10GbE	Cisco UCS FI6248 2	1/15	85
	FEX2208 B Port 8	10GbE	Cisco UCS FI6248 2	1/16	86

Table 12) Cisco UCS Fabric Interconnect 6248UP A cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS Fabric Interconnect 6248UP A	1/1	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 1	1
	1/2	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 2	2
	1/3	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 3	3
	1/4	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 4	4
	1/5	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 5	5
	1/6	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 6	6
	1/7	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 7	7
	1/8	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 A Port 8	8
	1/9	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 1	9
	1/10	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 2	10
	1/11	10GbE	Cisco UCS 5108 Chassis 2	FEX2208	11

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
				A Port 3	
	1/12	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 4	12
	1/13	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 5	13
	1/14	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 6	14
	1/15	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 7	15
	1/16	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 A Port 8	16
	1/17	10GbE	Cisco Nexus 5548UP 1	Eth1/1	17
	1/18	10GbE	Cisco Nexus 5548UP 1	Eth1/2	18
	1/19	10GbE	Cisco Nexus 5548UP 1	Eth1/3	19
	1/20	10GbE	Cisco Nexus 5548UP 1	Eth1/4	20
	1/21	10GbE	Cisco Nexus 5548UP 1	Eth1/5	21
	1/22	10GbE	Cisco Nexus 5548UP 1	Eth1/6	22
	1/23	10GbE	Cisco Nexus 5548UP 1	Eth1/7	23
	1/24	10GbE	Cisco Nexus 5548UP 1	Eth1/8	24
	1/25	10GbE	Cisco Nexus 5548UP 1	Eth1/9	87
	1/26	10GbE	Cisco Nexus 5548UP 1	Eth1/10	88
	1/27	10GbE	Cisco Nexus 5548UP 1	Eth1/11	89
	1/28	10GbE	Cisco Nexus 5548UP 1	Eth1/12	90
	1/29	10GbE	Cisco Nexus 5548UP 1	Eth1/13	91
	1/30	10GbE	Cisco Nexus 5548UP 1	Eth1/14	92
	1/31	10GbE	Cisco Nexus 5548UP 1	Eth2/1	25
	1/32	10GbE	Cisco Nexus 5548UP 2	Eth2/1	26
	MGMT0	GbE	GbE management switch (not shown)	Any	Not shown
	L1	GbE	Cisco UCS FI6248 2	L1	Not shown
	L2	GbE	Cisco UCS FI6248 2	L2	Not shown

Table 13) Cisco UCS Fabric Interconnect 6248UP B cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco UCS Fabric Interconnect 6248UP B	1/1	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 1	71
	1/2	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 2	72
	1/3	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 3	73
	1/4	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 4	74
	1/5	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 5	75
	1/6	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 6	76
	1/7	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 7	77
	1/8	10GbE	Cisco UCS 5108 Chassis 1	FEX2208 B Port 8	78
	1/9	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 1	79
	1/10	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 2	80
	1/11	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 3	81
	1/12	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 4	82
	1/13	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 5	83
	1/14	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 6	84
	1/15	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 7	85
	1/16	10GbE	Cisco UCS 5108 Chassis 2	FEX2208 B Port 8	86
	1/17	10GbE	Cisco Nexus 5548UP 2	Eth1/1	27
	1/18	10GbE	Cisco Nexus 5548UP 2	Eth1/2	28
	1/19	10GbE	Cisco Nexus 5548UP 2	Eth1/3	29
	1/20	10GbE	Cisco Nexus 5548UP 2	Eth1/4	30
	1/21	10GbE	Cisco Nexus 5548UP 2	Eth1/5	31

Local Device	Local Port	Connection	Remote Device	Remote Port	Cablin g Code
	1/22	10GbE	Cisco Nexus 5548UP 2	Eth1/6	32
	1/23	10GbE	Cisco Nexus 5548UP 2	Eth1/7	33
	1/24	10GbE	Cisco Nexus 5548UP 2	Eth1/8	34
	1/25	10GbE	Cisco Nexus 5548UP 2	Eth1/9	93
	1/26	10GbE	Cisco Nexus 5548UP 2	Eth1/10	94
	1/27	10GbE	Cisco Nexus 5548UP 2	Eth1/11	95
	1/28	10GbE	Cisco Nexus 5548UP 2	Eth1/12	96
	1/29	10GbE	Cisco Nexus 5548UP 2	Eth1/13	97
	1/30	10GbE	Cisco Nexus 5548UP 2	Eth1/14	98
	31	10GbE	Cisco Nexus 5548UP 1	Eth2/2	35
	32	10GbE	Cisco Nexus 5548UP 2	Eth2/2	36
	MGMT0	GbE	GbE management switch (not shown)	Any	Not shown
	L1	GbE	Cisco UCS FI6248 1	L1	Not shown
	L2	GbE	Cisco UCS FI6248 2	L2	Not shown

Table 14) Cisco Nexus 5548UP Switch A cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5548UP Switch A	Eth1/1	10GbE	Cisco UCS FI6248 1	1/17	17
	Eth 1/2	10GbE	Cisco UCS FI6248 1	1/18	18
	Eth 1/3	10GbE	Cisco UCS FI6248 1	1/19	19
	Eth 1/4	10GbE	Cisco UCS FI6248 1	1/20	20
	Eth 1/5	10GbE	Cisco UCS FI6248 1	1/21	21
	Eth 1/6	10GbE	Cisco UCS FI6248 1	1/22	22
	Eth 1/7	10GbE	Cisco UCS FI6248 1	1/23	23
	Eth 1/8	10GbE	Cisco UCS FI6248 1	1/24	24
	Eth 1/9	10GbE	Cisco UCS FI6248 1	1/25	87
	Eth 1/10	10GbE	Cisco UCS FI6248 1	1/26	88
	Eth 1/11	10GbE	Cisco UCS FI6248 1	1/27	89
	Eth 1/12	10GbE	Cisco UCS FI6248 1	1/28	90

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	Eth 1/13	10GbE	Cisco UCS FI6248 1	1/29	91
	Eth 1/14	10GbE	Cisco UCS FI6248 1	1/30	92
	FC 1/17	8Gb FC	EF560 Array 1	HIC 1 Port 1	37
	FC 1/18	8Gb FC	EF560 Array 1	HIC 1 Port 2	38
	FC 1/19	8Gb FC	EF560 Array 1	HIC 2 Port 1	41
	FC 1/20	8Gb FC	EF560 Array 1	HIC 2 Port 2	42
	FC 1/21	8Gb FC	EF560 Array 2	HIC 1 Port 1	45
	FC 1/22	8Gb FC	EF560 Array 2	HIC 1 Port 2	46
	FC 1/23	8Gb FC	EF560 Array 2	HIC 2 Port 1	49
	FC 1/24	8Gb FC	EF560 Array 2	HIC 2 Port 2	50
	FC 1/25	8Gb FC	EF560 Array 3	HIC 1 Port 1	53
	FC 1/26	8Gb FC	EF560 Array 3	HIC 1 Port 2	54
	FC 1/27	8Gb FC	EF560 Array 3	HIC 2 Port 1	57
	FC 1/28	8Gb FC	EF560 Array 3	HIC 2 Port 2	58
	FC 1/29	8Gb FC	EF560 Array 4	HIC 1 Port 1	61
	FC 1/30	8Gb FC	EF560 Array 4	HIC 1 Port 2	62
	FC 1/31	8Gb FC	EF560 Array 4	HIC 2 Port 1	65
	FC 1/32	8Gb FC	EF560 Array 4	HIC 2 Port 2	66
	Eth 2/1	10GbE	Cisco UCS FI6248 1	1/31	25
	Eth 2/2	10GbE	Cisco UCS FI6248 2	1/31	35
	Eth 2/3	10GbE	Cisco Nexus 5548UP 2	Eth 2/3	69
	Eth 2/4	10GbE	Cisco Nexus 5548UP 2	Eth 2/4	70

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	Mgmt0	1GbE	GbE management switch (not shown)	Any	Not shown

Table 15) Cisco Nexus 5548UP Switch B cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
Cisco Nexus 5548UP Switch B	Eth1/1	10GbE	Cisco UCS FI6248 2	1/17	27
	Eth 1/2	10GbE	Cisco UCS FI6248 2	1/18	28
	Eth 1/3	10GbE	Cisco UCS FI6248 2	1/19	29
	Eth 1/4	10GbE	Cisco UCS FI6248 2	1/20	30
	Eth 1/5	10GbE	Cisco UCS FI6248 2	1/21	31
	Eth 1/6	10GbE	Cisco UCS FI6248 2	1/22	32
	Eth 1/7	10GbE	Cisco UCS FI6248 2	1/23	33
	Eth 1/8	10GbE	Cisco UCS FI6248 2	1/24	34
	Eth 1/9	10GbE	Cisco UCS FI6248 2	1/25	93
	Eth 1/10	10GbE	Cisco UCS FI6248 2	1/26	94
	Eth 1/11	10GbE	Cisco UCS FI6248 2	1/27	95
	Eth 1/12	10GbE	Cisco UCS FI6248 2	1/28	96
	Eth 1/13	10GbE	Cisco UCS FI6248 2	1/29	97
	Eth 1/14	10GbE	Cisco UCS FI6248 2	1/30	98
	FC 1/17	8Gb FC	EF560 Array 1	HIC 1 Port 3	39
	FC 1/18	8Gb FC	EF560 Array 1	HIC 1 Port 4	40
	FC 1/19	8Gb FC	EF560 Array 1	HIC 2 Port 3	43
	FC 1/20	8Gb FC	EF560 Array 1	HIC 2 Port 4	44
	FC 1/21	8Gb FC	EF560 Array 2	HIC 1 Port 3	47
	FC 1/22	8Gb FC	EF560 Array 2	HIC 1 Port 4	48
	FC 1/23	8Gb FC	EF560 Array 2	HIC 2 Port 3	51
	FC 1/24	8Gb FC	EF560 Array 2	HIC 2 Port 4	52

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	FC 1/25	8Gb FC	EF560 Array 3	HIC 1 Port 3	55
	FC 1/26	8Gb FC	EF560 Array 3	HIC 1 Port 4	56
	FC 1/27	8Gb FC	EF560 Array 3	HIC 2 Port 3	59
	FC 1/28	8Gb FC	EF560 Array 3	HIC 2 Port 4	60
	FC 1/29	8Gb FC	EF560 Array 4	HIC 1 Port 3	63
	FC 1/30	8Gb FC	EF560 Array 4	HIC 1 Port 4	64
	FC 1/31	8Gb FC	EF560 Array 4	HIC 2 Port 3	67
	FC 1/32	8Gb FC	EF560 Array 4	HIC 2 Port 4	68
	Eth 2/1	10GbE	Cisco UCS FI6248 1	1/32	26
	Eth 2/2	10GbE	Cisco UCS FI6248 2	1/32	36
	Eth 2/3	10GbE	Cisco Nexus 5548UP 1	Eth 2/3	69
	Eth 2/4	10GbE	Cisco Nexus 5548UP 1	Eth 2/4	70
	Mgmt0	1GbE	GbE management switch (not shown)	Any	Not shown

Table 16) NetApp EF560 1 cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp EF560 1	HIC 1 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/17	37
	HIC 1 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/18	38
	HIC 1 Port 3	8Gb FC	Cisco Nexus 5548UP 2	FC 1/17	39
	HIC 1 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/18	40
	HIC 2 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/19	41
	HIC 2 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/20	42
	HIC 2 Port	8Gb FC	Cisco Nexus 5548UP 2	FC 1/19	43

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	3				
	HIC 2 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/20	44
	Controller 1 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown
	Controller 2 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown

Table 17) NetApp EF560 2 cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp EF560 2	HIC 1 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/21	37
	HIC 1 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/22	38
	HIC 1 Port 3	8Gb FC	Cisco Nexus 5548UP 2	FC 1/21	39
	HIC 1 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/22	40
	HIC 2 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/23	41
	HIC 2 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/24	42
	HIC 2 Port 3	8Gb FC	Cisco Nexus 5548UP 2	FC 1/23	43
	HIC 2 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/24	44
	Controller 1 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown
	Controller 2 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown

Table 18) NetApp EF560 3 cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp EF560 3	HIC 1 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/25	53
	HIC 1 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/26	54
	HIC 1 Port	8Gb FC	Cisco Nexus 5548UP 2	FC 1/25	55

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
	3				
	HIC 1 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/26	56
	HIC 2 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/27	57
	HIC 2 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/28	58
	HIC 2 Port 3	8Gb FC	Cisco Nexus 5548UP 2	FC 1/27	59
	HIC 2 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/28	60
	Controller 1 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown
	Controller 2 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown

Table 19) NetApp EF560 4 cabling.

Local Device	Local Port	Connection	Remote Device	Remote Port	Cabling Code
NetApp EF560 4	HIC 1 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/29	61
	HIC 1 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/30	62
	HIC 1 Port 3	8Gb FC	Cisco Nexus 5548UP 2	FC 1/29	63
	HIC 1 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/30	64
	HIC 2 Port 1	8Gb FC	Cisco Nexus 5548UP 1	FC 1/31	65
	HIC 2 Port 2	8Gb FC	Cisco Nexus 5548UP 1	FC 1/32	66
	HIC 2 Port 3	8Gb FC	Cisco Nexus 5548UP 2	FC 1/31	67
	HIC 2 Port 4	8Gb FC	Cisco Nexus 5548UP 2	FC 1/32	68
	Controller 1 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown
	Controller 2 Port 1	1 GbE	GbE management switch (not shown)	Any	Not shown

4.2 Cisco UCS Configuration

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

This deployment guide assumes that the Cisco UCS environment has the correct code version for the Cisco UCS management interface and that the correct firmware for the solution has been installed.

This section does not require the completion of any other sections of this document, but it does require that the equipment has been racked and cabled as per section 4.1, “Physical Infrastructure Layout.”

Perform Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments

Cisco UCS Fabric Interconnect 6248 1

To configure the Cisco UCS environment for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6248 Fabric Interconnect.

```
Enter the configuration method. (console/gui) ? console
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]: y
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]: y
Enter the switch fabric (A/B) []: A
Enter the system name:
Physical Switch Mgmt0 IP address :
Physical Switch Mgmt0 IPv4 netmask :
IPv4 address of the default gateway :
Cluster IPv4 address :
Configure the DNS Server IP address? (yes/no) [n]: y
DNS IP address :
Configure the default domain name? (yes/no) [n]:
Join centralized management environment (UCS Central)? (yes/no) [n]:
Following configurations will be applied:
Switch Fabric=A
System Name=
Enforced Strong Password=yes
Physical Switch Mgmt0 IP Address=
Physical Switch Mgmt0 IP Netmask=
Default Gateway=
Ipv6 value=0
DNS Server=

Cluster Enabled=yes
Cluster IP Address=
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized

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

Configuration file - Ok
```

2. Wait for the login prompt to verify that the configuration has been saved.

Cisco UCS Fabric Interconnect 6248 2

To configure the Cisco UCS environment for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the second Cisco UCS 6248 Fabric Interconnect.


```
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) ? y

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:
Peer Fabric interconnect Mgmt0 IPv4 Netmask:
Cluster IPv4 address          :

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

Physical Switch Mgmt0 IP address :


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

Configuration file - Ok
```

2. Wait for the login prompt to confirm that the configuration has been saved.

Log in to Cisco UCS Manager

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6248 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.
4. When prompted, enter `admin` as the user name and enter the Cisco UCS administrator password.
5. To log in to Cisco UCS Manager, click Login. Upon successful login, the following screen should be displayed:

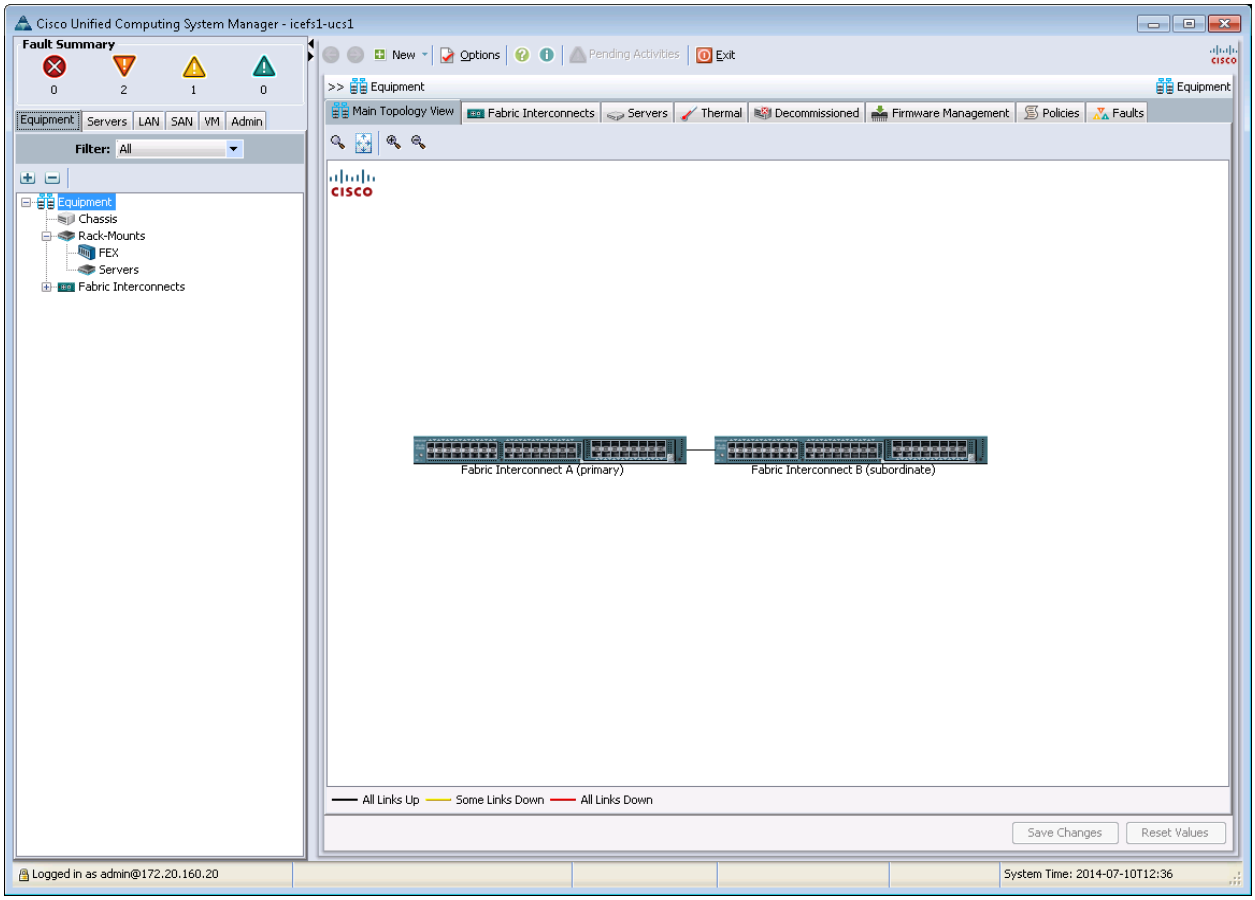

Upgrade Cisco UCS Manager to UCSM Release 2.2(5b)

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. In the list on the left, select Equipment.
3. On the right, select the Firmware Management tab and then the Installed Firmware tab.
4. If the Cisco UCS Manager Running Version in the center pane is not 2.2(5b), refer to http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html and select the appropriate Upgrade Guide to upgrade the Cisco UCS system to Release 2.2(5b). Install the version 2.2(5b) Infrastructure, B-Series, and C-Series software bundles.

Add Block of IP Addresses for Out-of-Band KVM Access

To create a block of IP addresses for server keyboard, video, mouse (KVM) access in the Cisco UCS environment, complete the following steps:

Note: This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Pools > Root > IP Pools > IP Pool Ext-Mgmt.

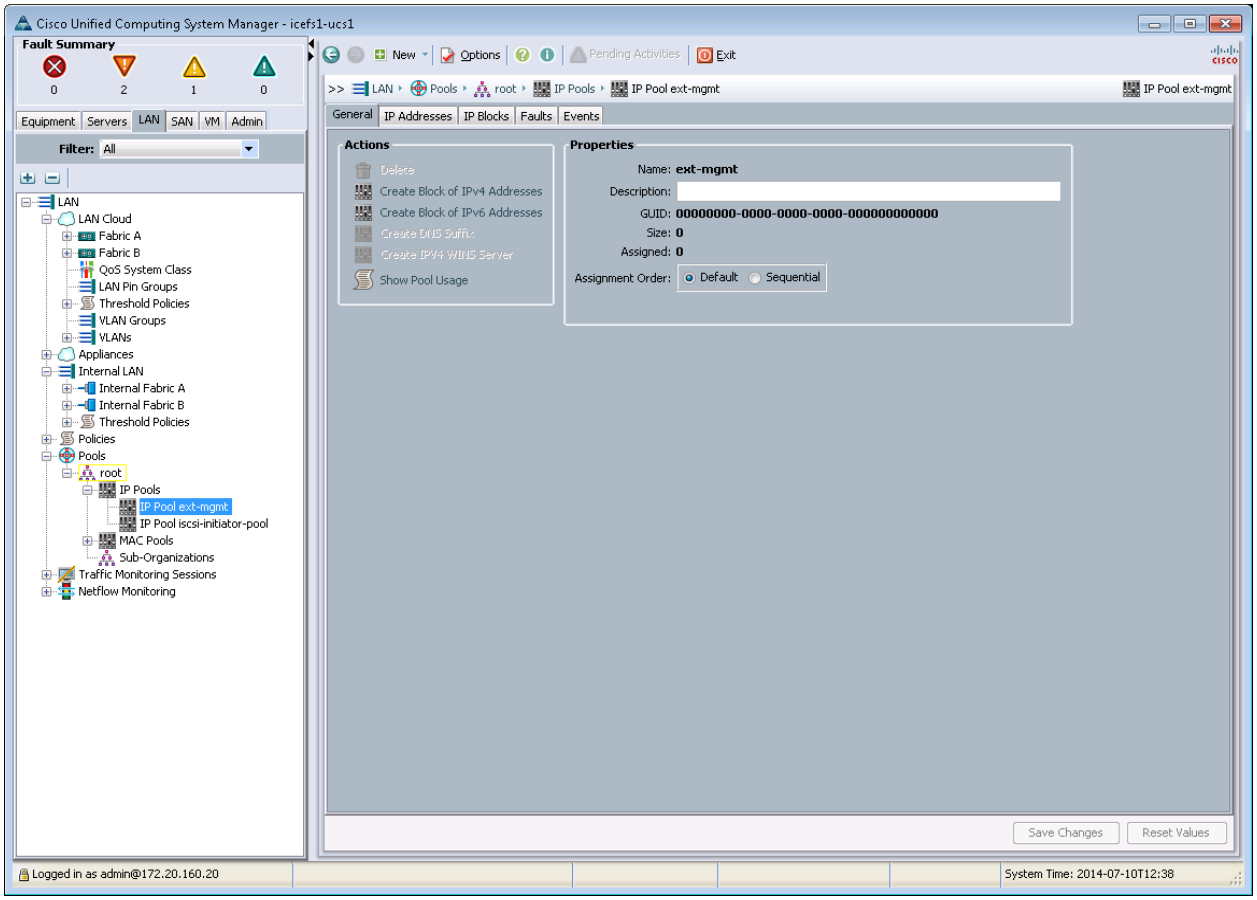

3. In the Actions pane, select Create Block of IP Addresses.
4. Enter the starting IP address of the block, the number of IP addresses required, and the subnet and gateway information.

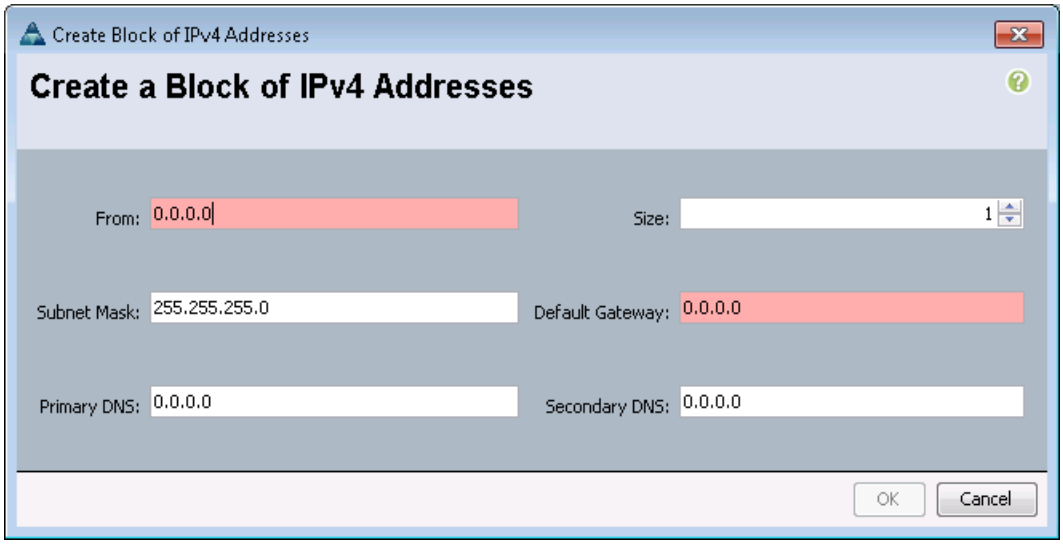

5. Click OK to create the IP block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

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

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of Cisco UCS B-series chassis and C-series servers. To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab and select Equipment in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to 8-link or set it to match the number of uplink ports that are cabled between the chassis and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.

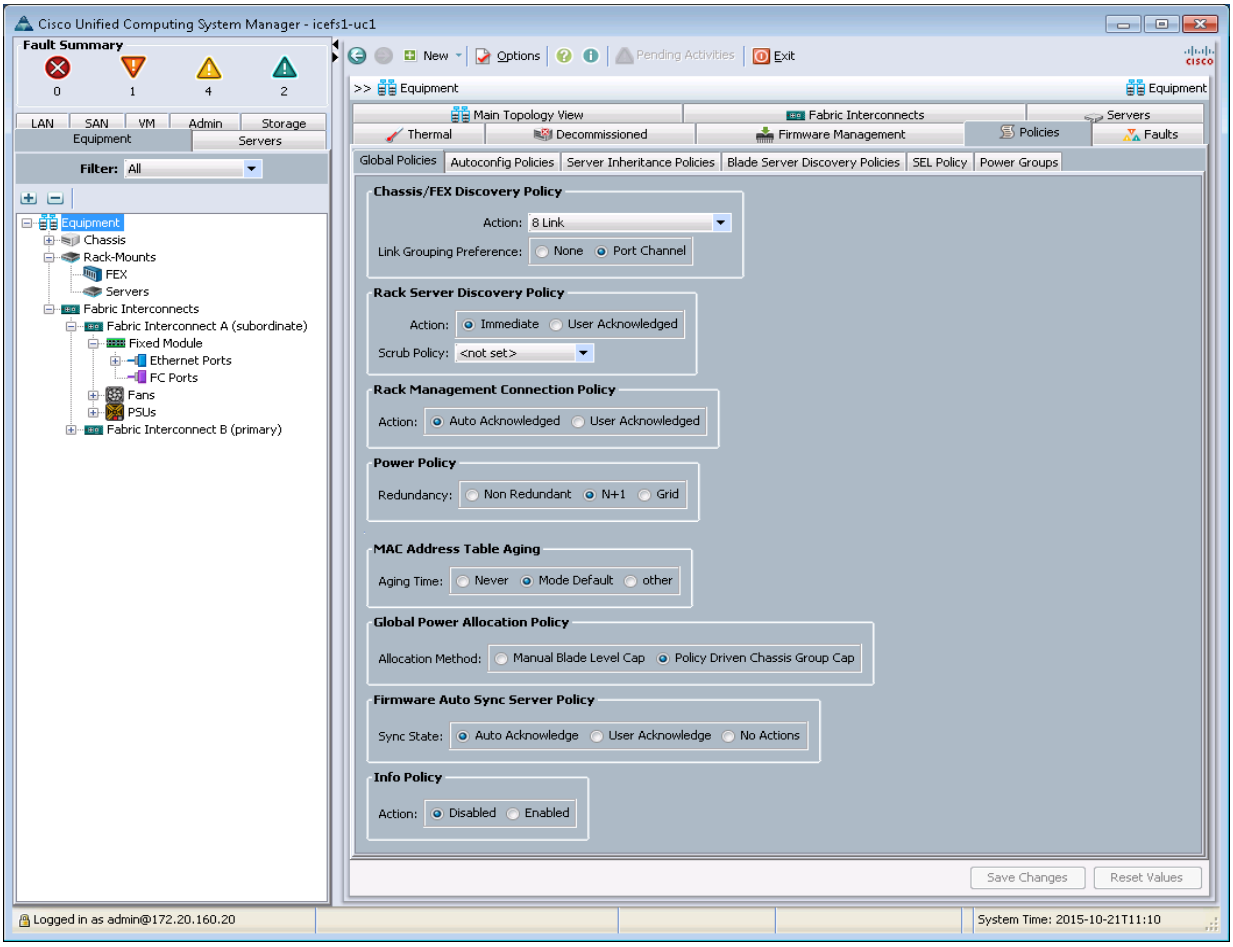

5. Click Save Changes.
6. Click OK.

Enable Server and Uplink Ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports 1 through 16 that are connected to the chassis, right-click them, and select Configure as Server Port.
5. Click Yes to confirm server ports and click OK.
6. Select ports 31 and 32 that are connected to the Cisco Nexus 5548 switches, right-click them, and select Configure as Uplink Ports. These ports carry non-SAN Ethernet traffic between the switches and the fabric interconnect ports.
7. Click Yes to confirm uplink ports and click OK.
8. Select ports 17 through 30, which serve as FCoE uplinks to the Cisco Nexus 5548 switch 1. Right-click them and select Configure as FCoE Uplink Port.
9. Click Yes to confirm FCoE uplink ports and click OK.
10. In the left pane, navigate to Fabric Interconnect A. In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that the ports in the If Role column have been configured correctly.

    Note: You might have to scroll to the top of the main pane to see the Physical Ports tab.

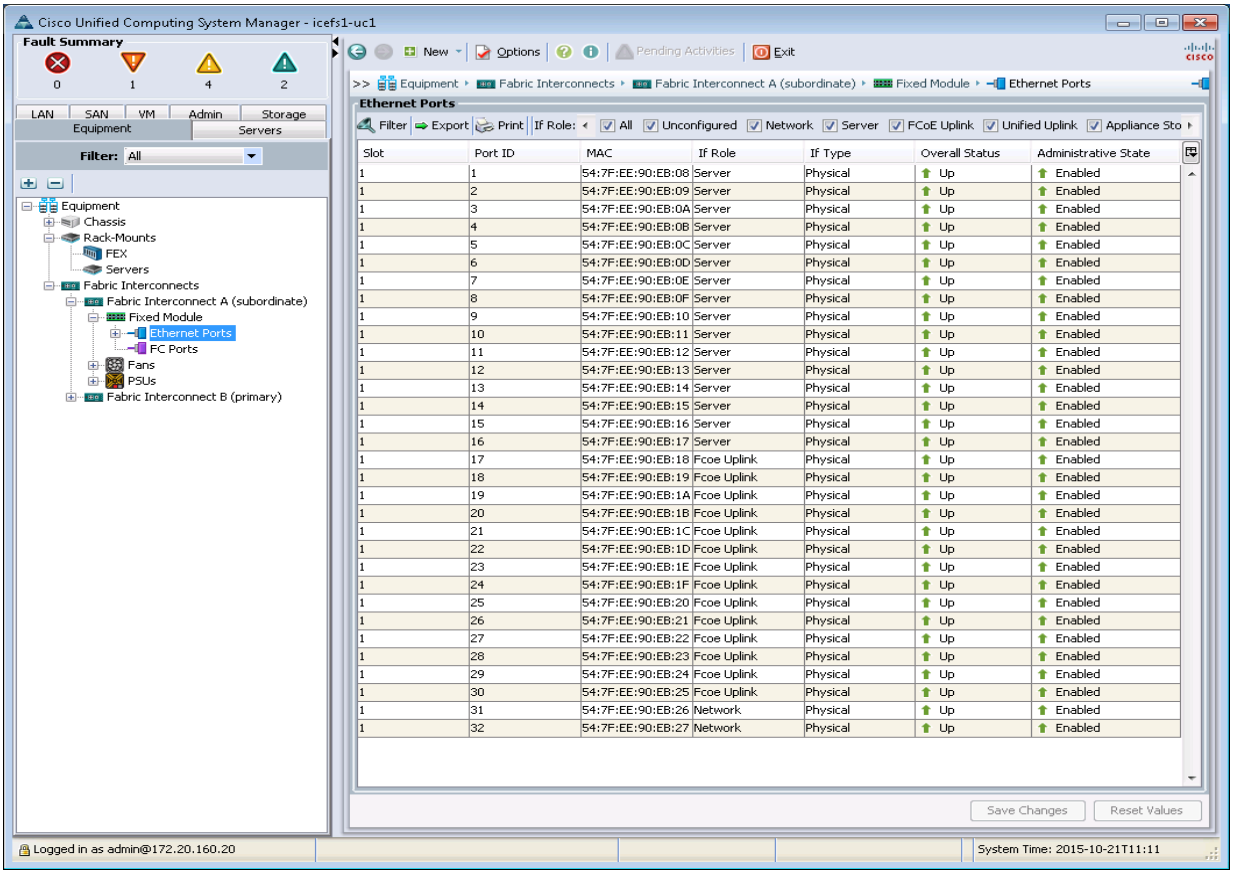

11. Repeat the preceding steps for Fabric Interconnect B (subordinate).
12. In the left pane, navigate to Fabric Interconnect B (subordinate). In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that the ports in the If Role column have been configured correctly.

    Note: You might have to scroll to the top of the main pane to see the Physical Ports tab.

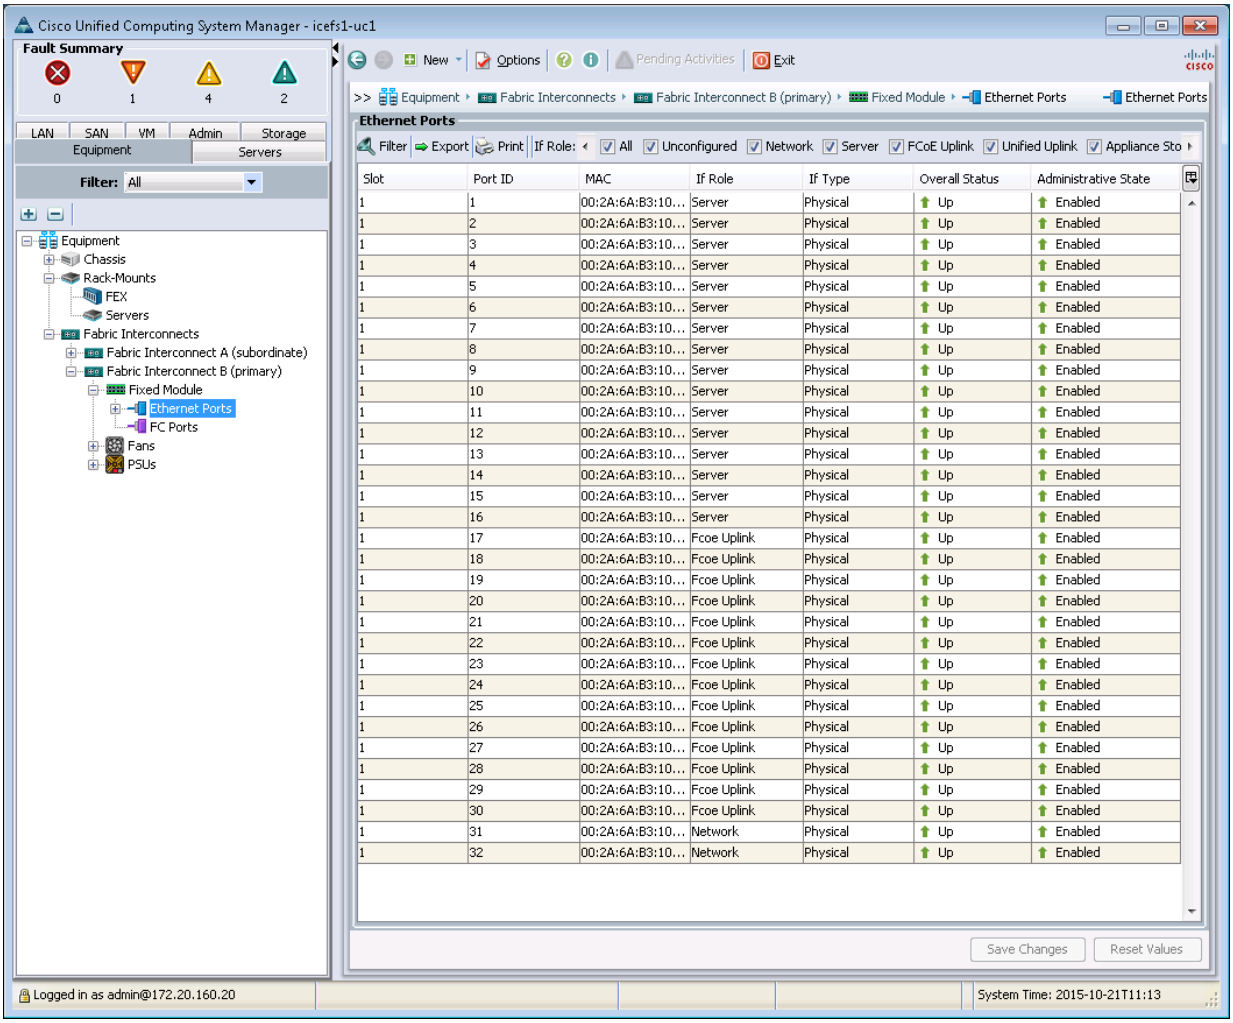

Acknowledge Cisco UCS Chassis

To acknowledge all Cisco UCS chassis, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.

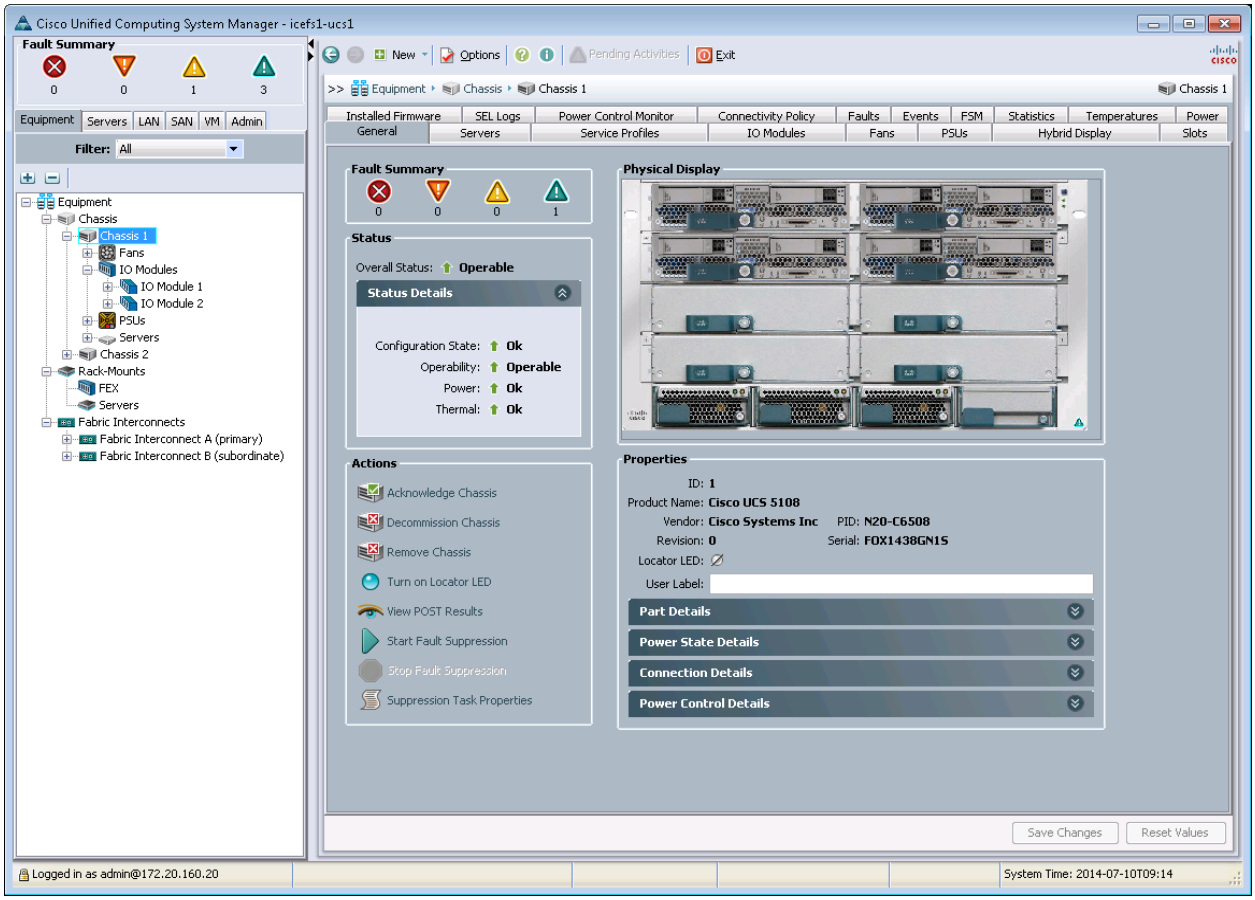

4. Click Yes and then click OK to complete acknowledging the chassis.

Create Uplink Port Channels to Cisco Nexus 5548UP Switches

To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

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

   Note: In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 5548UP switches and one from Fabric B to both Cisco Nexus 5548UP switches.

2. Under LAN > LAN Cloud, expand the Fabric A node.

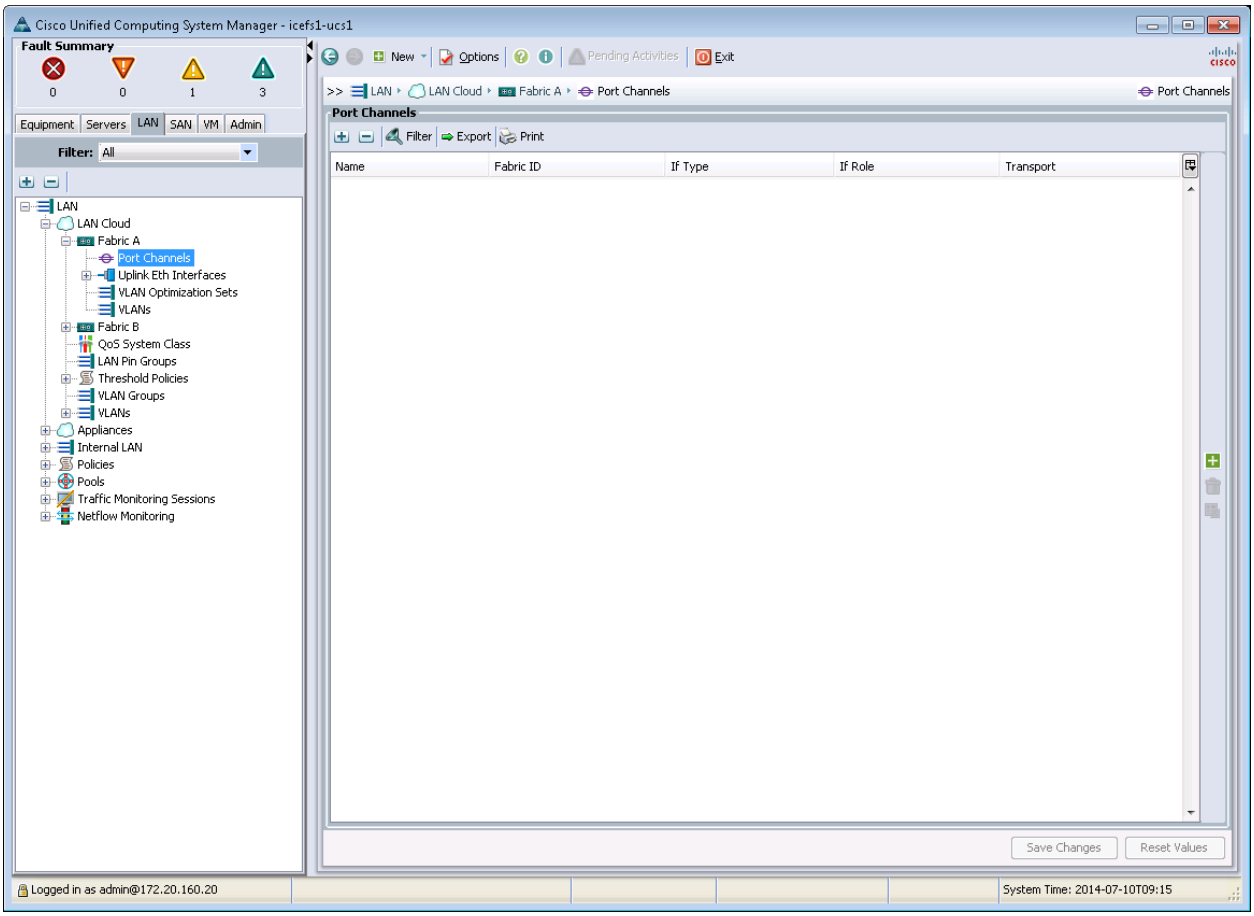

3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter `15` as the unique ID of the port channel.
6. Enter `vPC-15` as the name for the port channel.

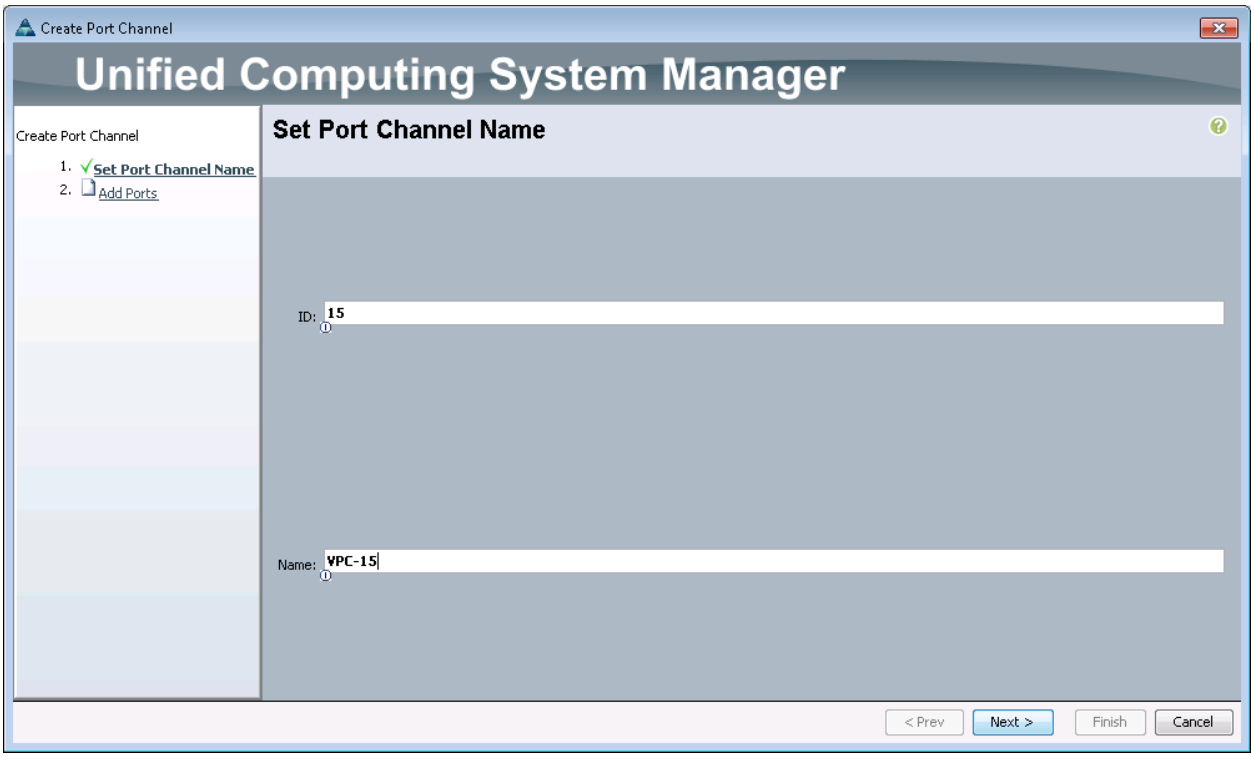

7. Click Next.
8. Select the following ports to be added to the port channel:
   - Slot ID 1 and port 31
   - Slot ID 1 and port 32
9. Click the double right-arrow button (>>) to add the ports to the port channel.

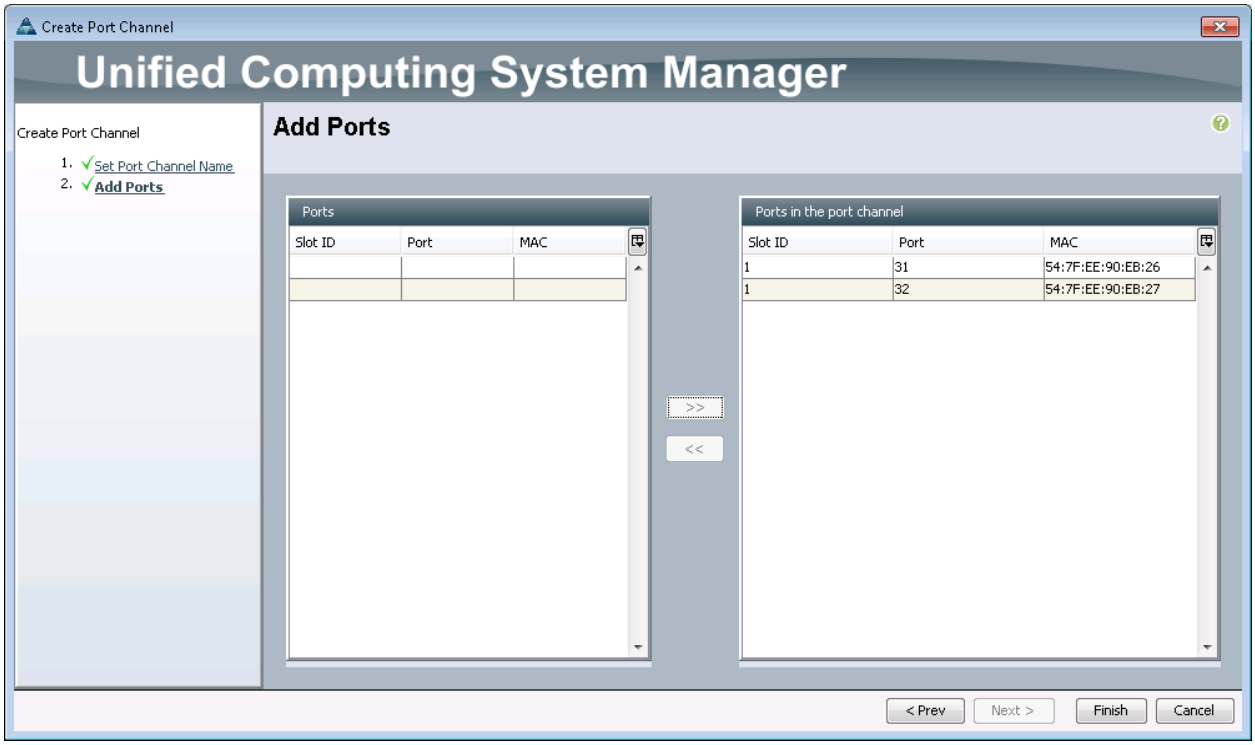

10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B node.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter `16` as the unique ID of the port channel.
16. Enter `vPC-16` as the name for the port channel.

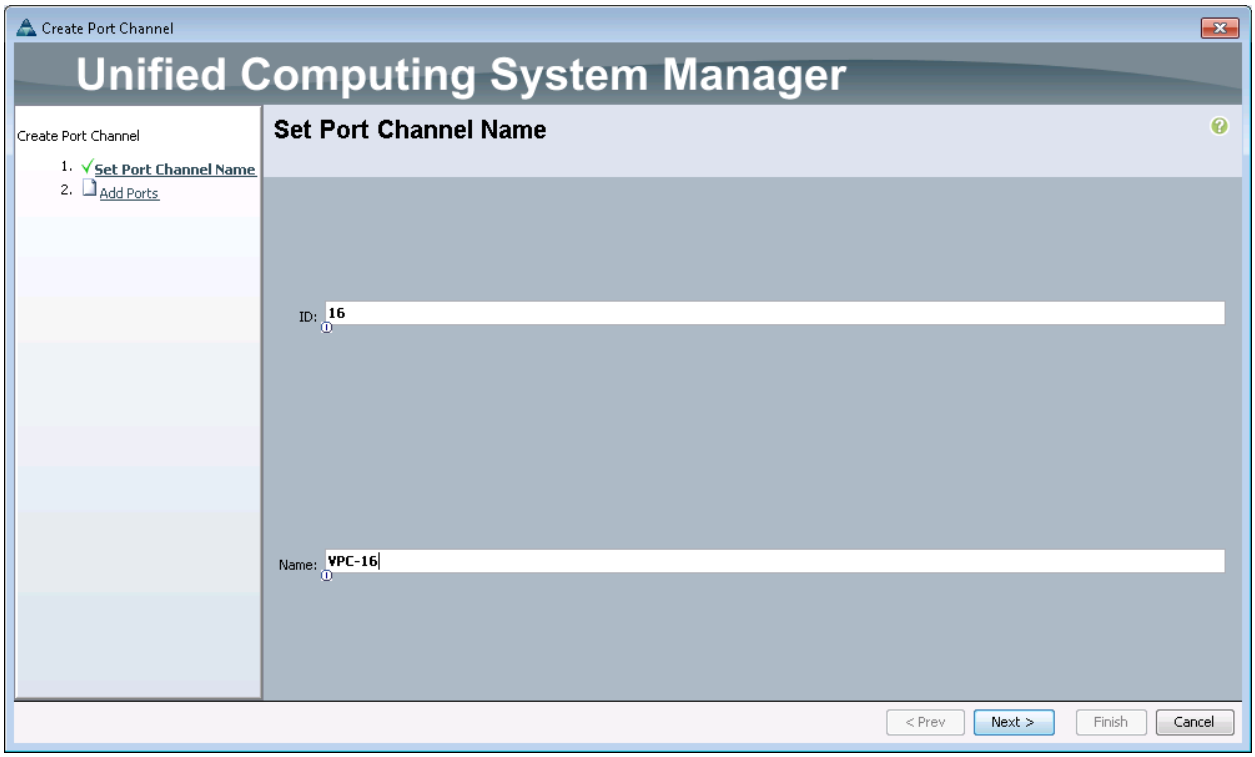

17. Click Next.
18. Select the following ports to be added to the port channel:
    - Slot ID 1 and port 31
    - Slot ID 1 and port 32
19. Click the double left-arrow button (>>) to add the ports to the port channel.

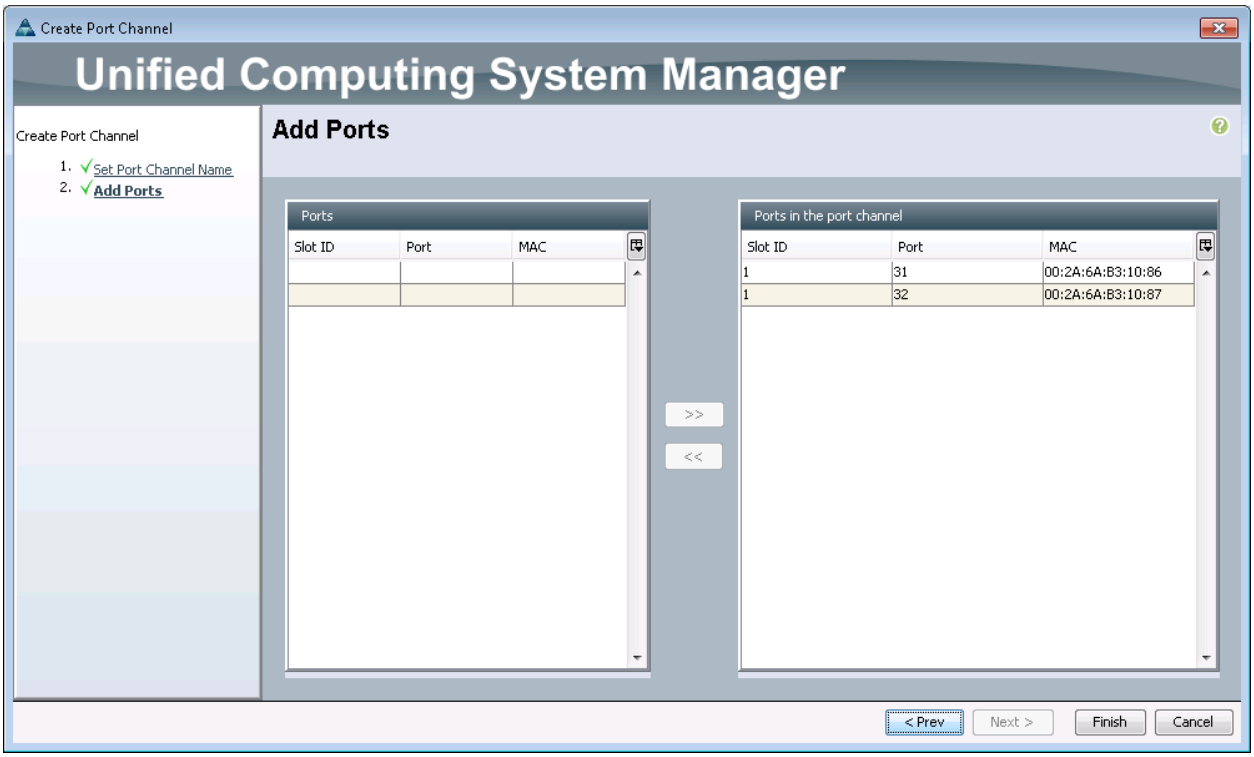

20. Click Finish to create the port channel.
21. Click OK.

Create an Organization

Organizations are used to organize resources and restrict access to various groups within the IT organization, thereby enabling multi-tenancy of the compute resources.

Note: Although this document does not assume the use of organizations, this procedure provides instructions for creating one.

To configure an organization in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Optional: Enter a description for the organization.
4. Click OK.
5. Click OK in the confirmation message.

Create MAC Address Pools

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > Root.

   Note: In this procedure, two MAC address pools are created: one for each switching fabric.
3. Right-click MAC Pools under the root organization.

4. Select Create MAC Pool to create the MAC address pool.
5. Enter `MAC-POOL-A` as the name for the MAC pool.
6. Optional: Enter a description for the MAC pool.

   Note: Keep the Assignment Order value as Default.

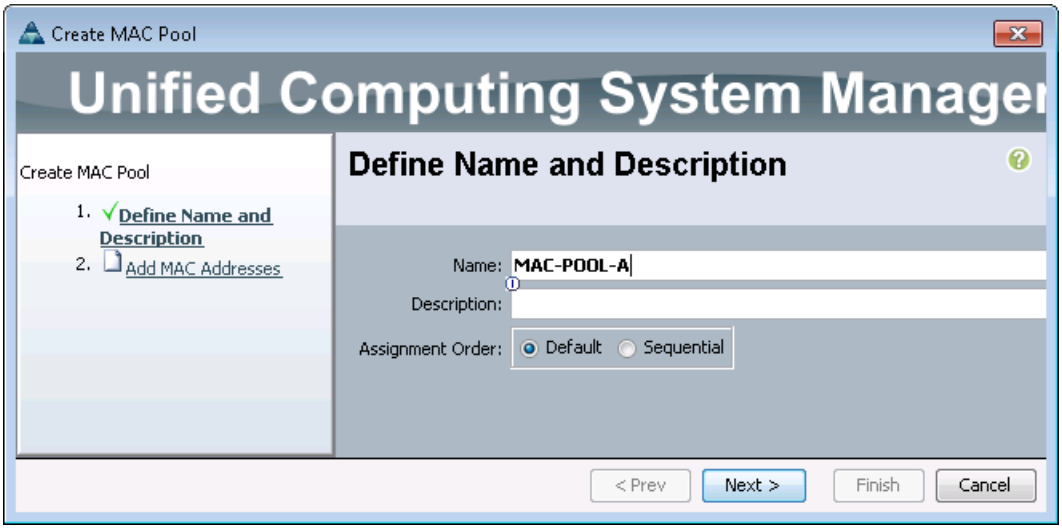

7. Click Next.

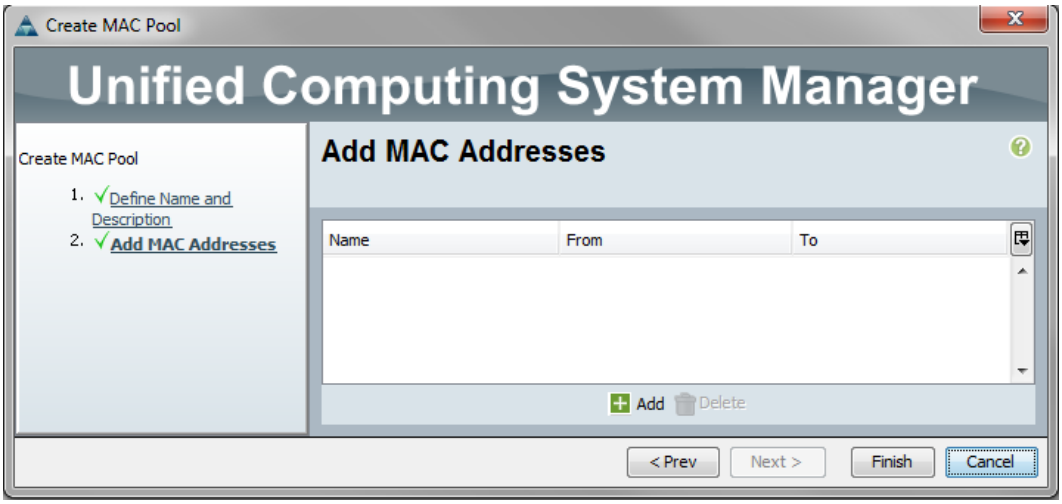

8. Click Add.
9. Specify a starting MAC address.

   Note: For the FlexPod solution, the NetApp recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as Fabric A addresses.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.

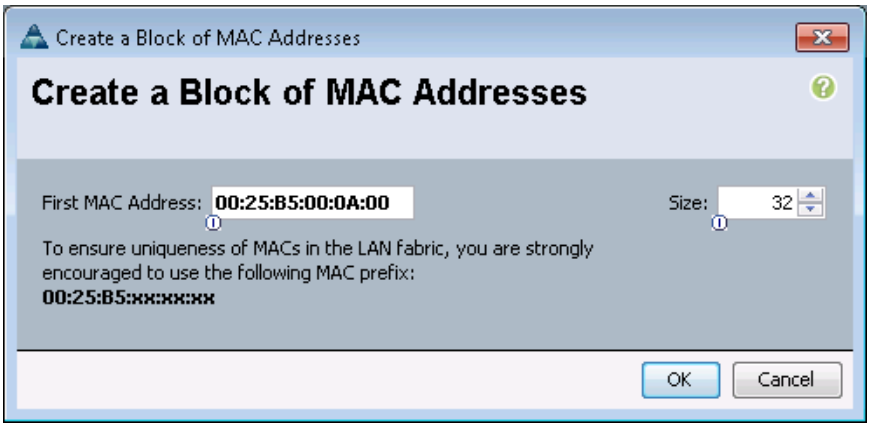

11. Click OK.
12. Click Finish.
13. In the confirmation message, click OK.
14. Right-click MAC Pools under the root organization.
15. Select Create MAC Pool to create the MAC address pool.
16. Enter `MAC-POOL-B` as the name for the MAC pool.
17. Optional: Enter a description for the MAC pool.

    Note: Select Default for the Assignment Order.

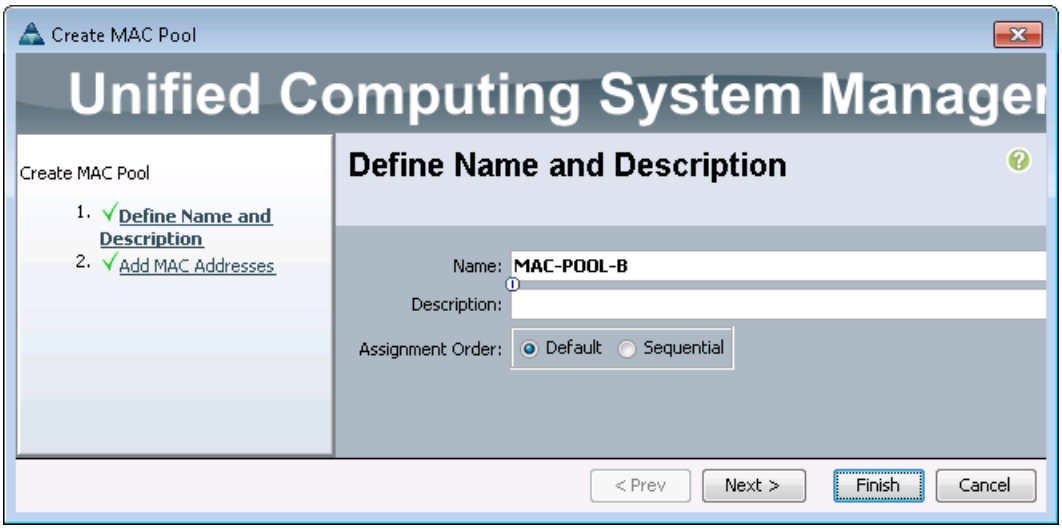

18. Click Next.

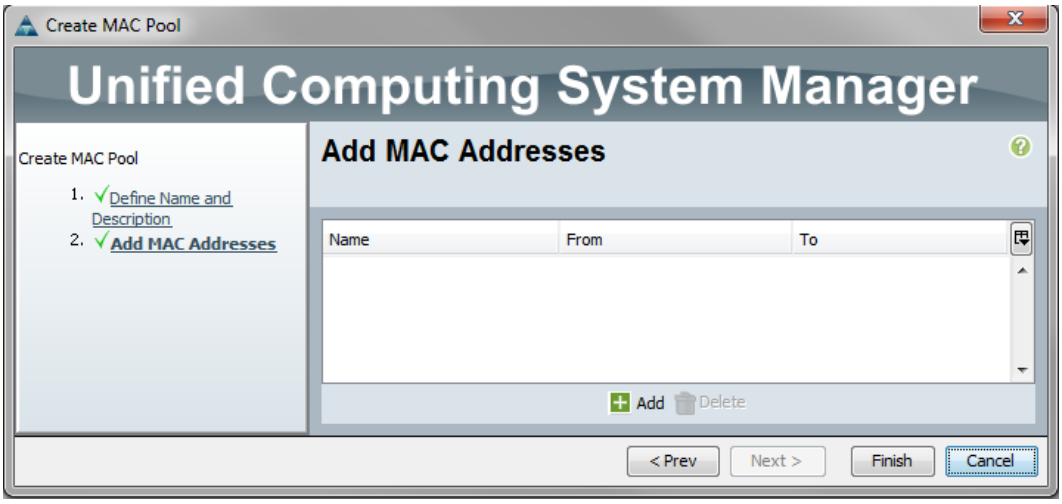

19. Click Add.
20. Specify a starting MAC address.

    Note: For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting MAC address to identify all the MAC addresses in this pool as Fabric B addresses.

21. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.

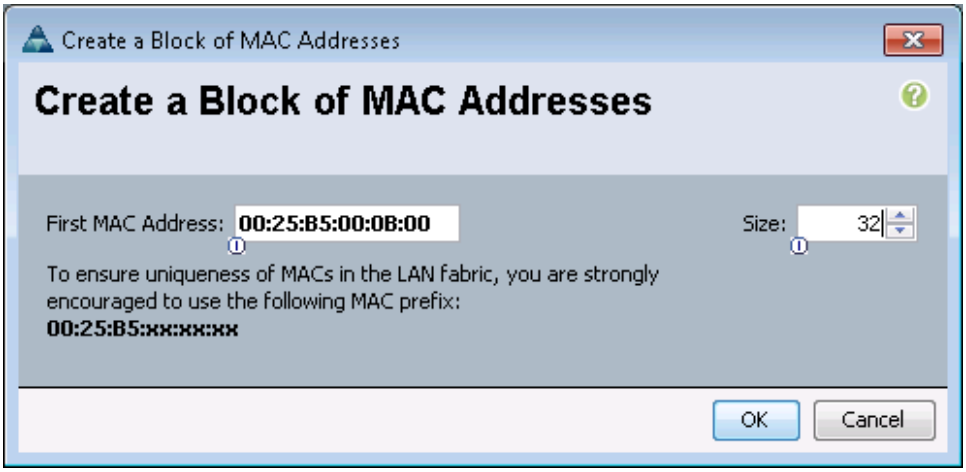

22. Click OK.

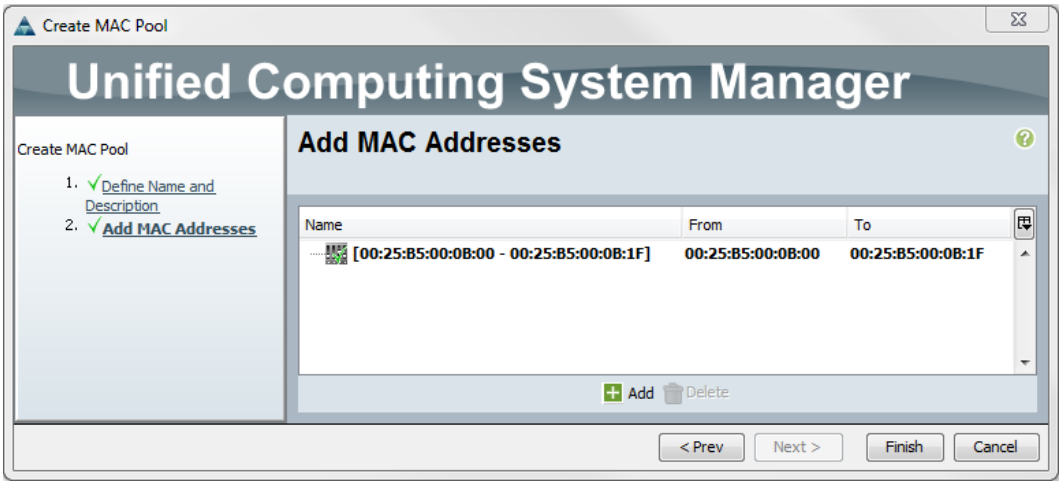

23. Click Finish.
24. In the confirmation message, click OK.

Create WWNN Pools

To configure the necessary worldwide node name (WWNN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN node in the navigation pane.
2. Select Pools > Root.

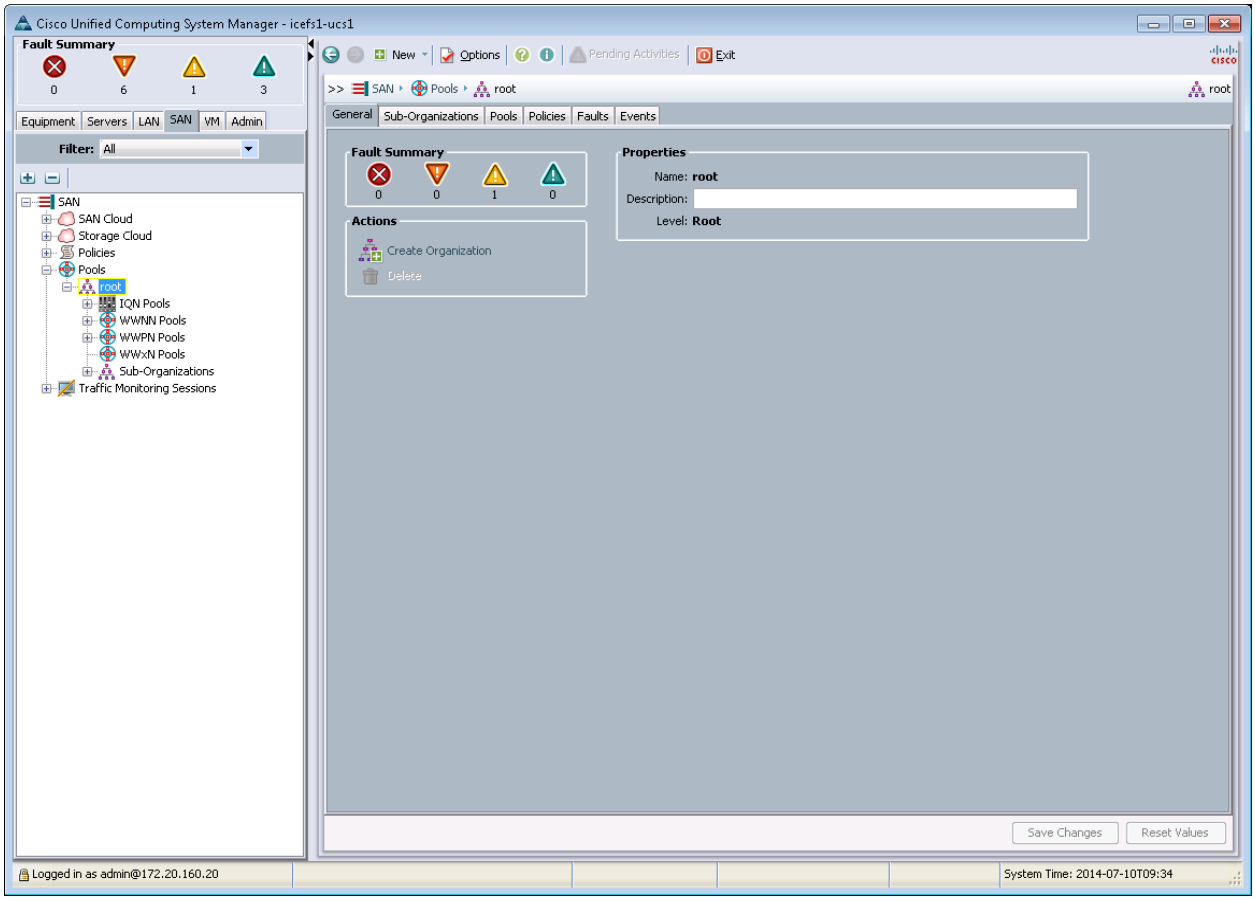

3. Right-click WWNN Pools.
4. Select Create WWNN Pool.
5. Enter `WWNN-POOL` as the name for WWNN pool.
6. Optional: Add a description for the WWNN pool.
7. Select Default for the Assignment Order.

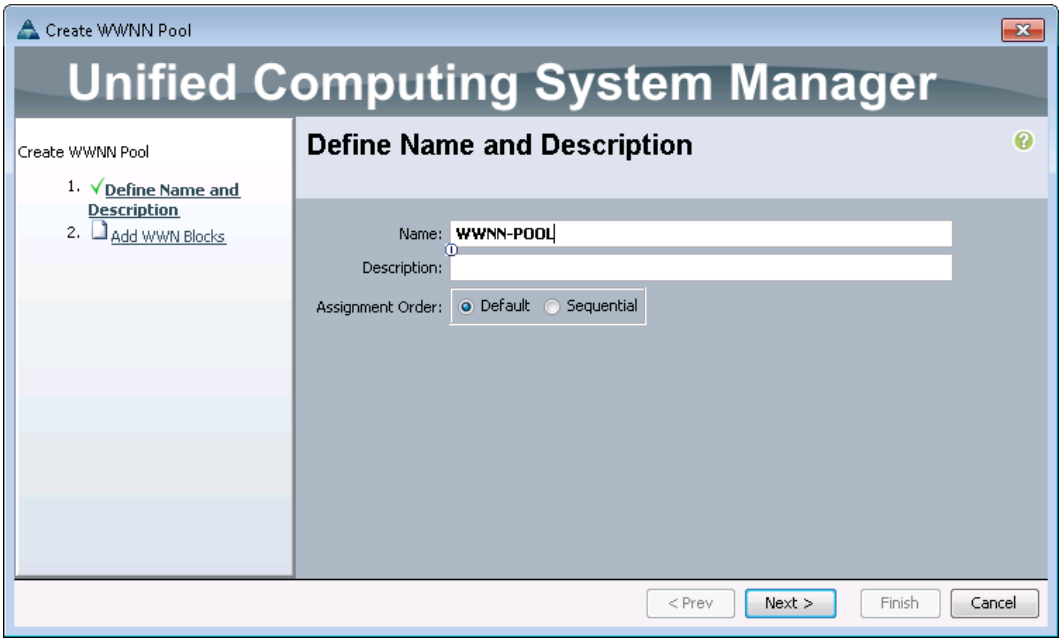

8. Click Next.

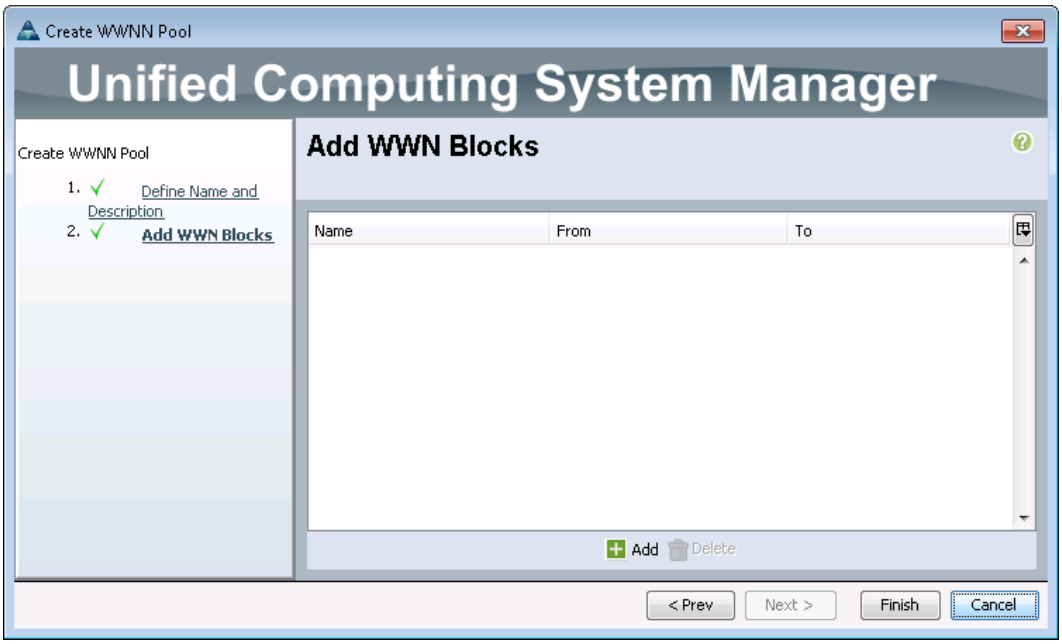

9. Click Add to add a block of WWNNs.
10. Either retain the default block of WWNNs or specify a base WWNN.
11. Specify a size for the WWNN block that is sufficient to support the available blade or server resources.

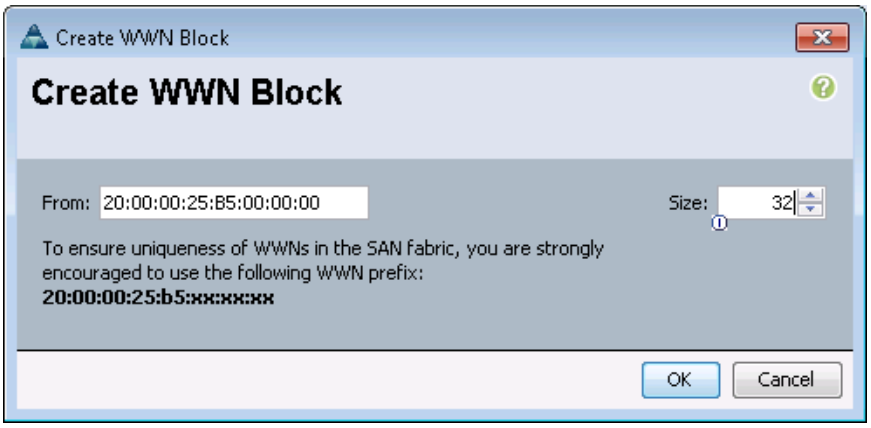

12. Click OK.

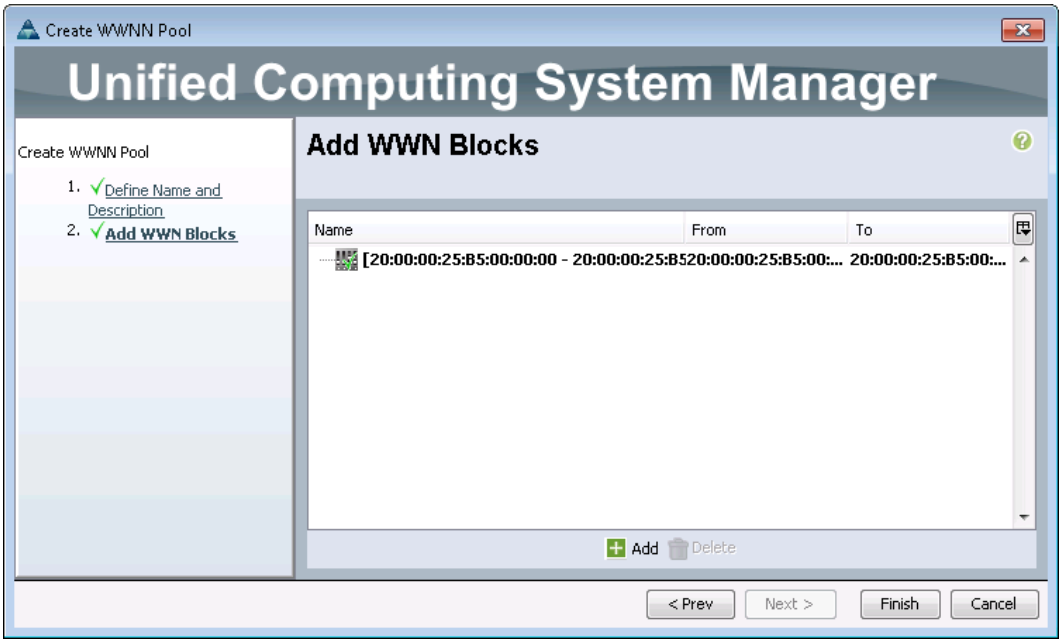

13. Click Finish.
14. Click OK.

Create WWPN Pools

To configure the necessary worldwide port name (WWPN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > Root.

   Note: In this procedure, two WWPN pools are created: one for Fabric A and one for Fabric B.
3. Right-click WWPN Pools.
4. Select Create WWPN Pool.
5. Enter `WWPN-POOL-A` as the name for WWPN pool for Fabric A.
6. Optional: Enter a description for this WWPN pool.

7. Select Default for Assignment Order.

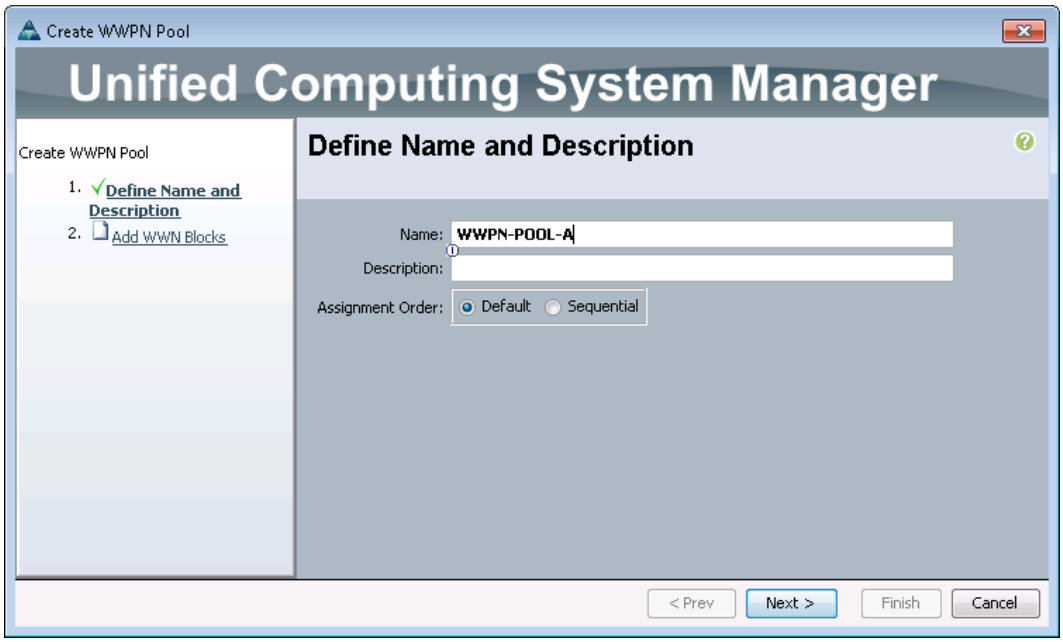

8. Click Next.

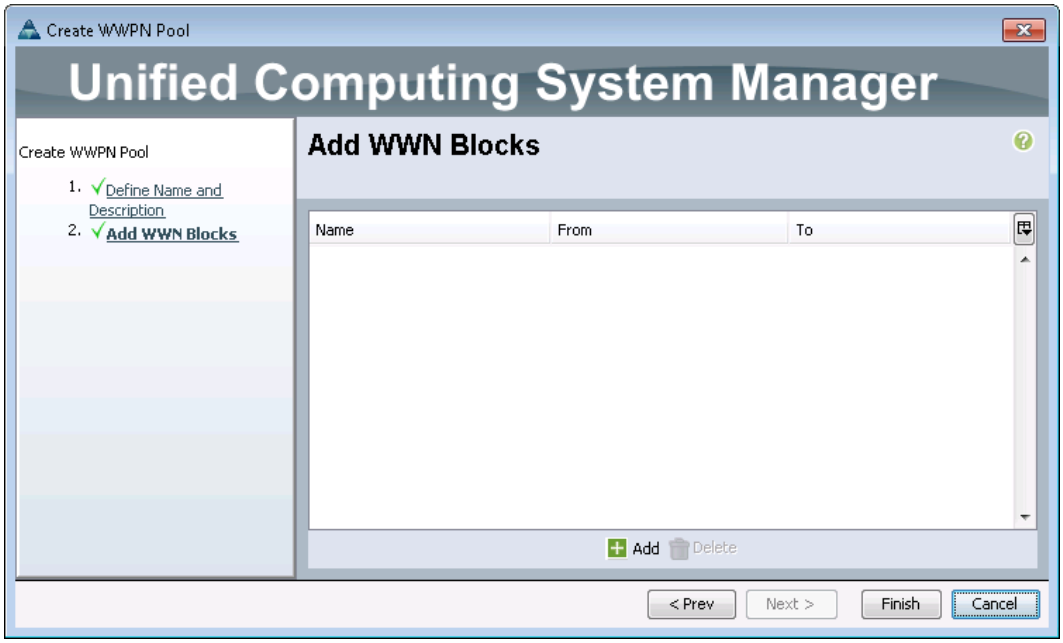

9. Click Add to add a block of WWPNs.
10. Specify the starting WWPN in the block for Fabric A.

    Note: For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting WWPN to identify all the WWPNs in this pool as Fabric A addresses.

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

12. Click OK.

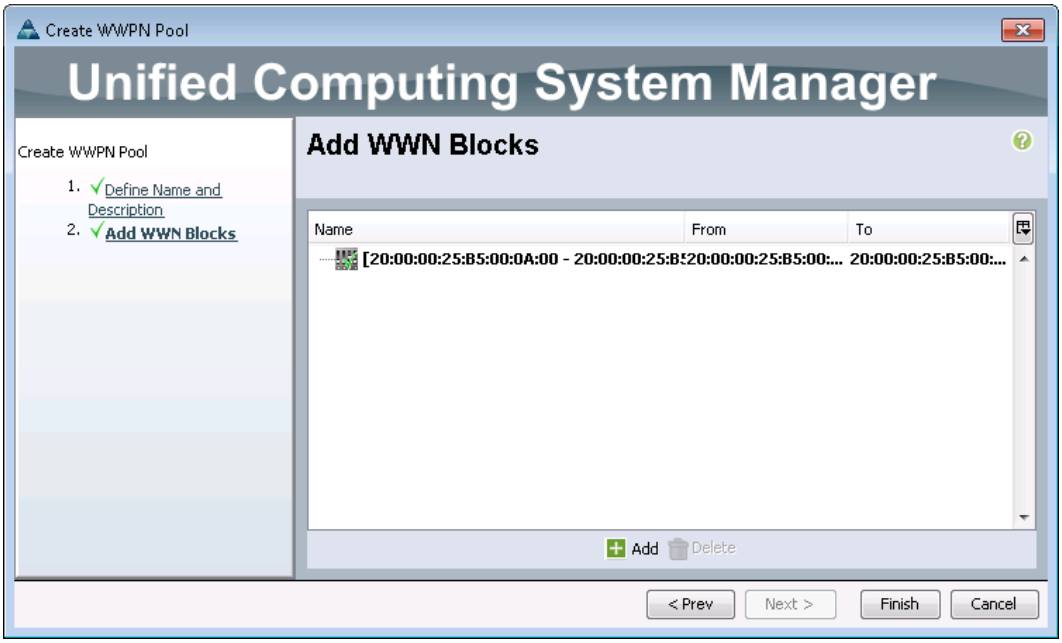

13. Click Finish to create the WWPN pool.
14. Click OK.
15. Right-click WWPN Pools.
16. Select Create WWPN Pool.
17. Enter `WWPN-POOL-B` as the name for the WWPN pool for Fabric B.
18. Optional: Enter a description for this WWPN pool.
19. Select Default for the Assignment Order.

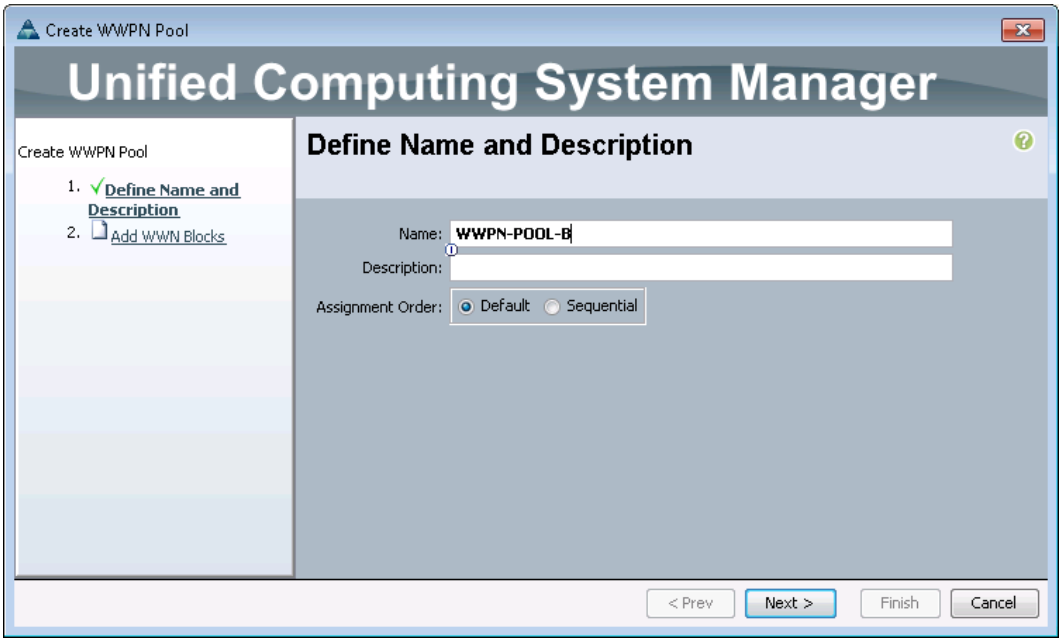

20. Click Next.

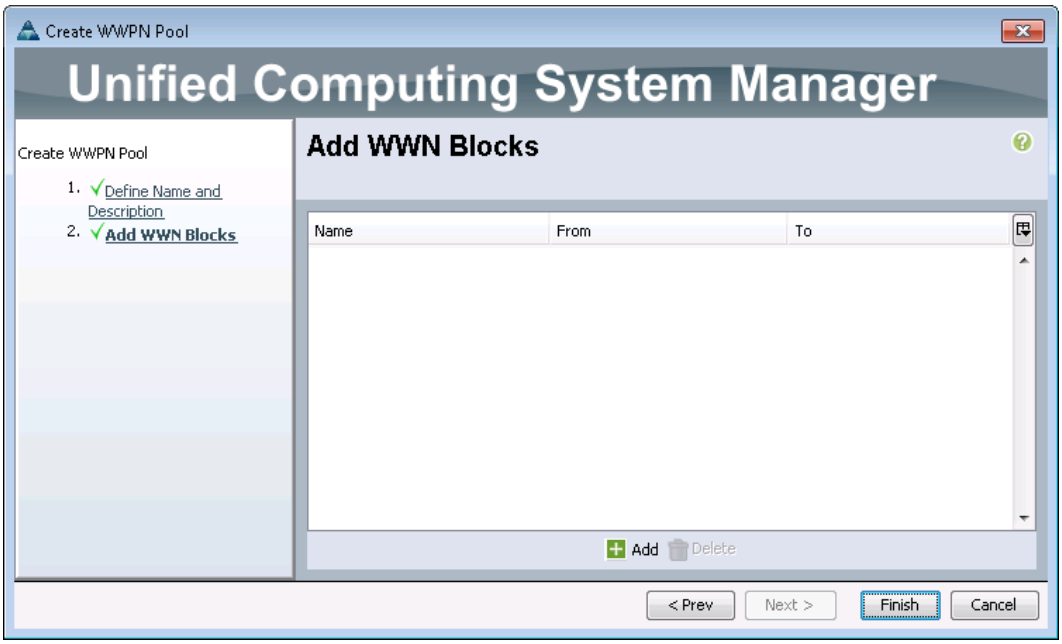

21. Click Add to add a block of WWPNs.
22. Enter the starting WWPN address in the block for Fabric B.

    Note: For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting WWPN to identify all the WWPNs in this pool as Fabric B addresses.

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

24. Click OK.

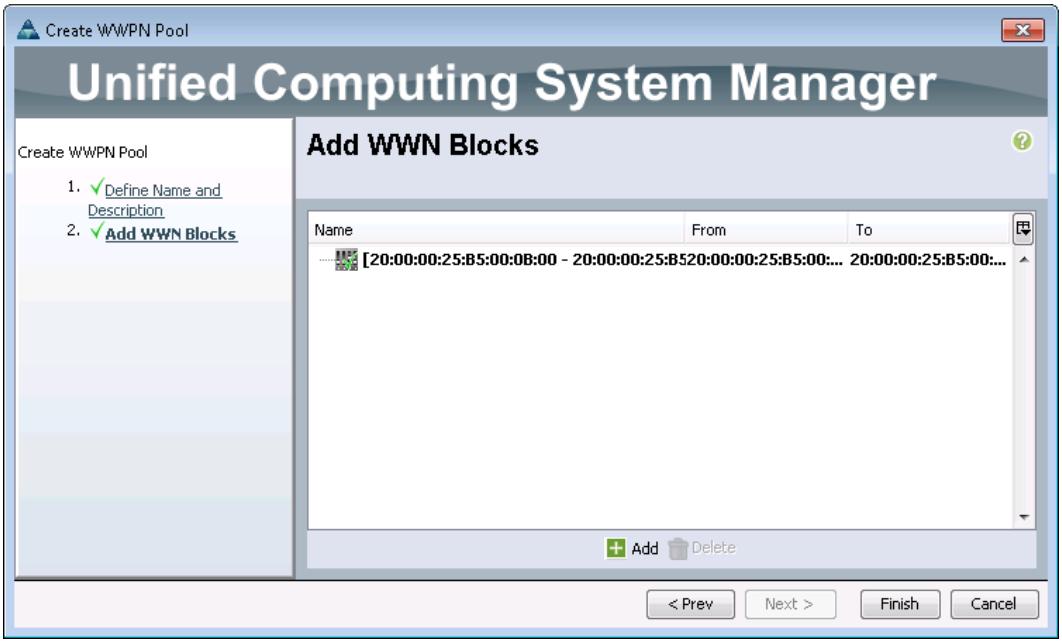

25. Click Finish.
26. Click OK.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > Root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID-POOL` as the name for the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Select the Derived option for Prefix.

8. Select Default for the Assignment Order.

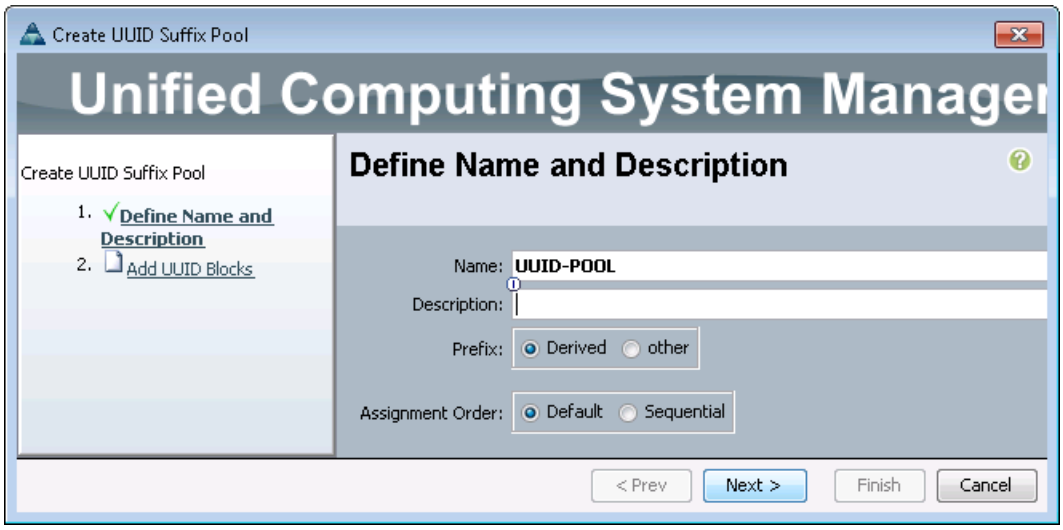

9. Click Next.

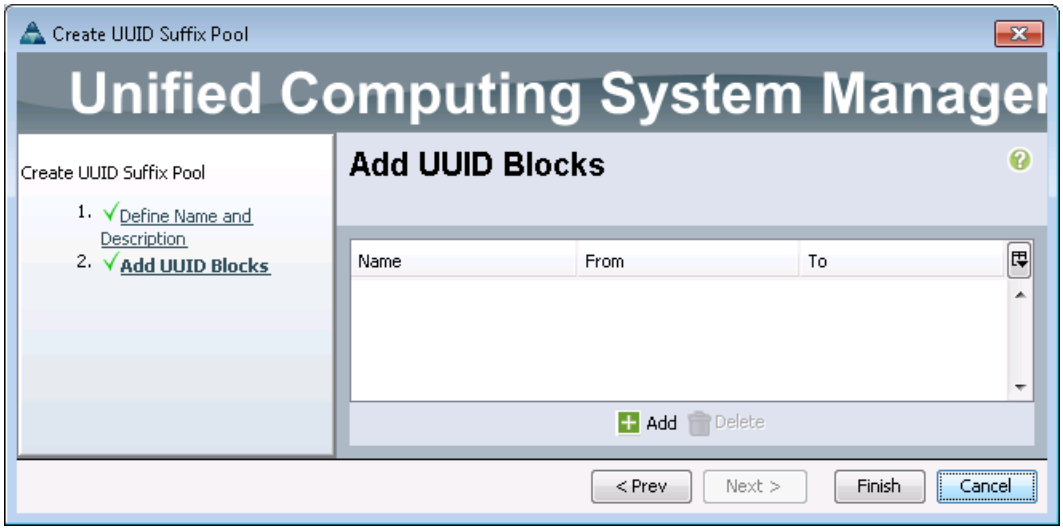

10. Click Add to add a block of UUIDs.
11. Select the From option as the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

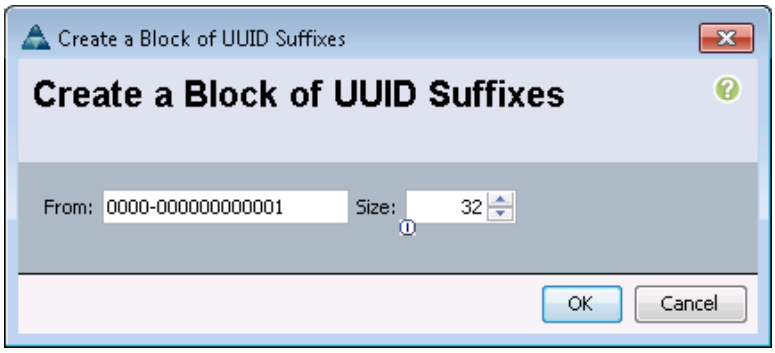

13. Click OK.

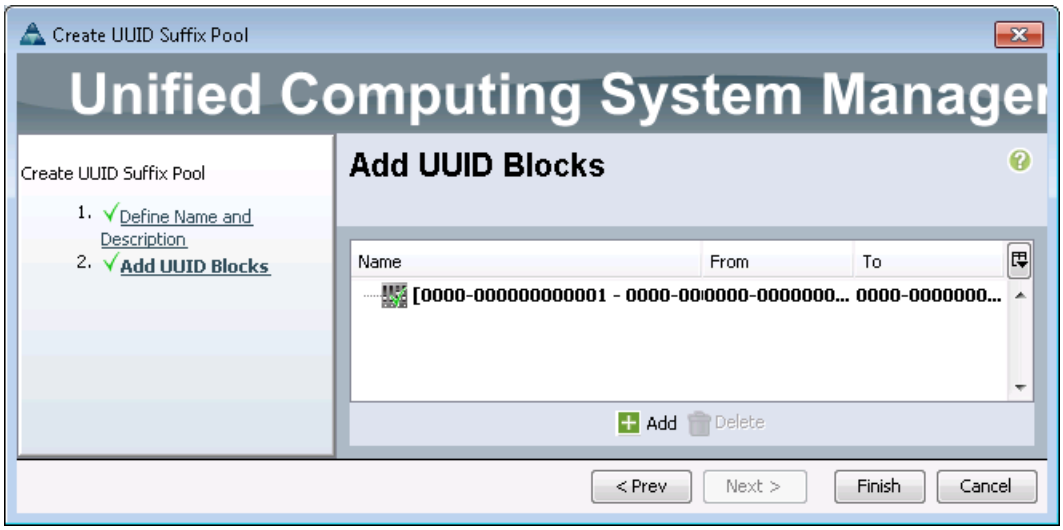

14. Click Finish.
15. Click OK.

Create Server Pool

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > Root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `ORACLE-POOL` as the name for server pool.
6. Optional: Enter a description for the server pool.

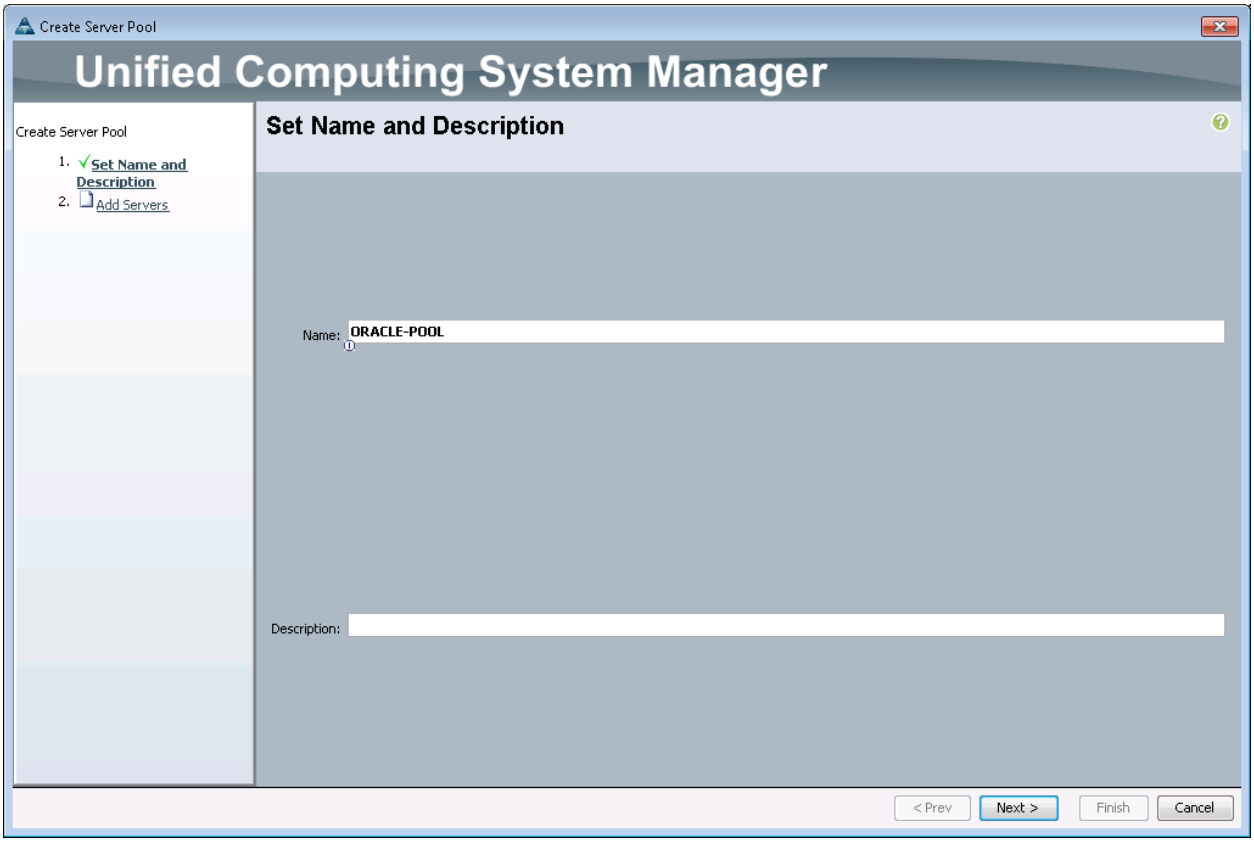

7. Click Next.
8. Select eight servers to be used for the Oracle RAC cluster and click the double right-arrow button (>>) to add them to the `ORACLE-POOL` server pool.

9. Click Finish.
10. Click OK.

Create VLANs

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

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

   Note: In this procedure, five VLANs are created.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `MGMT-VLAN` as the name for the VLAN to be used for management traffic.
6. Retain the Common/Global option selected for the scope of the VLAN.
7. Enter the ID of the management VLAN.
8. Retain the Sharing Type as None.

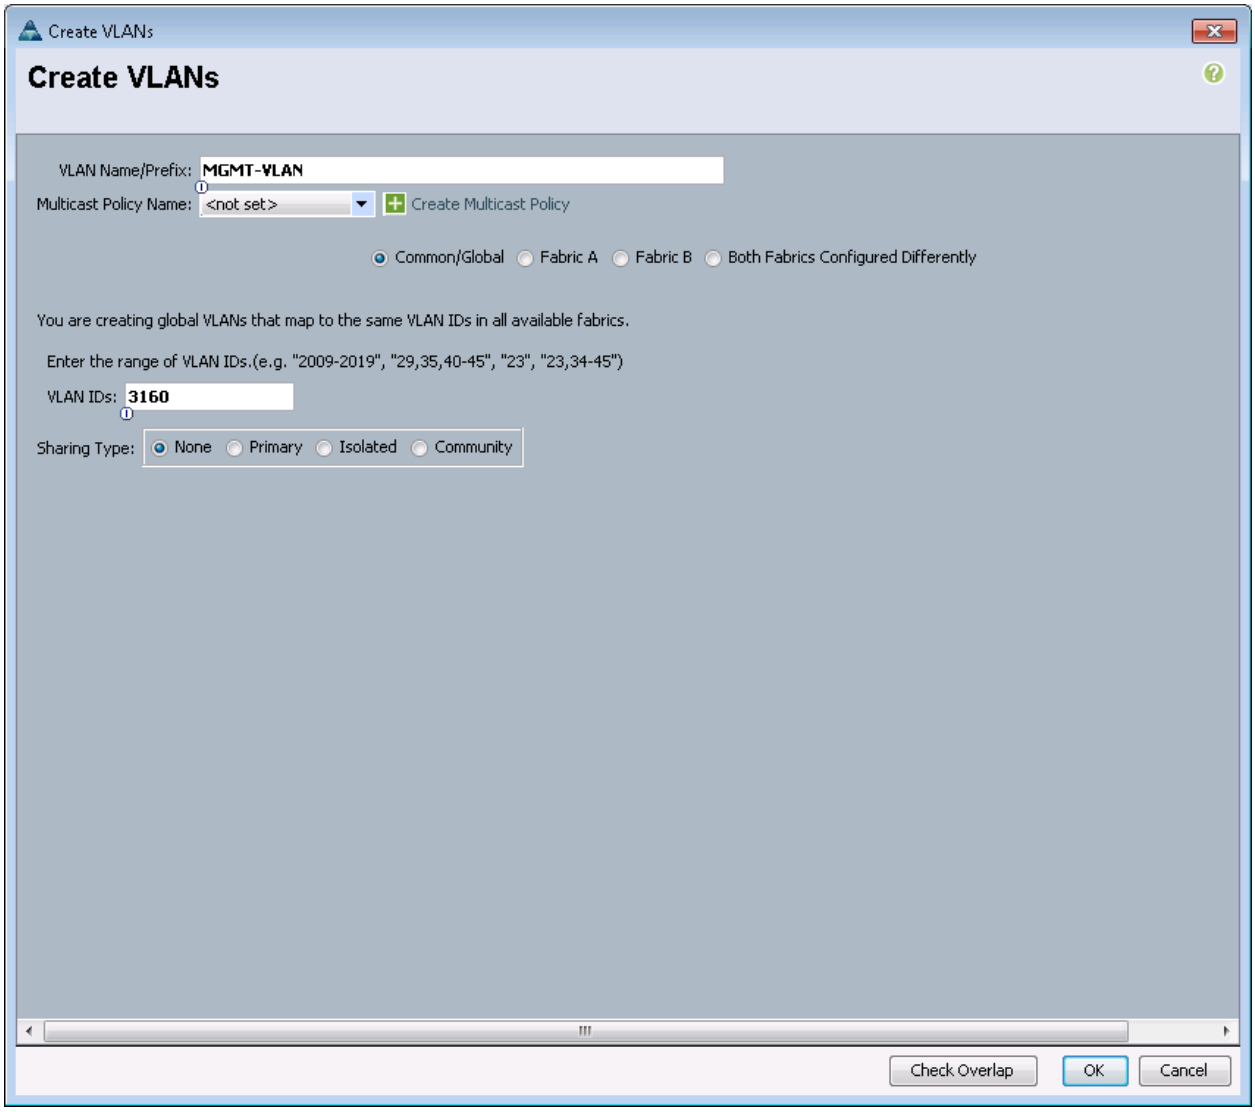

9. Click OK and then click OK again.
10. Right-click VLANs.
11. Select Create VLANs.
12. Enter `RAC-VIP-VLAN` as the name for the VLAN to be used for RAC intercluster traffic.
13. Retain the Common/Global option selected for the scope of the VLAN.
14. Enter the ID for the VLAN.
15. Retain the Sharing Type as None.

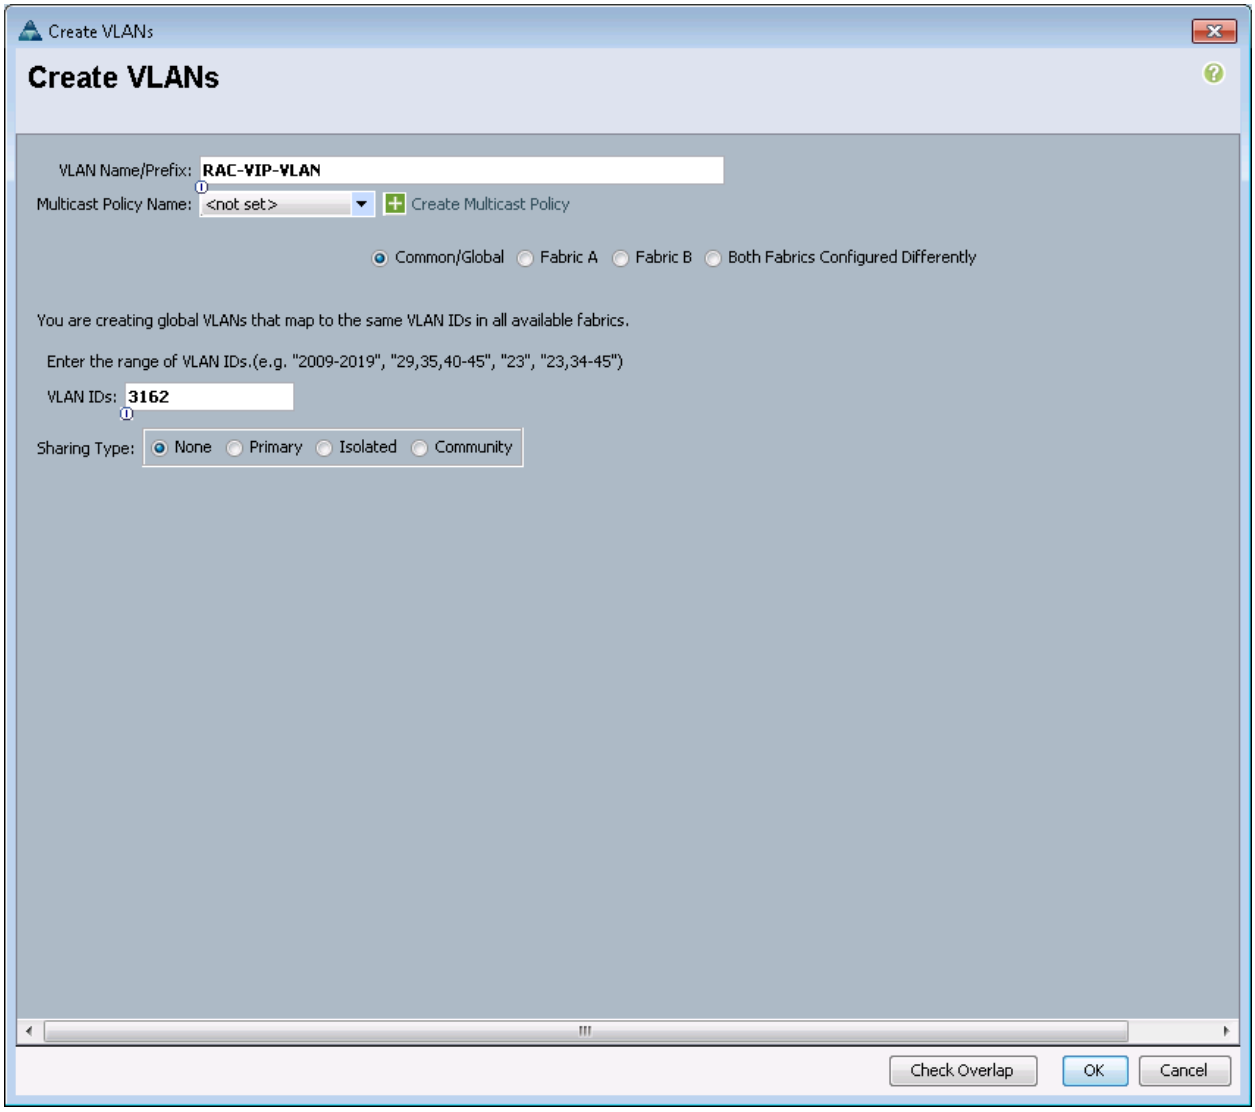

16. Click OK and then click OK again.

Create VSANs and FCoE Port Channels

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Expand the SAN > SAN Cloud tree.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter `VSAN-A` as the name for the VSAN for Fabric A.
6. Keep the Disabled option selected for FC Zoning.
7. Select Fabric A.
8. Enter the VSAN ID for Fabric A.
9. Enter the FCoE VLAN ID for Fabric A.

Note: For the FlexPod solution, it is recommended to use the same ID for the VSAN and the FCoE VLAN required for Fabric A.

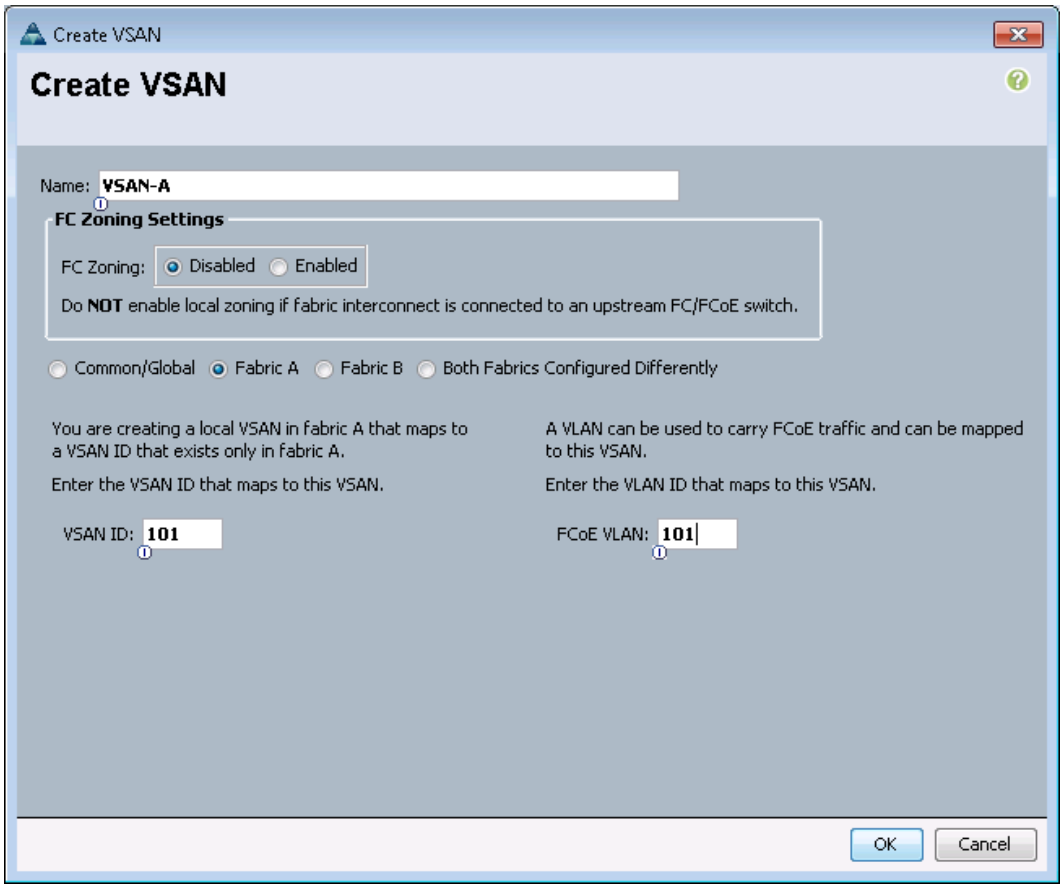

10. Click OK and then click OK again to create the VSAN.
11. Right-click VSANs.
12. Select Create VSAN.
13. Enter `VSAN-B` as the name for the VSAN for Fabric B.
14. Retain the Disabled option selected for FC Zoning.
15. Select Fabric B.
16. Enter the VSAN ID for Fabric B.
17. Enter the FCoE VLAN ID for Fabric B.

Note: NetApp recommends using the same ID for the VSAN and the FCoE VLAN required for Fabric B.

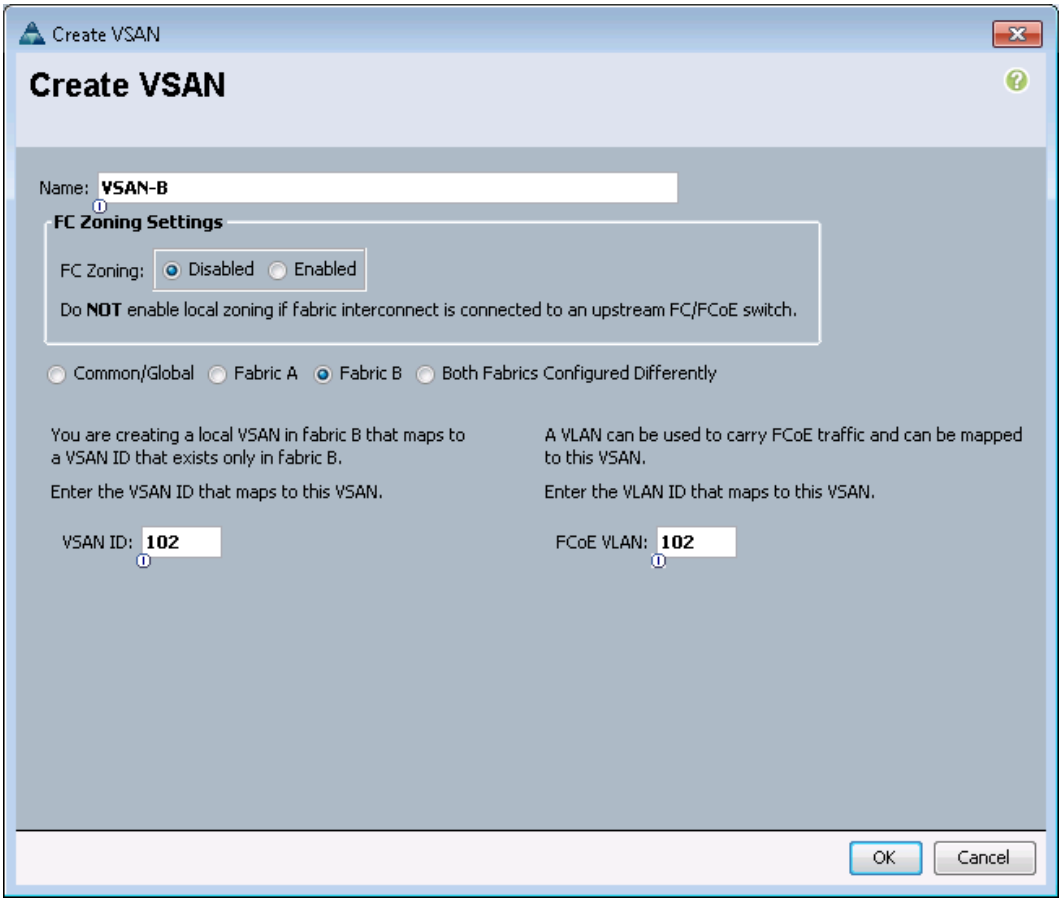

18. Click OK and then click OK again to create the VSAN.
19. In the navigation pane, under SAN > SAN Cloud, expand the Fabric A tree.
20. Right-click FCoE Port Channels.
21. Select Create FCoE Port Channel.
22. Enter `101` for the port channel ID and `Po101` for the port channel name.

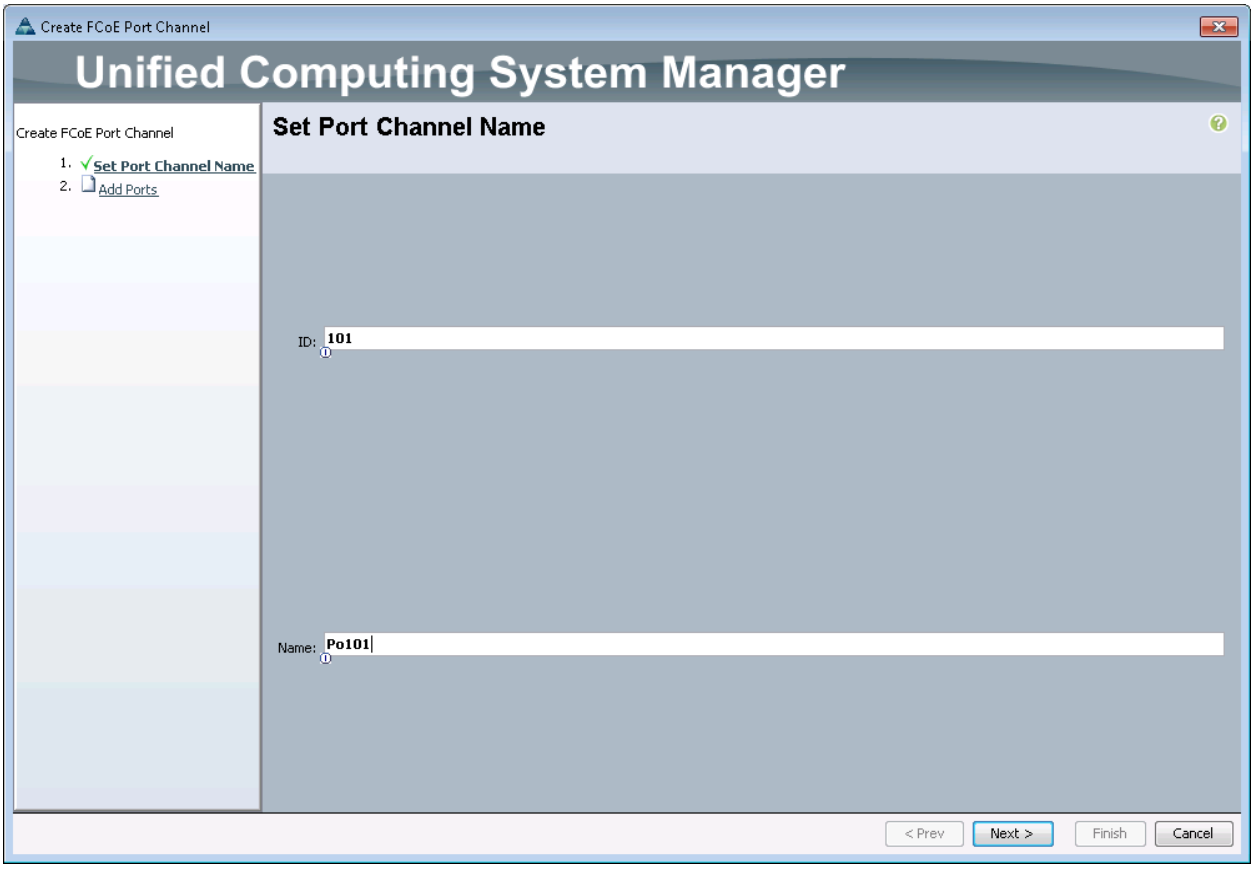

23. Click Next.
24. Select ports 17 through 30 and click the double right-arrow button (>>) to add the ports to the port channel.

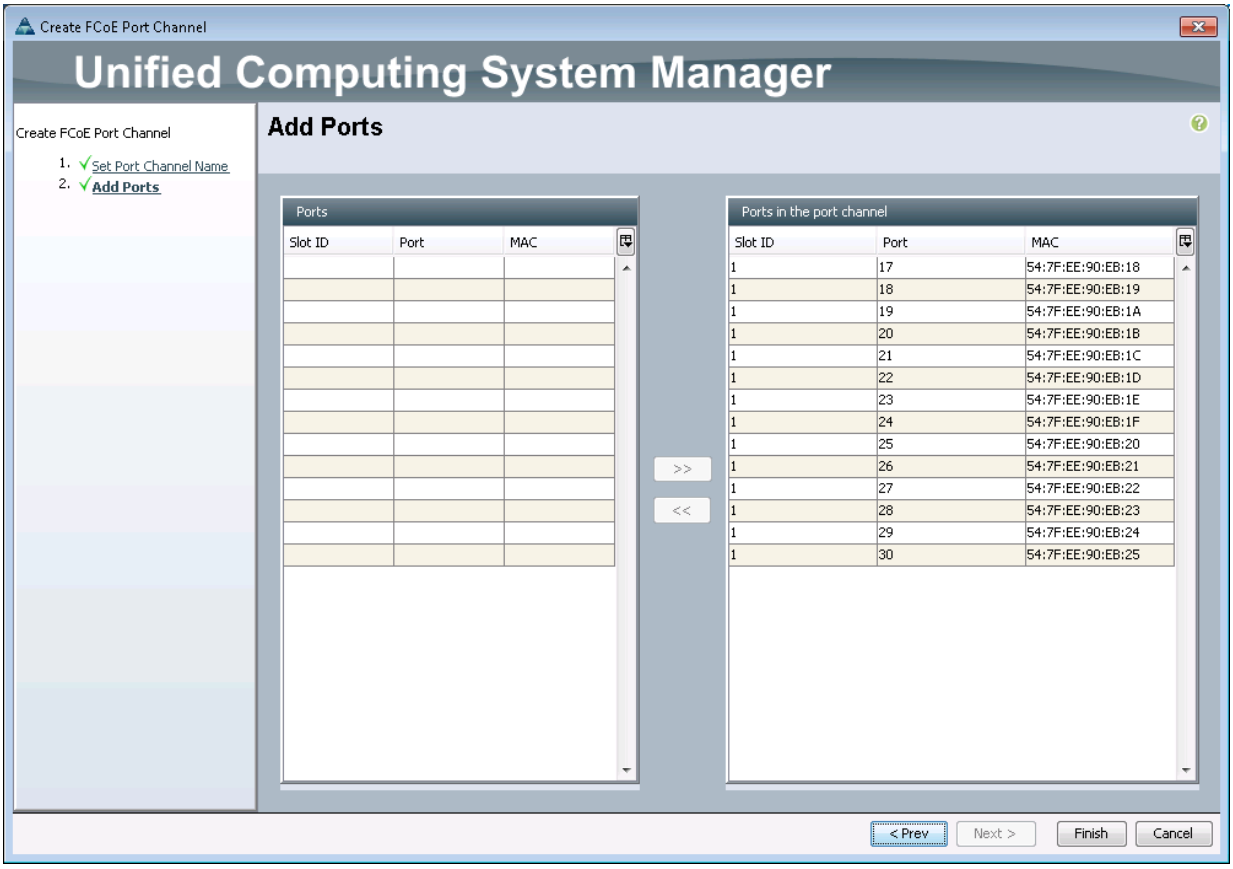

25. Click Finish.
26. Select the checkbox for Show Navigator for FCoE Port-Channel 101 (Fabric A).
27. Click OK to create the port channel.
28. Click OK to close the navigator.
29. In the navigation pane, under SAN > SAN Cloud, expand the Fabric B tree.
30. Right-click FCoE Port Channels.
31. Select Create FCoE Port Channel.
32. Enter `102` for the port channel ID and `Po102` for the port channel name.

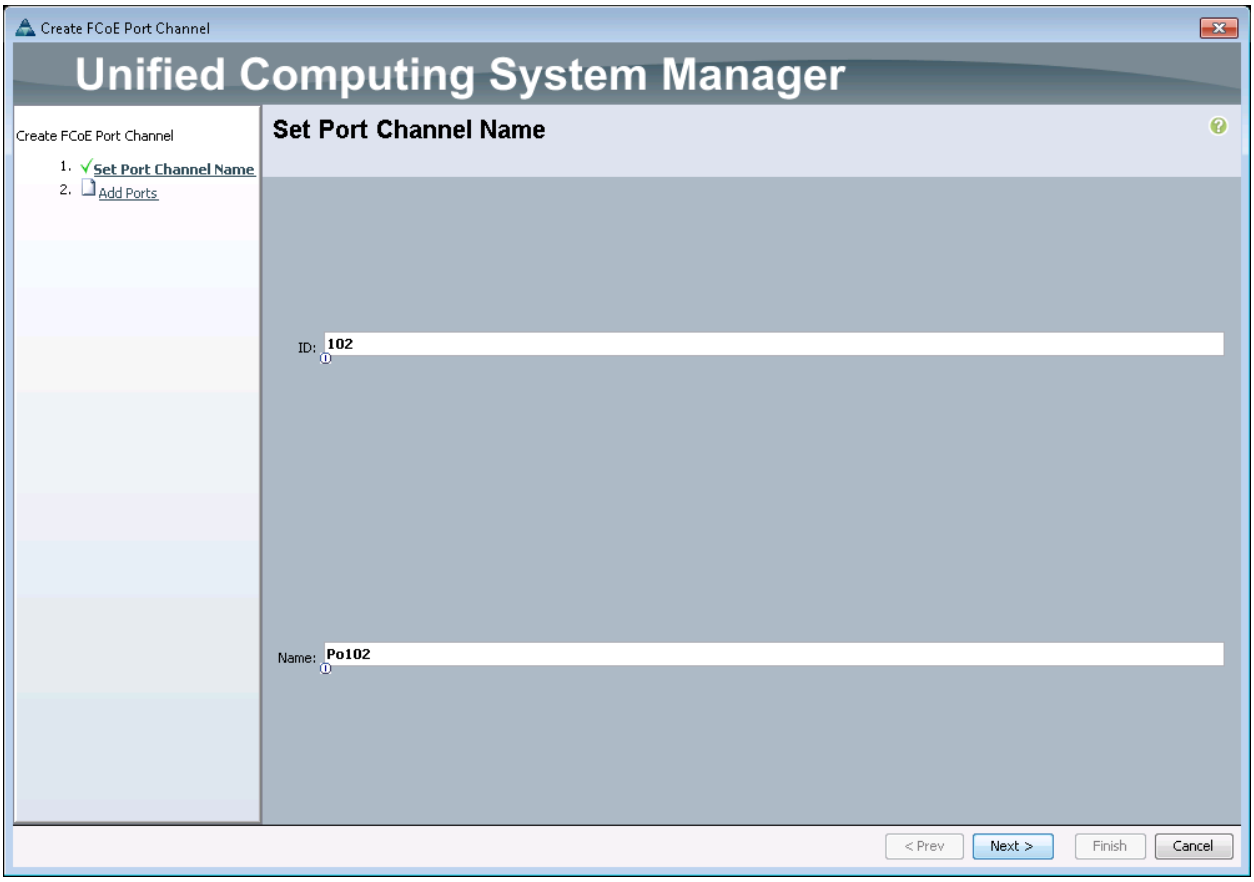

33. Click Next.
34. Select ports 17 through 30 and click the double right-arrow button (>>) to add the ports to the port channel.

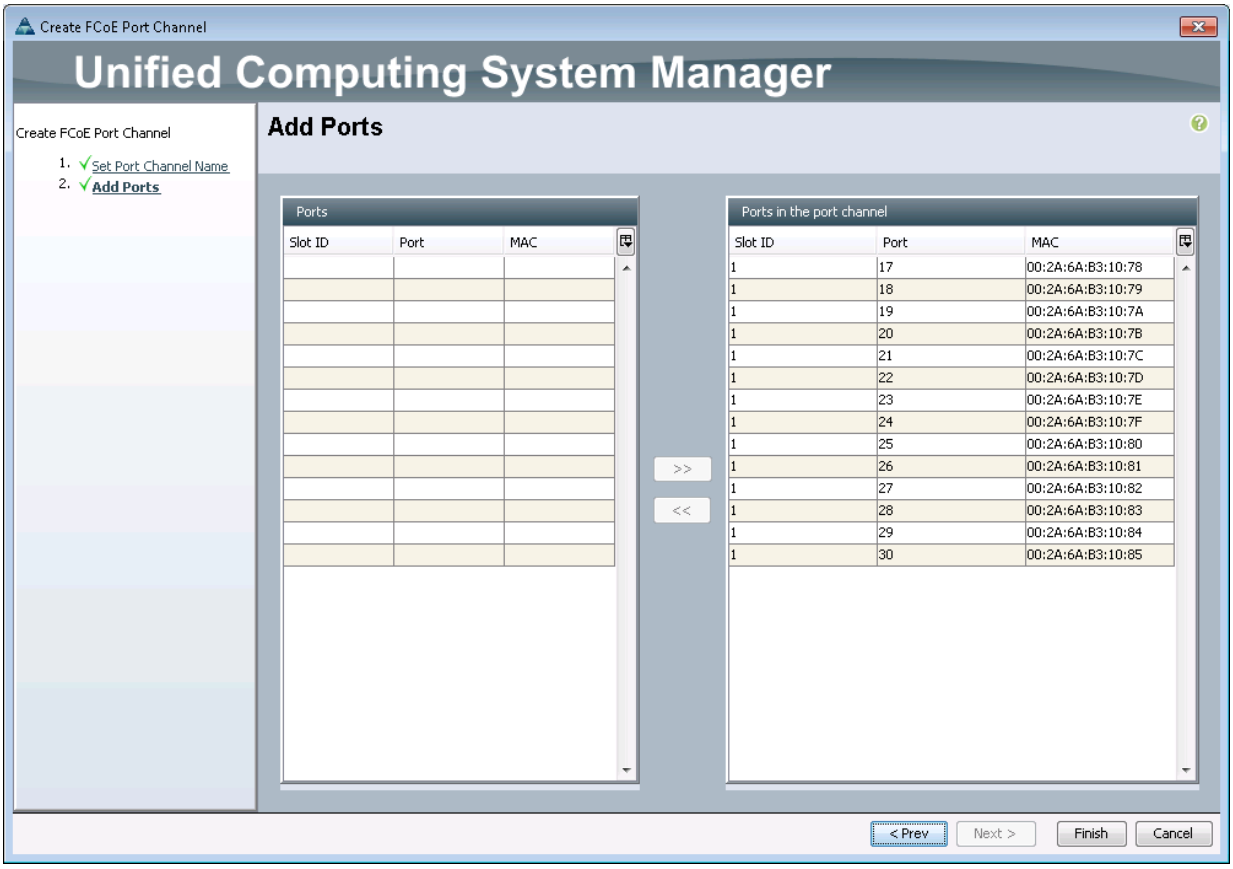

35. Click Finish.
36. Select the checkbox for Show Navigator for FCoE Port-Channel 102 (Fabric B).
37. Click OK to create the port channel.
38. Click OK to close the navigator.

Create Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapters, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > Root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter `RAC-HOST` as the name for the host firmware package.
6. Leave Simple selected.
7. Select the version 2.2(5b) for both the Blade and Rack packages.

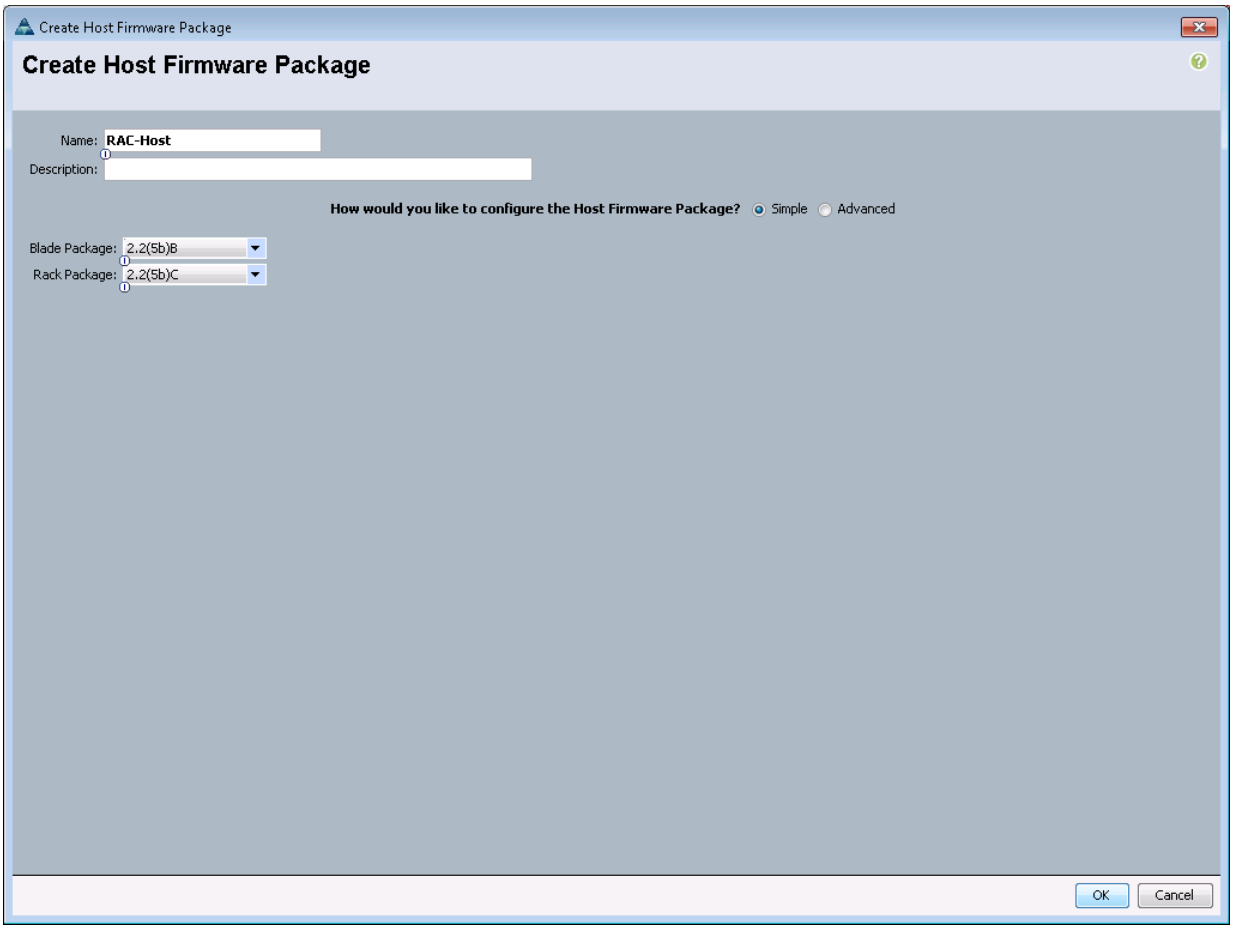

8. Click OK to create the host firmware package.
9. Click OK.

Set Jumbo Frames in Cisco UCS Fabric

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

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

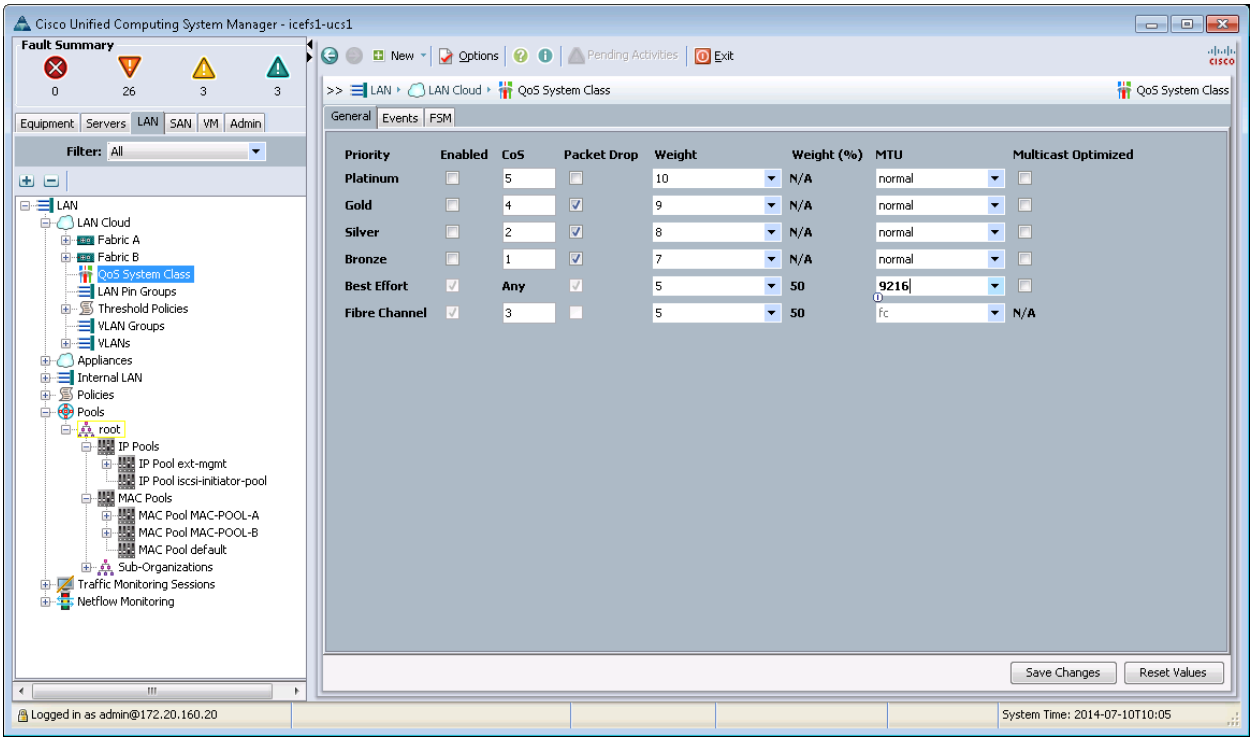

5. Click Save Changes.
6. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Policies > Root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `ENABLE-CDP` as the policy name.
6. For CDP, select the Enabled option.

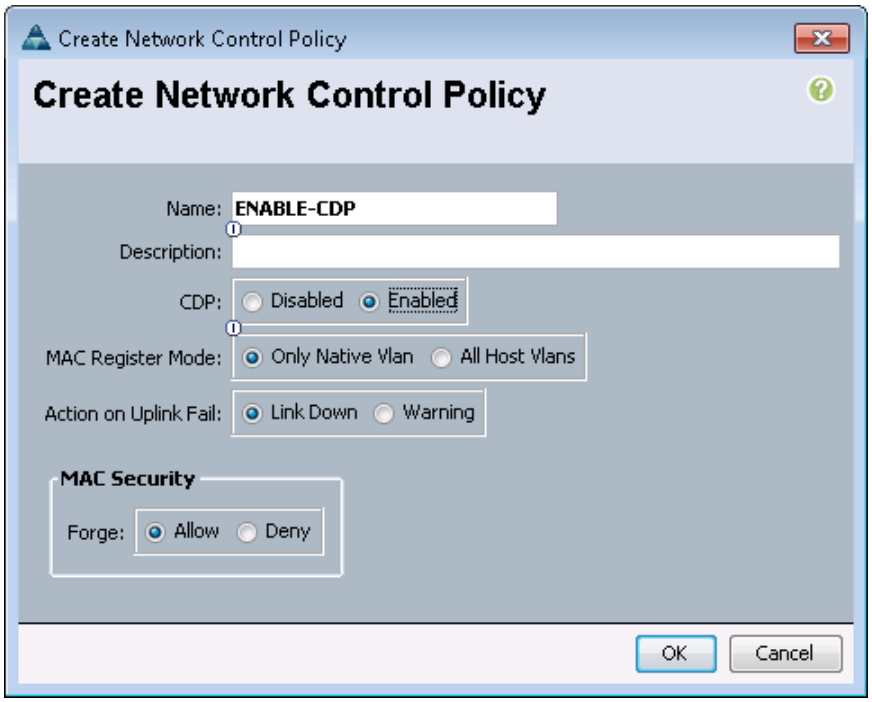

7. Click OK to create the network control policy.
8. Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > Root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter `NO-POWER-CAP` as the power control policy name.
6. Change the power capping setting to No Cap.

7. Click OK to create the power control policy.
8. Click OK.

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > Root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `RAC-HOST` as the BIOS policy name.
6. Change the Quiet Boot setting to Disabled.

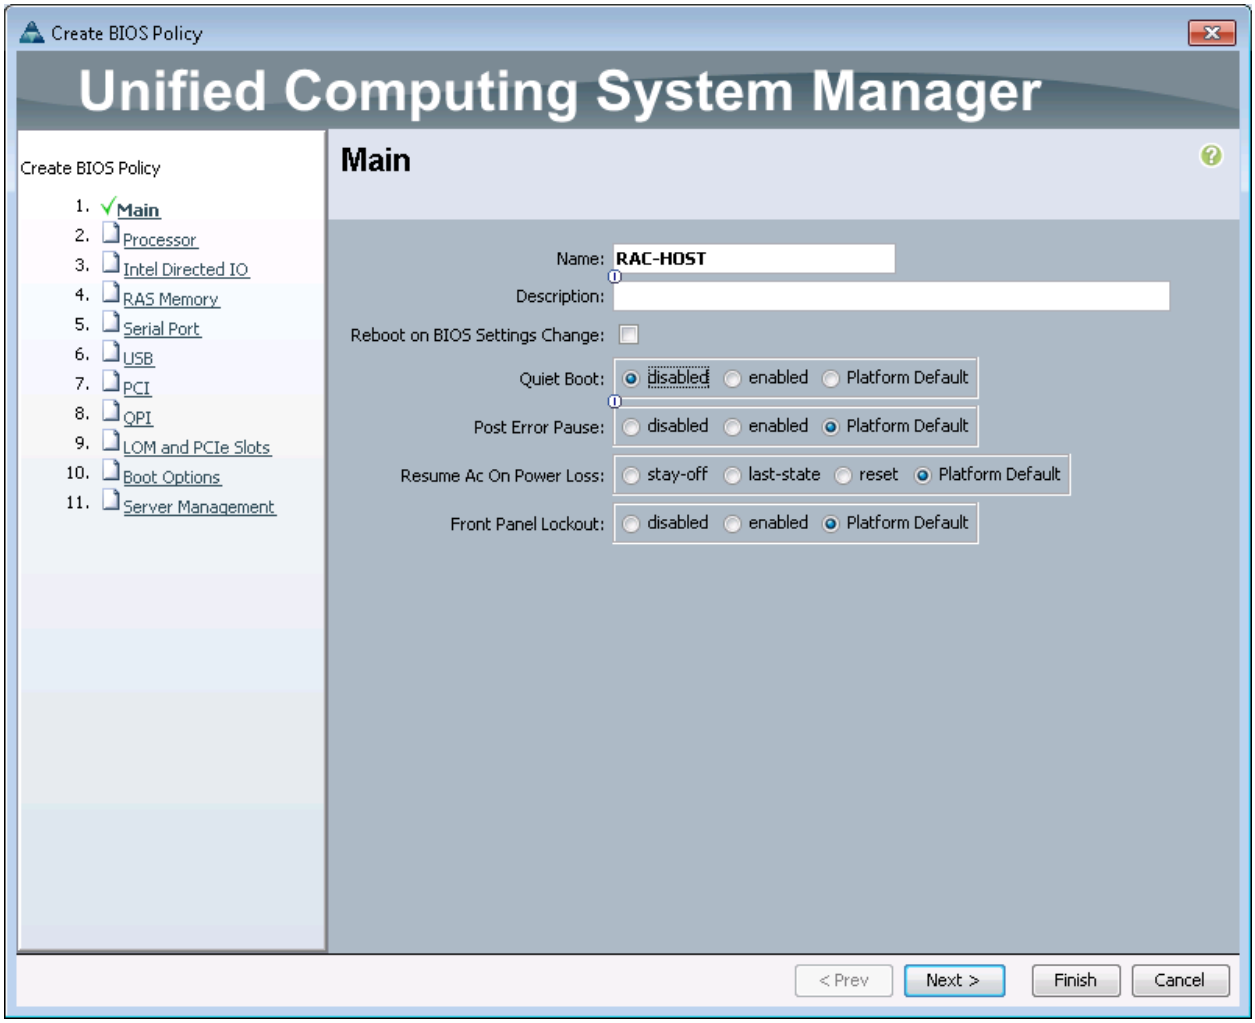

7. Click Finish to create the BIOS policy.
8. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

To create a vNIC/vHBA placement policy for the infrastructure hosts, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > Root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter `RAC-HOST` as the name for the placement policy.
6. Click 1 and, under the Selection Preference, select Assigned Only.

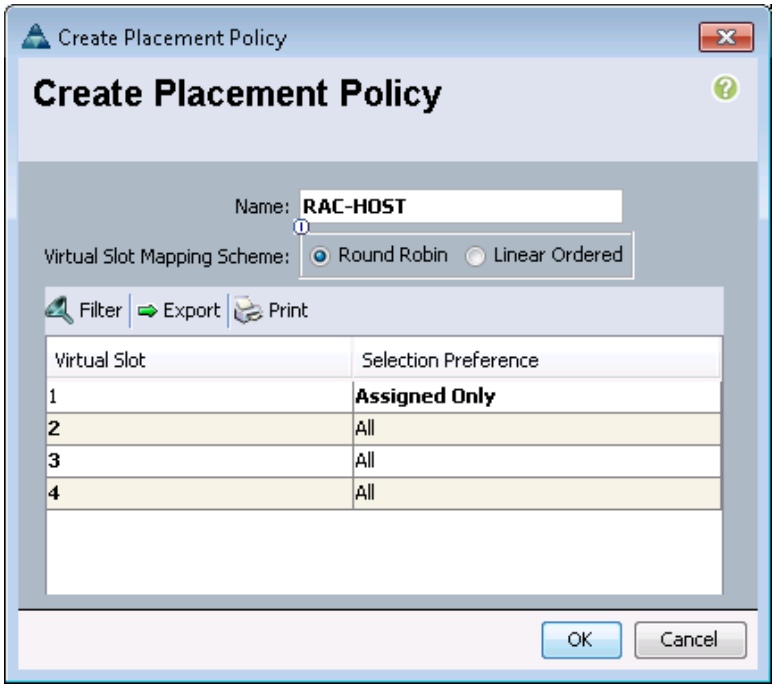

7. Click OK and then click OK again.

Update Default Maintenance Policy

To update the default maintenance policy, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Choose Policies > Root.
3. Choose Maintenance Policies > Default.
4. Change the Reboot Policy to User Ack.

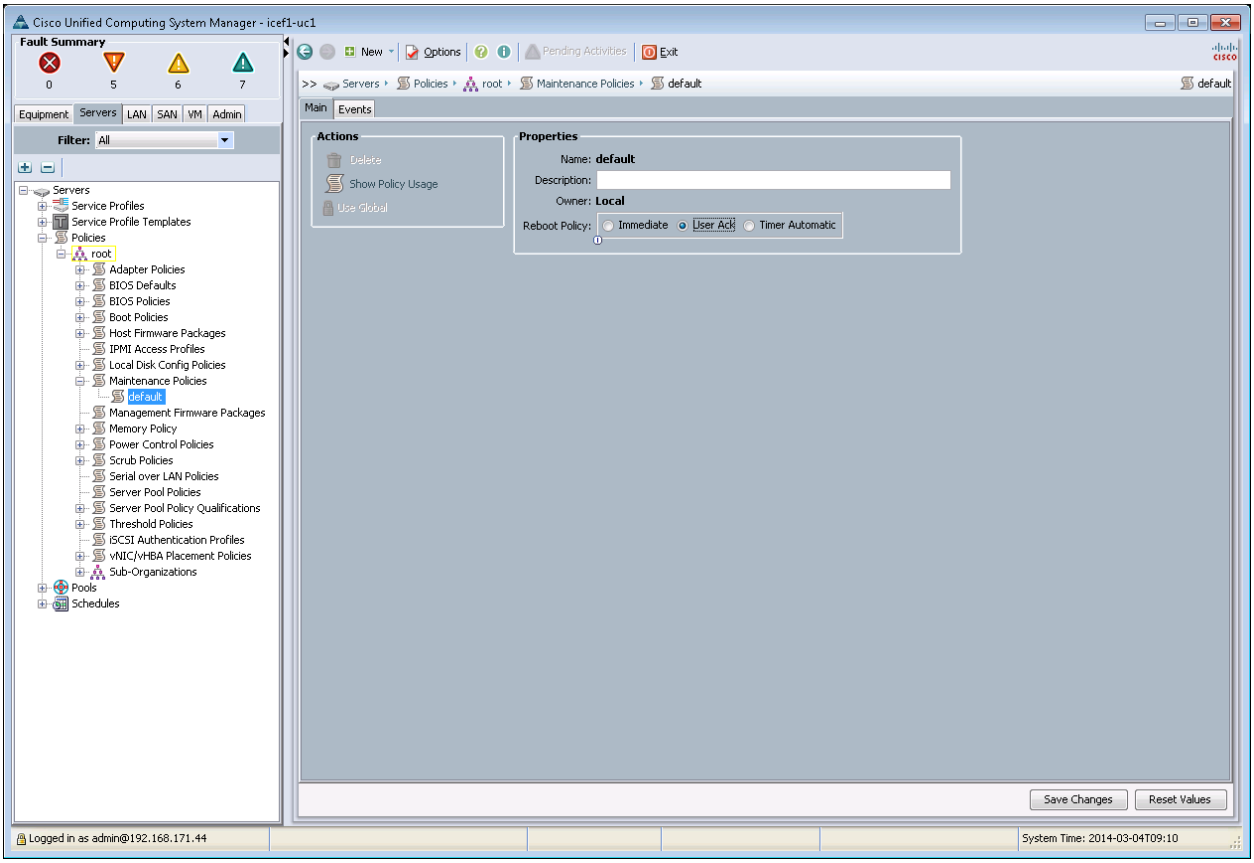

5. Click Save Changes.
6. Click OK to acknowledge the change.

Create vNIC Templates

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

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Policies > Root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `MGMT` as the vNIC template name.
6. For Fabric ID, select Fabric A.
7. Select the Enable Failover checkbox.

   Note: Under Target, do not select the VM checkbox.
8. Select Updating Template as the Template Type.
9. Under VLANs, select the checkboxes for `MGMT-VLAN`.
10. Set `MGMT-VLAN` as the native VLAN.
11. For MTU, enter `1500`.
12. In the MAC Pool list, select `MAC-POOL-A`.
13. In the Network Control Policy list, select `ENABLE-CDP`.

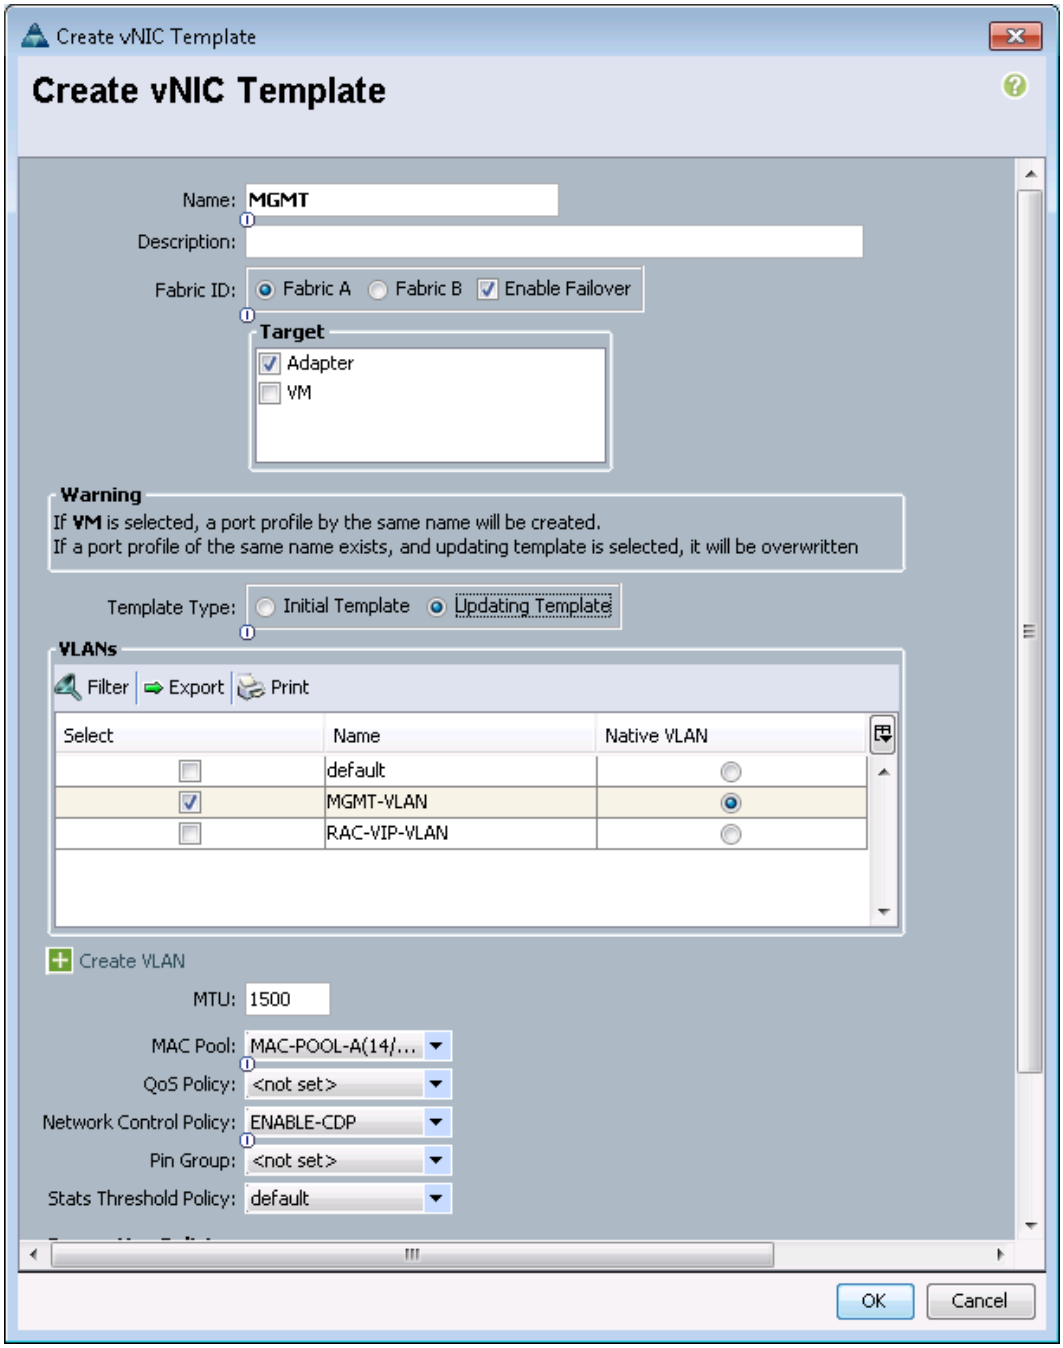

14. Click OK to create the vNIC template.
15. Click OK.
16. In the navigation pane, select the LAN tab.
17. Select Policies > Root.
18. Right-click vNIC Templates.
19. Select Create vNIC Template.
20. Enter `RAC-CLUSTER` as the vNIC template name.
21. Select Fabric B.

22. Select the Enable Failover checkbox.

    Note: Under Target, do not select the VM checkbox.

23. Select Updating Template as the template type.
24. Under VLANs, select the checkboxes for `RAC-VIP-VLAN`.
25. Set `RAC-VIP-VLAN` as the native VLAN.
26. For MTU, enter `9000`.
27. In the MAC Pool list, select `MAC_Pool_B`.
28. In the Network Control Policy list, select `Enable_CDP`.

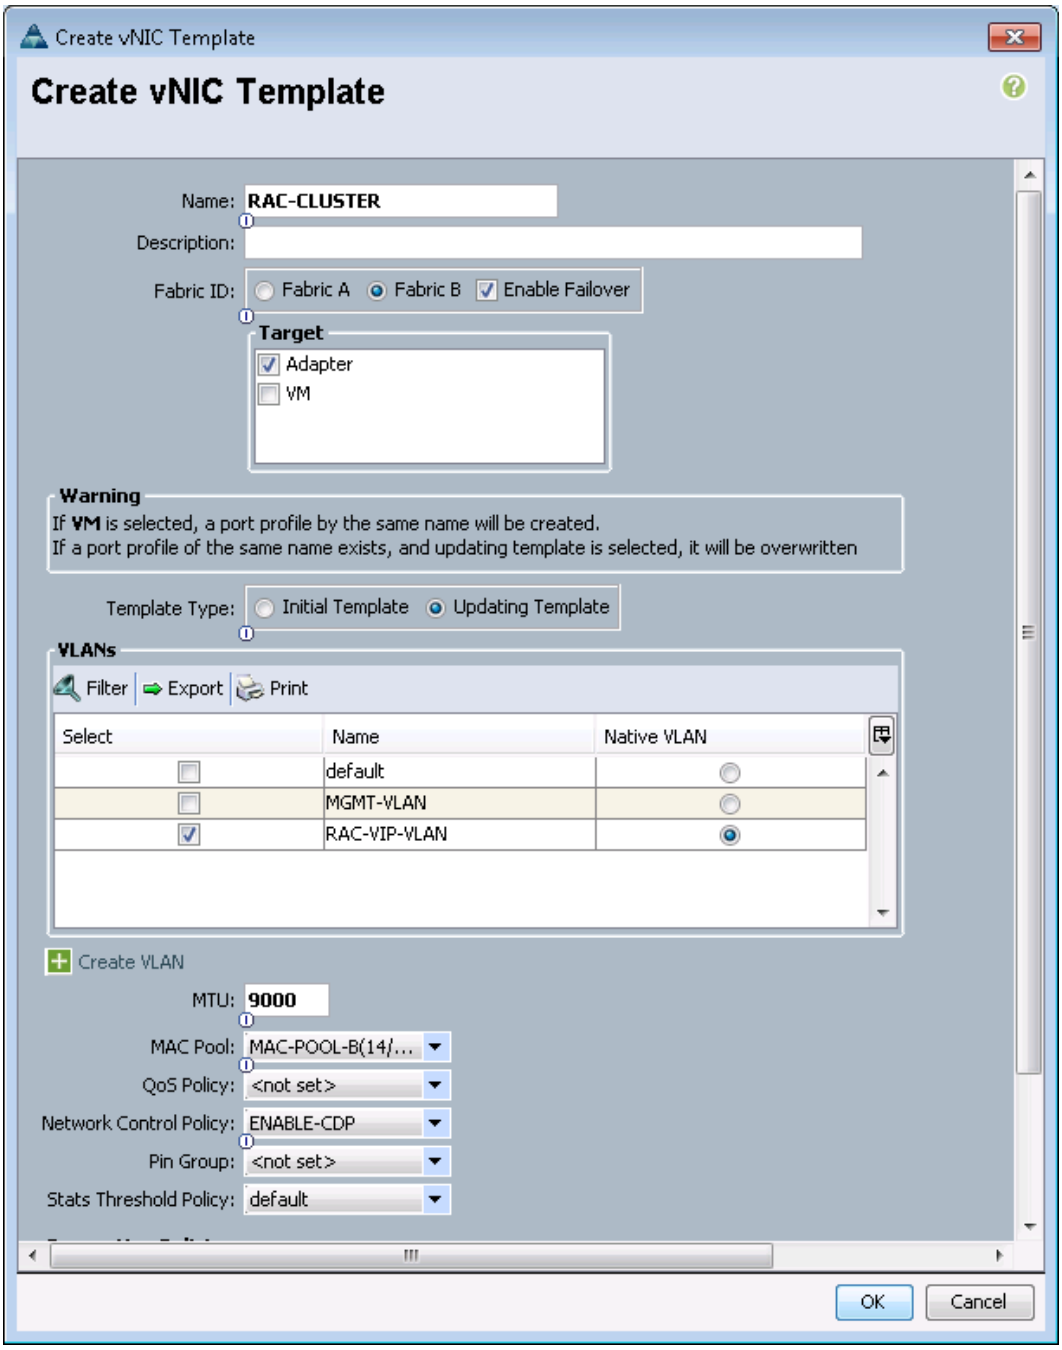

29. Click OK to create the vNIC template.
30. Click OK.

Create vHBA Templates for Fabric A and Fabric B

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

1. In Cisco UCS Manager, in the navigation pane, click the SAN tab.
2. Select Policies > Root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter `VSAN_A` as the vHBA template name.
6. Select A for Fabric ID.
7. In the Select VSAN list, select `VSAN-A`.
8. In the WWPN Pool list, select `WWPN-POOL-A`.

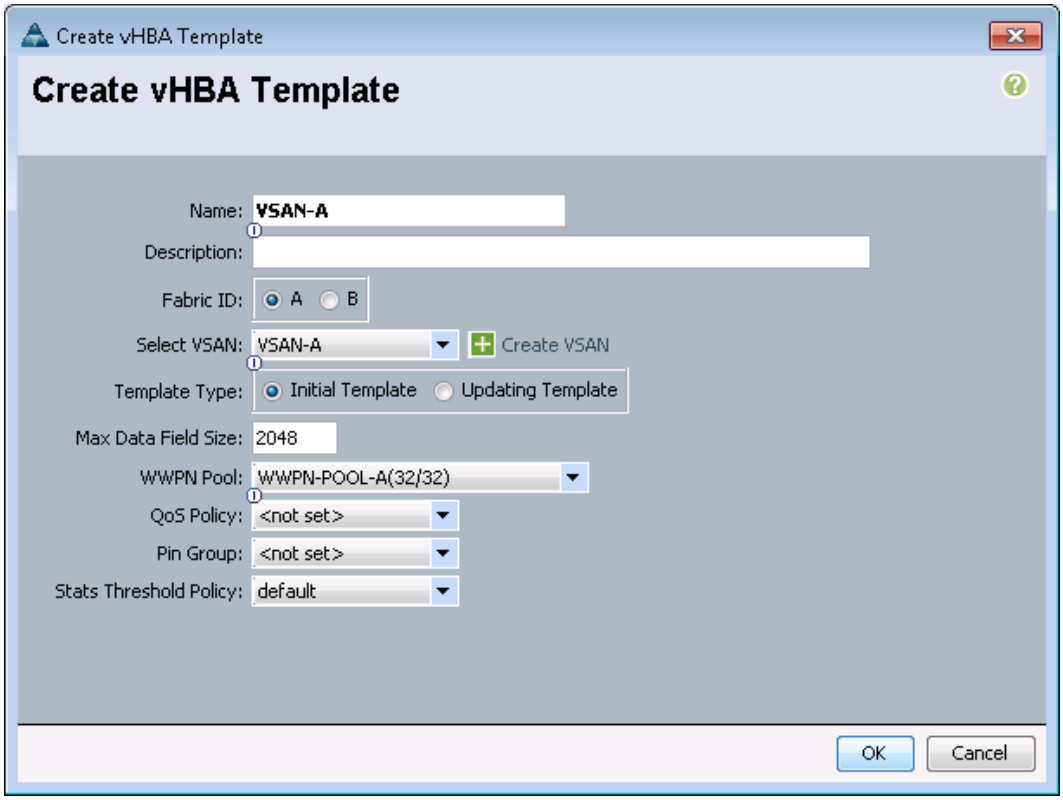

9. Click OK to create the vHBA template.
10. Click OK.
11. In the navigation pane, click the SAN tab.
12. Select Policies > Root.
13. Right-click vHBA Templates.
14. Select Create vHBA Template.
15. Enter `VSAN_B` as the vHBA template name.

16. Select B for Fabric ID.
17. In the Select VSAN list, select `VSAN-B`.
18. In the WWPN Pool, select `WWPN-POOL-B`.

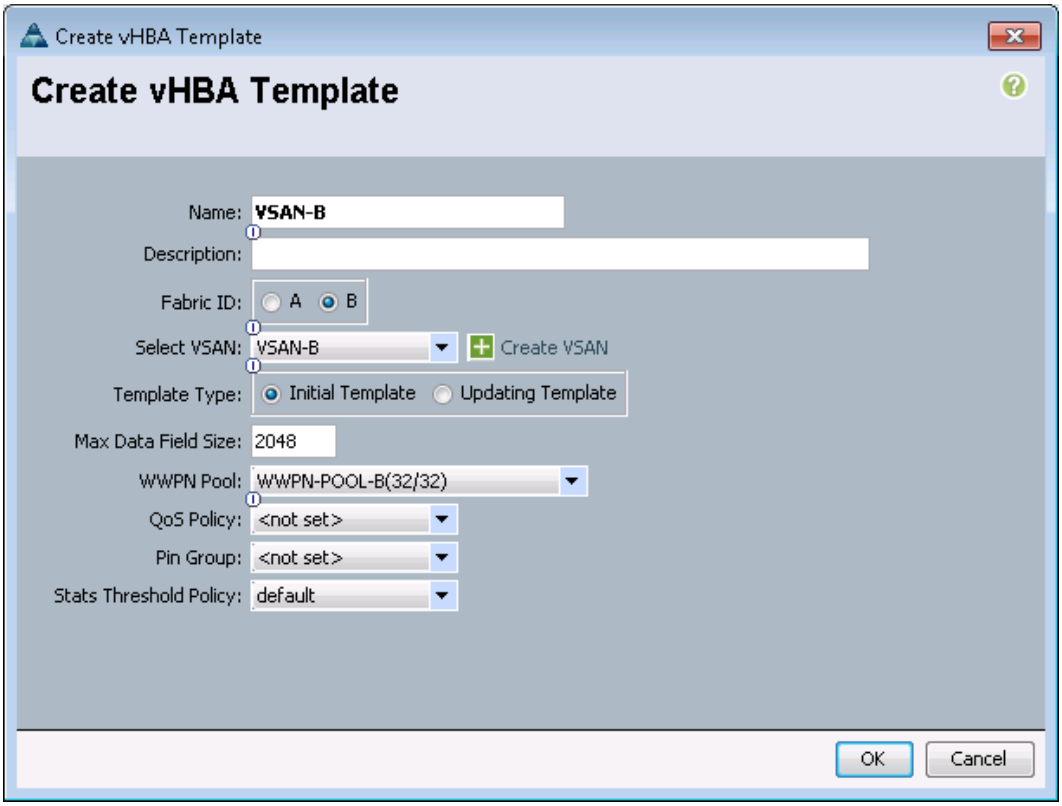

19. Click OK to create the vHBA template.
20. Click OK.

Create Boot Policies

This solution provides SAN boot volumes from an EF560 storage array using the two Fibre Channel host interface cards (HICs). Each HIC has four ports: two ports connected to each of the two Cisco Nexus 5548UP switches in the solution. Each Cisco Nexus switch corresponds to a unique SAN fabric. The boot policy uses one port per HIC per fabric (four ports total) to provide access to storage during boot operations.

For detailed port connectivity, see the Physical Cabling section of this document.

Two boot policies are configured in this procedure. The first policy configures fabric A as the primary path, and the second boot policy configures fabric B as the primary path.

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

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Policies > Root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `BOOT-FABRIC-A` as the name for the boot policy.
6. Optional: Enter a description for the boot policy.

Note: Do not select the Reboot on Boot Order Change checkbox.

7. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.

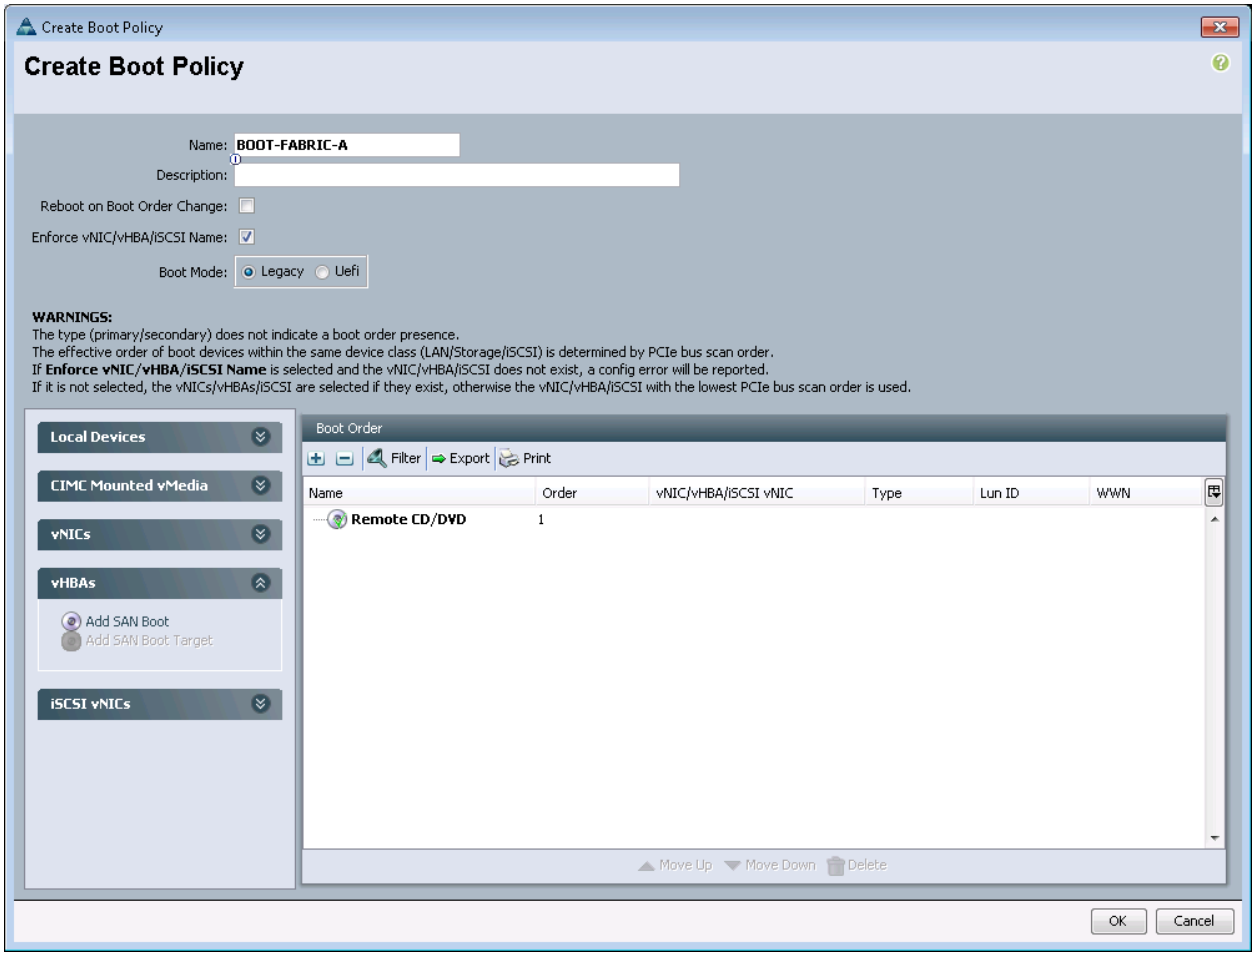

8. Expand the vHBAs drop-down menu and select Add SAN Boot.
9. In the Add SAN Boot dialog box, enter `VSAN-A` in the vHBA field.
10. Select Primary, or confirm that it is selected, for the Type option.

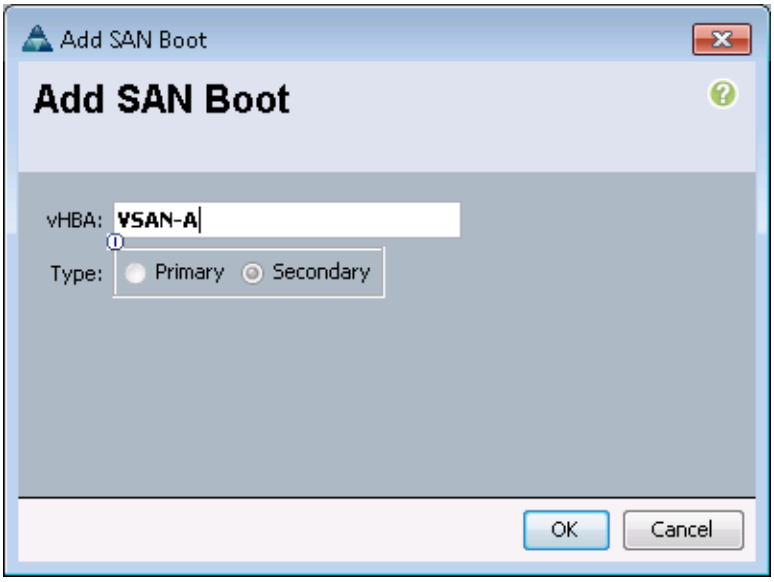

11. Click OK to add the SAN boot initiator.
12. From the vHBA drop-down menu, select Add SAN Boot Target.
13. Keep 0 as the value for Boot Target LUN.
14. Enter the WWPN for channel 1 of the first EF560 controller.

    Note: To obtain this information, see the Appendix of this document.
15. Select Primary for the SAN boot target type.

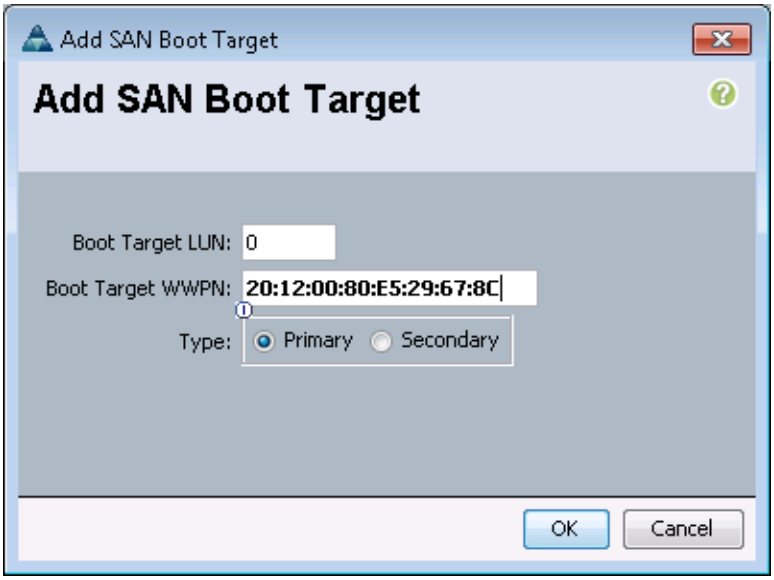

16. Click OK to add the SAN boot target.
17. From the vHBA drop-down menu, select Add SAN Boot Target.
18. Enter 0 as the value for Boot Target LUN.
19. Enter the WWPN for channel 1 of the second EF560 controller.

    Note: To obtain this information, see the Appendix of this document.

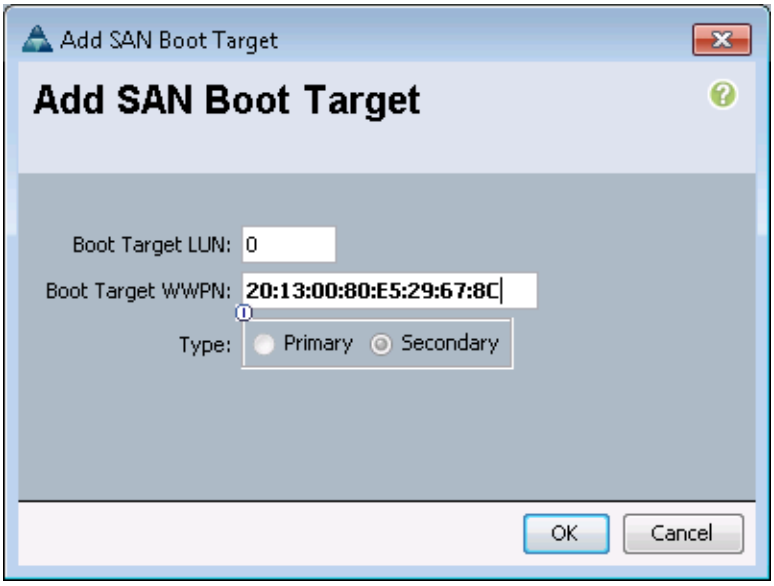

20. Click OK to add the SAN boot target.
21. From the vHBA drop-down menu, select Add SAN Boot.
22. In the Add SAN Boot dialog box, enter `VSAN-B` in the vHBA box.
23. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.

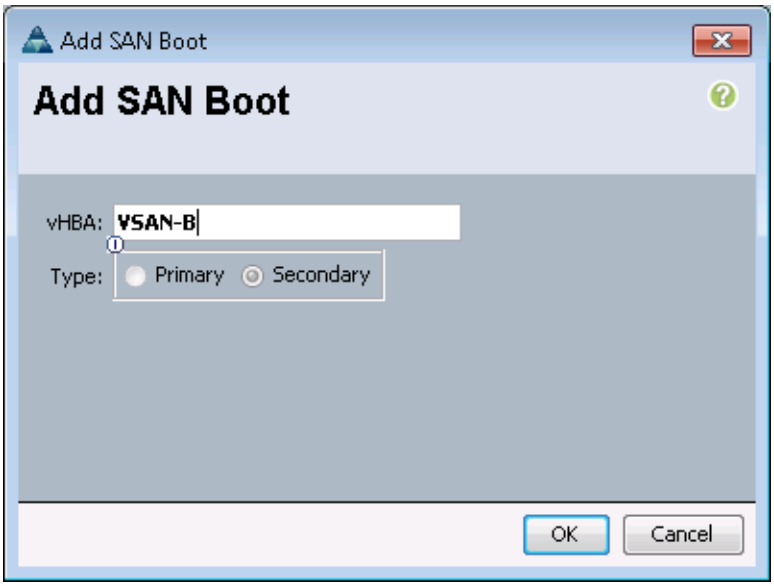

24. Click OK to add the SAN boot initiator.
25. From the vHBA drop-down menu, select Add SAN Boot Target.
26. Keep `0` as the value for Boot Target LUN.
27. Enter the WWPN for channel 3 of the first EF560 controller.

    Note: To obtain this information, see the Appendix of this document.
28. Select Primary for the SAN boot target type.

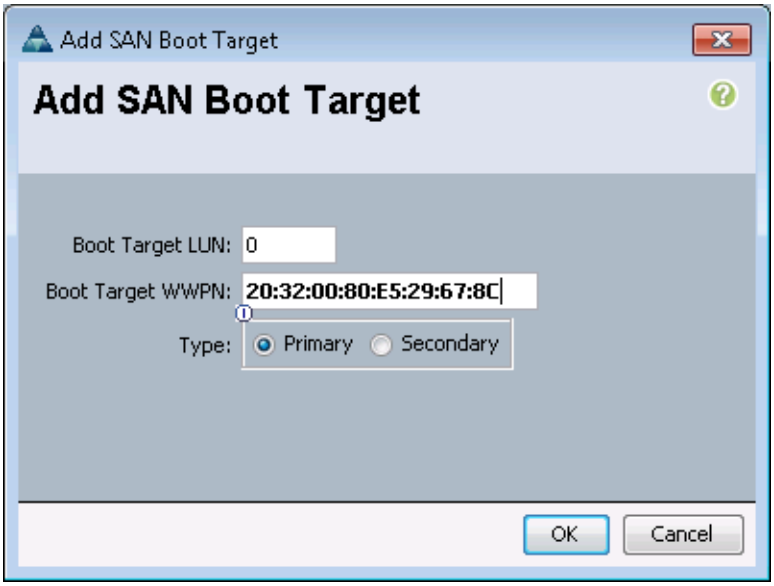

29. Click OK to add the SAN boot target.
30. From the vHBA drop-down menu, select Add SAN Boot Target.
31. Keep 0 as the value for Boot Target LUN.
32. Enter the WWPN for channel 3 of the second EF560 controller.

    Note: To obtain this information, see the Appendix of this document.

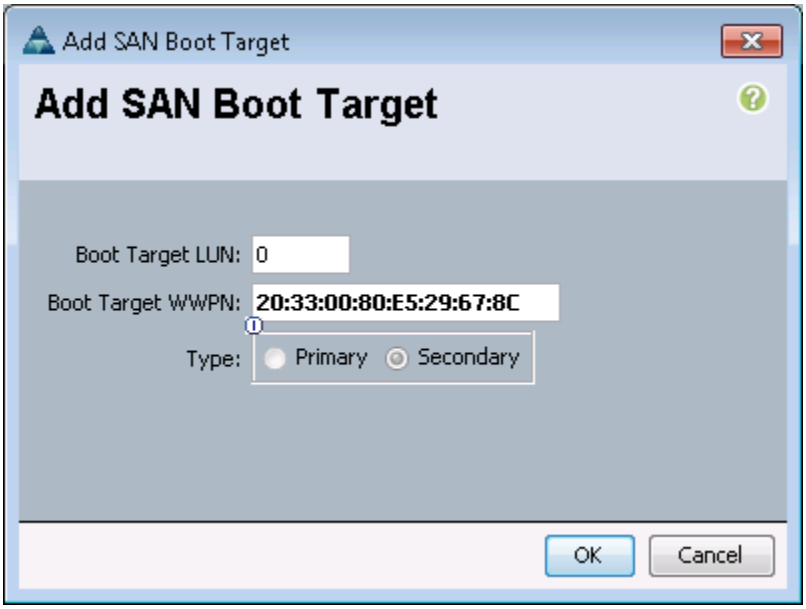

33. Click OK to add the SAN boot target.

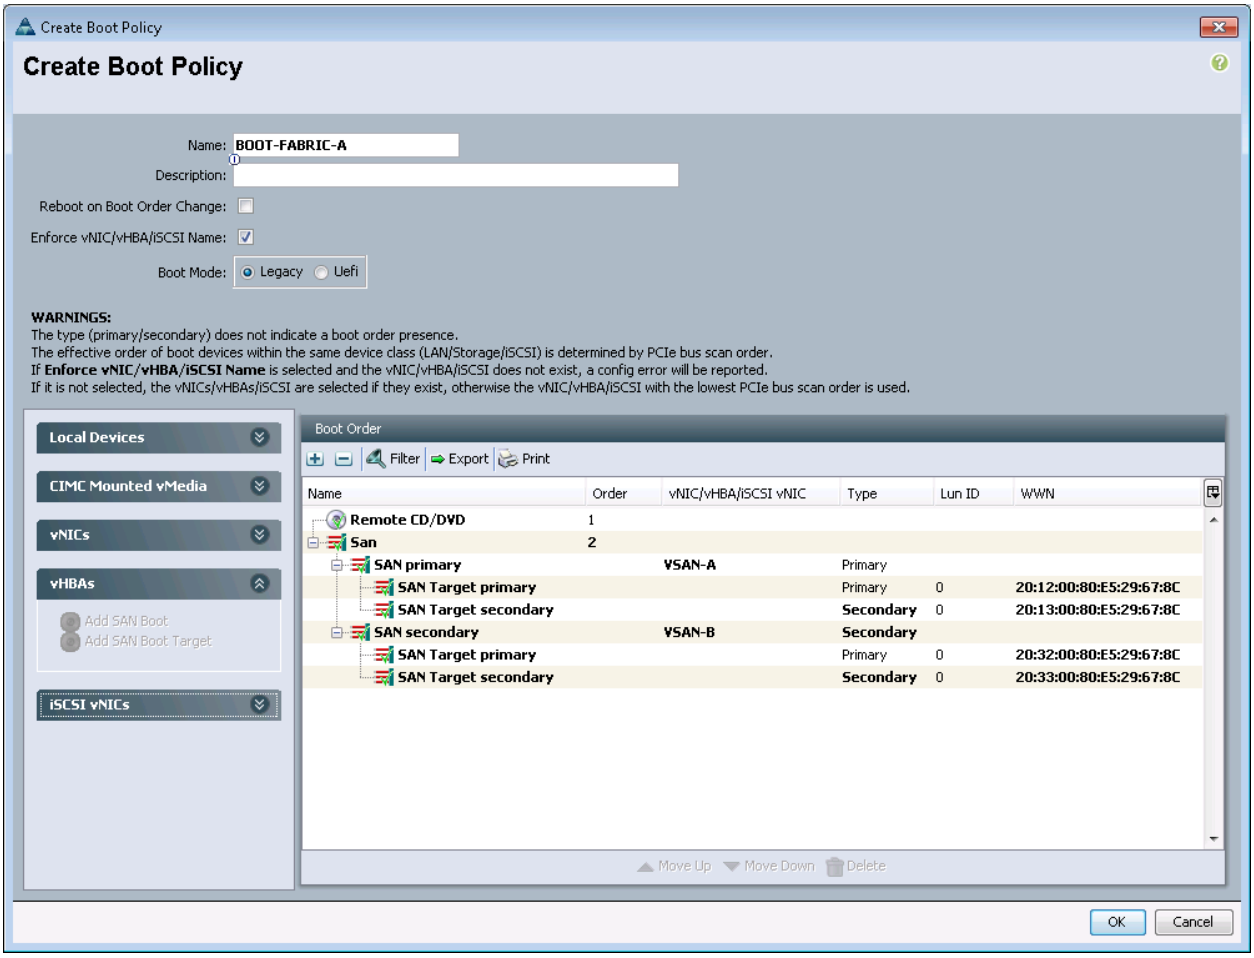

34. Click OK, then click OK again to create the boot policy.
35. Right-click Boot Policies.
36. Select Create Boot Policy.
37. Enter `BOOT-FABRIC-B` as the name for the boot policy.
38. Optional: Enter a description of the boot policy.

    Note: Do not select the Reboot on Boot Order Change option.
39. From the Local Devices drop-down menu, select Add Remote CD/DVD.

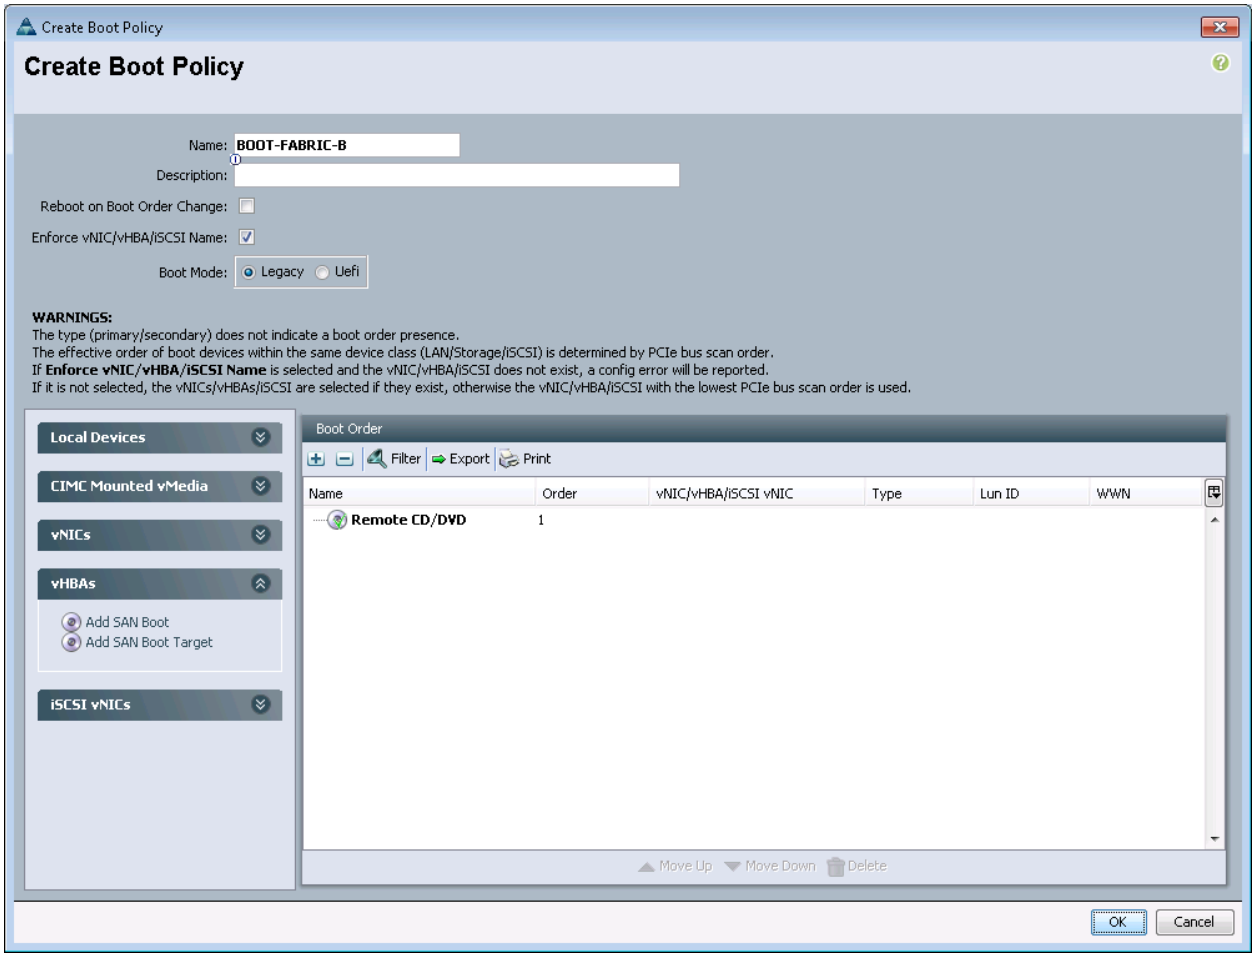

40. From the vHBA drop-down menu, select Add SAN Boot.
41. In the Add SAN Boot dialog box, enter `VSAN-B` in the vHBA box.
42. Confirm that Primary option is selected for the SAN boot type.

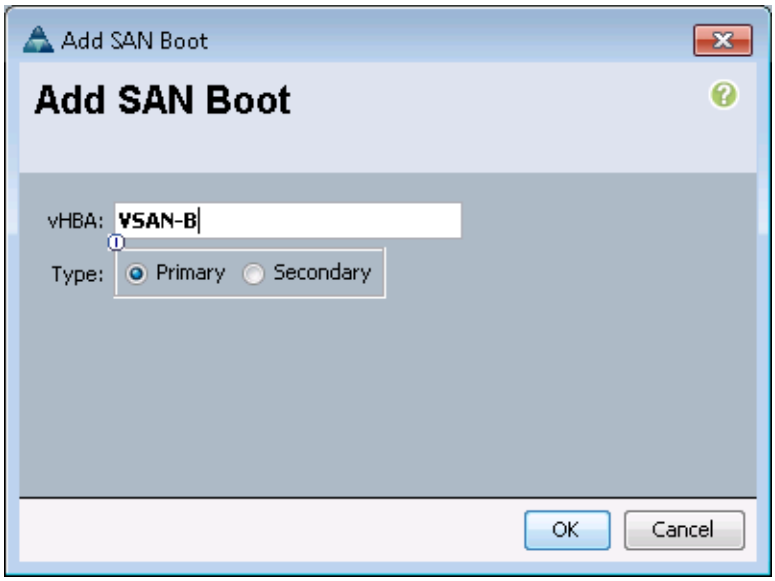

43. Click OK to add the SAN boot initiator.
44. From the vHBA drop-down menu, select Add SAN Boot Target.
45. Enter `0` as the value for Boot Target LUN.
46. Enter the WWPN for channel 3 of the first EF560 controller.

    Note: To obtain this information, see the Appendix of this document.
47. Select Primary option for the SAN boot target type.

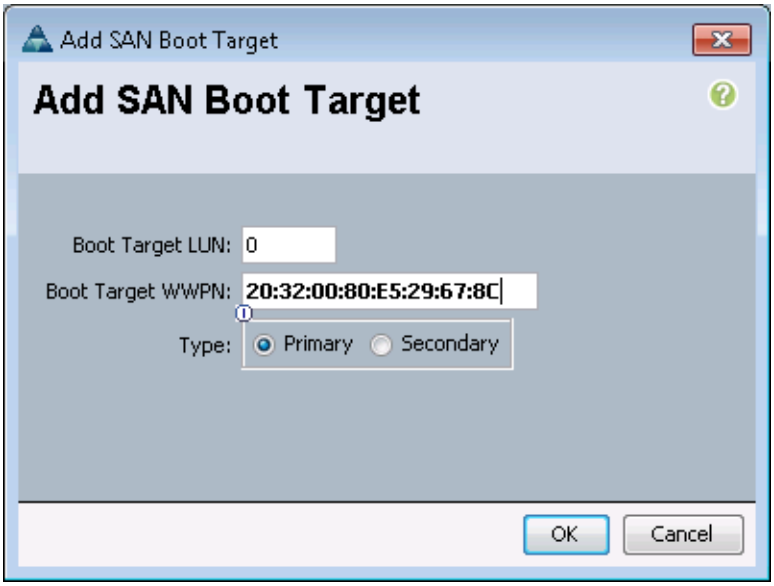

48. Click OK to add the SAN boot target.
49. From the `vHBA` drop-down menu, select Add SAN Boot Target.
50. Enter `0` as the value for Boot Target LUN.
51. Enter the WWPN for channel 3 of the second EF560 controller.

    Note: To obtain this information, see the Appendix of this document.

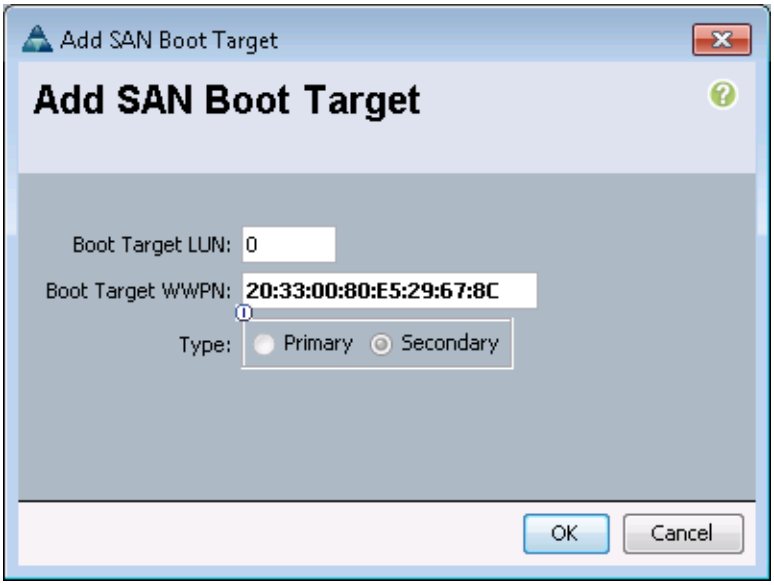

52. Click OK to add the SAN boot target.
53. From the vHBA menu, select Add SAN Boot.
54. In the Add SAN Boot dialog box, enter `VSAN-A` in the vHBA box.
55. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.

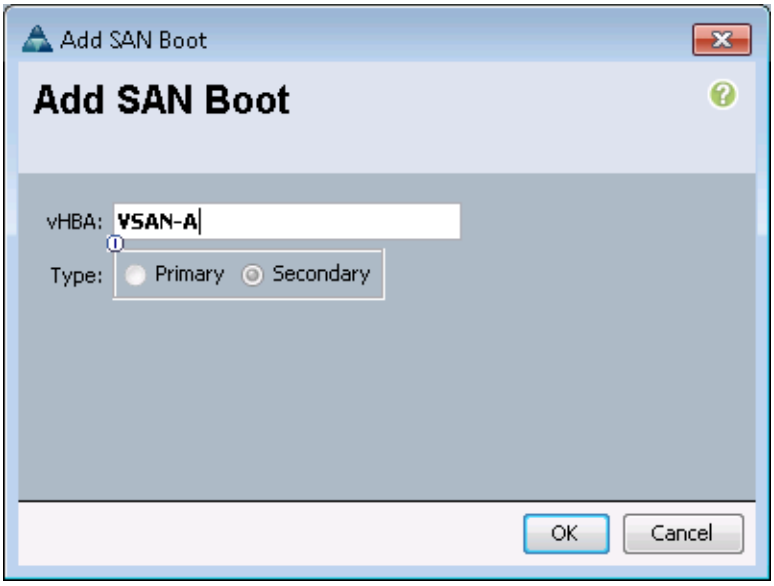

56. Click OK to add the SAN boot initiator.
57. From the vHBA menu, select Add SAN Boot Target.
58. Enter `0` as the value for Boot Target LUN.
59. Enter the WWPN for channel 1 of the first EF560 controller.

    Note: To obtain this information, see the Appendix of this document.

60. Select the Primary option for the SAN boot target type.

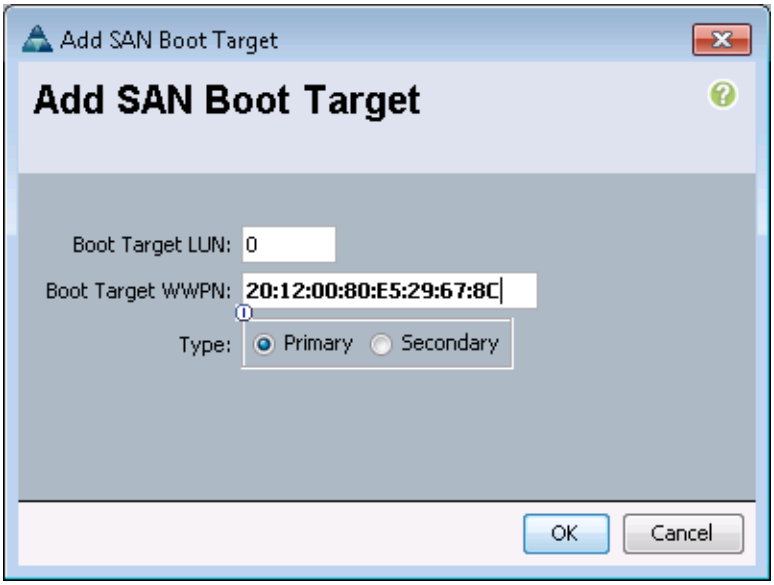

61. Click OK to add the SAN boot target.
62. From the vHBA drop-down menu, select Add SAN Boot Target.
63. Enter `0` as the value for Boot Target LUN.
64. Enter the WWPN for channel 1 of the second EF560 controller.

    Note: To obtain this information, see the Appendix of this document.

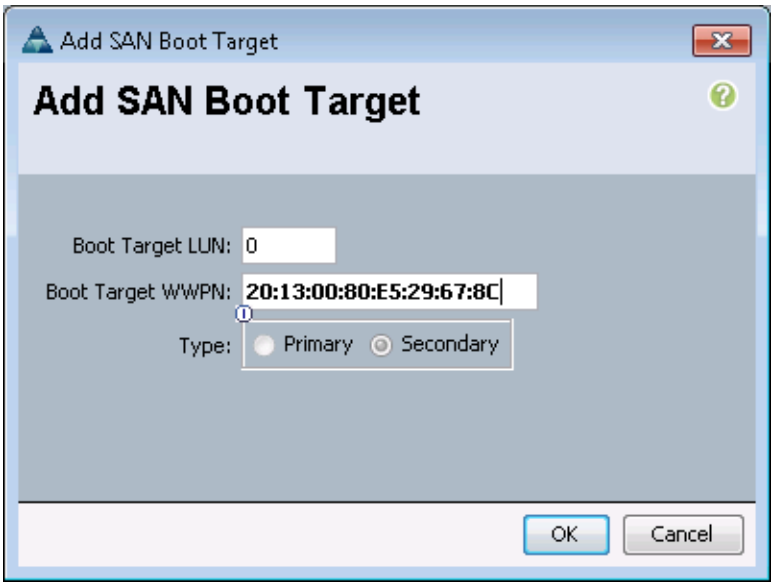

65. Click OK to add the SAN boot target.

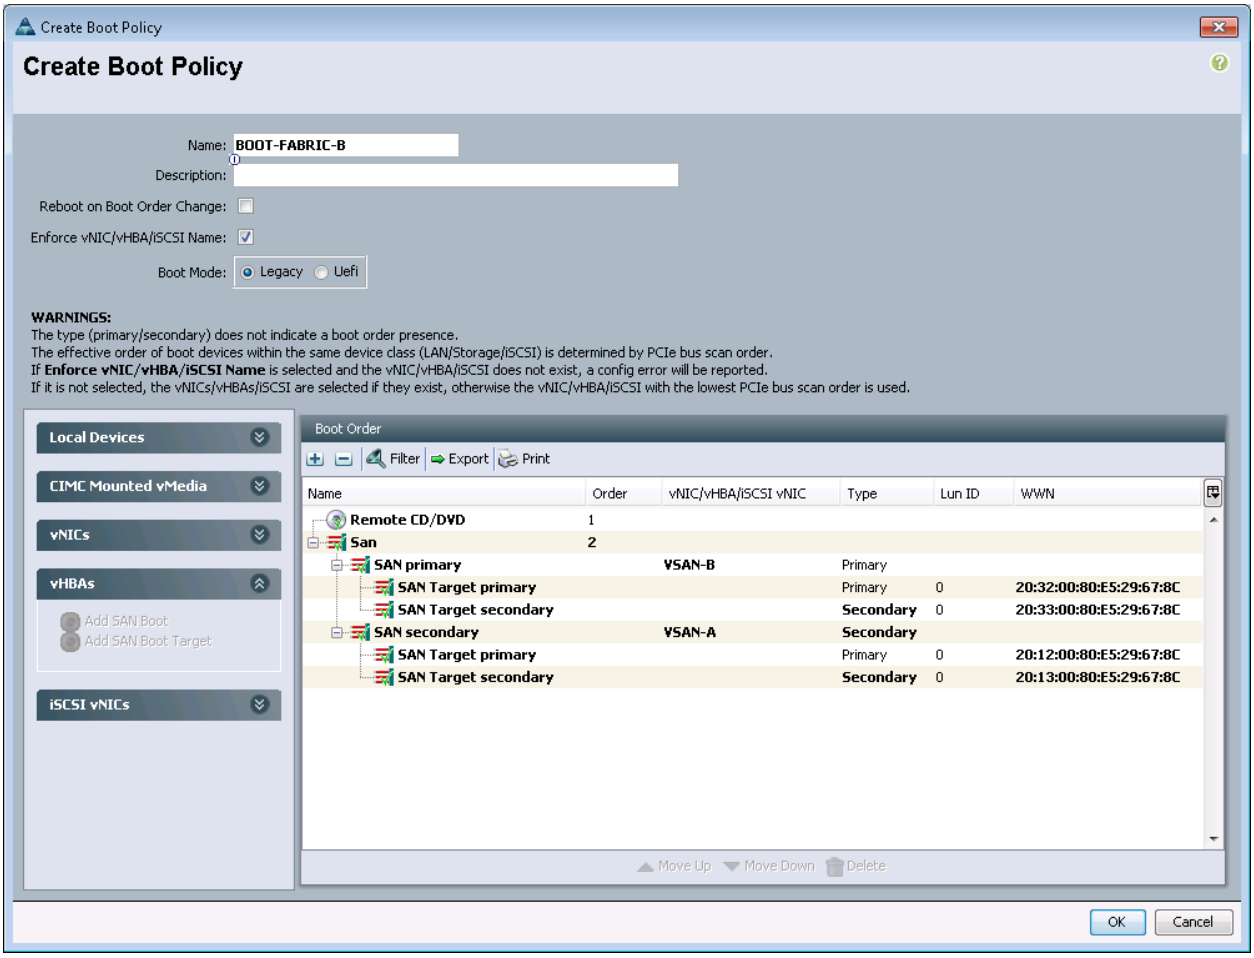

66. Click OK and then click OK again to create the boot policy.

Create Service Profile Templates

In this procedure, two service profile templates are created: one that boots primarily from fabric A and another that boots primarily from fabric B.

To create service profile templates, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Servers tab.
2. Select Service Profile Templates > Root.
3. Right-click Root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the Service Profile Template:
   a. Enter `RAC-SERVER-BOOT-FABRIC-A` as the name for the service profile template. This service profile template is configured to boot from Node 1 on Fabric A.
   b. Select the Updating Template option.
   c. Under UUID, select `UUID-POOL` as the UUID pool.

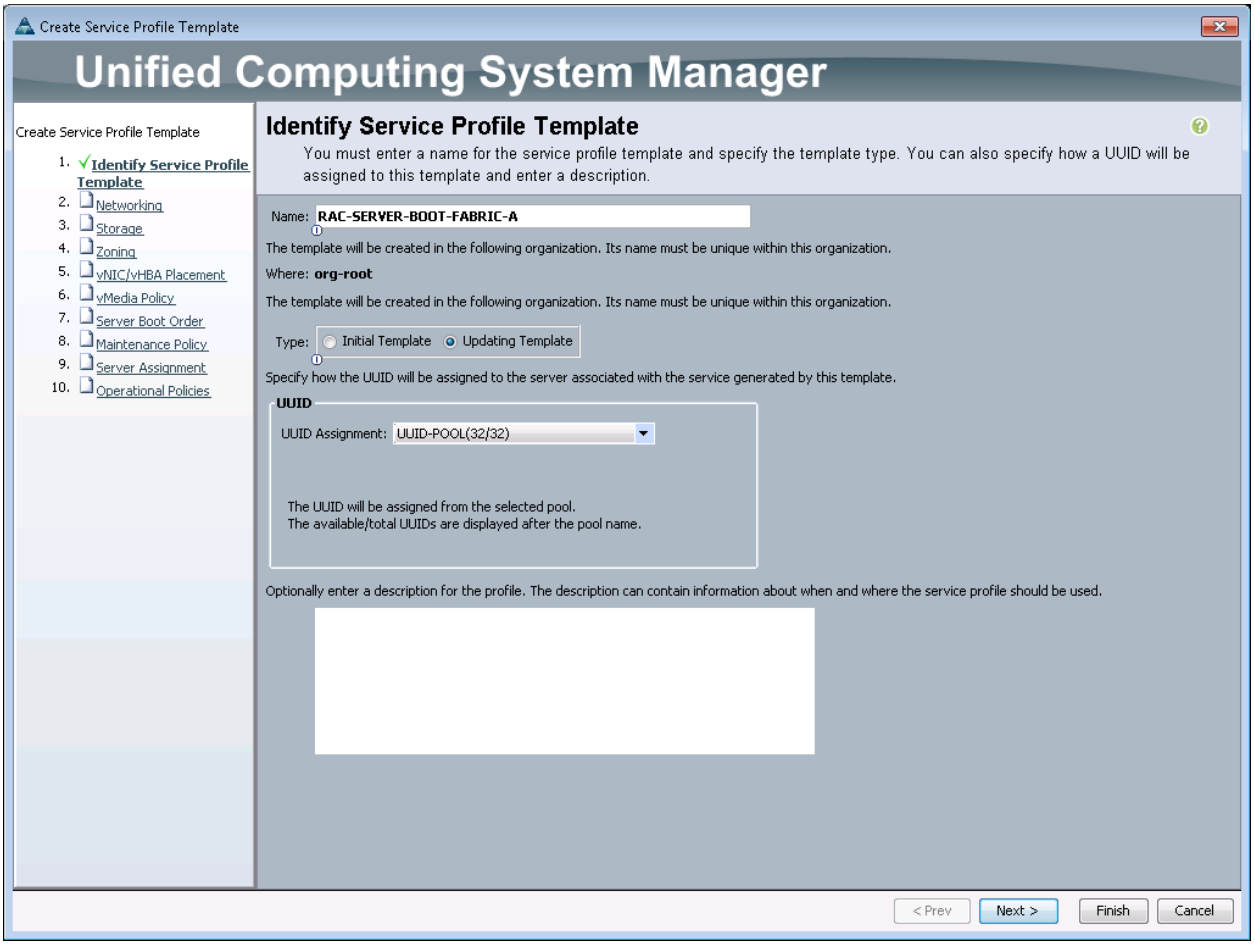

d. Click Next.

6. Configure the following Networking options:
   a. Retain the default setting for Dynamic vNIC Connection Policy.
   b. Select the Expert option to configure the LAN connectivity.

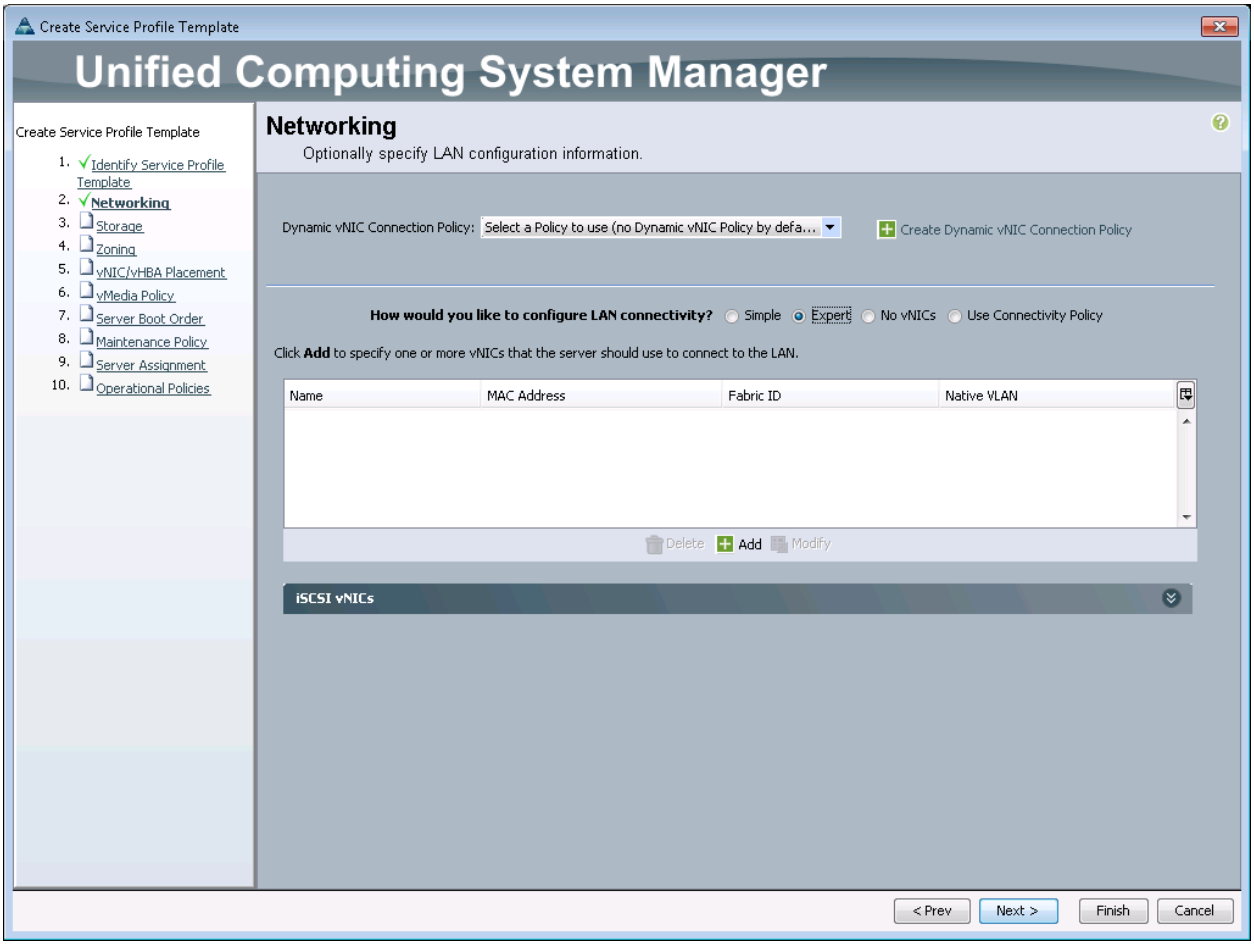

c. Click the upper Add button to add a vNIC to the template.

d. In the Create vNIC dialog box, enter `VNIC-MGMT` as the name for vNIC.

e. Select the Use vNIC Template checkbox.

f. In the vNIC Template list, select `MGMT`.

g. In the Adapter Policy list, select Linux.

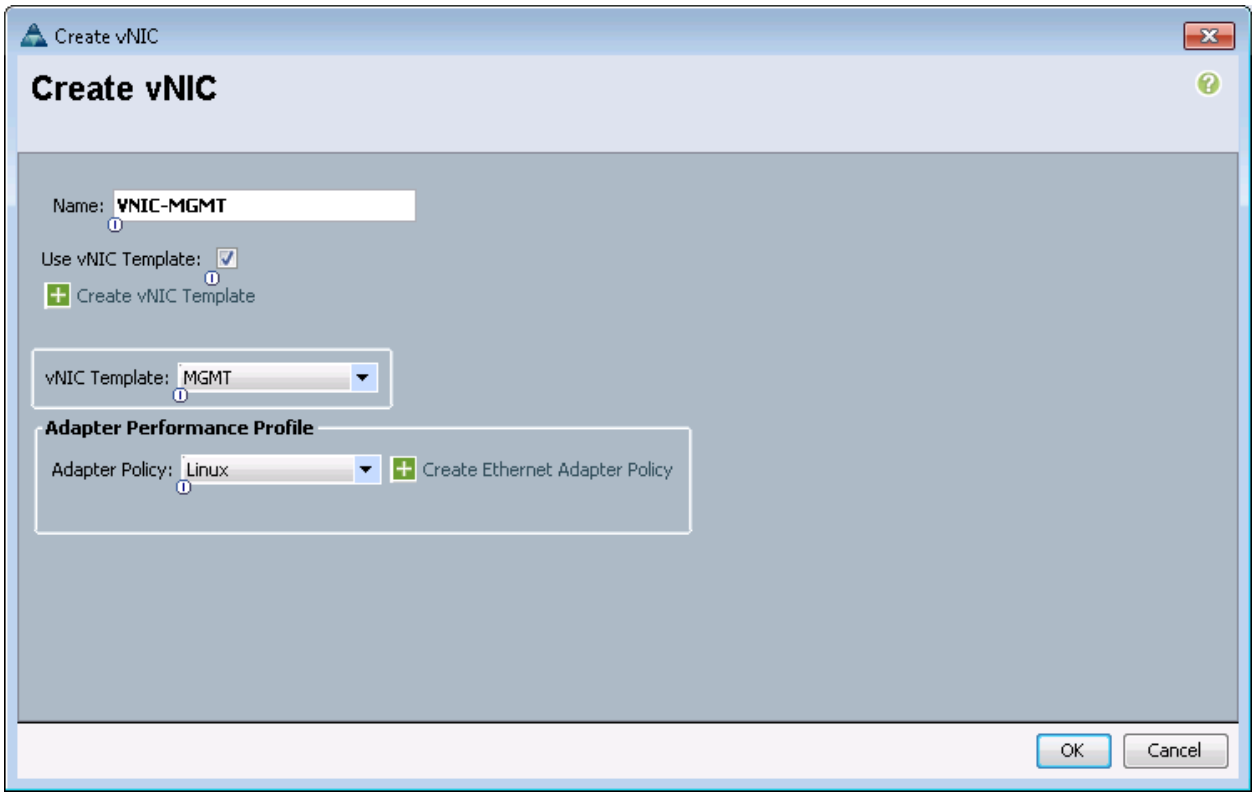

h. Click OK to add this vNIC to the template.

i. On the Networking page of the wizard, click the Add button to add another vNIC to the template.

j. In the Create vNIC box, enter `VNIC-RAC-CLUSTER` as the name for vNIC.

k. Select the Use vNIC Template checkbox.

l. In the vNIC Template list, select `RAC-CLUSTER`.

m. In the Adapter Policy list, select Linux.

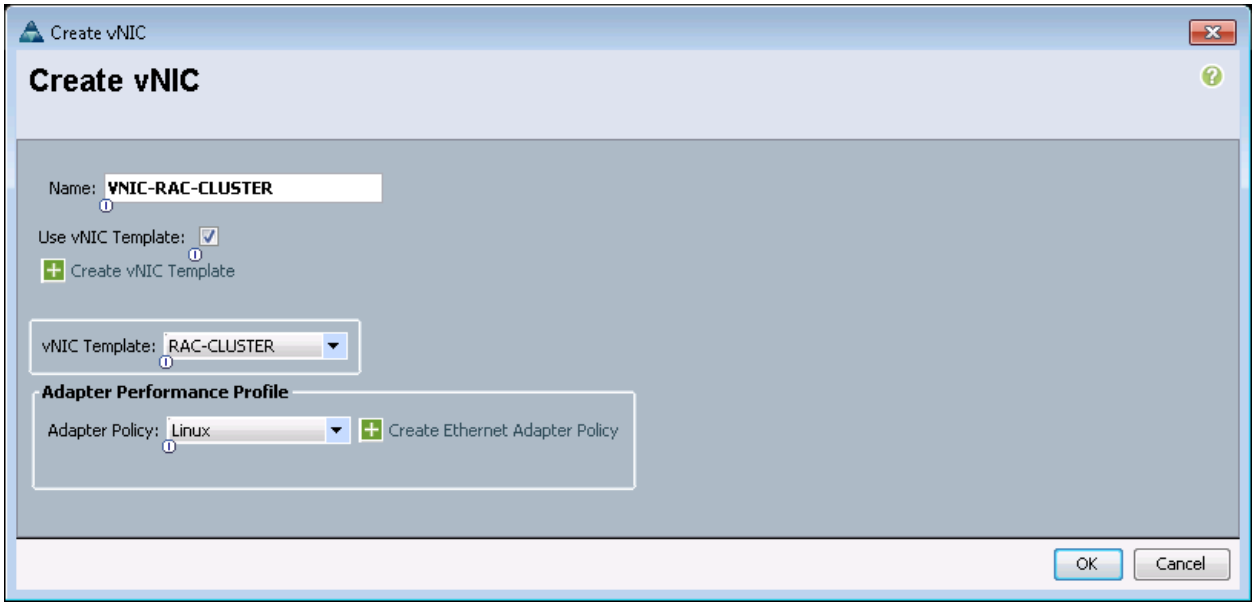

n. Click OK to add the vNIC to the template.

o. Review the table in the Networking page to confirm that both vNICs were created.

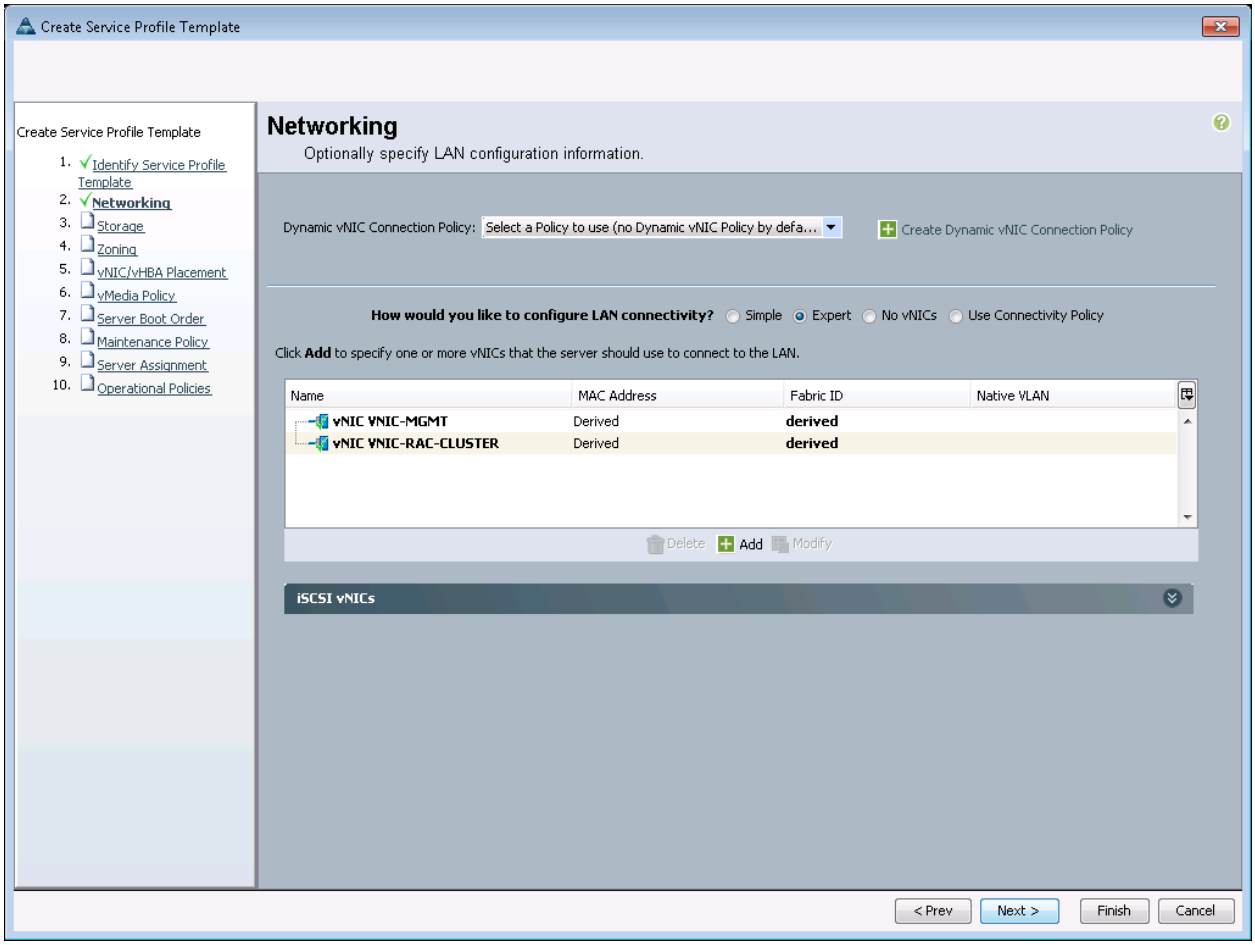

p. Click Next.

7. Configure the Storage options:

a. Do not alter the Local Storage option.

b. Select the Expert option to configure the SAN connectivity.

c. In the WWNN Assignment list, select `WWNN-POOL`.

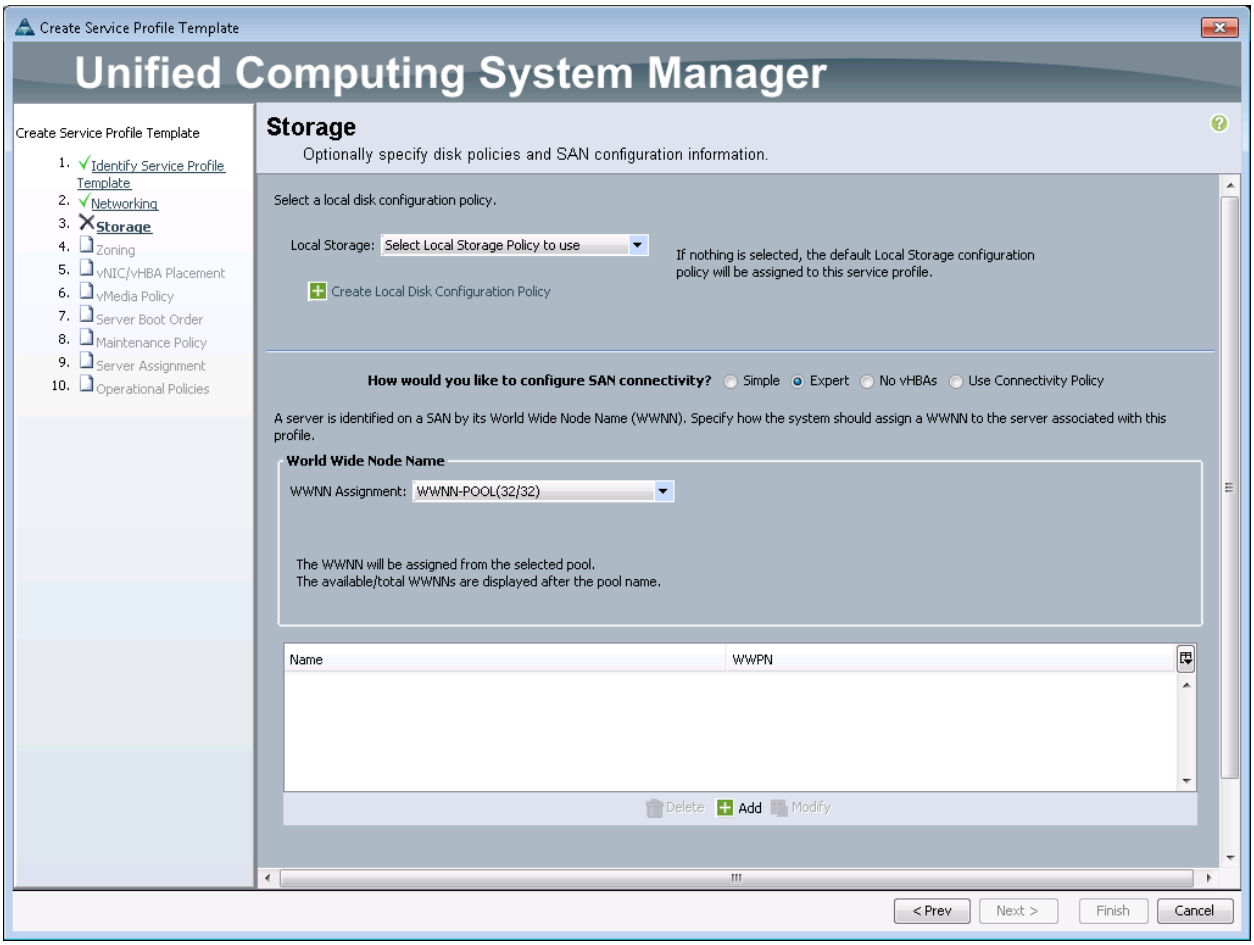

d. Click the Add button to add a vHBA to the template.

e. In the Create vHBA dialog box, enter `VSAN-A` as the name for vHBA.

f. Select the Use vHBA Template checkbox.

g. In the vHBA Template list, select `VSAN_A`.

h. In the Adapter Policy list, select Linux.

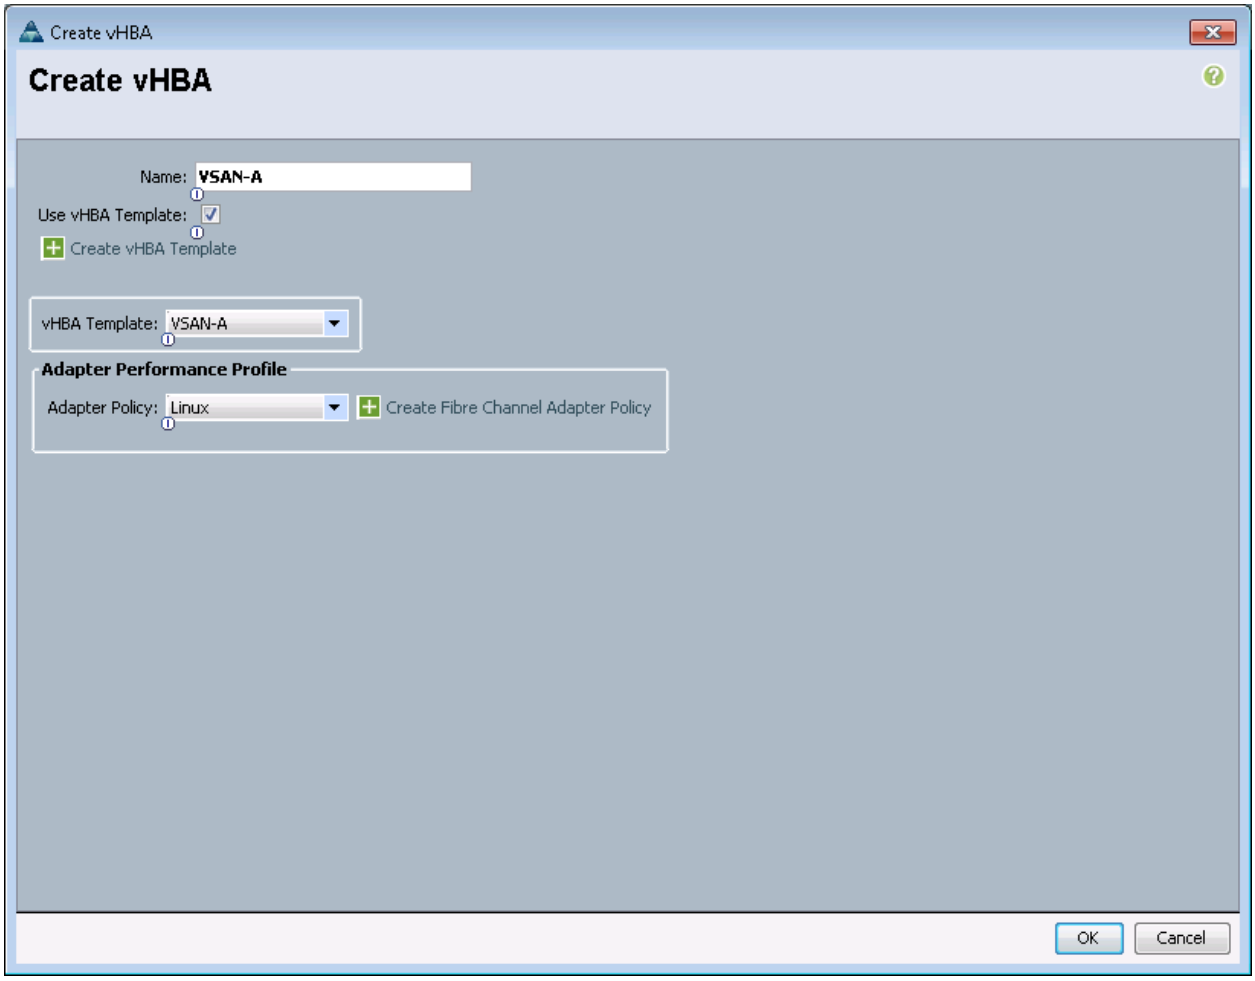

i. Click OK to add this vHBA to the template.
j. On the Storage page of the wizard, click Add to add another vHBA to the template.
k. In the Create vHBA dialog box, enter `VSAN-B` as the name for vHBA.
l. Select the checkbox for Use HBA Template.
m. In the vHBA Template list, select `VSAN-B`.
n. In the Adapter Policy list, select Linux.

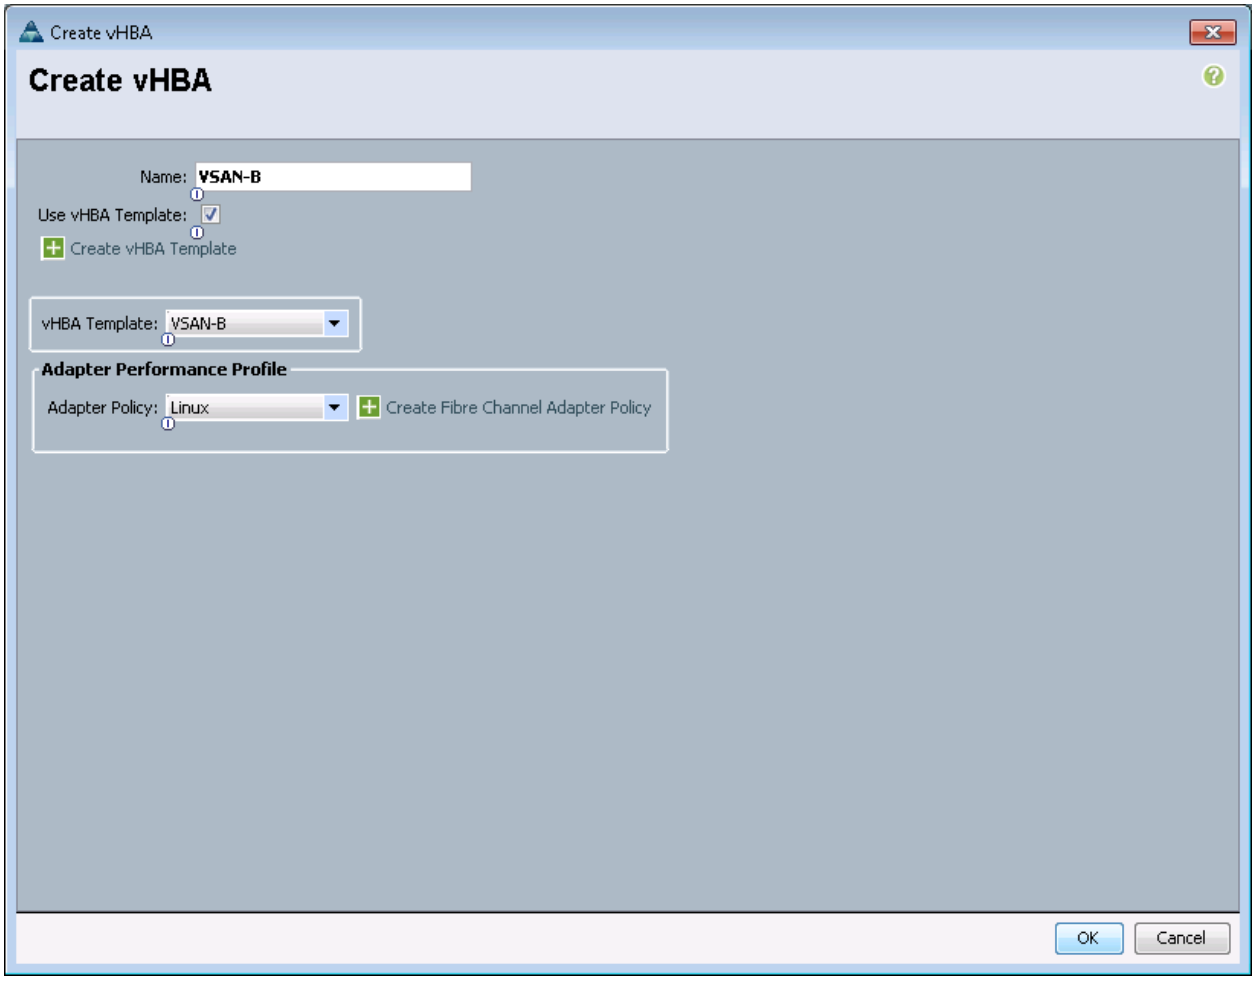

o. Click OK to add the vHBA to the template.

p. Review the table on the Storage page to verify that both vHBAs were created.

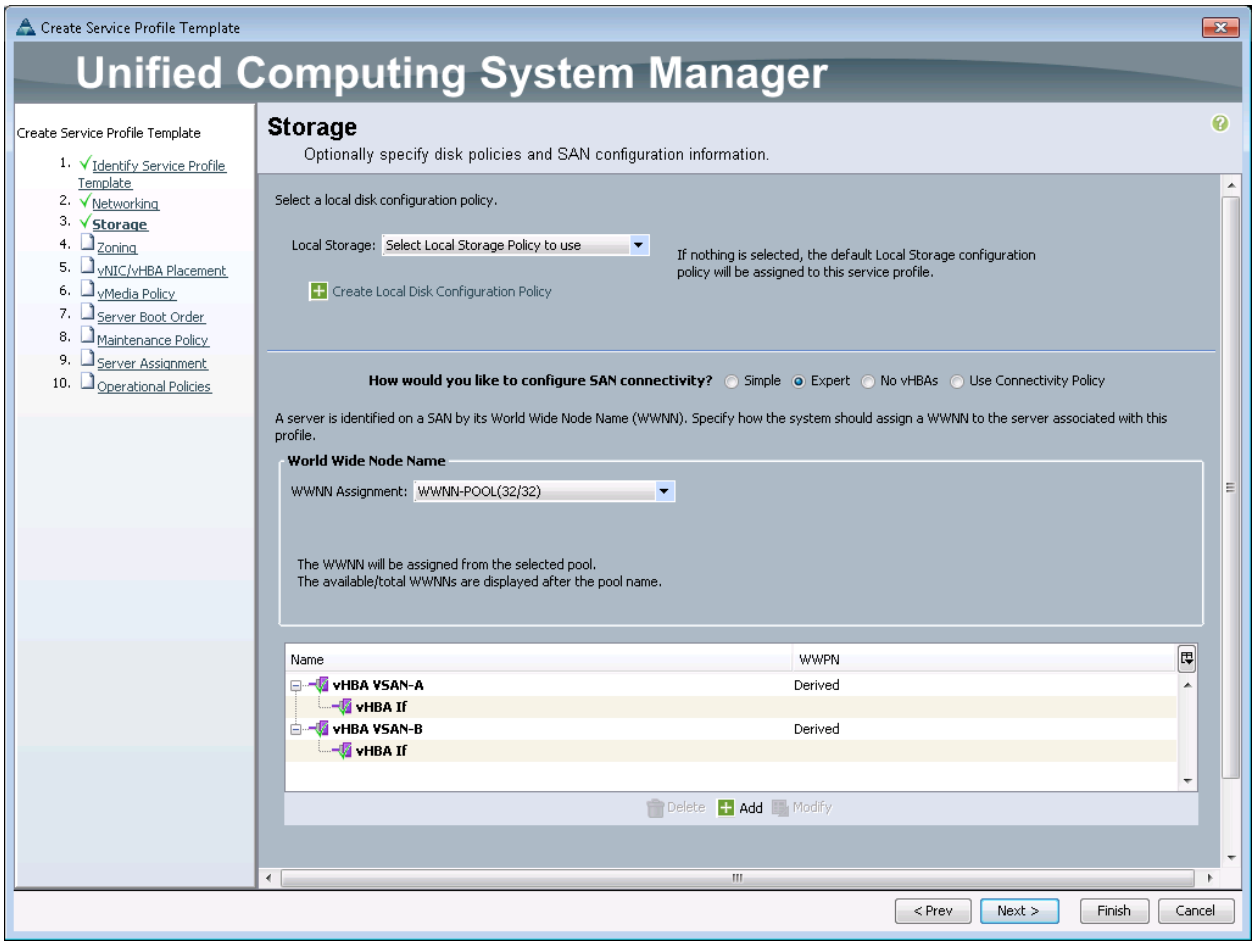

q. Click Next.

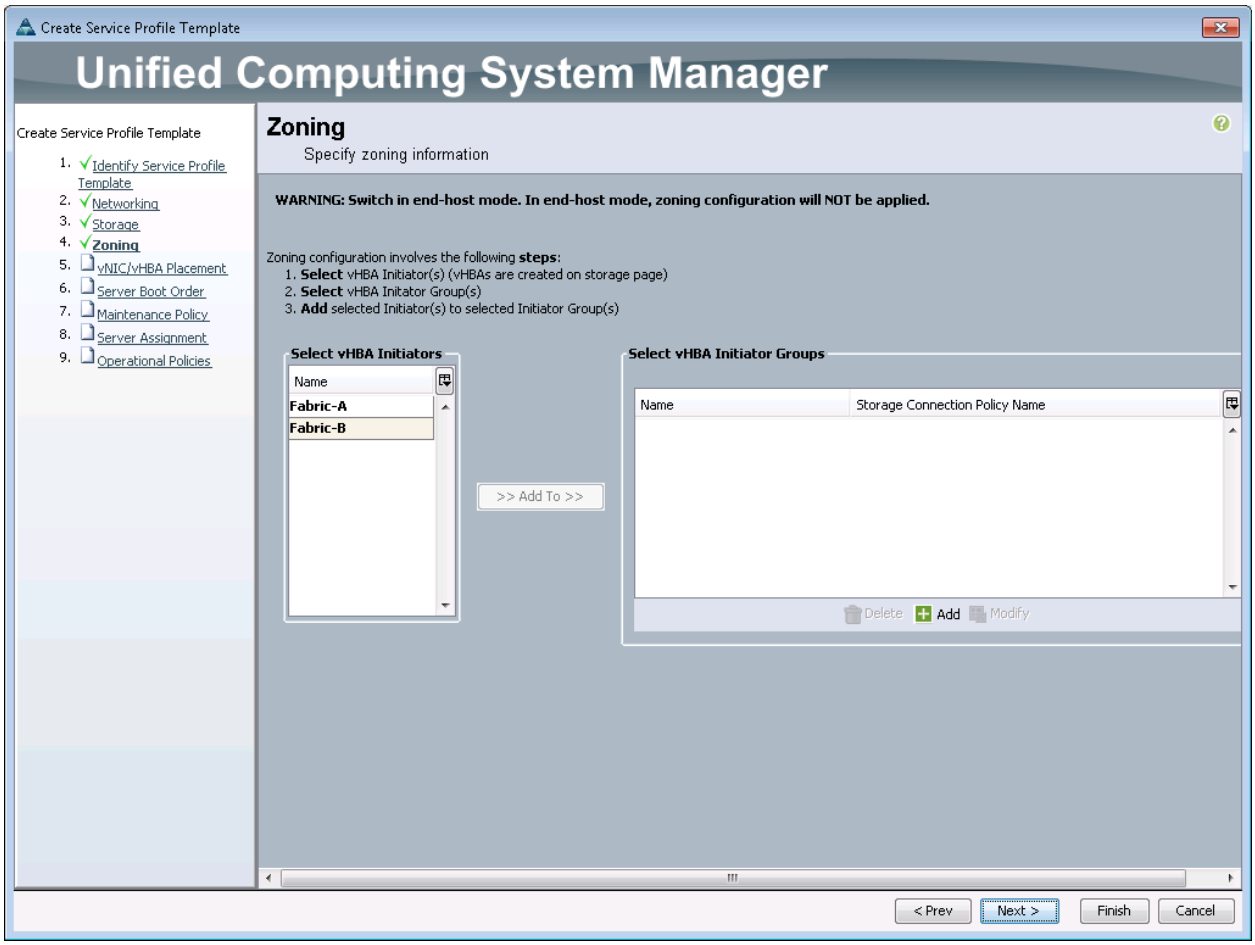

8. Set no Zoning options and click Next.
9. Set the following vNIC/vHBA placement options:
   a. In the Select Placement list, select the `RAC-HOST` placement policy.
   b. Select `vCon1` and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
      - `VHBA VSAN-A`
      - `VHBA VSAN-B`
      - `VNIC-MGMT`
      - `VNIC-RAC-CLUSTER`
   c. Review the table to verify that all of the vNICs and vHBAs were assigned to the policy in the appropriate order.

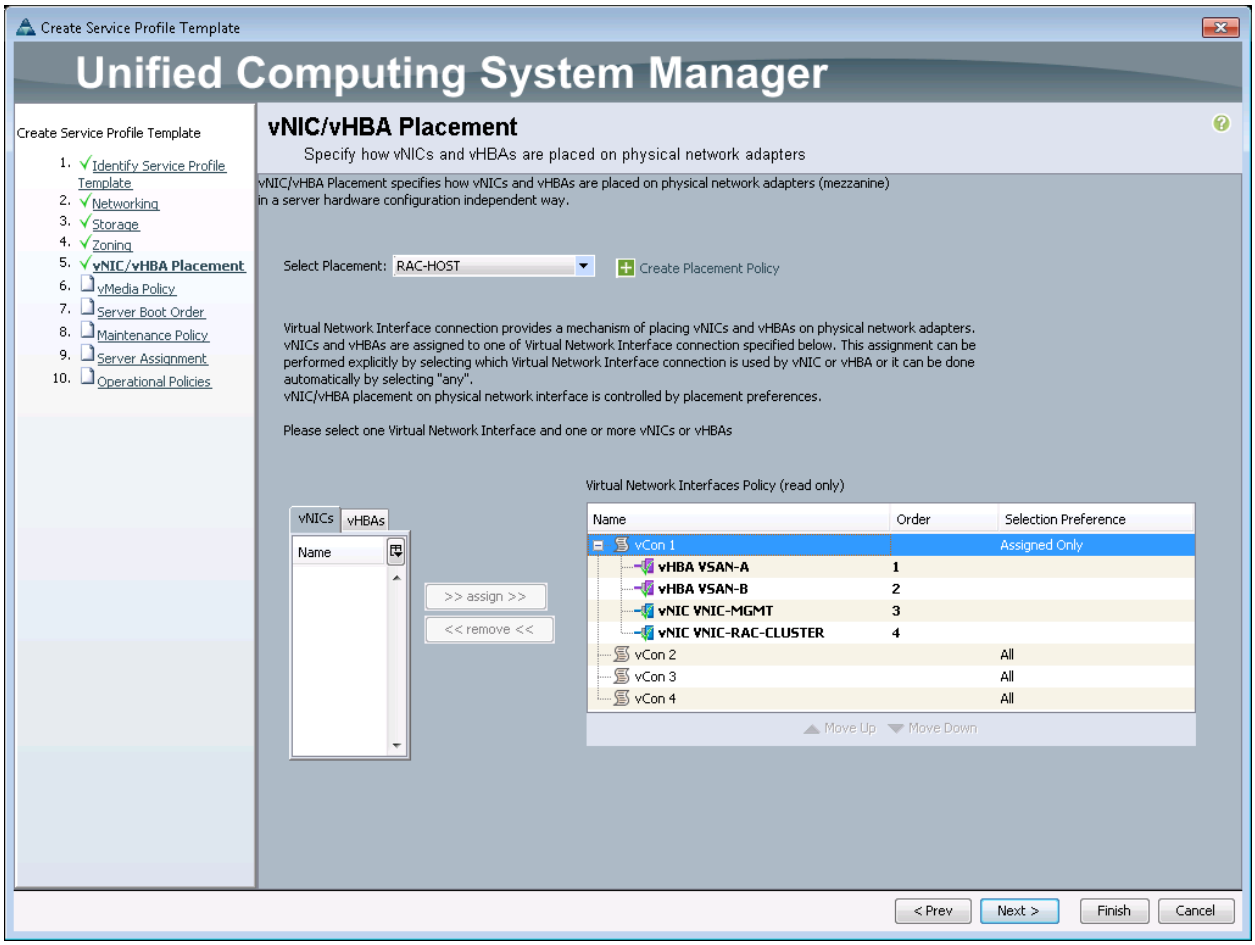

d. Click Next.

10. Set the Server Boot Order:

a. In the Boot Policy list, select `BOOT-FABRIC-A`.

b. Review the table to verify that all of the boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

c. Click Next.

11. Add a Maintenance Policy:

a. Confirm that maintenance policy is set to default.

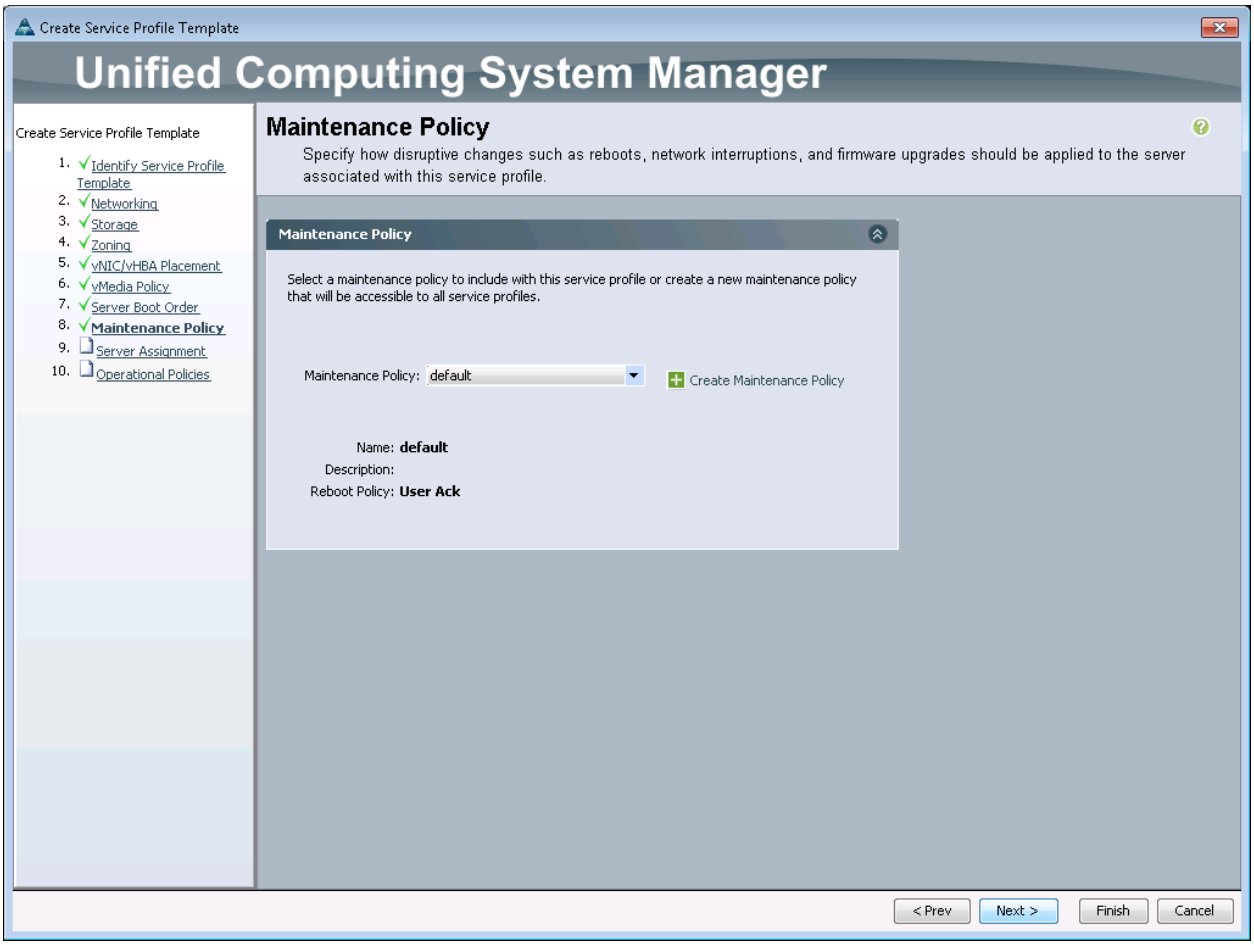

b. Click Next.

12. Specify the Server Assignment:
    a. In the Pool Assignment list, select `ORACLE-POOL`.
    b. Optional: Select a Server Pool Qualification policy.
    c. Select Down as the power state to be applied when the profile is associated with the server.
    d. Expand Firmware Management and select `RAC-HOST` from the Host Firmware list.

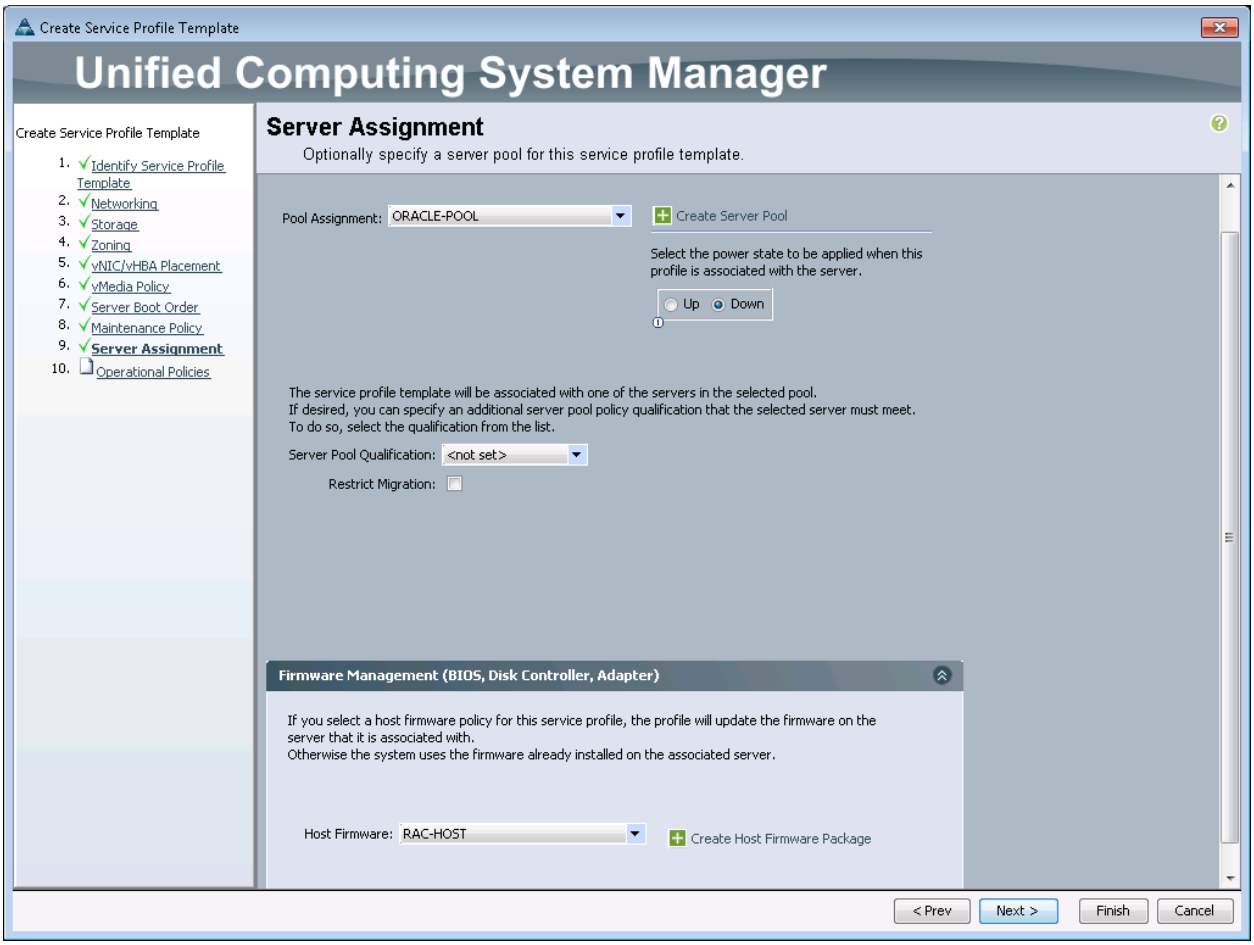

e. Click Next.

13. Add Operational Policies:

a. In the BIOS Policy list, select `RAC-HOST`.

b. Expand Power Control Policy Configuration and select `NO-POWER-CAP` in the Power Control Policy list.

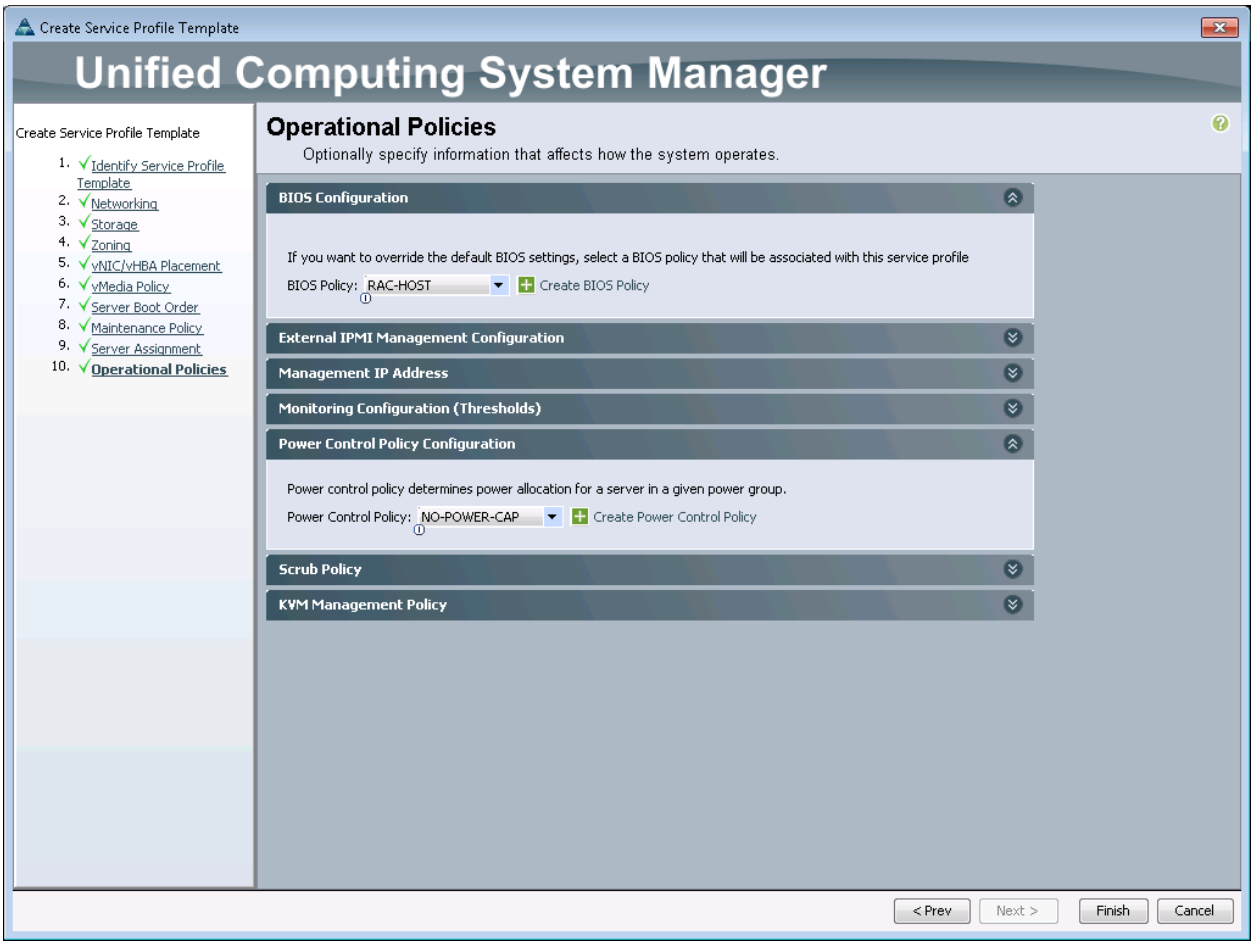

14. Click Finish to create the service profile template.
15. Click OK in the confirmation message.
16. Click the Servers tab in the navigation pane.
17. Select Service Profile Templates > Root.

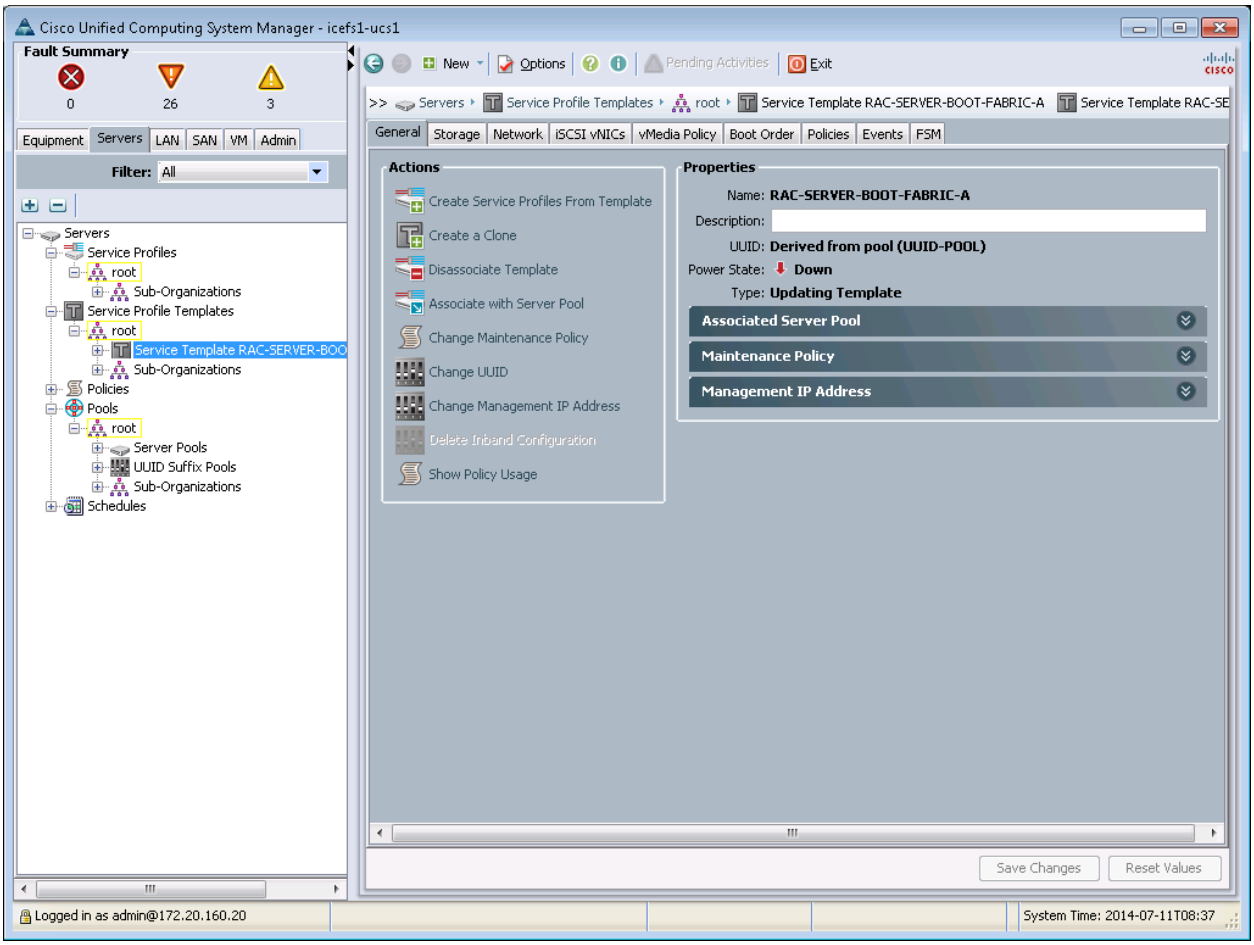

18. Right-click the previously created `RAC-SERVER-BOOT-FABRIC-A` template.
19. Select Create a Clone.
20. In the dialog box, enter `RAC-SERVER-BOOT-FABRIC-B` as the name for clone, select the Root Org, and click OK.

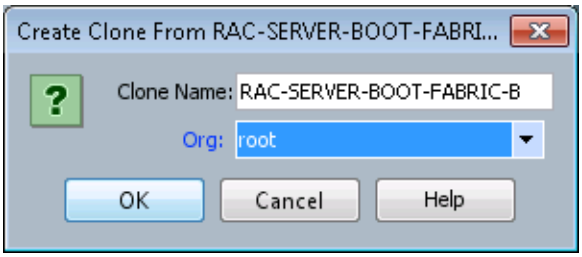

21. Click OK.
22. Select the newly cloned service profile template and click the Boot Order tab.

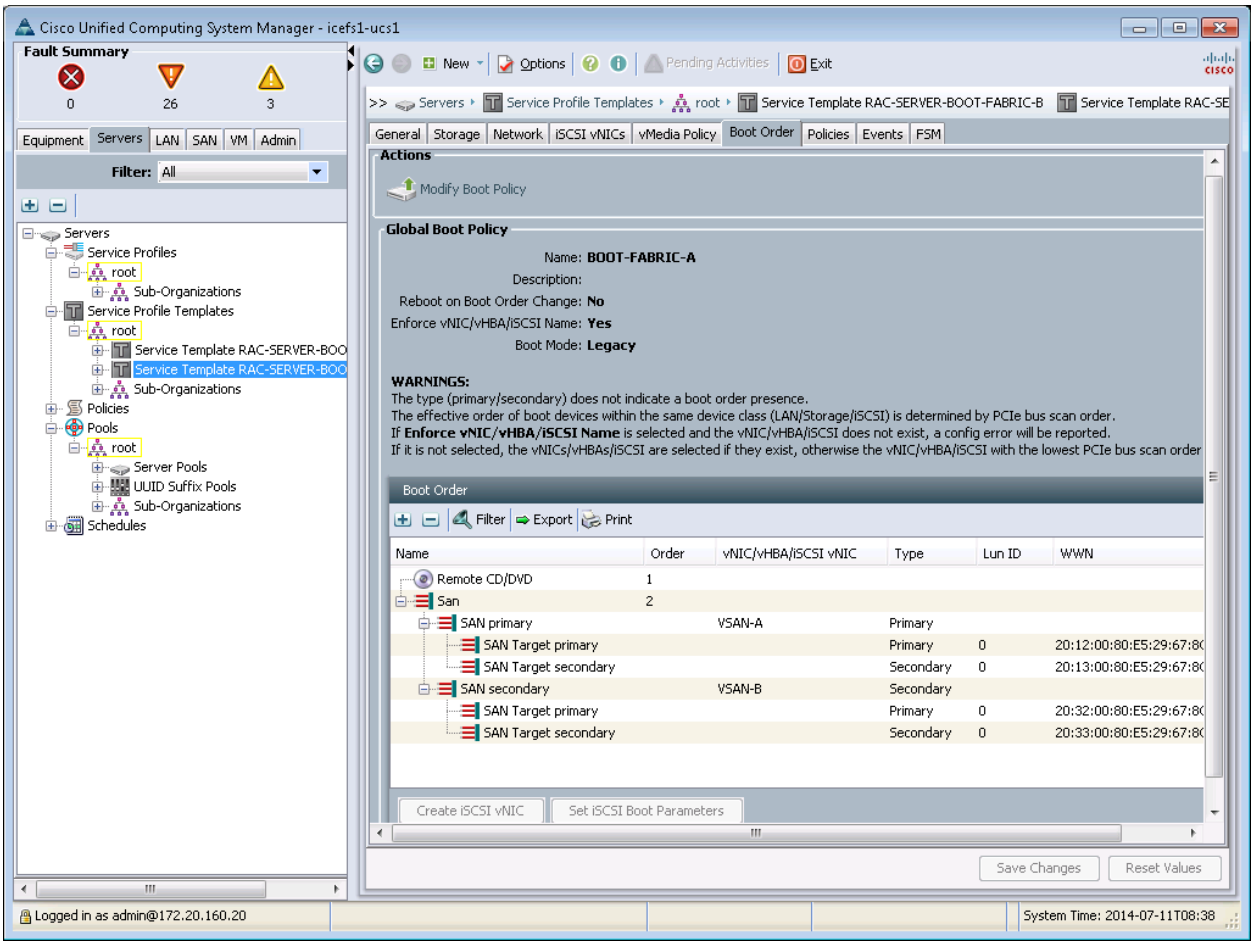

23. Click Modify Boot Policy.
24. In the Boot Policy list, select `BOOT-FABRIC-B`.

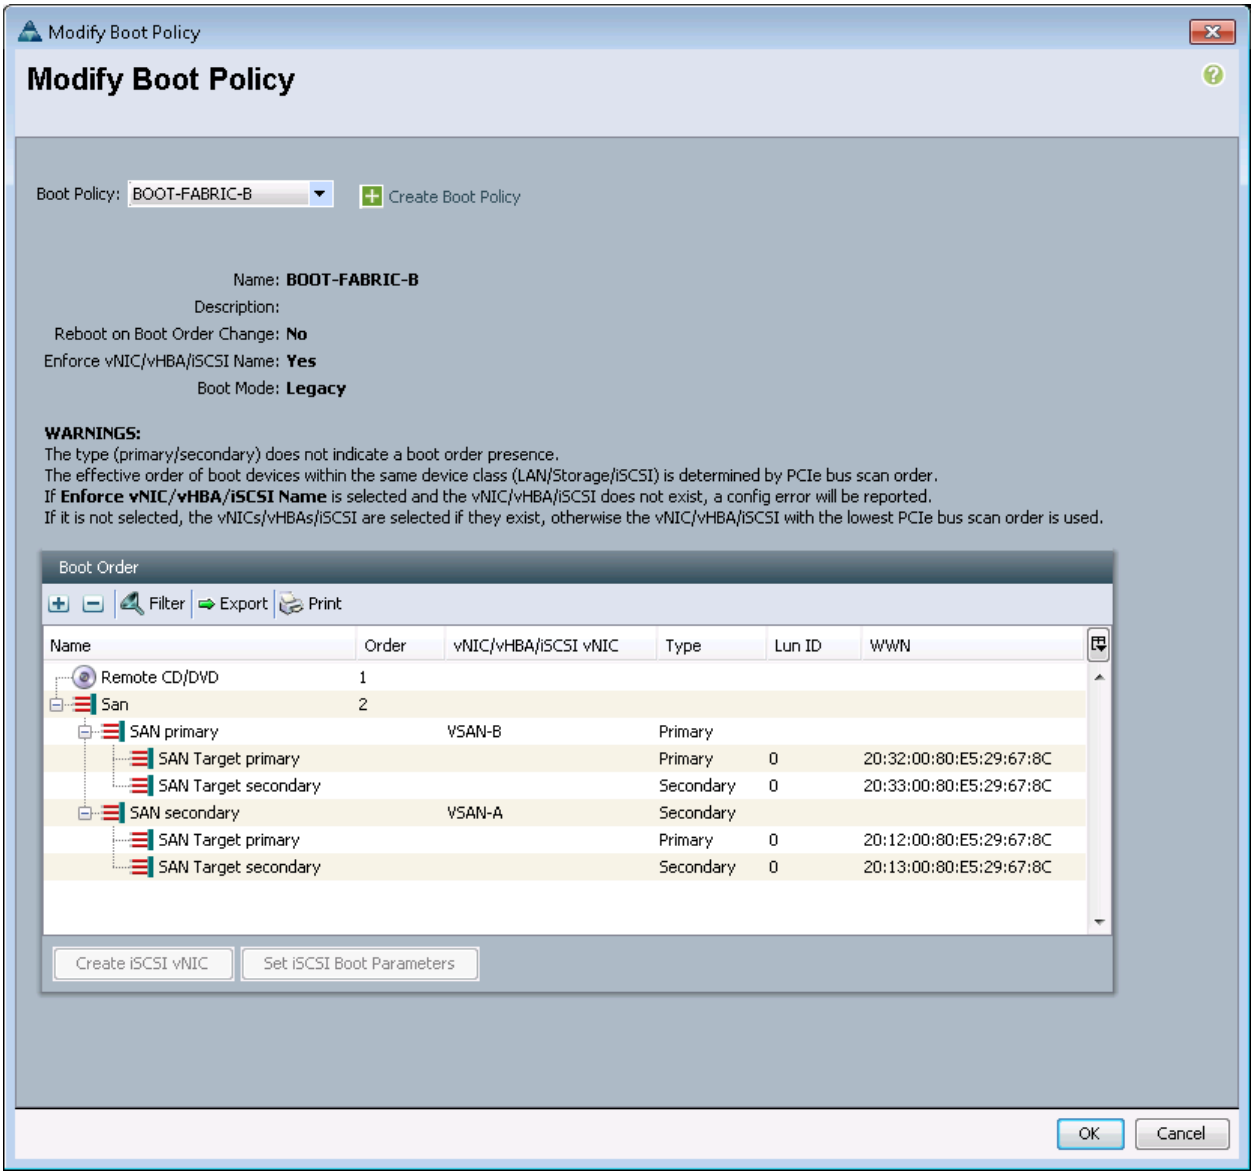

25. Click OK and then click OK again to close the confirmation window.
26. In the right pane, click the Network tab and then click Modify vNIC/HBA Placement.
27. Expand `vCon 1` and move `vHBA VSAN-B` ahead of `vHBA VSAN-A` in the placement order.
28. Click OK and then click OK again.

Create Service Profiles

To create service profiles from the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > Root > Service Template RAC-SERVER-BOOT-FABRIC-A.

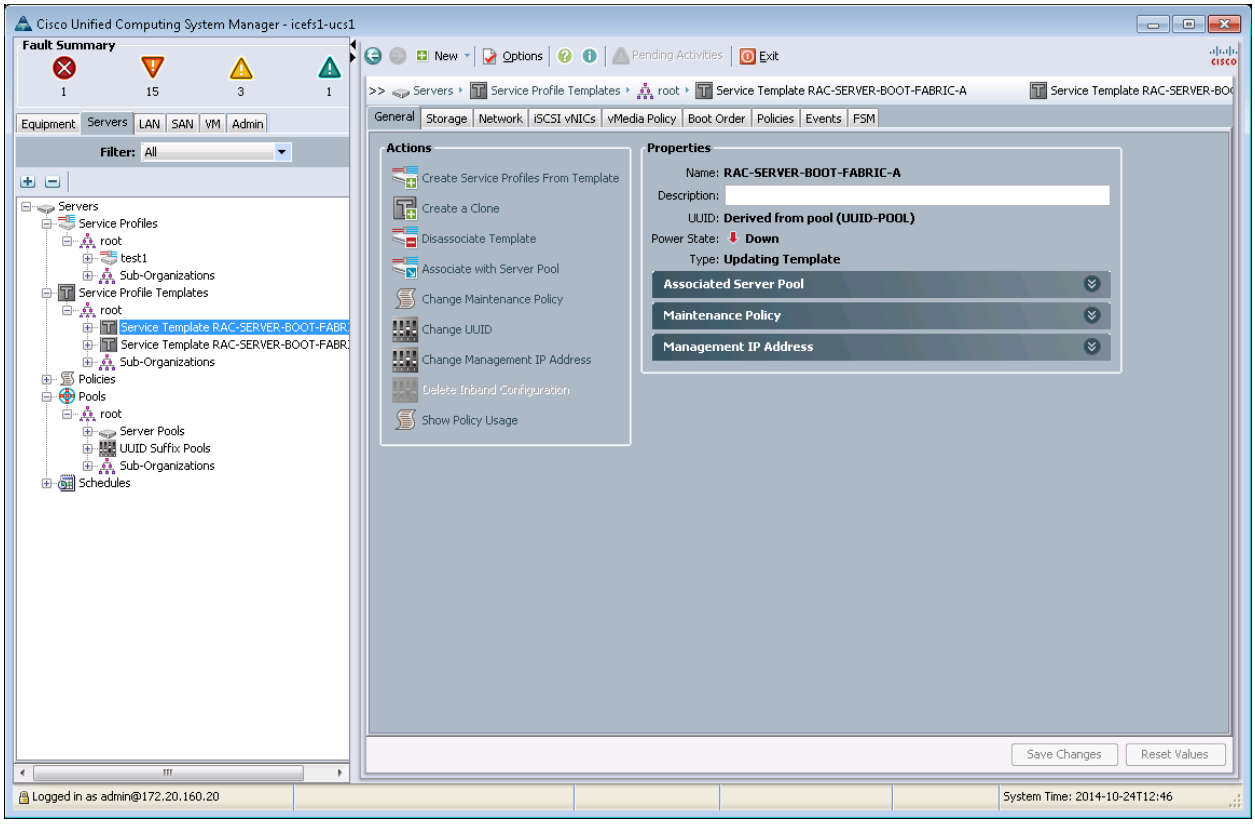

3. Right-click RAC-SERVER-BOOT-FABRIC-A and select Create Service Profiles from Template.
4. Enter `RAC-SERVER-BOOT-FABRIC-A-0` as the Naming Prefix.
5. Enter `1` as the Suffix Starting Number.
6. Enter `4` as the Number of Instances to create.

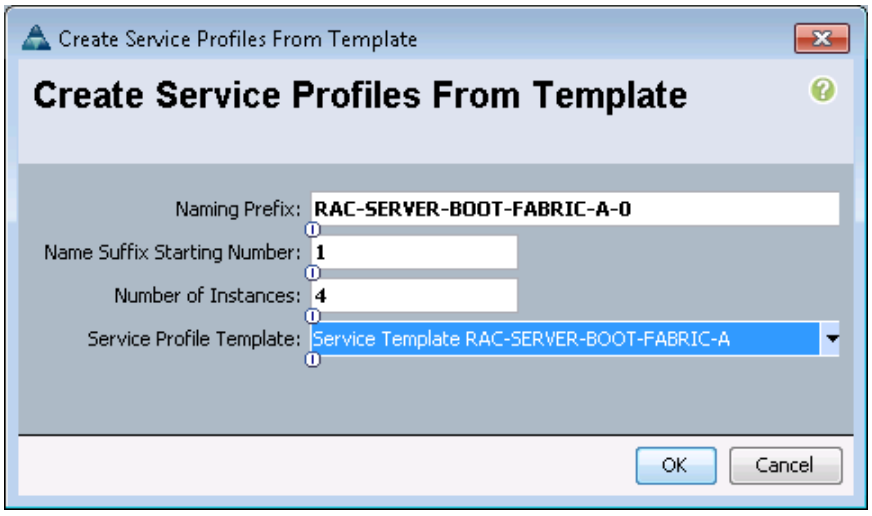

7. Click OK to create the service profile.
8. Click OK in the confirmation message.
9. Select Service Profile Templates > Root > Service Template RAC-SERVER-BOOT-FABRIC-B.

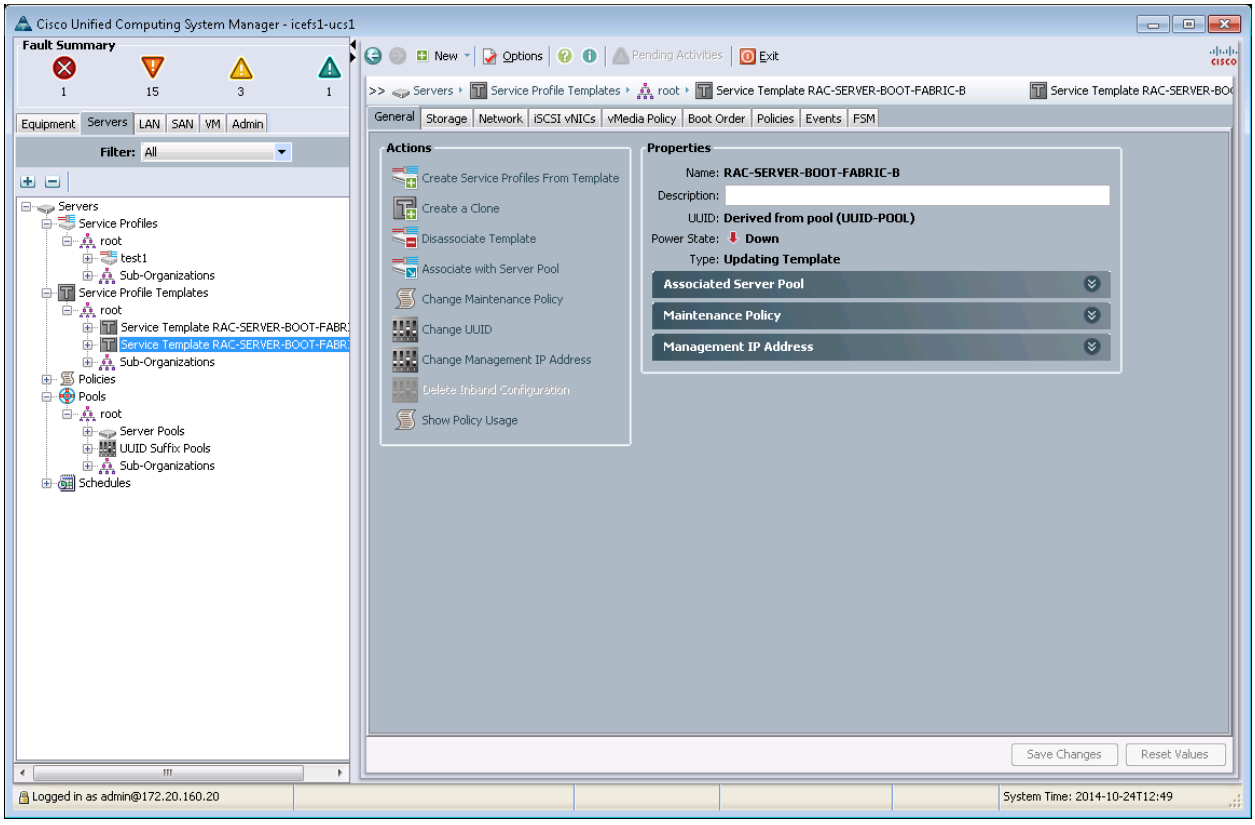

10. Right-click RAC-SERVER-BOOT-FABRIC-B and select Create Service Profiles from Template.
11. Enter `RAC-SERVER-BOOT-FABRIC-B-0` as the Naming Prefix.
12. Enter `1` for Name Suffix Starting Number.
13. Enter `4` for Number of Instances.

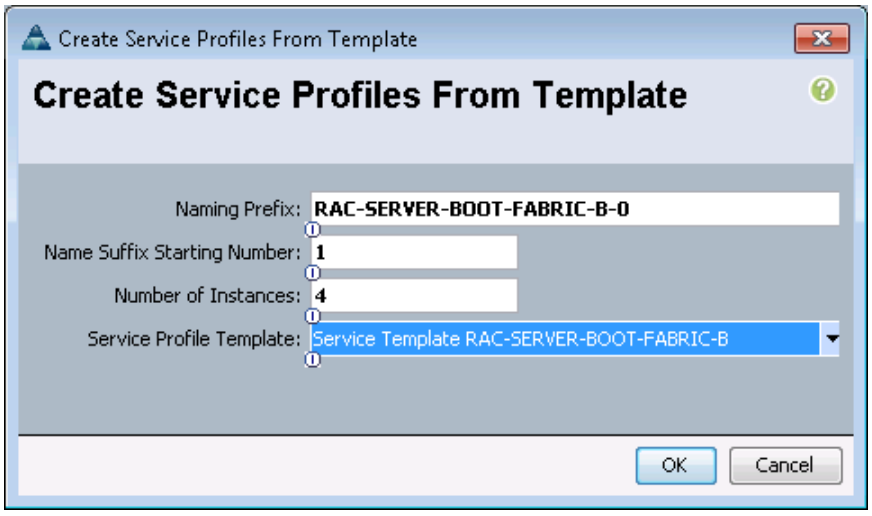

14. Click OK to create the service profile.
15. In the confirmation message, click OK.
16. Verify that all eight service profiles have been created. The service profiles are automatically associated with the servers from the servers assigned to the ORACLE-POOL server pool.

4.3 NetApp EF-Series Array Configuration

Table 20 provides requirements for configuring the NetApp EF-Series array, along with references and comments related to the configuration.

Table 20) Requirements, references, and comments for configuring NetApp EF-Series array.

Requirement	Reference	Comments
Storage system management software installation	NetApp E-Series Storage Systems Initial Configuration Guide and Software Installation for SANtricity Storage Manager 11.10	Refer to step 2 of the section "Installing the SANtricity Storage Manager Software" in the referenced document.
Storage system out-of-band management configuration	NetApp E-Series Storage Systems Initial Configuration Guide and Software Installation for SANtricity Storage Manager 11.10	Refer to step 6, option 2 of the referenced document.

Configure NetApp EF-Series Array

This section demonstrates using SANtricity with the EF560 to create the storage for the Oracle RAC database and map each of the RAC database nodes to the EF560 storage controllers. This requires the following procedures, which are described in detail in the following subsections:

- Create host groups in which to manage the Oracle RAC database servers.
- Create volume groups used to manage the EF560 storage provisioned for the Oracle RAC database nodes to use.
- Create the storage volumes used to boot the Oracle RAC database servers.
- Create the storage volumes used to house the Oracle RAC database.
- Provide access to the storage volumes from the Oracle RAC database nodes.

Table 21 shows the names and sizes of the volume groups, volumes, and host groups that were created on one of the four EF560 all-flash arrays. Comparable configurations were provisioned on the remaining three EF560 all-flash arrays to enable the Oracle RAC database to be provisioned evenly across all of the storage arrays.

Note: Refer to Table 24 in the appendix for the specific names that were used in this configuration. Table 21 is a sample from Table 24 to illustrate how each EF560 should be configured.

Table 21) Storage layout for one EF560 all-flash array.

Storage Array	Type	Volume Group RAID 10	Number of Physical Disks	Volume/LUN Name	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name
EF560-1	Data	A1VG1	10	DATALUN1	450	3.6TB	2	RAC
				DATALUN2	450			
				DATALUN3	450			
				DATALUN4	450			
		A2VG2	10	DATALUN5	450			
				DATALUN6	450			
				DATALUN7	450			

Storage Array	Type	Volume Group RAID 10	Number of Physical Disks	Volume/LUN Name	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name
				DATALUN8	450			
	Redo logs and boot LUNs	A1LOGVG	2	RAC1BOOT	78	256GB		
				RAC2BOOT	78			
				LOG1	50			
				LOG2	50			

Create Host and Host Group

In order to map the storage devices to an Oracle RAC database node, the first step is to create a host group that includes the associated FC HBA WWPNs for all of the Oracle RAC database server nodes. For this solution we created a single host group called “RAC” and included a total of eight hosts (one for each of the RAC nodes). All the storage devices are mapped to the RAC host group so that all the RAC nodes have access to those devices.

Follow these steps to create the host group and the hosts:

1. Launch the E-Series SANtricity management interface for one of the four EF560 storage arrays and select the Host Mappings tab.

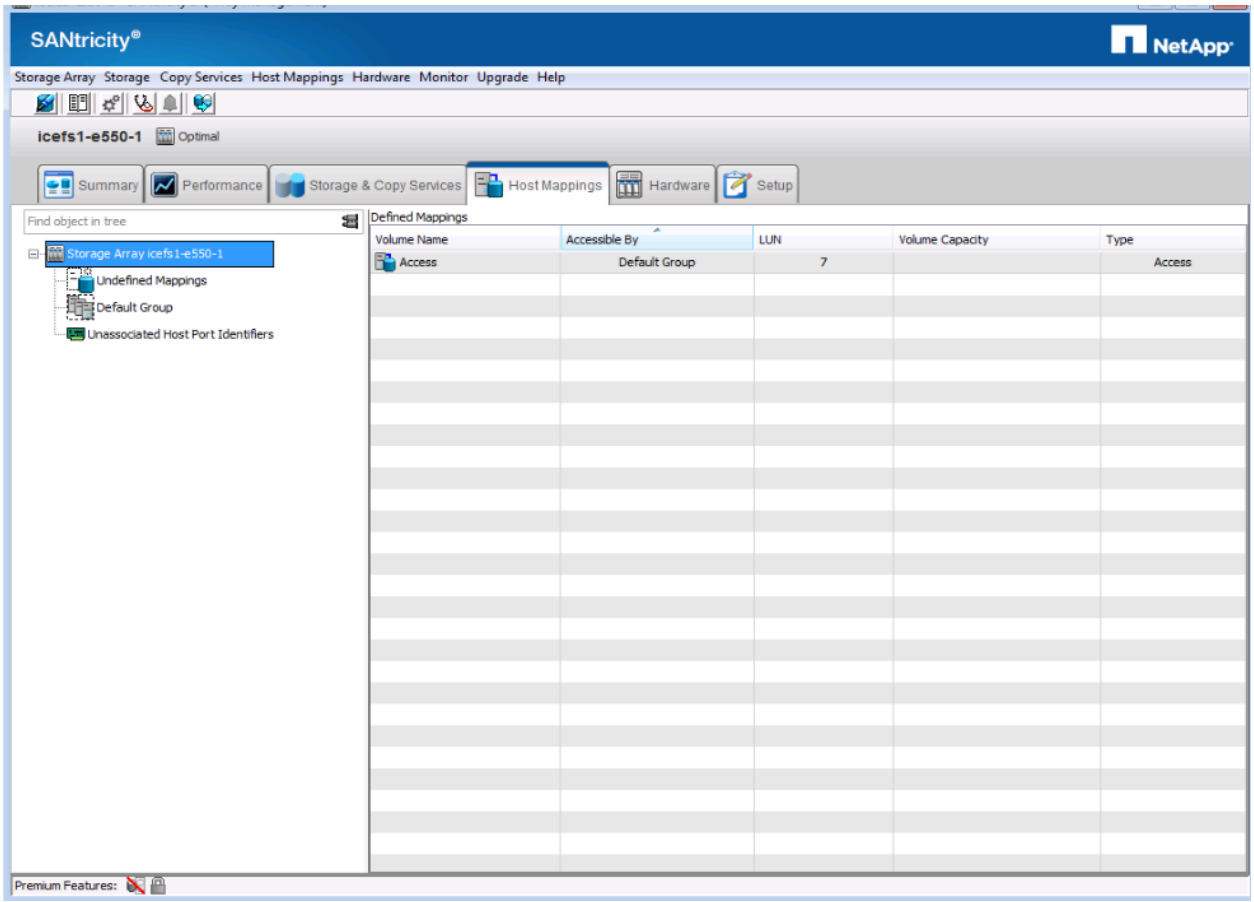

2. On the left navigation bar, right-click the Storage array.

3. Select Define.
4. Select Host Group.
5. In the menu to enter the host group name, give the host group name as `RAC`.

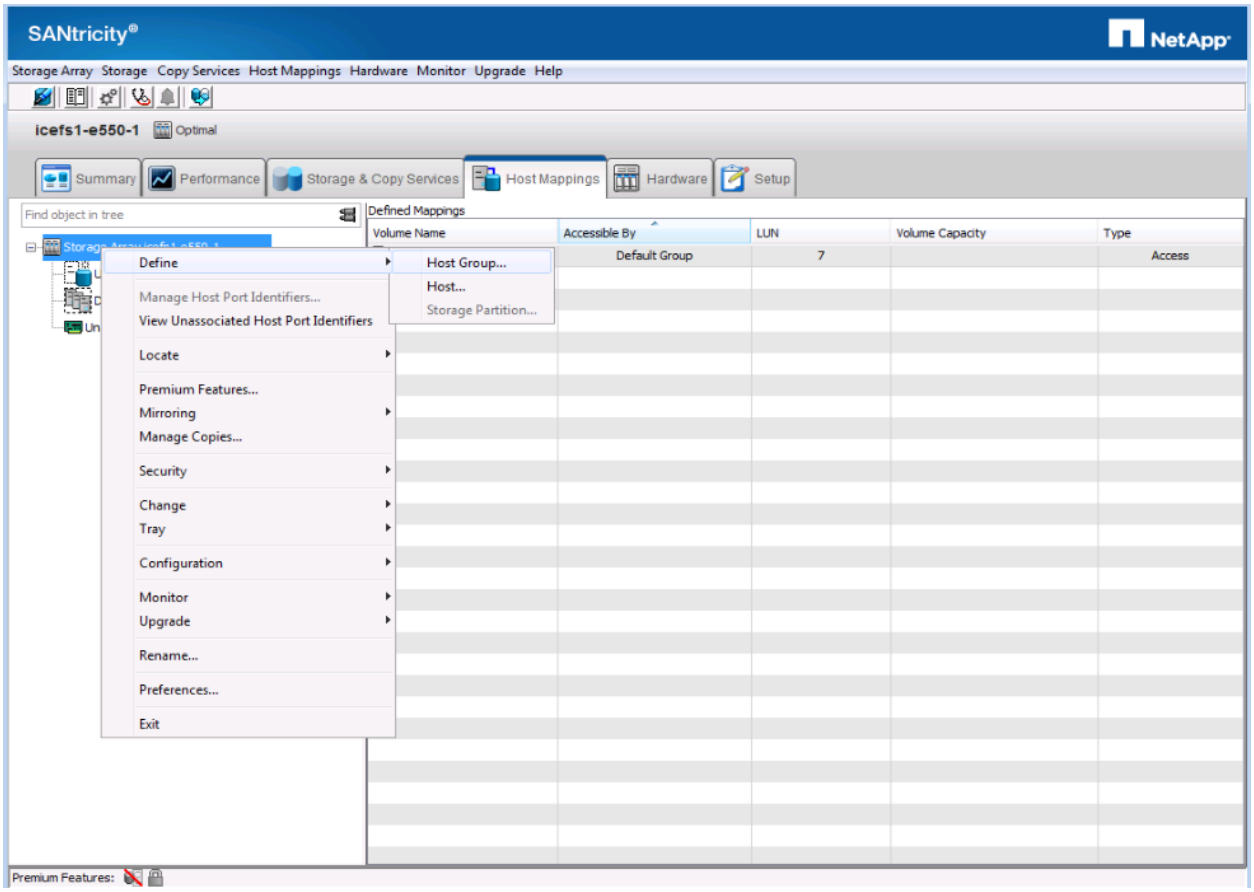

6. Click OK to complete the host group creation.

Create Volume Group

Now that you have created the host groups and added the hosts associated with the Oracle RAC nodes, the next step is to provision the storage on the EF560 for use by the Oracle RAC database. Complete the following steps to create the volume groups and associated LUNs used to provide storage for the Oracle RAC database.

To create a volume group from unconfigured capacity in the storage system using SANtricity, complete the following steps:

1. In the SANtricity Array Management window (AMW), click the Hardware tab and verify that the required number of hot spare drives have been allocated. For this setup, each EF560 system was provisioned with two hot spares.

   Note: For more information about configuring hot spare drives, refer to the section “Using Hot Spare Drives” in the SANtricity online help.

2. In the AMW, select the Storage and Copy Services tab, right-click Total Unconfigured Capacity, and select Create Volume Group.

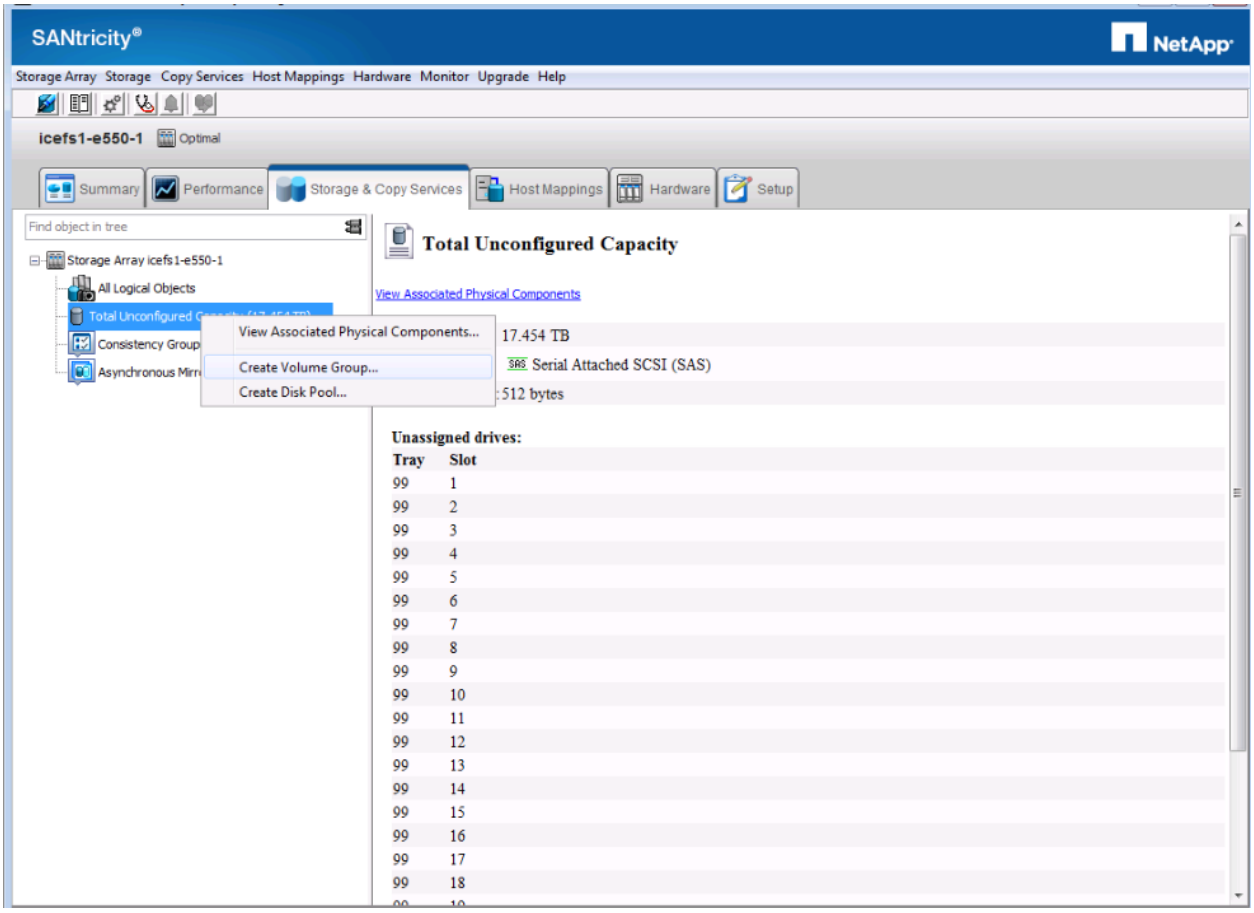

3. Click Next to proceed to the next screen.

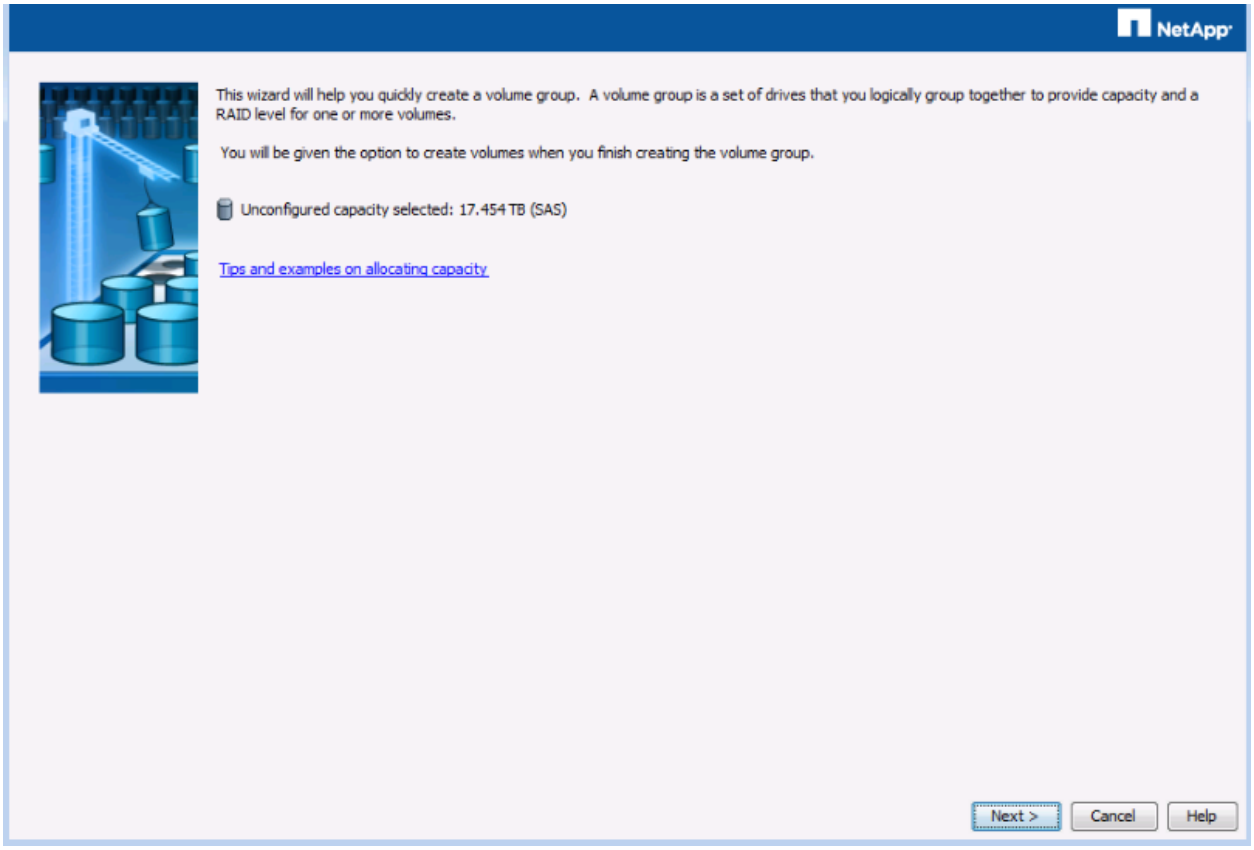

4. Provide a name for the new volume group. In this configuration, the volume group name is `A1VG1`.

   Note: Volume group names must not exceed 30 characters and must not contain spaces. The name string can contain letters, numbers, underscores (_), dashes (-), and pound signs (#).

5. To create a volume group automatically, select Automatic (Recommended) and click Next.

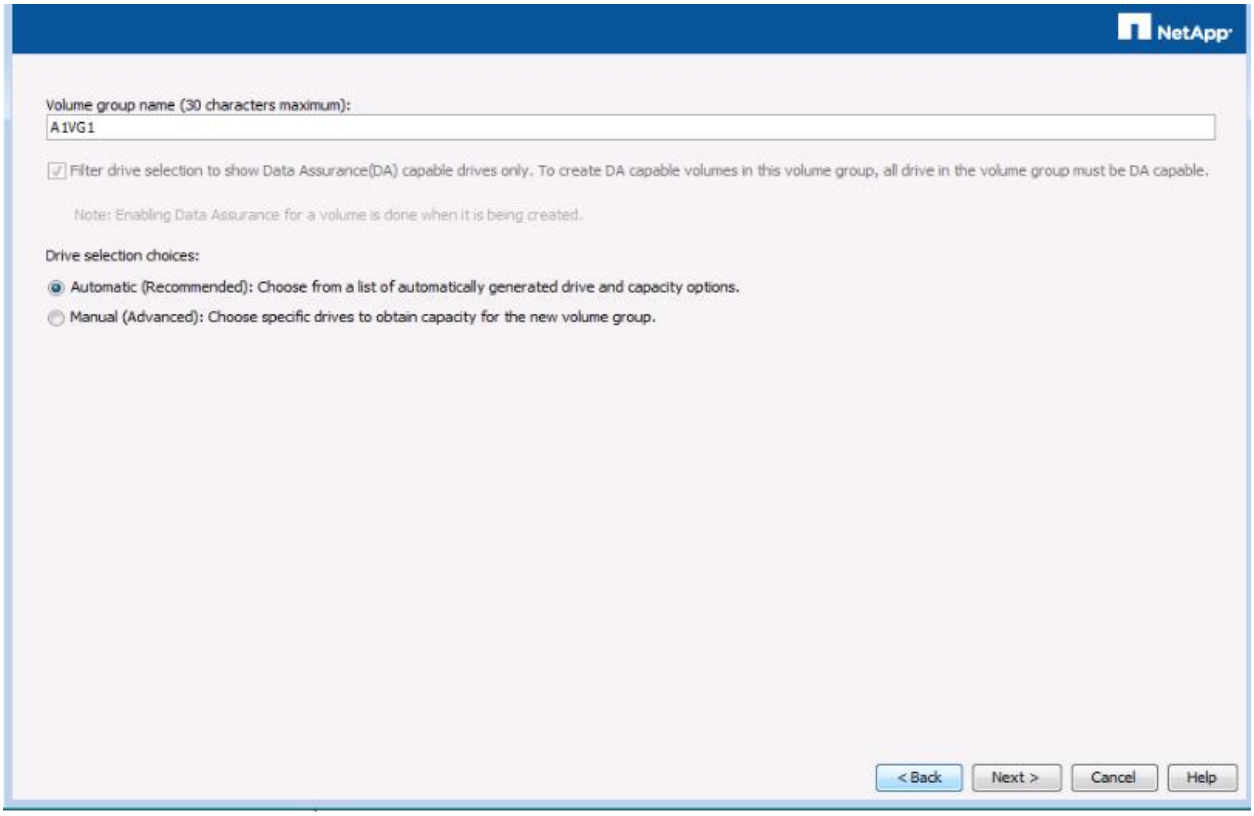

6. Select RAID level from the drop down. In this case we are selecting RAID10.
7. Select the number of drives. We are selecting 10 drives in this case.
8. Click finish to create the RAID group.
9. Repeat step 2 through step 8 to create all the required volume groups defined in Table 24. For this solution, a total of two volume groups were created with the RAID type as RAID 10, and one volume group was created with the RAID type as RAID 1, on each of the four EF560 systems:
   - Two volume groups were created with 10 drives for the Oracle data files.
   - One volume group was created with 2 drives for the redo logs and RAC-node boot LUNs.

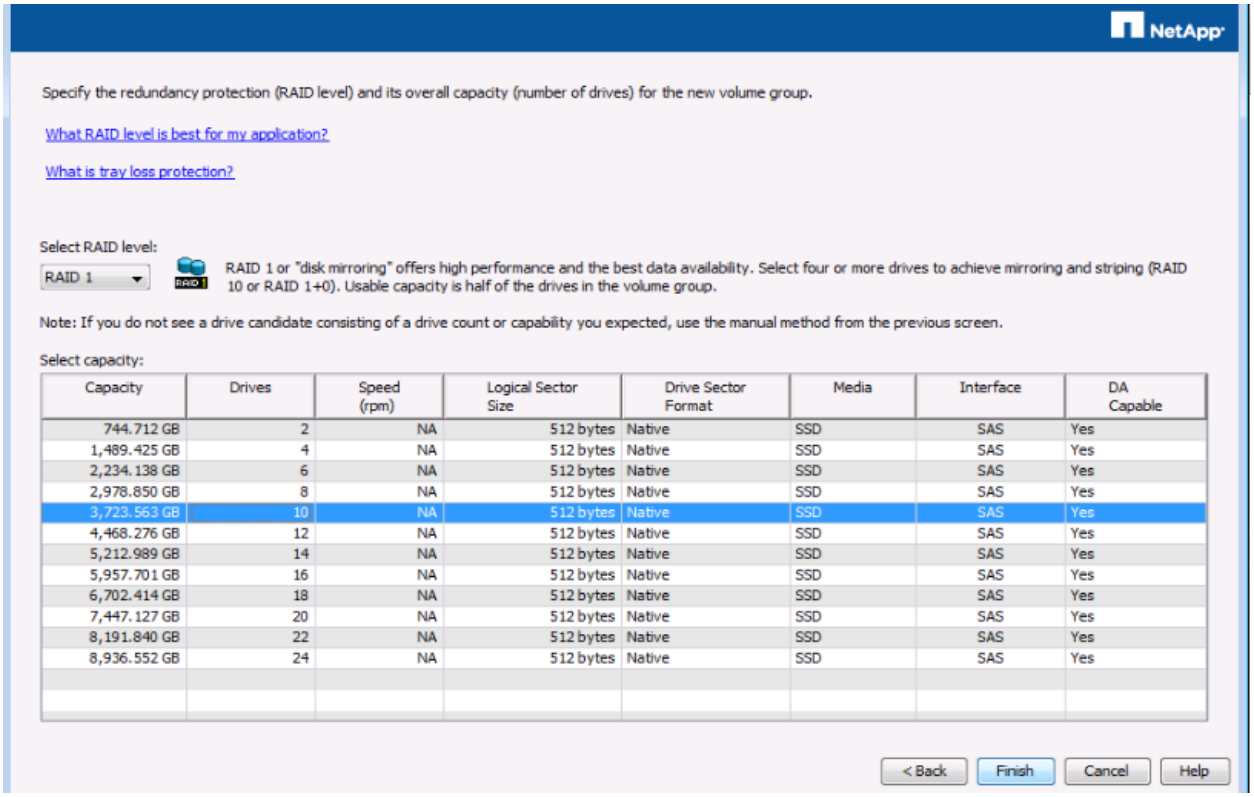

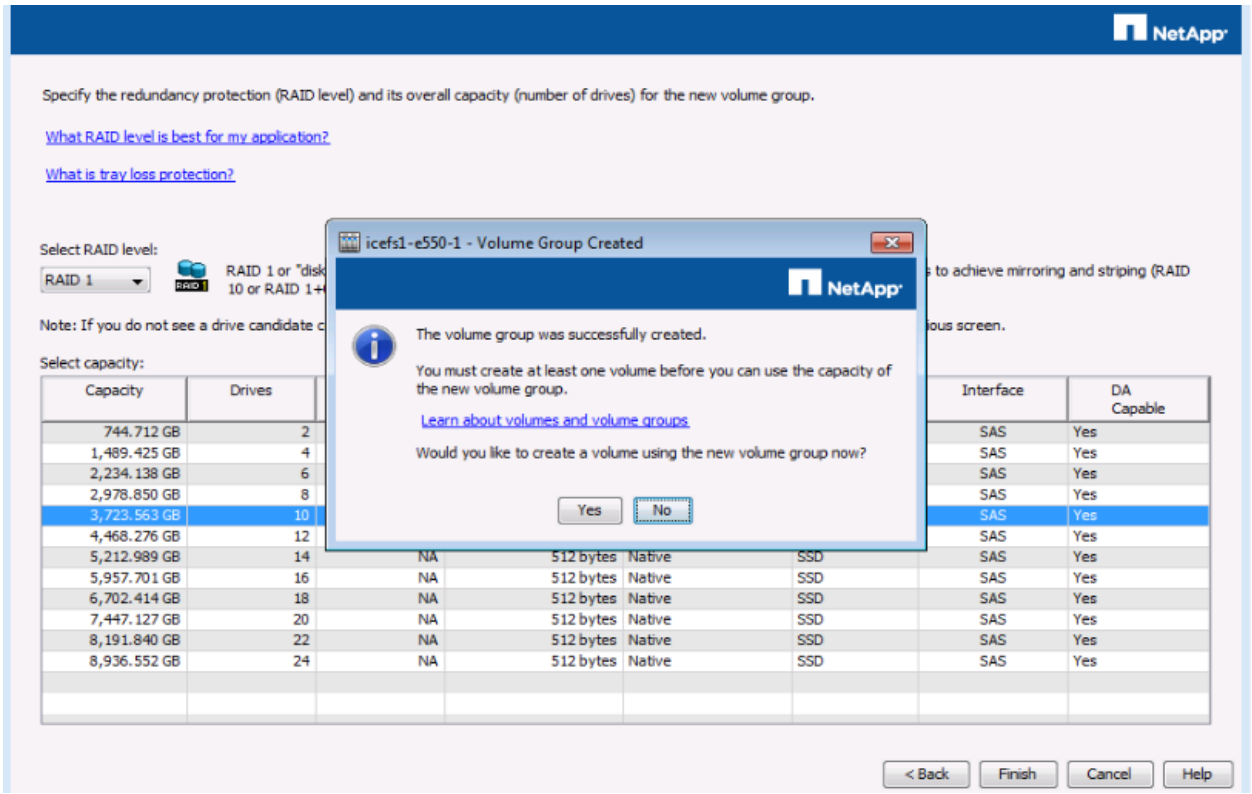

Create Volume for Boot LUN

This section describes using the A1LOGVG volume group to provision storage for the Oracle Linux boot LUNs used by the Oracle RAC nodes. This process creates a series of volumes (LUNs) and associates them with a specific volume group to facilitate manageability.

Note: Before the Oracle provisioning, NetApp recommends that you create the boot LUNs for the RAC database nodes.

To create a volume, complete the following steps:

1. Select the volume group that was created for the redo logs and boot LUNs. In this example, the A1LOGVG volume group was used to create the boot LUNs for two RAC nodes.
2. Right-click Free Capacity and select Create Volume.

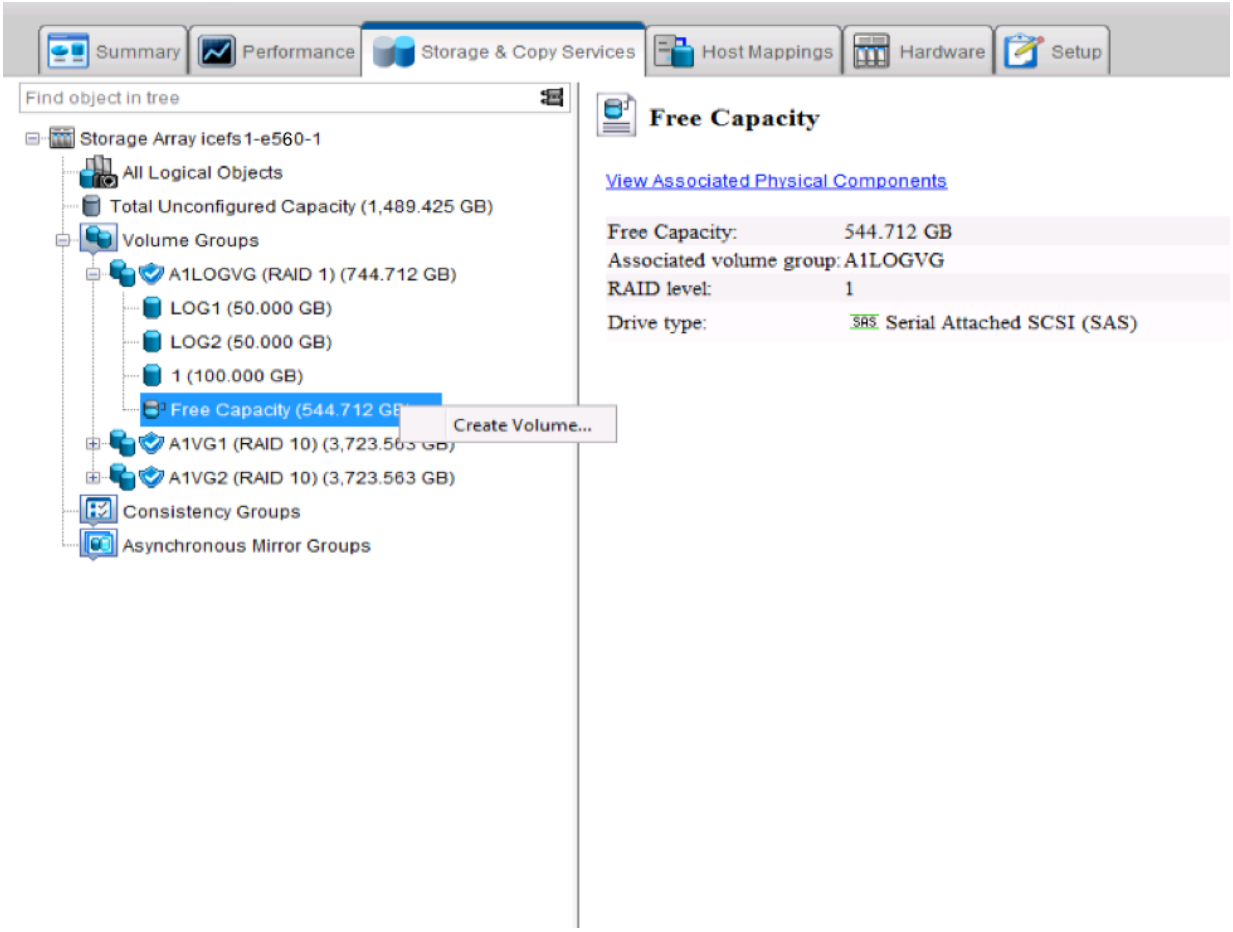

3. Enter the new volume capacity from the available capacity in volume group A1LOGVG. For this setup, we used 100GB for each boot LUN for the RAC nodes.
4. Enter a new volume name, in this case `RAC1BOOT`.
5. From the Map to Host drop-down list, select the host name.
6. Disable data assurance.
7. Select File System for the Volume I/O characteristics type.
8. Disable dynamic cache read prefetch.
9. Click Finish to create the new volume.

10. Repeat step 1 through step 9 for each RAC node listed in step 5 earlier so that each RAC node has a configured boot LUN. After the boot LUN creation, the setup used for this solution has entries as provided in Table 22.

    Note: Although it is not a requirement, NetApp recommends that you install the Oracle Linux operating system on all the RAC nodes immediately after creating the boot LUNs. This makes it easier to identify the correct LUNs on the EF560 that is provisioned for the process. If you would like to install the Oracle Linux OS now, perform the procedures in the section “Oracle Linux Installation on RAC Nodes.”

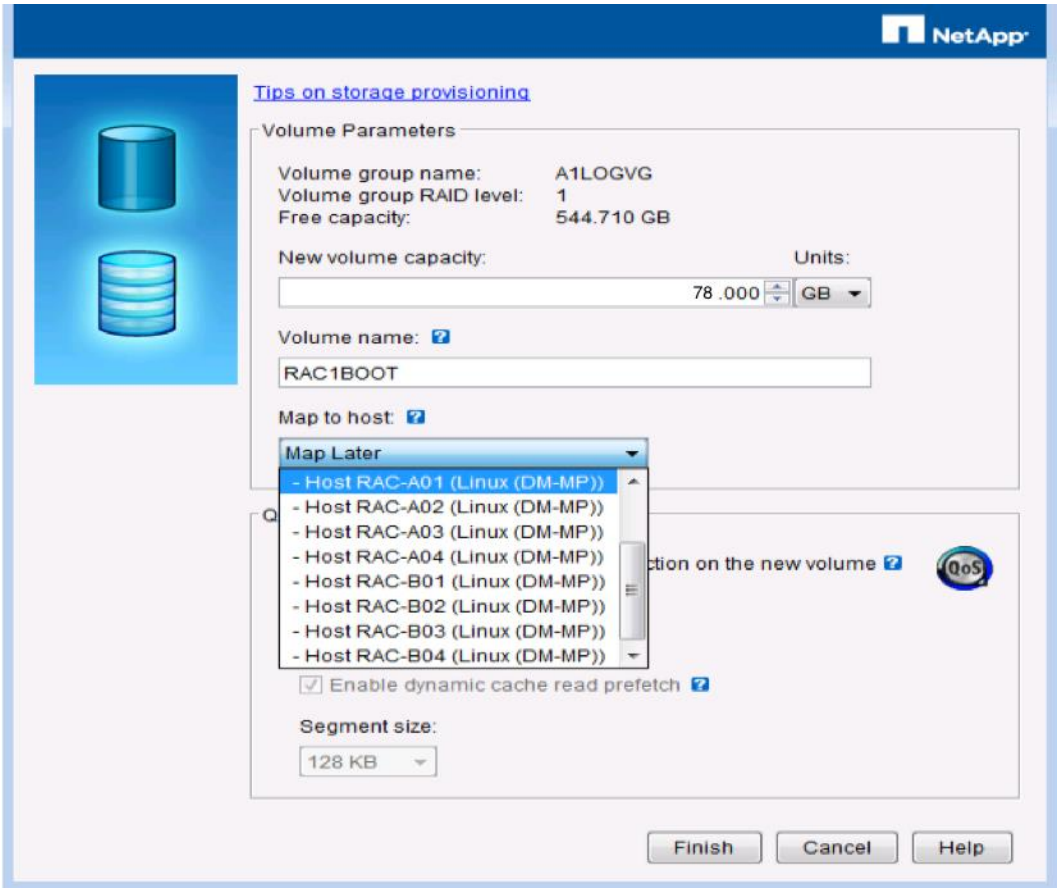

Table 22) Storage layout for boot LUNs for RAC nodes.

Storage Array	Type	Volume Group RAID 1	Number of Physical Disks	Volume/ LUN Name	Allocated Capacity in GB	Mapped Host Name	Total Allocated Capacity	Spare Disks
EF560-1	Boot LUNs for the RAC nodes	A1LOGVG	2	RAC1BOOT	78	RAC-A01	624GB	2
				RAC2BOOT	78	RAC-A02		
EF560-2		A2LOGVG		RAC3BOOT	78	RAC-A03		
				RAC4BOOT	78	RAC-A04		
EF560-3		A3LOGVG		RAC5BOOT	78	RAC-B01		
				RAC6BOOT	78	RAC-B02		

Storage Array	Type	Volume Group RAID 1	Number of Physical Disks	Volume/ LUN Name	Allocated Capacity in GB	Mapped Host Name	Total Allocated Capacity	Spare Disks
EF560-4		A4LOGVG		RAC7BOOT	78	RAC-B03		
				RAC8BOOT	78	RAC-B04		

Volume Creation for Oracle RAC Grid

This section explains how to create the disks for the grid cluster. A total of three disks are created for the grid infrastructure. For this setup, the second EF560 array out of the four is used to house the disks for the grid. Follow the steps to create the disks for grid infrastructure.

1. Launch the SANtricity GUI for the second storage array EF560-2.
2. Select one of the volume groups you created for the Oracle installation. This is A2VG1 in this example.
3. Select free capacity, right-click, and select Create Volume.
4. Enter the new volume capacity from the available capacity in the volume group A2VG1. For this setup we used 5GB for each Oracle Cluster Ready Services (CRS) disk.
5. Enter a new volume name, in this case CRS1.
6. From the Map to host drop-down list, select Host Group RAC.
7. Disable Data Assurance.
8. Select File System in the Volume I/O characteristics type.
9. Disable dynamic cache read prefetch.
10. Click Finish to create the new volume.
11. You are prompted to create another volume within the same volume group.
12. Click Yes.
13. Repeat step 4 through step 10 to create a total of three 5GB CRS LUNs (CRS1 through CRS3) for the volume group A2VG1. Table 23 shows the results of the CRS disk setup after all three disks are created.

Table 23) Storage layout for CRS disks for RAC cluster.

Storage Array	Type	Volume Group RAID 10	Number of Physical Disks	Volume/ LUN Name	Allocated Capacity in GB	Mapped Host Group Name	Total Allocated Capacity	Spare Disks
EF560-2	CRS disk	A2LOGVG	2	GRID1	5	RAC	5GB	2
				GRID2	5			
				GRID3	5			

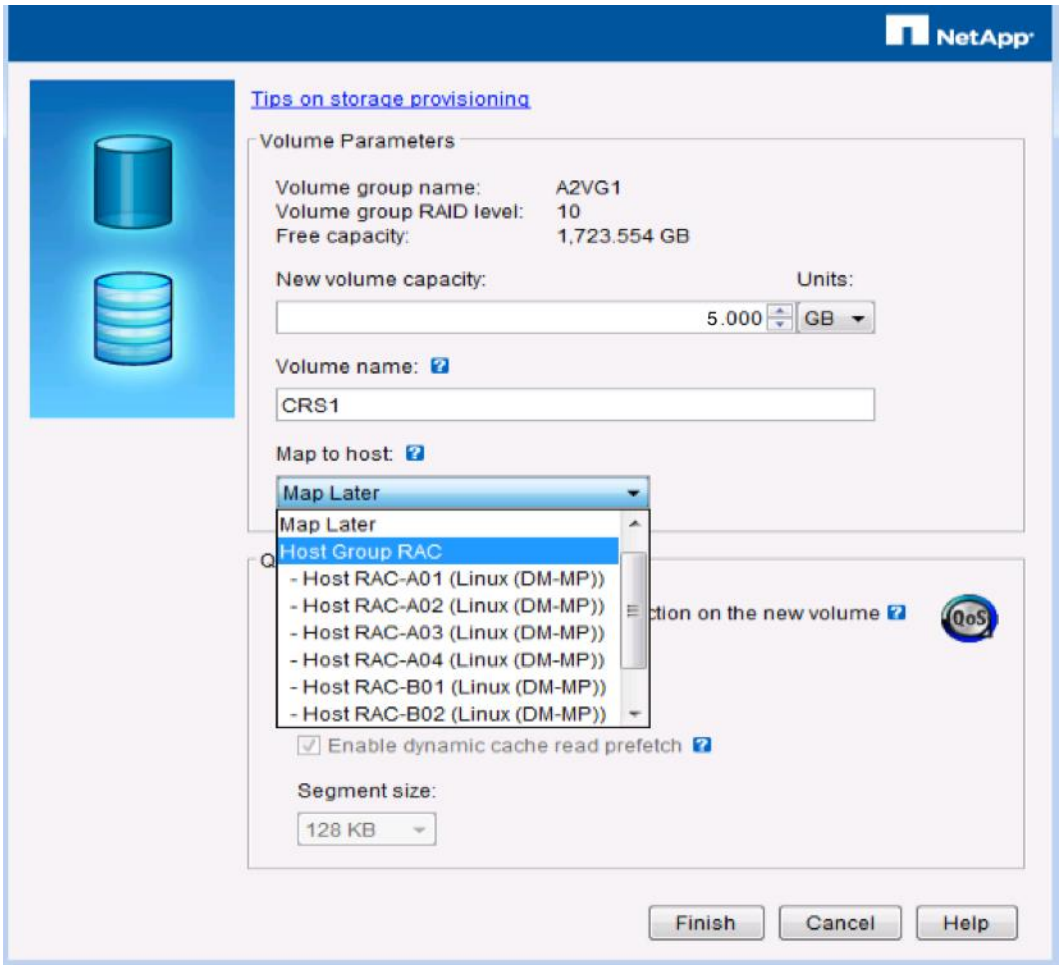

Volume Creation for Oracle RAC Database

This section describes using the A1VG1 volume group to provision storage for the Oracle RAC database accessed by the Oracle RAC nodes. This process creates a series of volumes (LUNs) and associates them with a specific volume group to facilitate manageability. To create a volume, complete the following steps:

1. Select one of the volume groups you created for the Oracle installation. This volume group in this example is A1VG1.
2. Select Free Capacity, right-click, and select Create Volume.

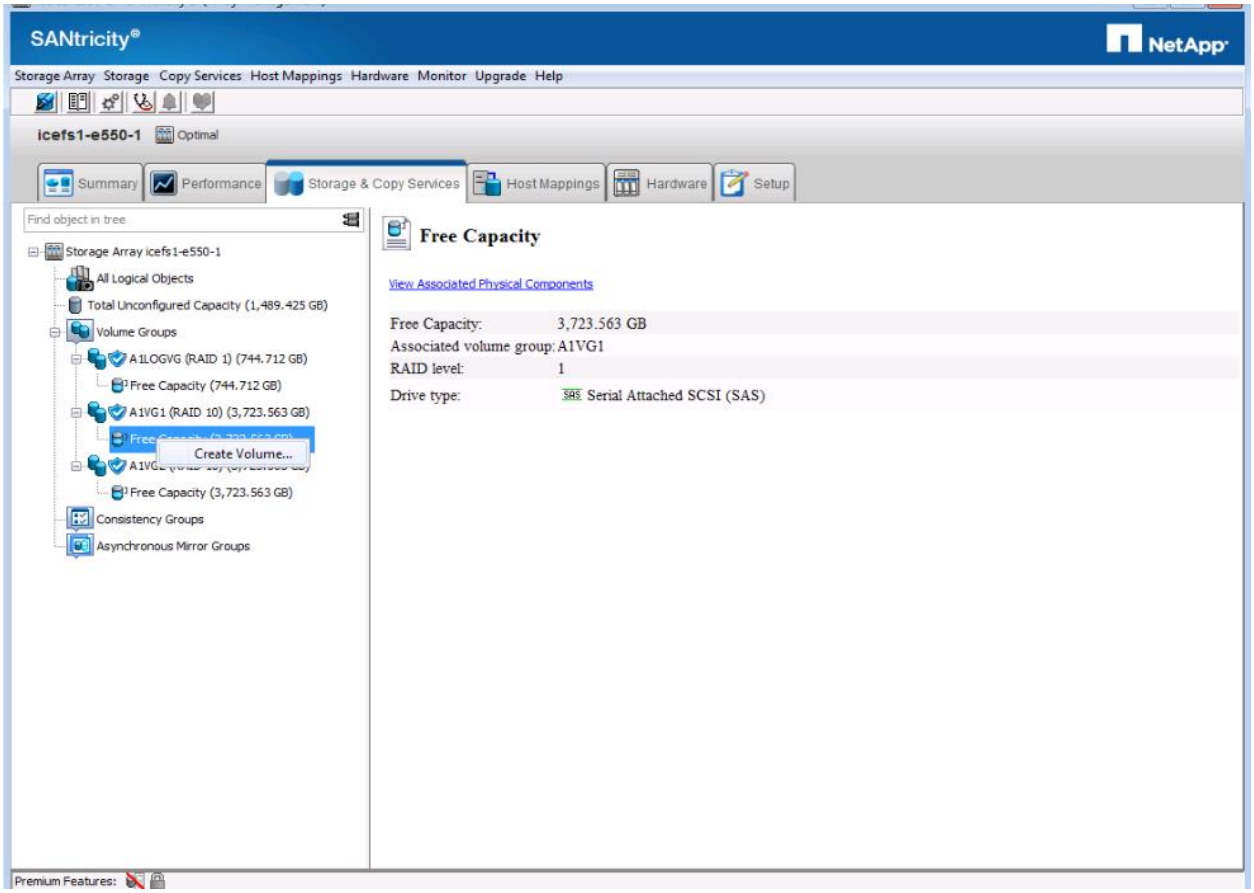

3. Enter the new volume capacity from the available capacity in the volume group A1VG1. For this setup, we used 450GB for each volume.
4. Enter a new volume name, in this case `DATALUN1`.

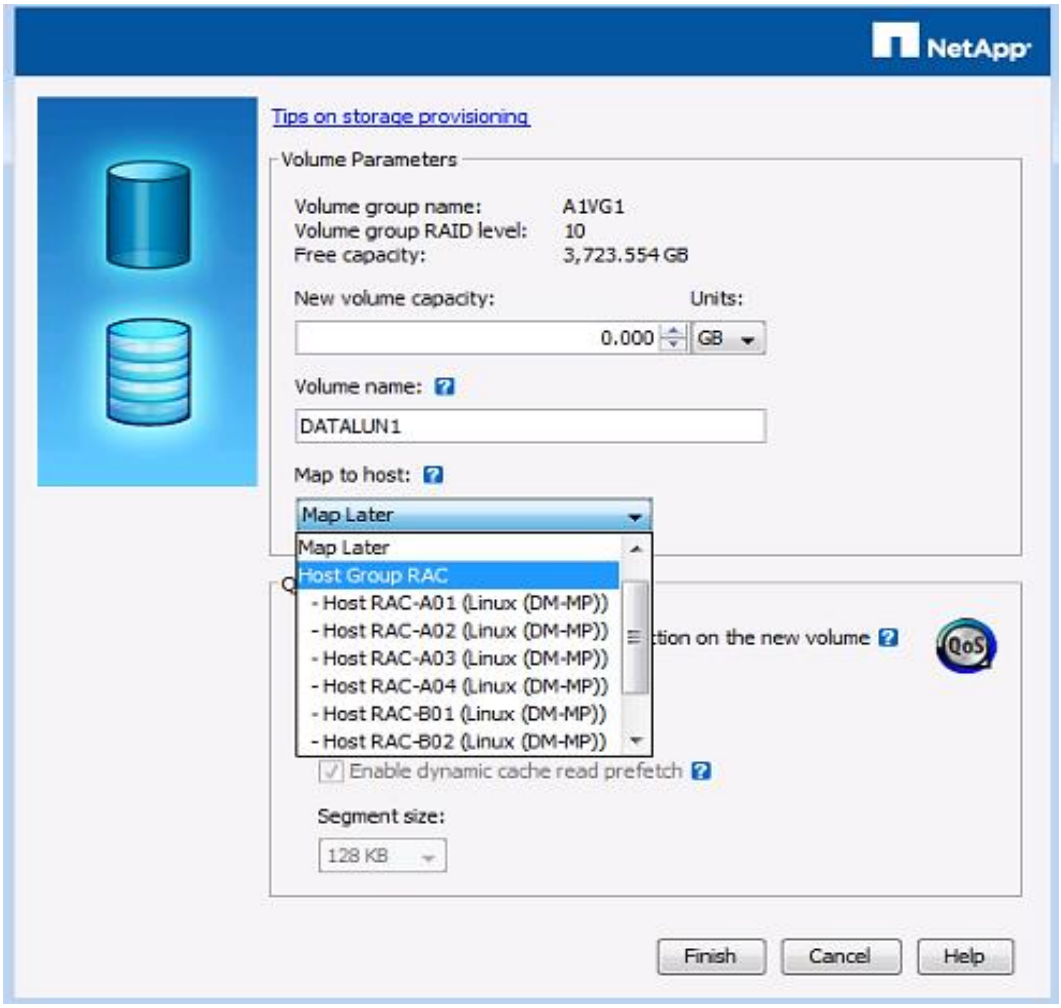

5. From the Map to Host drop-down list, select Host Group RAC.
6. Disable Data Assurance.
7. Select Database in the Volume I/O characteristics type.
8. Disable dynamic cache read prefetch.
9. Click Finish to create the new volume.
10. You are prompted to create another volume within the same volume group.
11. Click Yes.
12. Repeat step 3 through step 9 to create a total of four 450GB data LUNs (DATALUN1 through DATALUN4) for the volume group A1VG1.
13. Repeat step 1 and select second volume group A1VG2 to create the rest of the data LUNs.
14. Repeat step 3 through step 9 to create a total of four 450GB data LUNs (DATALUN5 through DATALUN8) for the volume group A1VG2.
15. Repeat this process to provision the rest of the EF560 storage arrays as part of the solution.

Note: For this setup, each data volume group was created with four 450GB LUNs, which makes eight LUNs per EF560 system. A total of 32 LUNs were created for all four EF560 systems together. See Table 24 in the appendix for details.

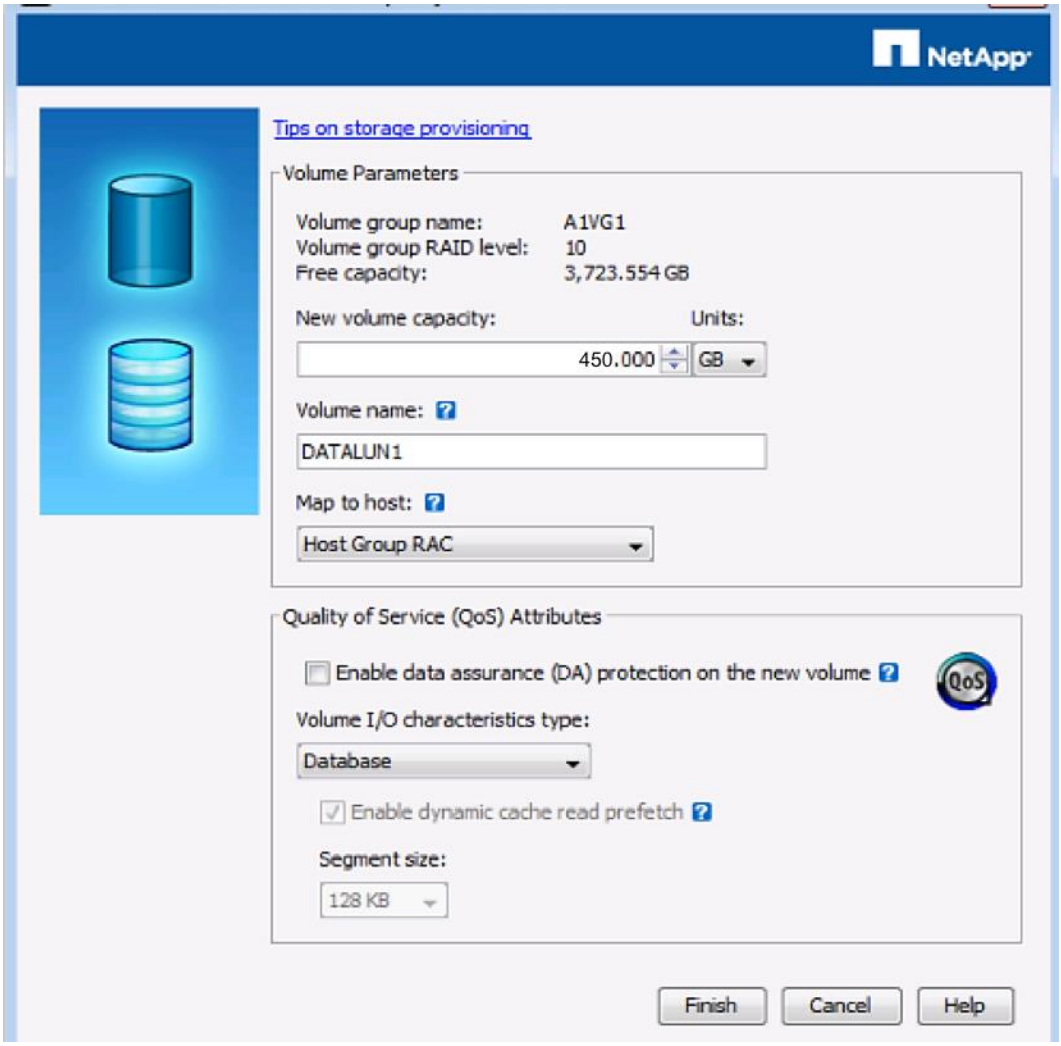

4.4 Cisco Nexus 5548UP Network Configuration

This section describes the installation and configuration of the Cisco Nexus 5548UP switches used in this solution. Section 4.2, “Cisco UCS Configuration,” must be completed before the procedure in this section can be performed.

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment with Cisco Nexus NX-OS 7.0(0)N1(1) or later.

Initial Cisco Nexus 5548UP Switch A Configuration

To set up the initial configuration for the Cisco Nexus switch A, complete the following steps:

1. Configure the switch.

   Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter power-on autoprovisioning.

   Note: Review the configuration summary before enabling the configuration.

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



```
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin":
Confirm the password for "admin":
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name:
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address:
Mgmt0 IPv4 netmask:
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway:
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa): rsa
Number of rsa key bits <1024-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address:
Enter basic FC configurations (yes/no) [n]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

Initial Cisco Nexus 5548UP Switch B Configuration

To set up the initial configuration for the Cisco Nexus switch B, complete the following steps:

1. Configure the switch.

   Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter power-on autoprovisioning.

   Note: Review the configuration summary before enabling the configuration.

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin":
Confirm the password for "admin":
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name:
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address:
Mgmt0 IPv4 netmask:
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway:
Enable the telnet service? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa): rsa
Number of rsa key bits <1024-2048> : 1024
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address:
Enter basic FC configurations (yes/no) [n]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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


FlexPod Cisco Nexus Fibre Channel (FC) Storage

This section describes the steps required to configure the FCoE and FC capabilities required for the Cisco Nexus 5548UP switches.

Enable Licenses

To license the Cisco Nexus switches, complete the following steps:

1. For the Cisco Nexus 5548UP switch A and Cisco Nexus 5548UP switch B, log in as `admin`.
2. Run the following commands:

```
config t
feature fcoe
feature npiv
feature lacp
feature vpc
feature lldp
```

Configure FC Ports

To configure the FC ports on the Cisco Nexus switches, complete the following step:

1. On the Cisco Nexus 5548UP switch A and Cisco Nexus 5548UP switch B, from the global configuration mode, run the following commands:

```
slot 1
port 17-32 type fc
exit
copy run start
reload
```

Set Global Configurations

To enable global configurations on the Cisco Nexus switches, complete the following step:

1. On the Cisco Nexus 5548UP switch A and Cisco Nexus 5548UP switch B, run the following commands to set global configurations and jumbo frames in QoS:

```
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
```

Create VLANs

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

1. On the Cisco Nexus 5548UP switch A, from the global configuration mode, run the following commands. VLAN IDs for the Native-VLAN, the Fabric A FCoE VLAN, the in-band management VLAN, and the RAC interconnect VLAN are entered:

```
vlan _____
```



```
name Native-VLAN
exit
vlan _____
name FCoE_Fabric_A
exit
vlan _____
name IB-MGMT
exit
vlan _____
name RAC-Interconnnect-1 VLAN
exit
```

2. On Cisco Nexus 5548UP switch B, from the global configuration mode, run the following commands. VLAN IDs for the Native-VLAN, the Fabric B FCoE VLAN, the in-band management VLAN, and the RAC interconnect VLAN are entered:

```
vlan _____
name Native-VLAN
exit
vlan _____
name FCoE_Fabric_B
exit
vlan _____
name IB-MGMT
exit
vlan _____
name RAC-Interconnnect-1 VLAN
exit
```

Add Individual Port Descriptions for Troubleshooting Switch A and Switch B

To add individual port descriptions for troubleshooting activity and verification for Cisco Nexus 5548UP switch A, complete the following step:

1. From the global configuration mode, run the following commands (which include example component names):

```
interface fc1/17
switchport description EF560 controller 01a-1
exit
interface fc1/18
switchport description EF560 controller 01a-2
exit
interface fc1/19
switchport description EF560 controller 01b-1
exit
interface fc1/20
switchport description EF560 controller 01b-2
exit
interface fc1/21
switchport description EF560 controller 02a-1
exit
interface fc1/22
switchport description EF560 controller 02a-2
exit
interface fc1/23
switchport description EF560 controller 02b-1
exit
interface fc1/24
switchport description EF560 controller 02b-2
exit
interface fc1/25
switchport description EF560 controller 03a-1
exit
interface fc1/26
switchport description EF560 controller 03a-2
exit
interface fc1/27
switchport description EF560 controller 03b-1
```



```
exit
interface fc1/28
switchport description EF560 controller 03b-2
exit
interface fc1/29
switchport description EF560 controller 04a-1
exit
interface fc1/30
switchport description EF560 controller 04a-2
exit
interface fc1/31
switchport description EF560 controller 04b-1
exit
interface fc1/32
switchport description EF560 controller 04b-2
exit
interface Eth1/1
description UCS_FlexPod-A:1/17
exit
interface Eth1/2
description UCS_FlexPod-A:1/18
exit
interface Eth1/3
description UCS_FlexPod-A:1/19
exit
interface Eth1/4
description UCS_FlexPod-A:1/20
exit
interface Eth1/5
description UCS_FlexPod-A:1/21
exit
interface Eth1/6
description UCS_FlexPod-A:1/22
exit
interface Eth1/7
description UCS_FlexPod-A:1/23
exit
interface Eth1/8
description UCS_FlexPod-A:1/24
exit
interface Eth1/9
description UCS_FlexPod-A:1/25
exit
interface Eth1/10
description UCS_FlexPod-A:1/26
exit
interface Eth1/11
description UCS_FlexPod-A:1/27
exit
interface Eth1/12
description UCS_FlexPod-A:1/28
exit
interface Eth1/13
description UCS_FlexPod-A:1/29
exit
interface Eth1/14
description UCS_FlexPod-A:1/30
exit
interface Eth2/3
description nexus_B:2/3
exit
interface Eth2/4
description nexus_B:2/4
exit
interface eth2/1
description UCS_FlexPod-A:1/31
exit
interface eth2/2
description UCS_FlexPod-B:1/31
exit
```


To add individual port descriptions for troubleshooting activity and verification for Cisco Nexus 5548UP switch B, complete the following step:

1. From the global configuration mode, run the following commands (which include example component names):

```
interface fc1/17
switchport description EF560 controller 01a-3
exit
interface fc1/18
switchport description EF560 controller 01a-4
exit
interface fc1/19
switchport description EF560 controller 01b-3
exit
interface fc1/20
switchport description EF560 controller 01b-4
exit
interface fc1/21
switchport description EF560 controller 02a-3
exit
interface fc1/22
switchport description EF560 controller 02a-4
exit
interface fc1/23
switchport description EF560 controller 02b-3
exit
interface fc1/24
switchport description EF560 controller 02b-4
exit
interface fc1/25
switchport description EF560 controller 03a-3
exit
interface fc1/26
switchport description EF560 controller 03a-4
exit
interface fc1/27
switchport description EF560 controller 03b-3
exit
interface fc1/28
switchport description EF560 controller 03b-4
exit
interface fc1/29
switchport description EF560 controller 04a-3
exit
interface fc1/30
switchport description EF560 controller 04a-4
exit
interface fc1/31
switchport description EF560 controller 04b-3
exit
interface fc1/32
switchport description EF560 controller 04b-4
exit
interface Eth1/1
description UCS_FlexPod-B:1/17
exit
interface Eth1/2
description UCS_FlexPod-B:1/18
exit
interface Eth1/3
description UCS_FlexPod-B:1/19
exit
interface Eth1/4
description UCS_FlexPod-B:1/20
exit
interface Eth1/5
description UCS_FlexPod-B:1/21
exit
interface Eth1/6
```



```
description UCS_FlexPod-B:1/22
exit
interface Eth1/7
description UCS_FlexPod-B:1/23
exit
interface Eth1/8
description UCS_FlexPod-B:1/24
exit
interface Eth1/9
description UCS_FlexPod-B:1/25
exit
interface Eth1/10
description UCS_FlexPod-B:1/26
exit
interface Eth1/11
description UCS_FlexPod-B:1/27
exit
interface Eth1/12
description UCS_FlexPod-B:1/28
exit
interface Eth1/13
description UCS_FlexPod-B:1/29
exit
interface Eth1/14
description UCS_FlexPod-B:1/30
exit
interface Eth2/3
description nexus_A:2/3
exit
interface Eth2/4
description nexus_A:2/4
exit
interface eth2/1
description UCS_FlexPod-A:1/32
exit
interface eth2/2
description UCS_FlexPod-B:1/32
exit
```

Create Port Profiles for Switch A and Switch B

Port profiles are used to simplify ongoing network administration and configuration. Ports with similar configurations can be grouped within port profiles. Configuration changes can then be made to the port profile and are applied to all members of the port profile. FlexPod recommends port profiles for the following port types:

- Cisco UCS Ethernet ports
- Cisco UCS FCoE ports
- Cisco Nexus VPC ports

To create the Ethernet traffic port profiles, complete the following step on both the switches:

1. On Cisco Nexus 5548UP switch A, from the global configuration mode, run the following commands. Text is entered instead of the actual VLAN IDs in the following example:

```
port-profile default max-ports 512
port-profile type port-channel UCS-Ethernet
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan in band management vlan id, RAC interconnect 1 vlan id
spanning-tree port type edge trunk
state enabled
port-profile type port-channel vPC-Peer-Link
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan in band management vlan id, RAC interconnect 1 vlan id
spanning-tree port type network
```



```
state enabled
port-profile type port-channel UCS-FCOE-FABRIC-A
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan FCoE fabric a vlan id
spanning-tree port type edge trunk
state enabled
exit
```

2. On the Cisco Nexus 5548UP switch B, from the global configuration mode, run the following commands:

```
port-profile default max-ports 512
port-profile type port-channel UCS-Ethernet
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan in band management vlan id, RAC interconnect 1 vlan id
spanning-tree port type edge trunk
state enabled
port-profile type port-channel vPC-Peer-Link
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan in band management vlan id, RAC interconnect 1 vlan id
spanning-tree port type network
state enabled
port-profile type port-channel UCS-FCOE-FABRIC-B
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan FCoE fabric b vlan id
spanning-tree port type edge trunk
state enabled
exit
```

Create Port Channels for Switch A and Switch B

To create the necessary port channels between devices, complete the following step on both the switches:

1. On the Cisco Nexus 5548UP switch A, from the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
exit
interface Eth2/3-4
channel-group 10 mode active
no shutdown
exit
interface Po13
description UCS_FlexPod-A
exit
interface Eth2/1
channel-group 13 mode active
no shutdown
exit
interface Po14
description UCS_FlexPod-B
exit
interface Eth2/2
channel-group 14 mode active
no shutdown
exit
interface Po15
description UCS_FlexPod-fabric-A-FCOE
exit
interface Eth1/1-14
channel-group 15 mode active
exit
copy run start
```


2. On the Cisco Nexus 5548UP switch B, from the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
exit
interface Eth2/3-4
channel-group 10 mode active
no shutdown
exit
interface Po13
description UCS_FlexPod-A
exit
interface Eth2/1
channel-group 13 mode active
no shutdown
exit
interface Po14
description UCS_FlexPod-B
exit
interface Eth2/2
channel-group 14 mode active
no shutdown
exit
interface Po15
description UCS_FlexPod-fabric-B-FCOE
exit
interface Eth1/1-14
channel-group 15 mode active
exit
copy run start
```

Add Port Profiles to Port Channels for Switch A and Switch B

Port channels and their member ports inherit their configuration from the previously configured port profiles.

To assign port profiles to the appropriate port channels, complete the following step on both the switches:

1. On the Cisco Nexus 5548UP switch A, from the global configuration mode, run the following commands:

```
interface Po10
inherit port-profile vPC-Peer-Link
exit
interface Po13
inherit port-profile UCS-Ethernet
exit
interface Po14
inherit port-profile UCS-Ethernet
exit
interface Po15
inherit port-profile UCS-FCOE-FABRIC-A
exit
copy run start
```

2. On the Cisco Nexus 5548UP switch B, from the global configuration mode, run the following commands:

```
interface Po10
inherit port-profile vPC-Peer-Link
exit
interface Po13
inherit port-profile UCS-Ethernet
exit
interface Po14
inherit port-profile UCS-Ethernet
exit
```



```
interface Po15
inherit port-profile UCS-FCOE-FABRIC-B
exit
copy run start
```

Configure Virtual Port Channels for Switch A and Switch B

To configure virtual port channels (vPCs), complete the following step on both switches:

1. On the Cisco Nexus 5548UP switch A, from the global configuration mode, run the following commands. Text is entered instead of the actual VPC domain ID and source and destination IP addresses of the Cisco Nexus switches in the following example:

```
vpc domain Nexus VPC Domain ID
role priority 10
peer-keepalive destination nexus B mgmt0 ip source nexus A mgmt0 ip
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

2. On the Cisco Nexus 5548UP switch B, from the global configuration mode, run the following commands. Text is entered instead of the actual VPC domain ID and source and destination IP addresses of the Cisco Nexus switches in the following example:

```
vpc domain Nexus VPC Domain ID
role priority 20
peer-keepalive destination nexus A mgmt0 ip source nexus B mgmt0 ip
auto-recovery
exit
interface Po10
vpc peer-link
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 5548UP switches included in the FlexPod environment into the infrastructure. The procedures described in section “FlexPod Cisco Nexus Base” can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after the configuration is complete.

Create VSANs, Assign and Enable Virtual Fibre Channel Ports

This procedure configures the Fibre Channel over Ethernet (FCoE) connections between the Cisco Nexus 5548UP switches, the Cisco UCS Fabric Interconnects, and the NetApp storage systems.

Cisco Nexus 5548UP Switch A

To configure virtual storage area networks (VSANs), create and update relevant port profiles, assign virtual Fibre Channel (vFC) ports, and enable vFC ports on switch A, complete the following step. You create a VLAN for FCoE traffic and a corresponding VSAN for fabric A:

1. From the global configuration mode, run the following commands:

```
vlan ____
name FCoE_Fabric_A
fcoe vsan ____
exit
interface vfc15
switchport description UCS_FlexPod-A:FCoE
bind interface po15
switchport trunk allowed vsan 101
no shutdown
vsan database
vsan 101name Fabric_A
vsan 101 interface vfc15
vsan 101 interface fc1/17
vsan 101 interface fc1/18
vsan 101 interface fc1/19
vsan 101 interface fc1/20
vsan 101 interface fc1/21
vsan 101 interface fc1/22
vsan 101 interface fc1/23
vsan 101 interface fc1/24
vsan 101 interface fc1/25
vsan 101 interface fc1/26
vsan 101 interface fc1/27
vsan 101 interface fc1/28
vsan 101 interface fc1/29
vsan 101 interface fc1/30
vsan 101 interface fc1/31
vsan 101 interface fc1/32
exit
copy run start
```

Cisco Nexus 5548UP B

To configure VSANs, create and update relevant port profiles, assign vFC ports, and enable vFC ports on switch B, complete the following step. You create a VLAN for FCoE traffic and a corresponding VSAN for fabric A. Examples are written below:

1. From the global configuration mode, run the following commands:

```
vlan ____
name FCoE_Fabric_B
fcoe vsan 102
exit
interface vfc15
switchport description UCS_FlexPod-B:FCoE
bind interface po15
switchport trunk allowed vsan 102
no shutdown
vsan database
vsan 102 name Fabric_B
vsan 102 interface vfc15
vsan 102 interface fc1/17
vsan 102 interface fc1/18
vsan 102 interface fc1/19
vsan 102 interface fc1/20
vsan 102 interface fc1/21
vsan 102 interface fc1/22
vsan 102 interface fc1/23
vsan 102 interface fc1/24
vsan 102 interface fc1/25
vsan 102 interface fc1/26
vsan 102 interface fc1/27
vsan 102 interface fc1/28
vsan 102 interface fc1/29
```



```
vsan 102 interface fc1/30
vsan 102 interface fc1/31
vsan 102 interface fc1/32
exit
copy run start
```

Create Device Aliases

To configure device aliases on both switches, complete the following steps.

Cisco Nexus 5548UP Switch A

1. From the global configuration mode, run the following commands. Enter the WWPN for the preceding device in each line of command. Examples are written below:

```
device-alias database
device-alias name RAC-SERVER-01-A pwwn _________________
device-alias name RAC-SERVER-01-A-2 pwwn _________________
device-alias name RAC-SERVER-02-A pwwn _________________
device-alias name RAC-SERVER-02-A-2 pwwn _________________
device-alias name RAC-SERVER-03-A pwwn _________________
device-alias name RAC-SERVER-03-A-2 pwwn _________________
device-alias name RAC-SERVER-04-A pwwn _________________
device-alias name RAC-SERVER-04-A-2 pwwn _________________
device-alias name RAC-SERVER-05-A pwwn _________________
device-alias name RAC-SERVER-05-A-2 pwwn _________________
device-alias name RAC-SERVER-06-A pwwn _________________
device-alias name RAC-SERVER-06-A-2 pwwn _________________
device-alias name RAC-SERVER-07-A pwwn _________________
device-alias name RAC-SERVER-07-A-2 pwwn _________________
device-alias name RAC-SERVER-08-A pwwn _________________
device-alias name RAC-SERVER-08-A-2 pwwn _________________
device-alias name EF560 controller 01a-1 pwwn _________________
device-alias name EF560 controller 01a-2 pwwn _________________
device-alias name EF560 controller 01b-1 pwwn _________________
device-alias name EF560 controller 01b-2 pwwn _________________
device-alias name EF560 controller 02a-1 pwwn _________________
device-alias name EF560 controller 02a-2 pwwn _________________
device-alias name EF560 controller 02b-1 pwwn _________________
device-alias name EF560 controller 02b-2 pwwn _________________
device-alias name EF560 controller 03a-1 pwwn _________________
device-alias name EF560 controller 03a-2 pwwn _________________
device-alias name EF560 controller 03b-1 pwwn _________________
device-alias name EF560 controller 03b-2 pwwn _________________
device-alias name EF560 controller 04a-1 pwwn _________________
device-alias name EF560 controller 04a-2 pwwn _________________
device-alias name EF560 controller 04b-1 pwwn _________________
device-alias name EF560 controller 04b-2 pwwn _________________
exit
device-alias commit
```

Cisco Nexus 5548UP Switch B

1. From the global configuration mode, run the following commands. Enter the WWPN for the preceding device in each line of command. Examples are written:

```
device-alias database
device-alias name RAC-SERVER-01-B pwwn _________________
device-alias name RAC-SERVER-01-B-2 pwwn _________________
device-alias name RAC-SERVER-02-B pwwn _________________
device-alias name RAC-SERVER-02-B-2 pwwn _________________
device-alias name RAC-SERVER-03-B pwwn _________________
device-alias name RAC-SERVER-03-B-2 pwwn _________________
device-alias name RAC-SERVER-04-B pwwn _________________
device-alias name RAC-SERVER-04-B-2 pwwn _________________
device-alias name RAC-SERVER-05-B pwwn _________________
device-alias name RAC-SERVER-05-B-2 pwwn _________________
device-alias name RAC-SERVER-06-B pwwn _________________
device-alias name RAC-SERVER-06-B-2 pwwn _________________
```



```
device-alias name RAC-SERVER-07-B pwwn ________________
device-alias name RAC-SERVER-07-B-2 pwwn ________________
device-alias name RAC-SERVER-08-B pwwn ________________
device-alias name RAC-SERVER-08-B-2 pwwn ________________
device-alias name EF560 controller 01a-3 pwwn ________________
device-alias name EF560 controller 01a-4 pwwn ________________
device-alias name EF560 controller 01b-3 pwwn ________________
device-alias name EF560 controller 01b-4 pwwn ________________
device-alias name EF560 controller 02a-3 pwwn ________________
device-alias name EF560 controller 02a-4 pwwn ________________
device-alias name EF560 controller 02b-3 pwwn ________________
device-alias name EF560 controller 02b-4 pwwn ________________
device-alias name EF560 controller 03a-3 pwwn ________________
device-alias name EF560 controller 03a-4 pwwn ________________
device-alias name EF560 controller 03b-3 pwwn ________________
device-alias name EF560 controller 03b-4 pwwn ________________
device-alias name EF560 controller 04a-3 pwwn ________________
device-alias name EF560 controller 04a-4 pwwn ________________
device-alias name EF560 controller 04b-3 pwwn ________________
device-alias name EF560 controller 04b-4 pwwn ________________
exit
device-alias commit
```

Create Zones

Cisco Nexus 5548UP Switch A

To create zones for the service profiles on switch A, complete the following steps:

1. Create a zone for each service profile.

   Note: A single host is demonstrated in the following example. Perform this procedure for all Oracle RAC hosts. Each host HBA WWPN is zoned with all EF560 targets on the fabric.

```
zone name RAC-SERVER-01-A vsan 101
member device-alias RAC-SERVER-01-A
member device-alias RAC-SERVER-01-A-2
member device-alias EF560-1a-1
member device-alias EF560-1a-2
member device-alias EF560-1b-1
member device-alias EF560-1b-2
member device-alias EF560-2a-1
member device-alias EF560-2a-2
member device-alias EF560-2b-1
member device-alias EF560-2b-2
member device-alias EF560-3a-1
member device-alias EF560-3a-2
member device-alias EF560-3b-1
member device-alias EF560-3b-2
member device-alias EF560-4a-1
member device-alias EF560-4a-2
member device-alias EF560-4b-1
member device-alias EF560-4b-2
exit
```

2. After the zones for the Cisco UCS service profiles have been created, create the zone set and add the necessary members.

```
zoneset name FlexPod vsan 101
member RAC-SERVER-01-A
member RAC-SERVER-02-A
member RAC-SERVER-03-A
member RAC-SERVER-04-A
member RAC-SERVER-05-A
member RAC-SERVER-06-A
member RAC-SERVER-07-A
member RAC-SERVER-08-A
exit
```


3. Activate the zone set.

```
zoneset activate name FlexPod vsan 101
exit
copy run start
```

Cisco Nexus 5548UP Switch B

To create zones for the service profiles on switch B, complete the following steps:

1. Create a zone for each service profile.

   Note: A single host is demonstrated in the following example. Replicate this procedure for all Oracle RAC hosts. Each host HBA WWPN is zoned with all EF560 targets on the fabric.

```
zone name RAC-SERVER-01-B vsan 102
member device-alias RAC-SERVER-01-B
member device-alias RAC-SERVER-01-B-2
member device-alias EF560-1a-3
member device-alias EF560-1a-4
member device-alias EF560-1b-3
member device-alias EF560-1b-4
member device-alias EF560-2a-3
member device-alias EF560-2a-4
member device-alias EF560-2b-3
member device-alias EF560-2b-4
member device-alias EF560-3a-3
member device-alias EF560-3a-4
member device-alias EF560-3b-3
member device-alias EF560-3b-4
member device-alias EF560-4a-3
member device-alias EF560-4a-4
member device-alias EF560-4b-3
member device-alias EF560-4b-4
exit
```

2. After the zones for the Cisco UCS service profiles have been created, create the zone set and add the necessary members.

```
zoneset name FlexPod vsan 102
member RAC-SERVER-01-B
member RAC-SERVER-02-B
member RAC-SERVER-03-B
member RAC-SERVER-04-B
member RAC-SERVER-05-B
member RAC-SERVER-06-B
member RAC-SERVER-07-B
member RAC-SERVER-08-B
exit
```

3. Activate the zone set.

```
zoneset activate name FlexPod vsan 102
exit
copy run start
```

4.5 Oracle Real Application Cluster Configuration

The following section describes the installation and configuration of Oracle RAC nodes into an eight-node cluster.

Note: The storage configuration of the boot LUNs must be complete before you start the Oracle Linux installation on the Oracle RAC nodes.

Oracle Linux Installation on RAC Nodes

This section provides details around how to install and configure Oracle Linux 6.6 for use in this solution. From the Cisco UCS director, launch the KVM console for the first RAC node.

1. After the console is up and running, activate the virtual device to map the virtual CD-ROM.

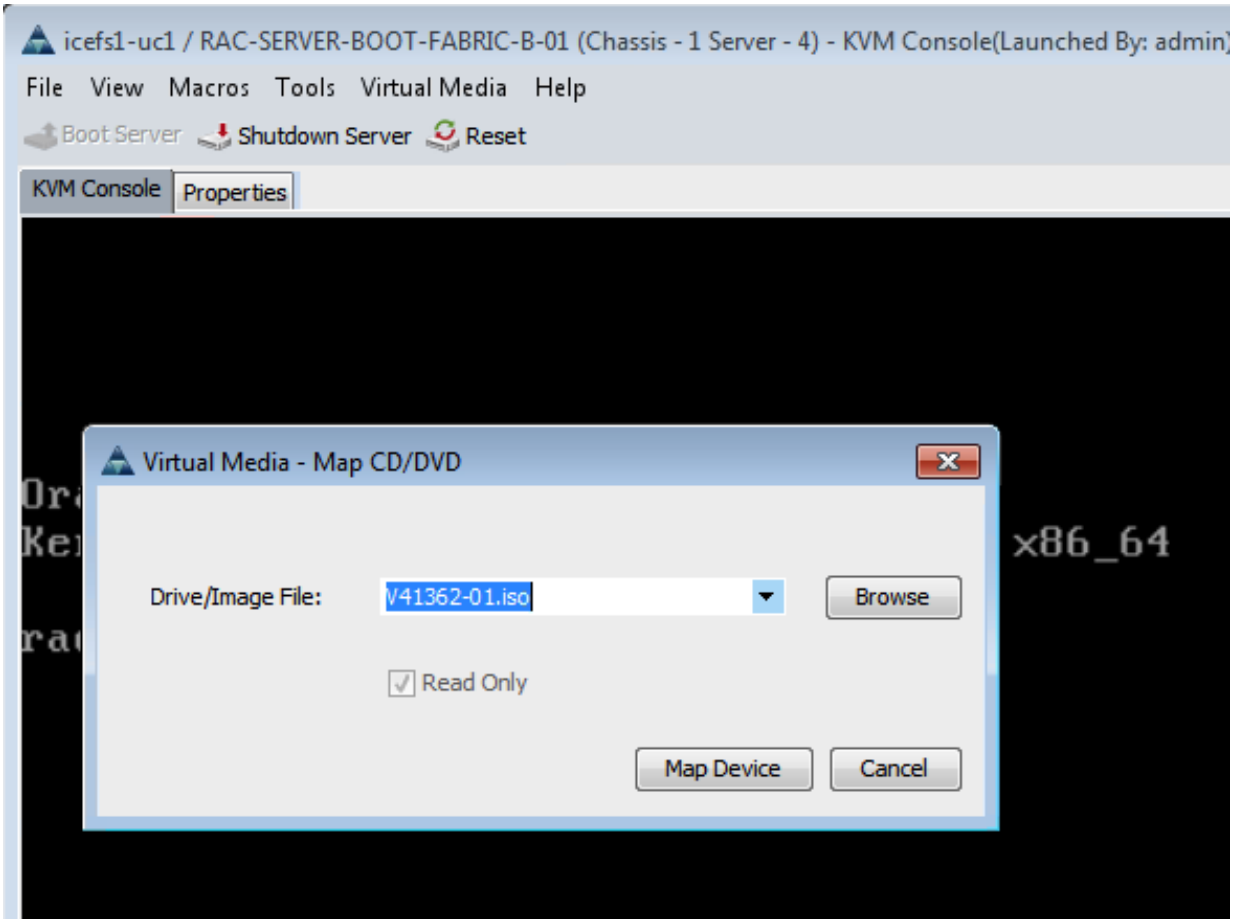

2. Map the virtual CD-ROM on the KVM console to pick the Oracle Linux Server 6.6 installation media. The install media should be attached to the Cisco UCS system manager.

Reboot the RAC node through the KVM console so that the installation media is recognized.

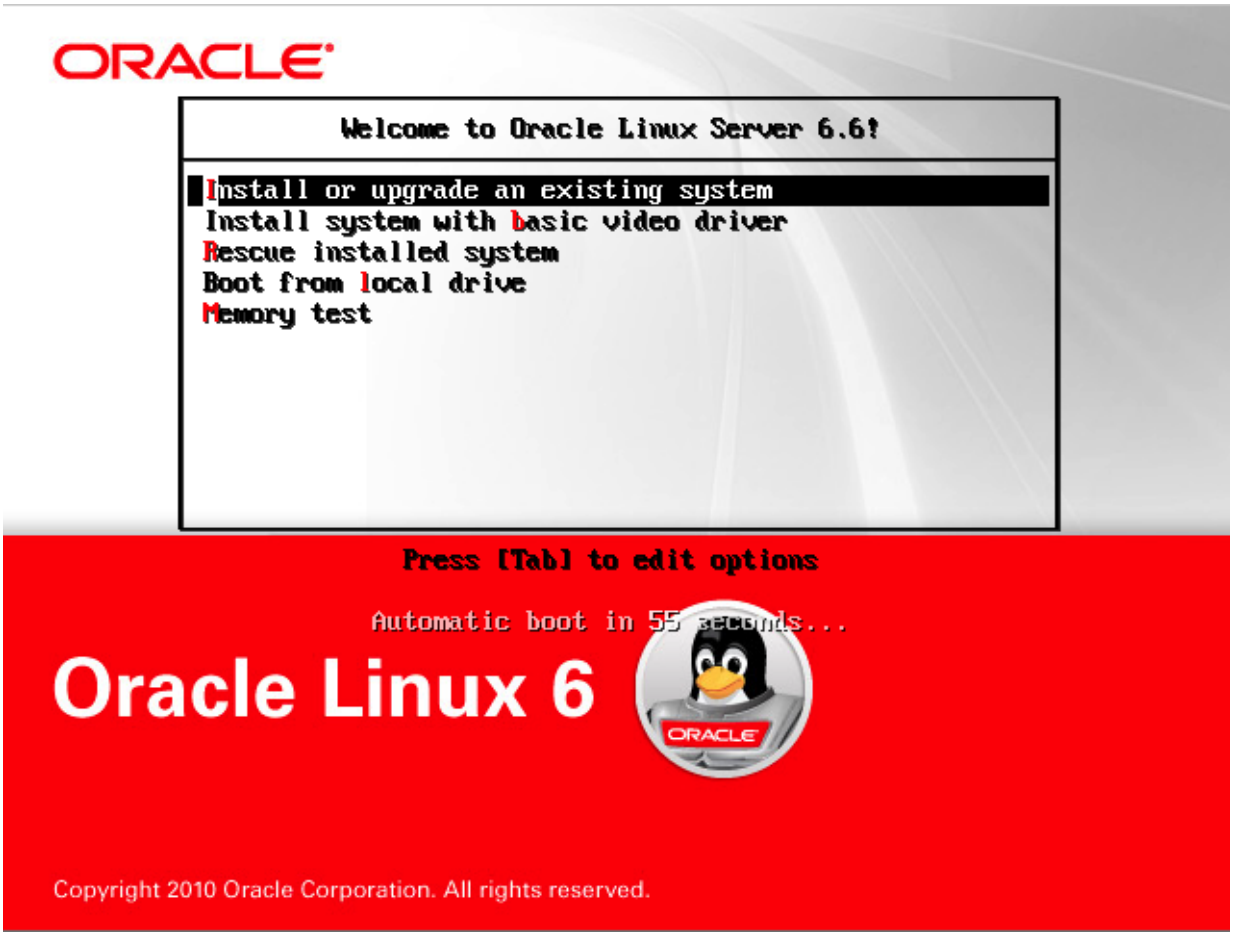

3. On the first splash screen, select Skip.

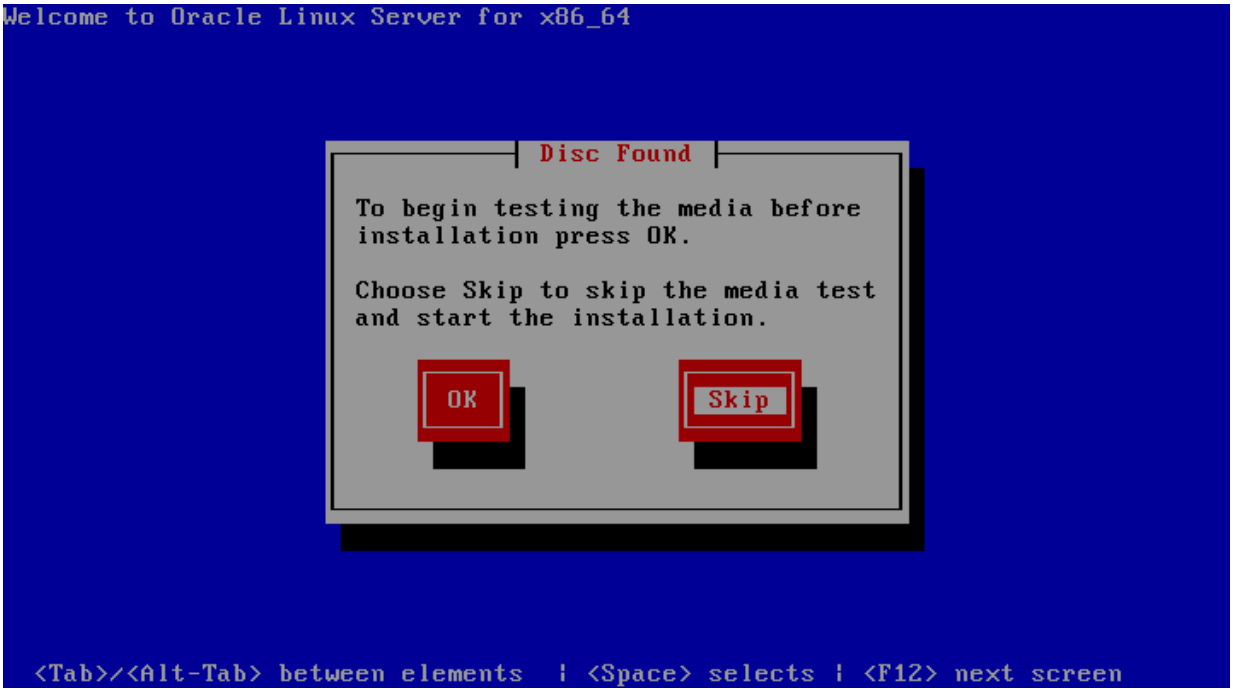

4. Click Next.

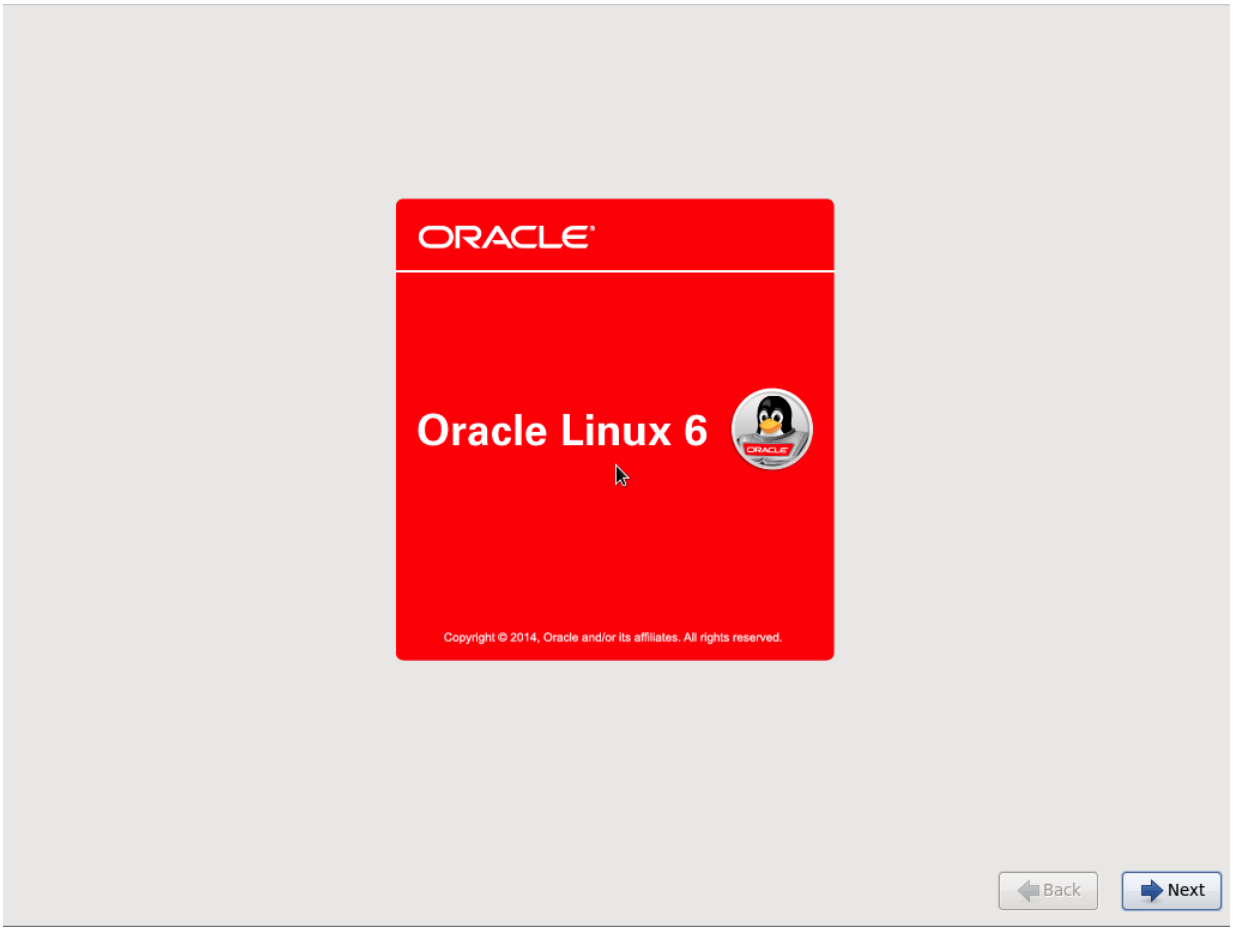

5. Click Next to select English as the language.

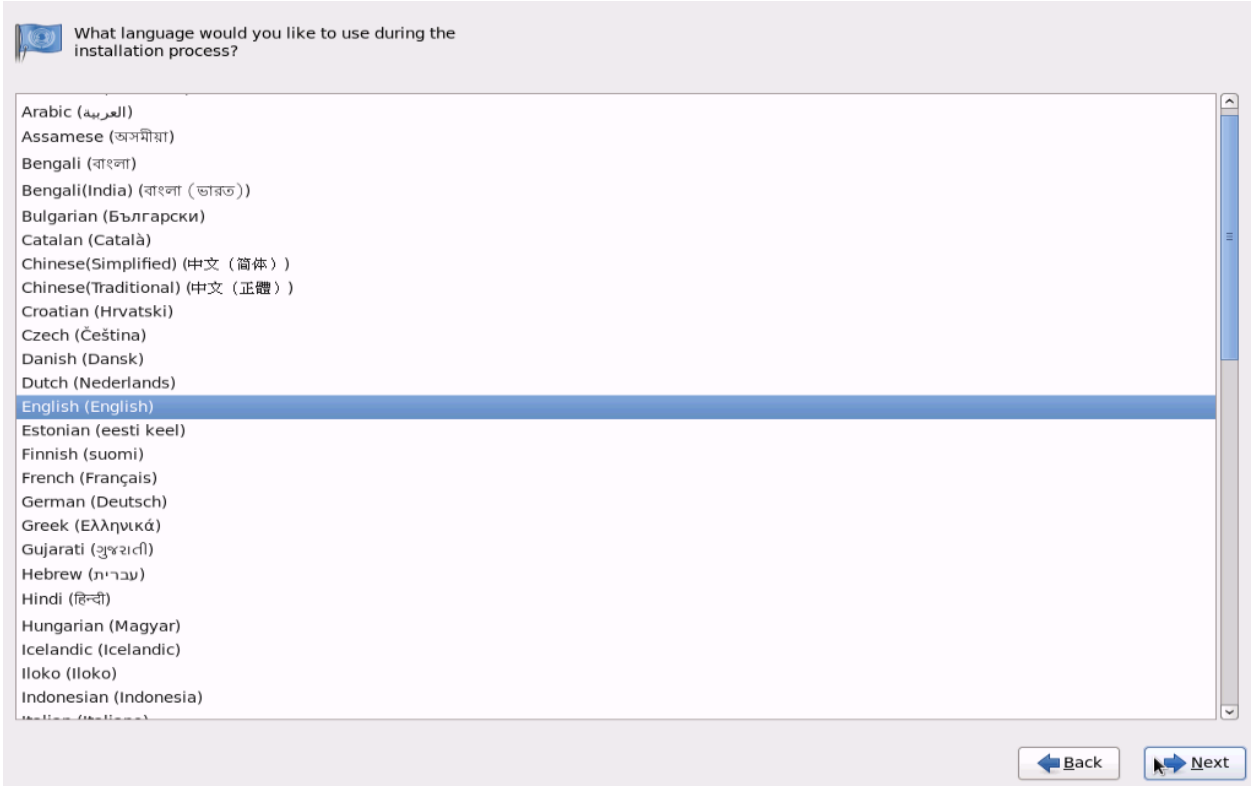

6. Click Next to use U.S. English for the keyboard layout.

7. Select Specialized Storage Devices so that the boot from the NetApp storage can be initialized.
8. Click Next.

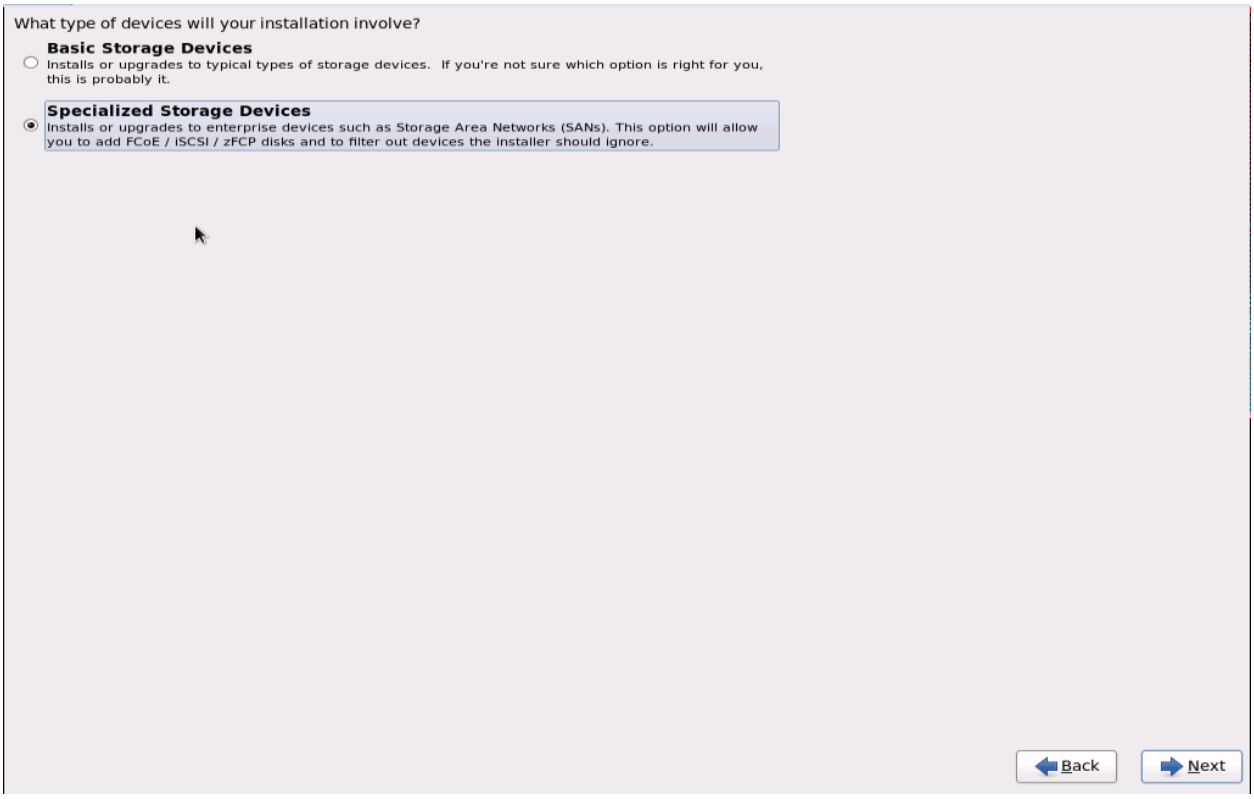

9. Select the listed NetApp disk. This example assumes that only the storage associated with the RAC node boot process has been provisioned at this time. As a result, there are no other NetApp storage devices available.
10. Click Next.

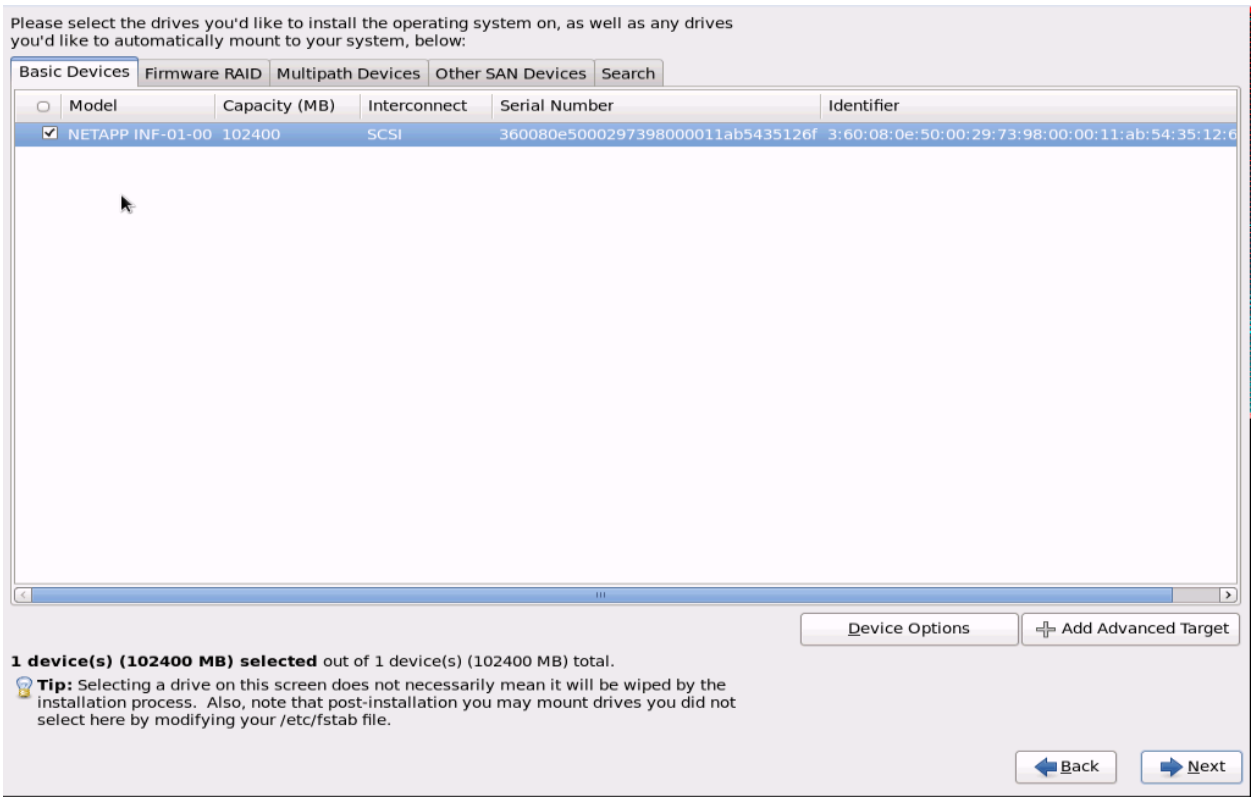

11. In the splash window select Yes, Discard Any Data.
12. Click Next.

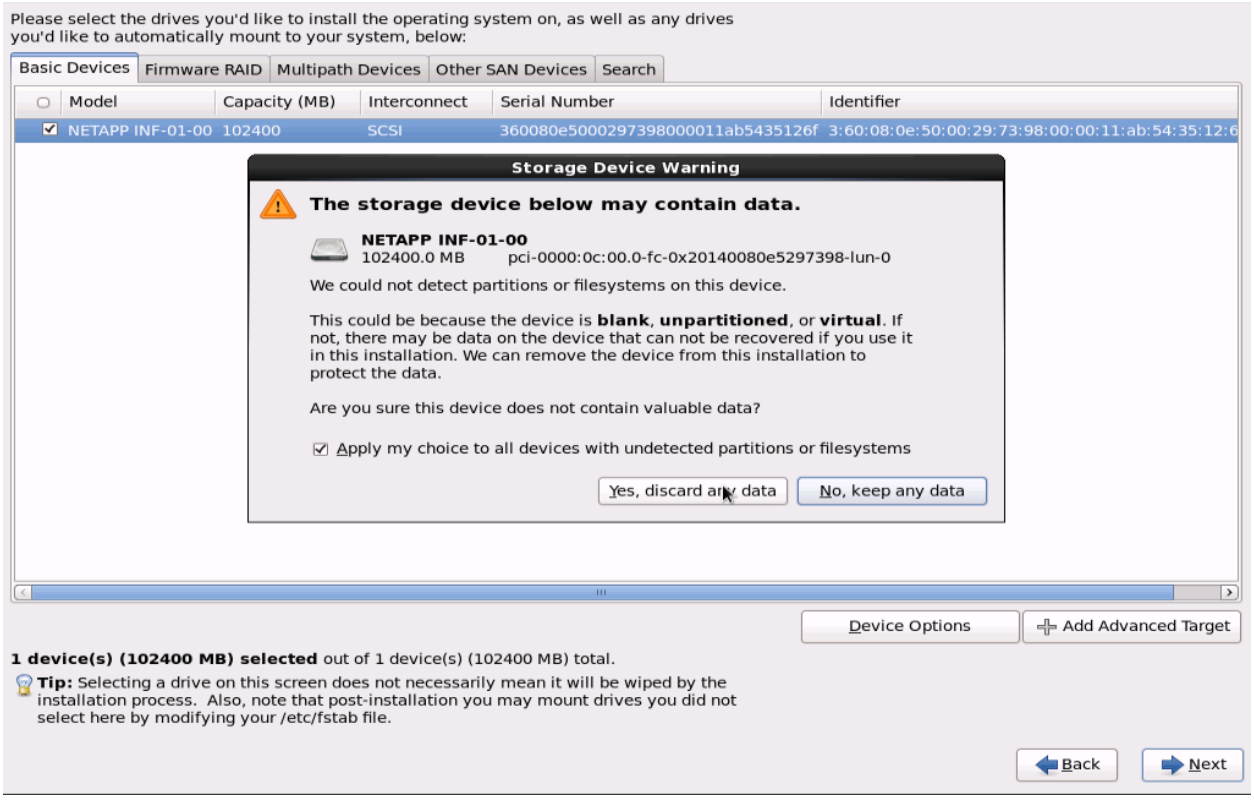

13. Enter the public host name; in this case, we used `rac-server-01`.
14. Click Next.

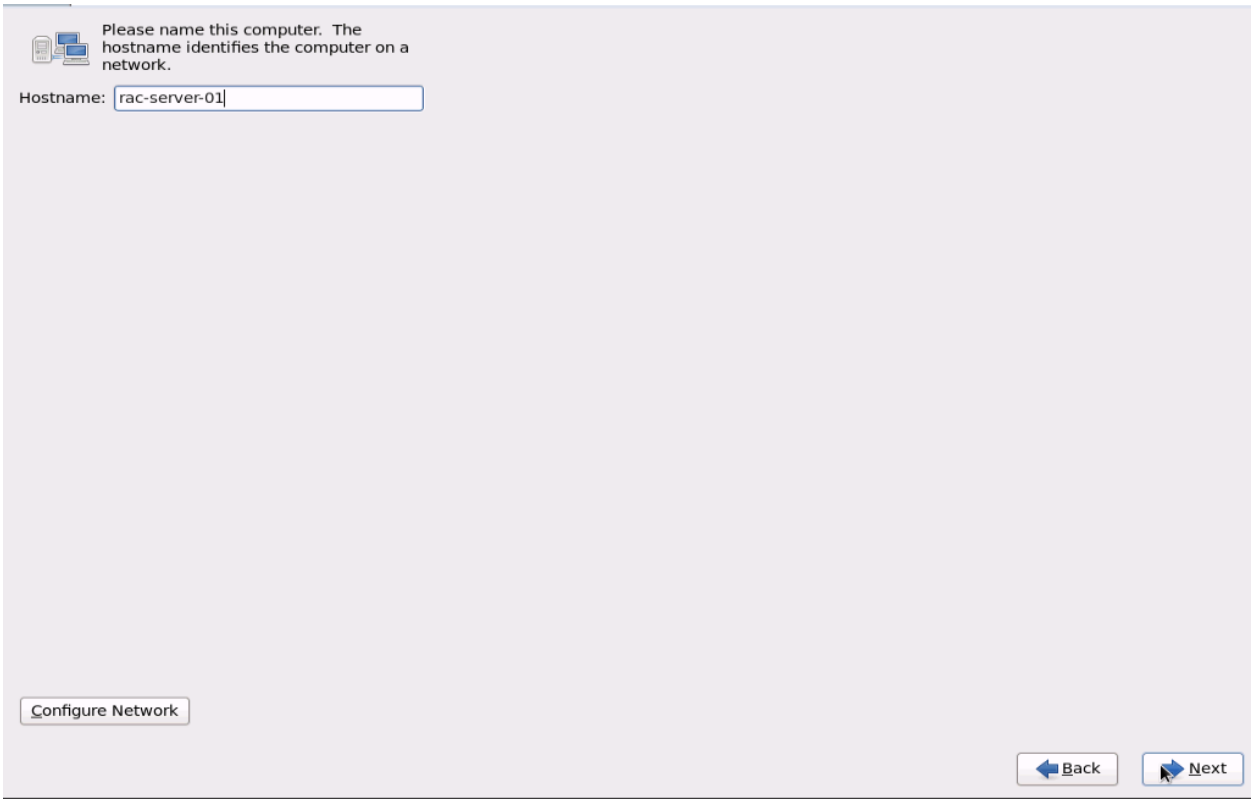

15. Select the time zone and click Next.

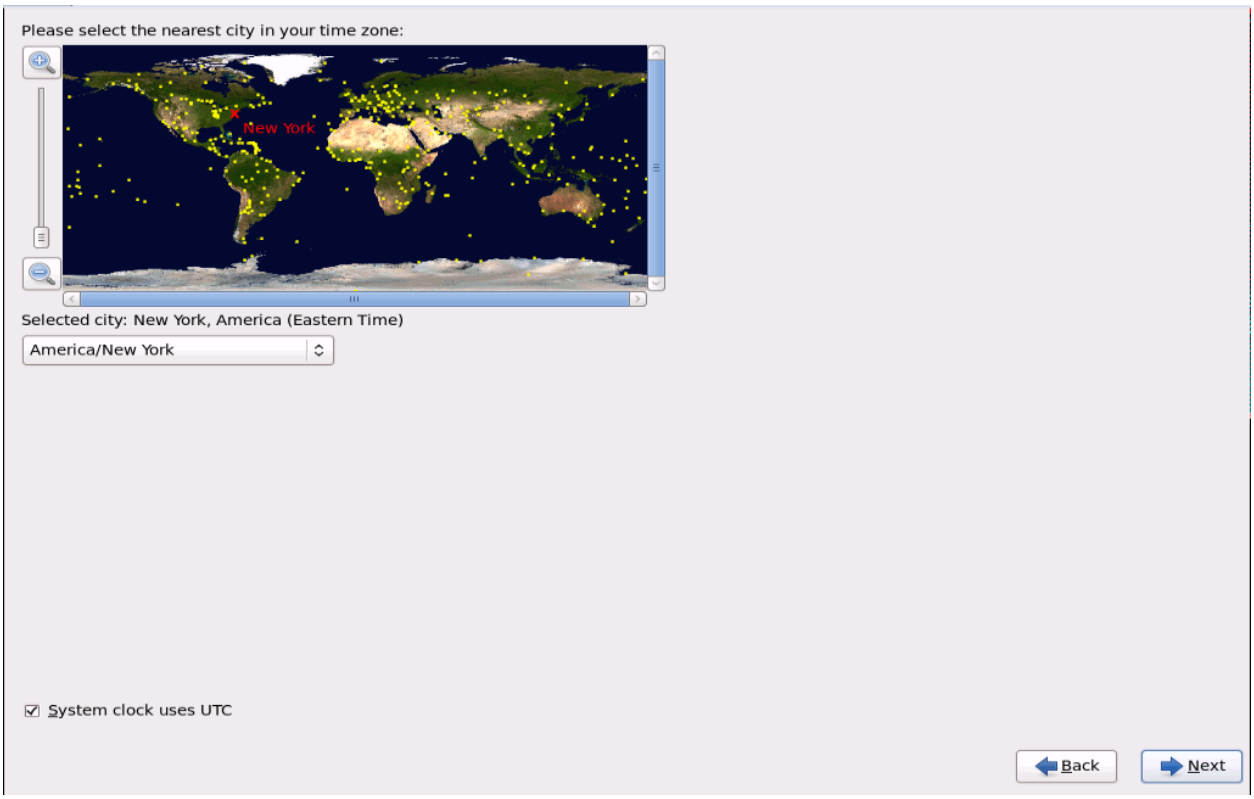

16. Because this is a completely new installation, select Use All Space to wipe out all the data.
17. Click Next.

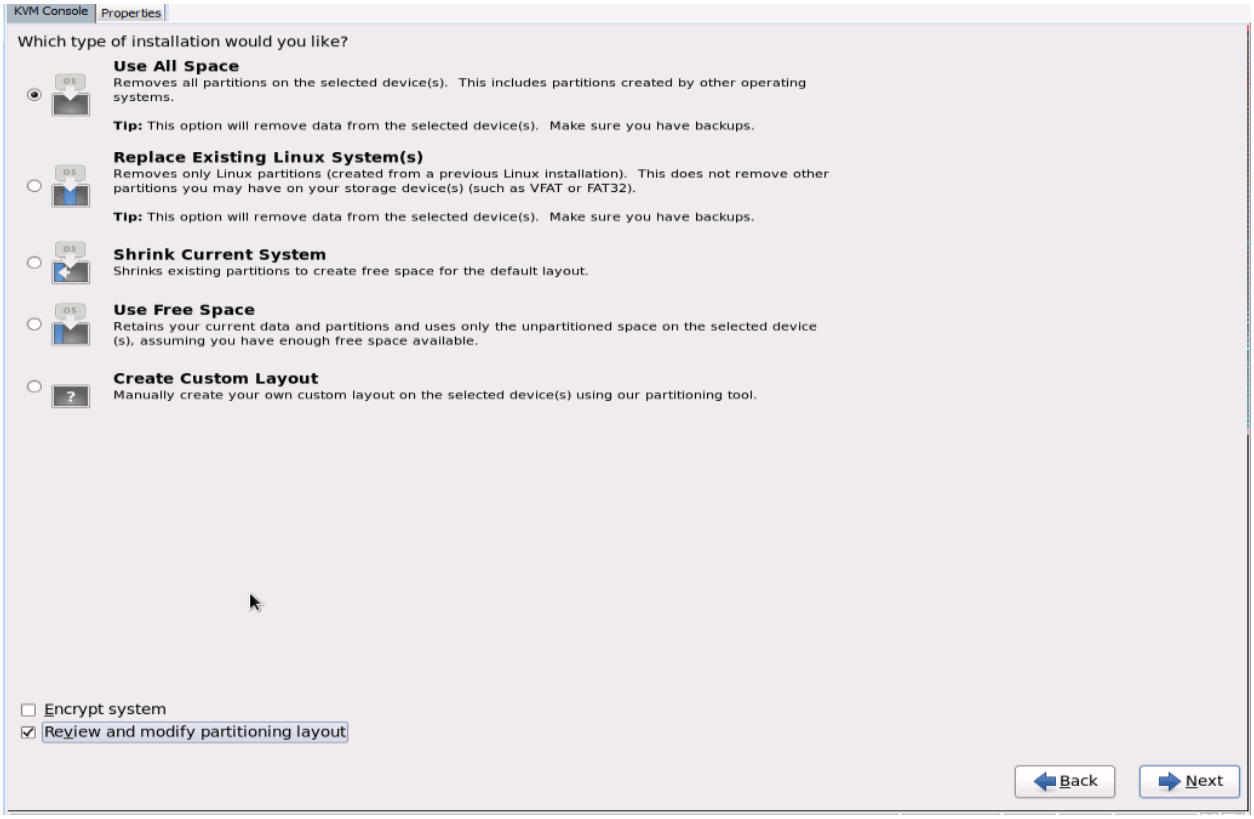

18. Select the physical volume and click Create.

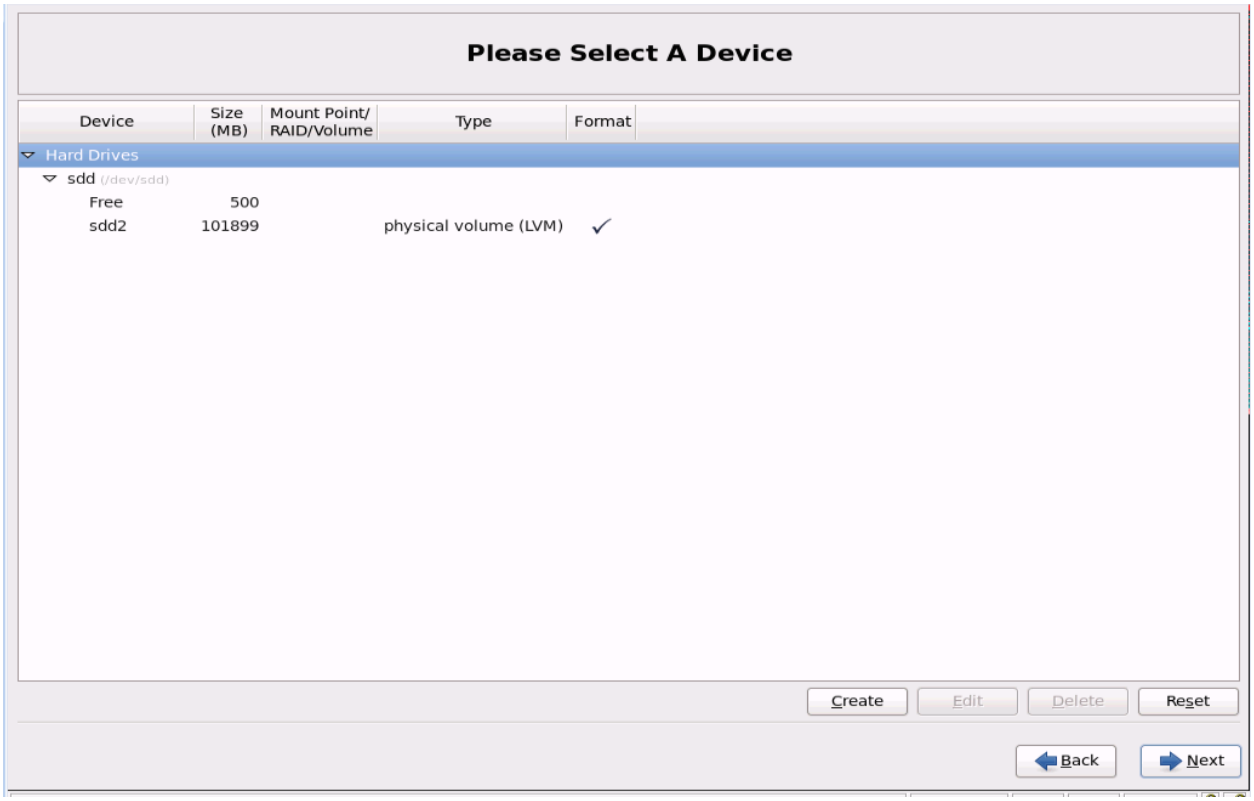

19. In the popup window, below Create Partition, select Standard Partition.
20. Click Create.

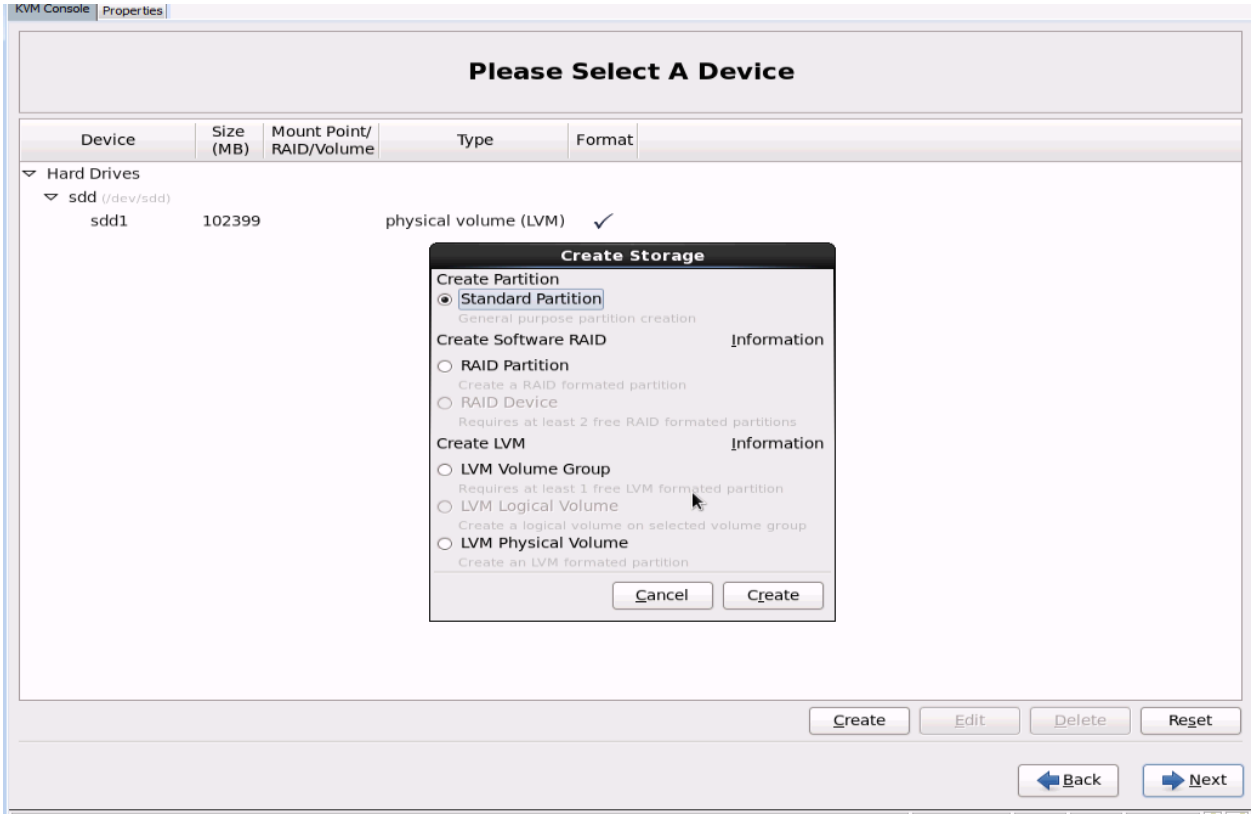

21. For the mount point, select /boot.
22. For File System Type, select ext4.
23. For the validation of this solution, we used a 5000MB size. This value can be a lower number depending on the user requirement.
24. From the additional options, select Fixed Size.
25. Click OK.

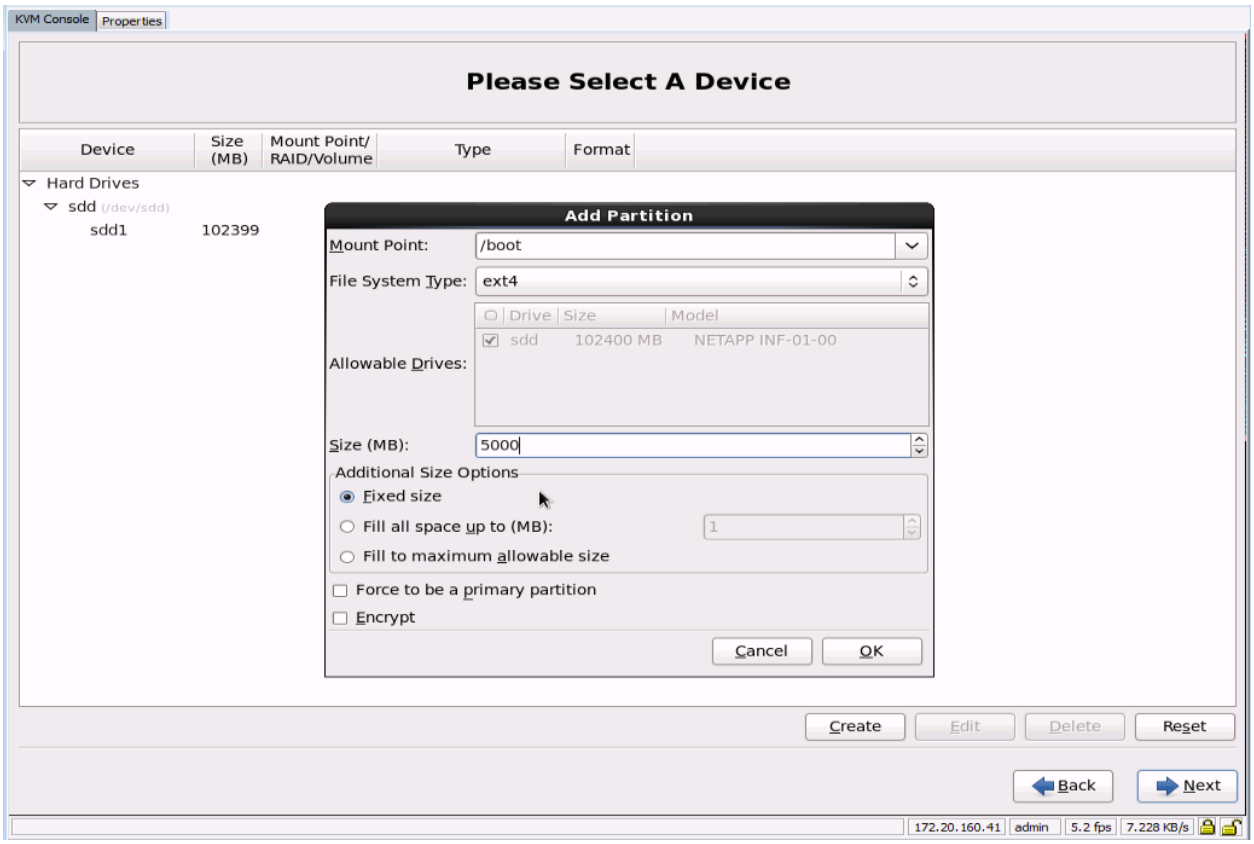

26. Select the physical volume again.
27. Click Create.
28. In the popup window, select Logical Volume Group.
29. Click Add on the Logical Volume Group.
30. In the new window, select File System Type as Swap.
31. Either use the default name for the logical volume or give the appropriate name for the swap logical volume.
32. Enter the volume size. This verified architecture used a size of 36GB.
33. Click OK.

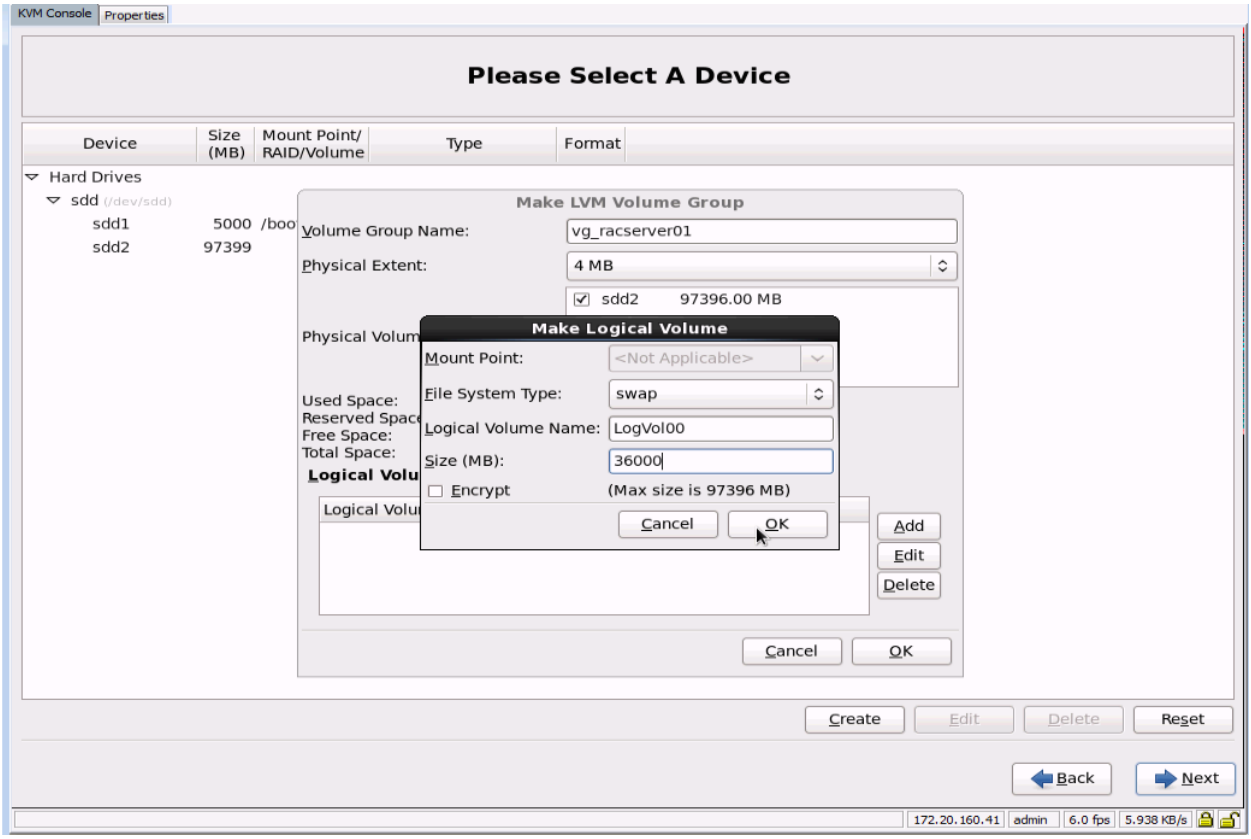

34. Click Add in the Logical Volume Group pane.
35. For Mount Point, select /.
36. For File System Type, select ext4.
37. Either use a default logical volume name or provide an appropriate name.
38. Use the rest of the space from the physical volume. This can be changed depending on specific user requirements.
39. Click OK.

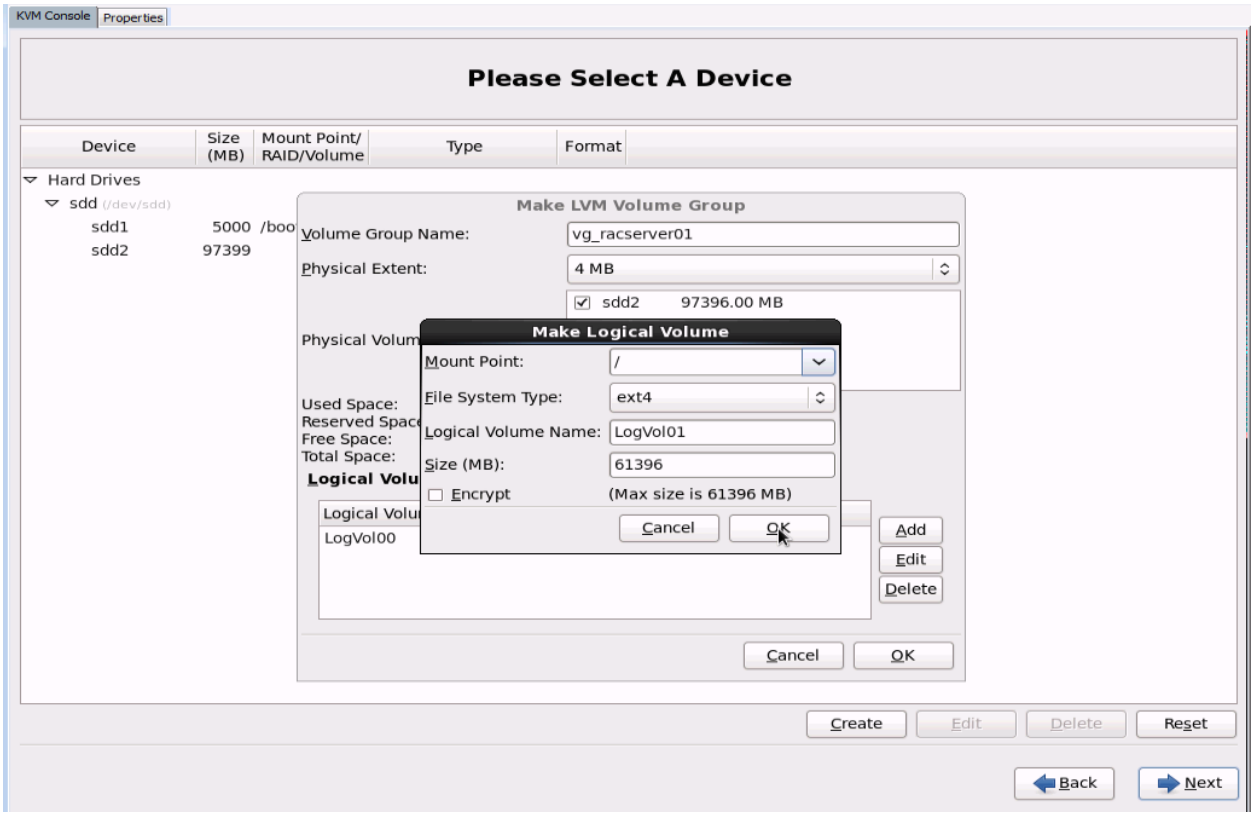

40. Click OK to create the volume group and the logical volumes.
41. Click Next.

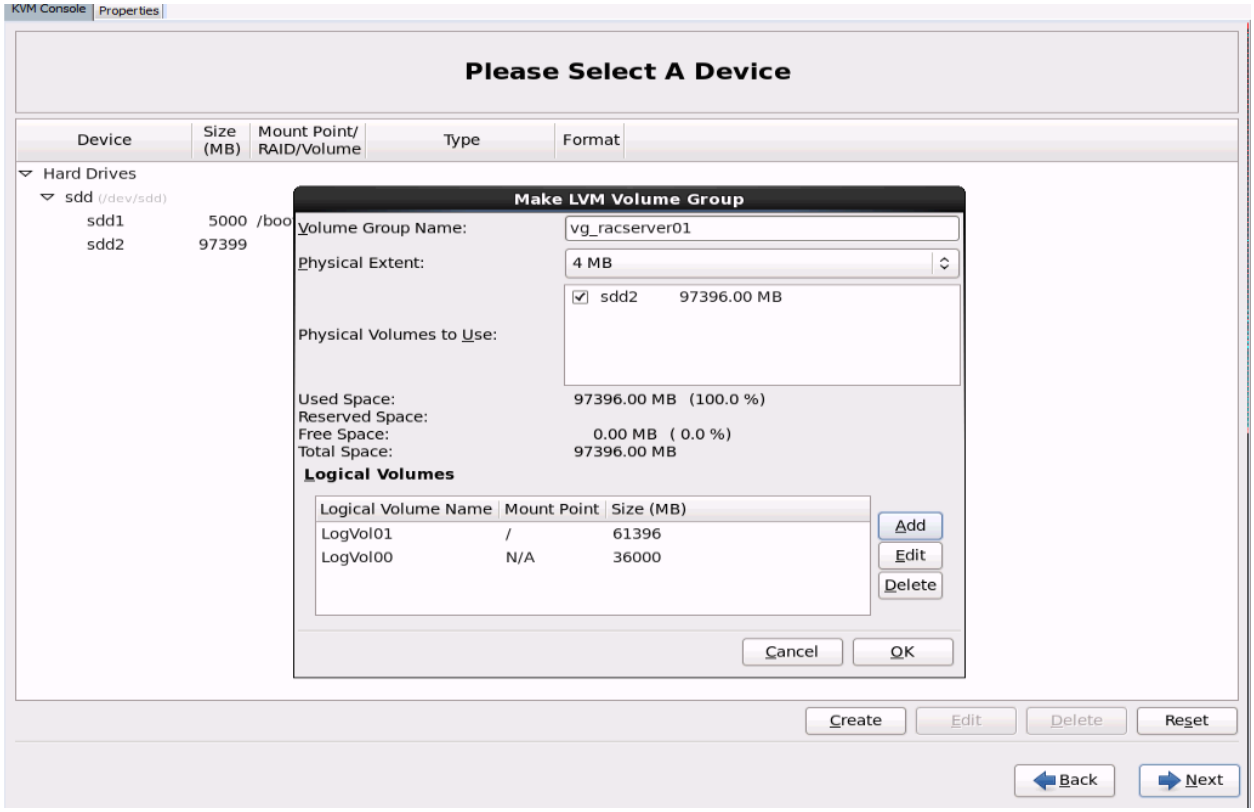

42. Click Next.

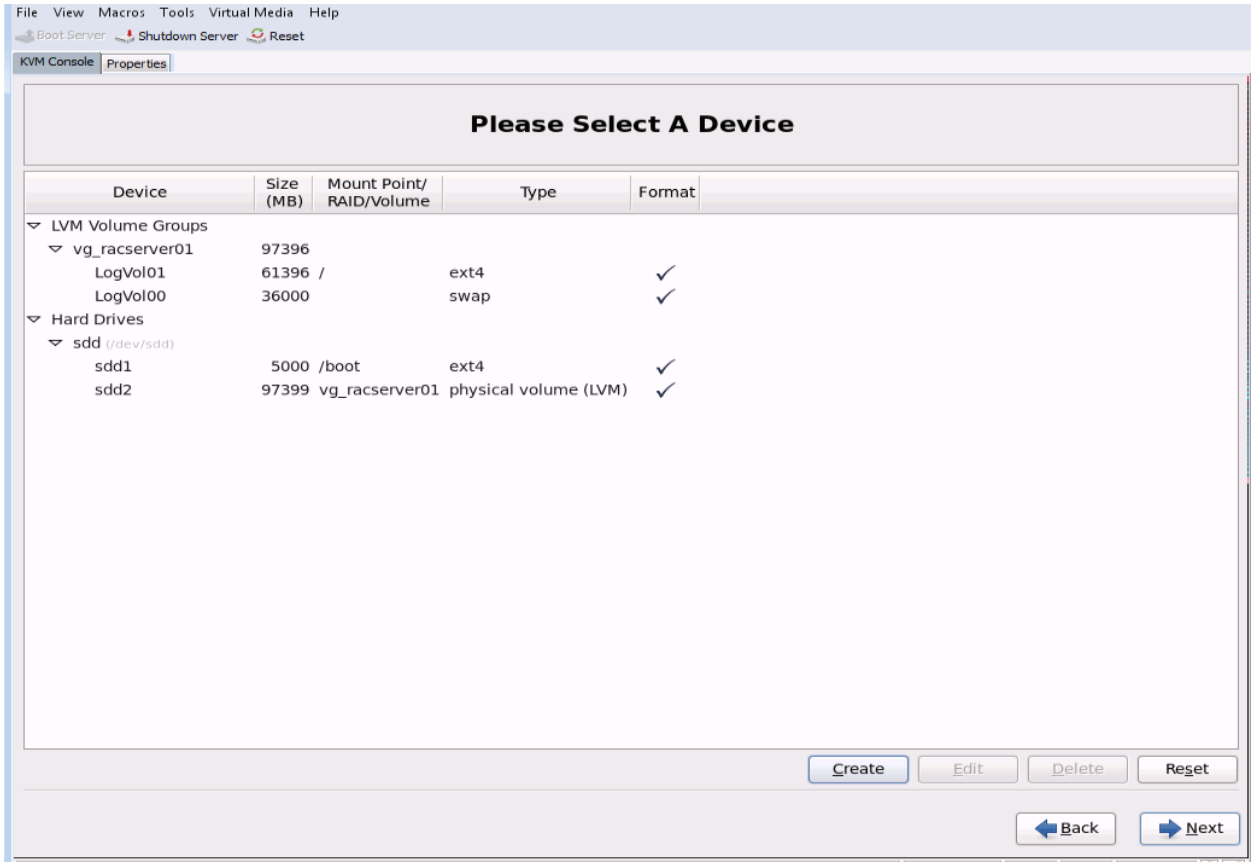

43. Select Write Changes to Disk in the message window.
44. Click Next.

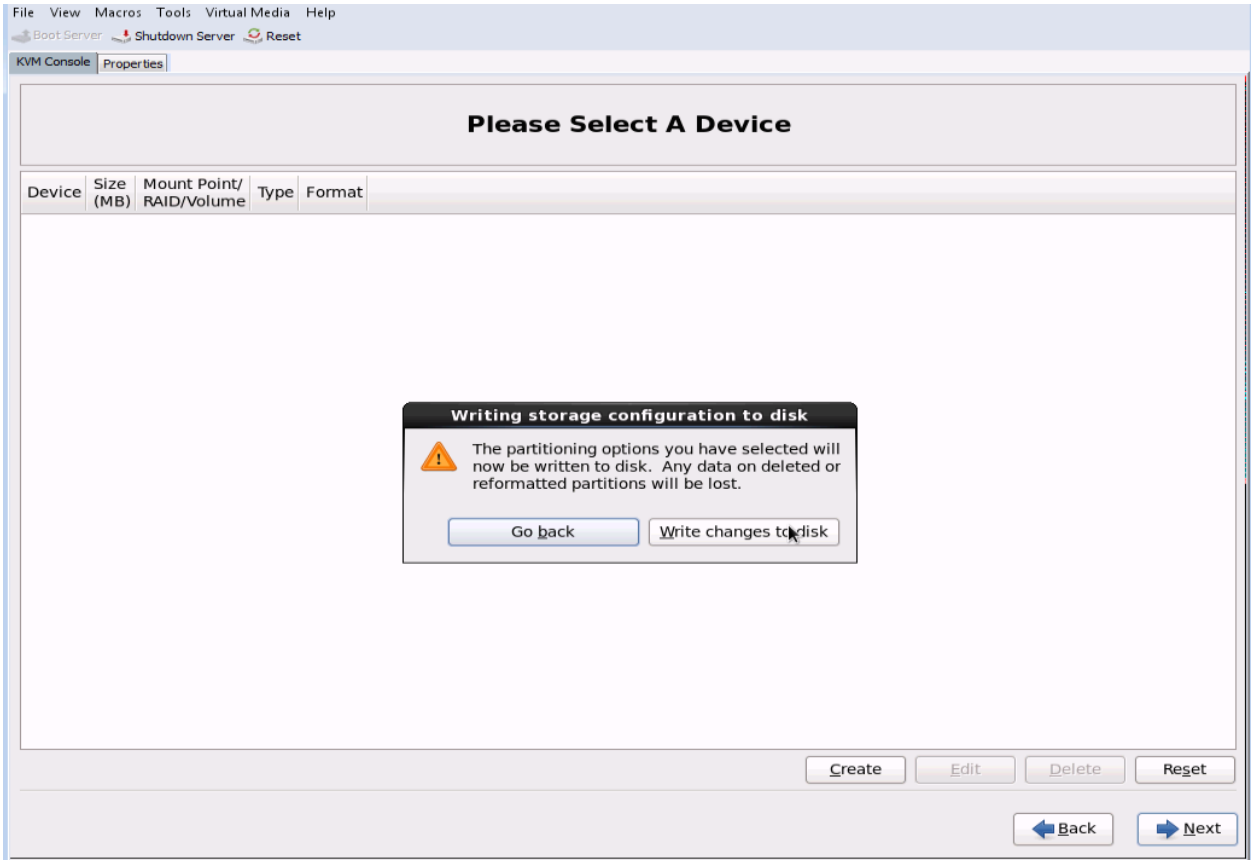

45. Click Next.

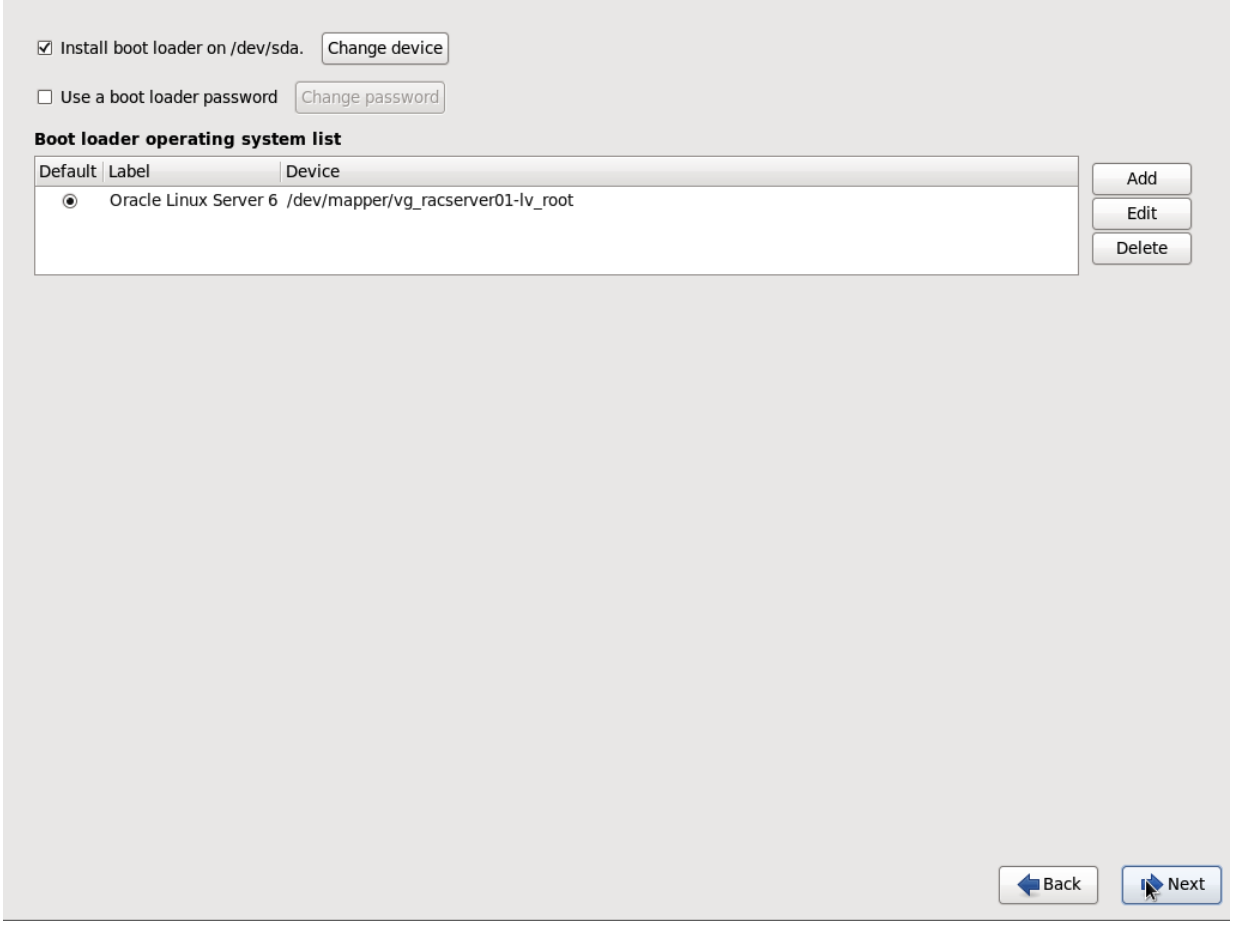

46. Select Database Server.
47. Click Next.

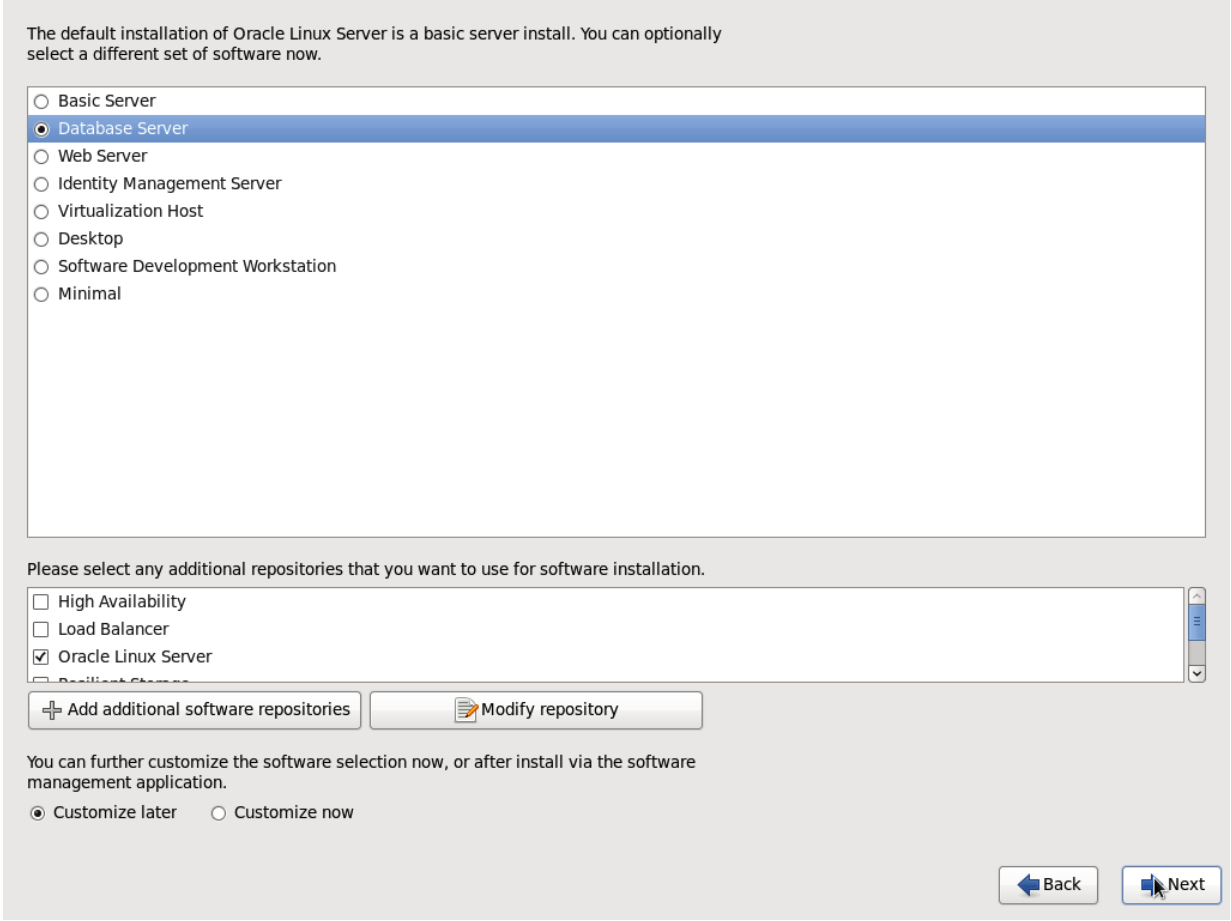

The installation of Oracle Linux starts.

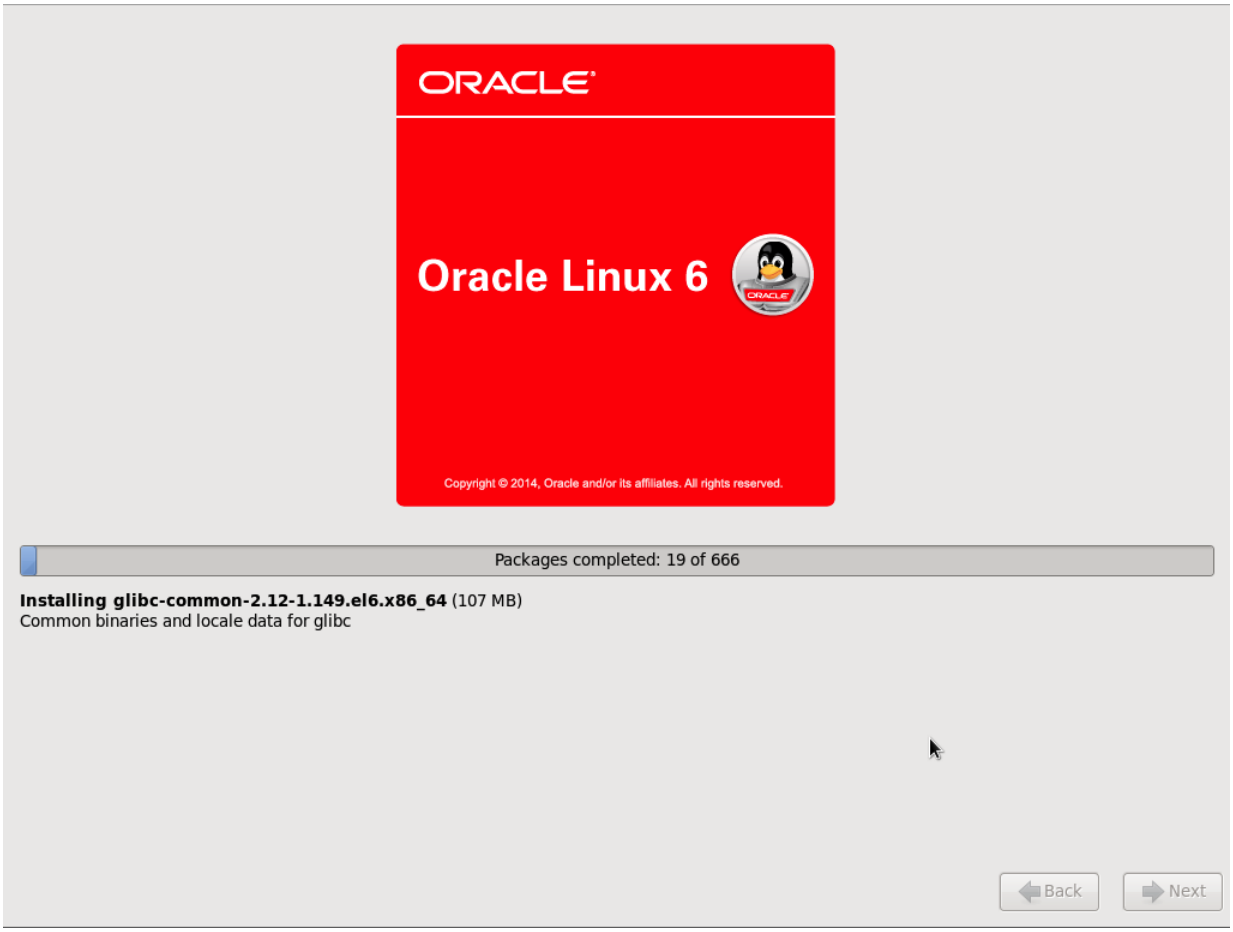

48. Click Reboot to reboot the server.
49. When the node starts to reboot, change the boot order from CD-ROM (the default) to HDD, which is the newly installed Oracle Linux image.
50. Using the KVM console, log in as a root user and set the network for public IP and the RAC interconnect. For this setup, eth0 was used as a public network, and eth1 was used as private network. For this setup, the following commands were used to set up the Ethernet interfaces:

```
rac-server-01]# cd /etc/sysconfig/network-scripts
rac-server-01]# vi ifcfg-eth0
DEVICE=eth0
HWADDR=00:25:B5:00:0A:1F
TYPE=Ethernet
ONBOOT=yes
NM_CONTROLLED=no
BOOTPROTO=no
IPADDR=172.20.160.50
NETMASK=255.255.255.0
GATEWAY=172.20.160.1
DNS1=10.61.186.19

Save the file to commit the changes

rac-server-01]# vi ifcfg-eth1
DEVICE=eth2
HWADDR=00:25:B5:00:0A:0F
TYPE=Ethernet
UUID=61a99cae-8cba-49bb-b528-58eb249ca10f
ONBOOT=yes
NM_CONTROLLED=no
```



```
BOOTPROTO=no
IPADDR=172.20.161.111
NETMASK=255.255.255.0
MTU=9000
```

51. Repeat the same for all the other RAC nodes.

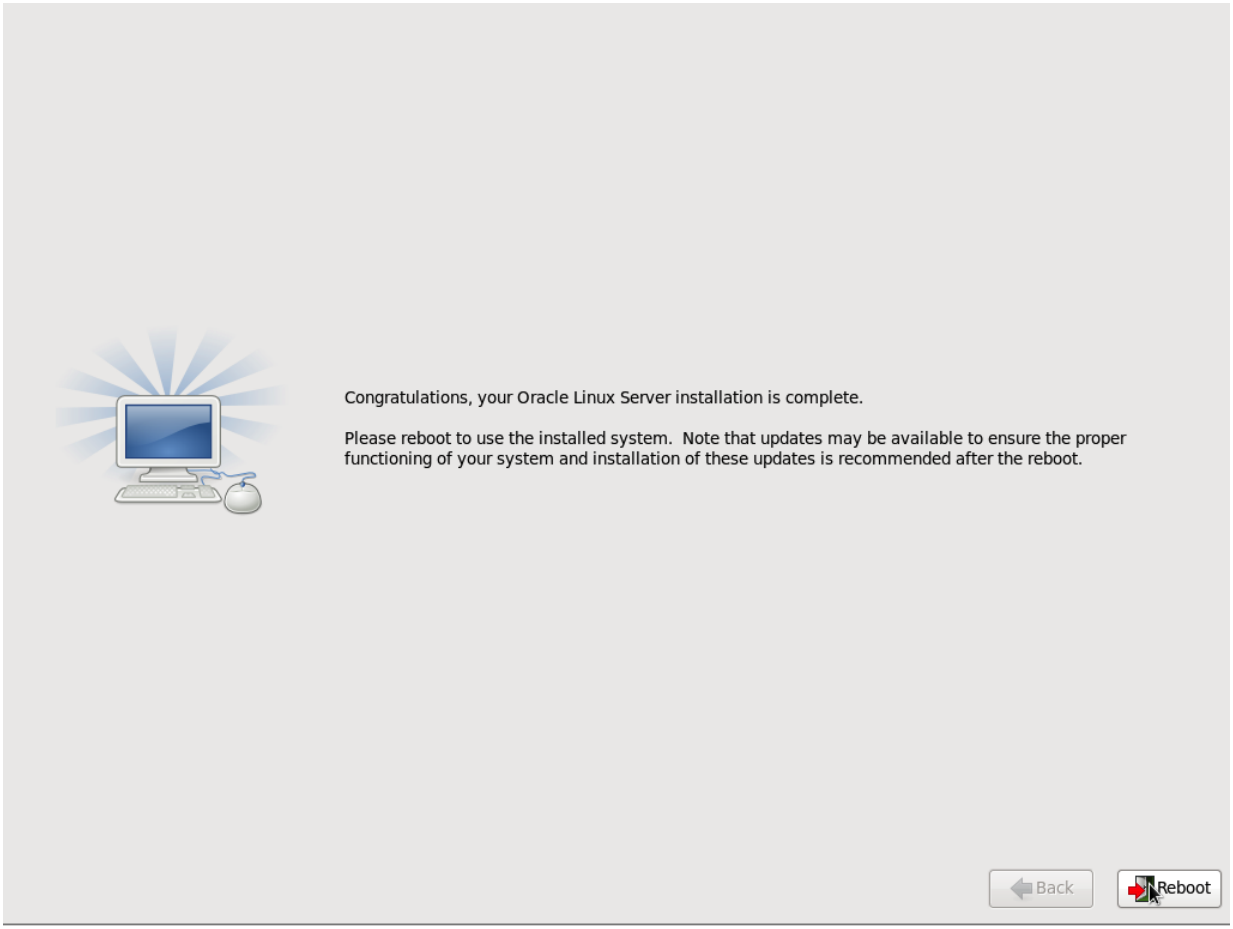

Assuming you have performed the steps to install Oracle Linux before provisioning the Oracle RAC database storage, refer back to sections “Volume Creation for Oracle RAC Grid” and “Volume Creation for Oracle RAC Database” to create the storage for the database. If you have already done this, proceed to section “Installing Cisco UCS Drivers on RAC Nodes.”

Installing Cisco UCS Drivers on RAC Nodes

Cisco UCS Virtual Interface Card (VIC) drivers facilitate communication between the Oracle Linux operating system running on the RAC nodes and the Cisco UCS VICs on the Cisco UCS blades or the RAC nodes itself. The Cisco UCS VIC driver ISO bundle includes an eNIC driver and an fNIC driver. The eNIC is the driver for the Cisco UCS VIC Ethernet NIC. The fNIC is the driver for the Cisco UCS VIC Fibre Channel over Ethernet HBA (FCoE). The drivers can be downloaded and installed from the Downloading Cisco UCS VIC Drivers page of the Cisco website. Driver version are specified in section 3.2, “Software Requirements.”

After the drivers are installed for all the RAC nodes, reboot all the RAC nodes before the Oracle RAC Node preparation.

Oracle RAC Node Preparation

Before you start the Oracle software installation, complete the following tasks to make the servers (RAC nodes) ready.

Setting Up User Accounts

Oracle recommends that you use different user accounts for the installation of the grid infrastructure (GI) and the Oracle RDBMS home. The GI is installed in a separate Oracle base, owned by user `grid`. After the grid install is done, the GI home ownership is changed to root so that it is inaccessible to unauthorized users.

1. Create the operating system group oinstall by running the following command as a root user:

   Note: For the verification of this solution, we created only one group. In a real production environment you have to add other groups to conform to security requirements.

```
groupadd -g 501 oinstall
```

2. Create the operating system users oracle and grid by running the following command as a root user:

   Note: For the verification of this solution, we created only these two users. This might vary in a production environment.

```
useradd -u 501 -c "Grid User"  -g oinstall
useradd –u 502 –c "Oracle User"  -g oinstall
```

3. Set the password for the preceding two users.

Setting Secure Shell (SSH) Passwordless Logins

In order to perform the Oracle installation, SSH passwordless login for the grid, oracle, and root users must be enabled. As part of this setup, all the RAC nodes were set up to enable the passwordless login for these three users: grid, oracle, and root. To setup passwordless login for Oracle installation, complete the following steps:

Note: The following example shows how to set up passwordless login for one user and two RAC nodes.

1. Log in to the first RAC node 1 (for example, rac-server-01) as a root user.
2. Generate a public/private key pair.

```
ssh-keygen
```

   Note: When you run the `ssh-keygen` command, you are prompted to answer several questions. Just press Enter each time until you are returned to a prompt. For example:

```
[root@rac-server-01 ~]# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
51:cc:1f:e2:b6:a5:44:f8:6f:da:3c:32:4f:5d:1d:19 root@stlrx600s6-5
The key's randomart image is:
+--[ RSA 2048]----+
|         +.    E |
|        ..= .   o|
|        .+ o . o |
|         .= o   o|
|        So =    o|
```



```
|          o o. . |
|           =. .  |
|          +.+    |
|           +..   |
+-----------------+
```

3. Copy the public key from the file `id_rsa.pub`, which is located in the `.ssh` directory of the `home` directory of the user, and paste it to a file called `authorized_keys` under the `.ssh` directory. Make sure that the text is all on one line.
4. Log in to first RAC node 2 (for example, rac-server-02) as a root user.
5. Generate a public/private key pair by typing this command:

```
ssh-keygen
```

6. When you run this command, you are prompted to answer several questions. Press Enter each time until you are returned to a prompt, just as you did for rac-server-01.
7. Copy the public key from the file `id_rsa.pub,` which is located in the `.ssh` directory of the `home` directory of the user, and paste it to a file called `authorized_keys` in the `.ssh` directory. Make sure the text is all on one line.
8. Copy the content of `authorized_keys` from rac-server-01 to the same file in rac-server-02.
9. Copy the content of `authorized_keys` from rac-server-02 to the same file in rac-server-01.
10. After you are done with the preceding two steps, the `authorized_keys` file has two lines, one from each server, and the file is identical on both servers (rac-server-01 and rac-server-02).
11. Test to make sure the secure shell communication is working without having to enter passwords by doing an SSH between two servers.
12. Follow the same steps for all the RAC nodes for the root, grid, oracle users. After you are done, the `authorized_keys` file for each user on each RAC node has eight entries. This enables passwordless logins between RAC nodes for all three users.

Networking

In order to set up RAC clustering, you need to set up the following names:

- A unique cluster name.
- Public host names for each of the RAC nodes.
- Public virtual host names for each of the RAC nodes. The virtual host is used to reroute client requests sent to the node if the node is down. Oracle recommends that the host name be in this format: public hostname-vip. Also, the virtual host name must be in the same subnet as your public IP address, must be registered with your DNS for each node, and must be capable of performing a reverse DNS lookup. The following command output shows how the reverse lookup has been set for one of the RAC nodes as an example:

```
[root@rac-server-01 oracle]# nslookup rac-server-vip-01
Server:         10.61.186.19
Address:        10.61.186.19#53

Name:   rac-server-vip-01.ice.rtp.netapp.com
Address: 172.20.160.60

[root@rac-server-01 oracle]# nslookup 172.20.160.60
Server:         10.61.186.19
Address:        10.61.186.19#53

60.160.20.172.in-addr.arpa      name = rac-server-vip-01.ice.rtp.netapp.com.
```

- A private host name for each node in the cluster. The private network should be on a dedicated switch or on a separate VLAN and preferably on a 10GbE network.

- All the public, virtual, and private IPs should be added to the `/etc/hosts` file on each of the RAC nodes. In order to avoid any confusion, all the RAC nodes were copied with the same `/etc/hosts` file for this setup, as shown in the following:

```
#--public
172.20.160.50 rac-server-01
172.20.160.51 rac-server-02
172.20.160.52 rac-server-03
172.20.160.53 rac-server-04
172.20.160.54 rac-server-05
172.20.160.55 rac-server-06
172.20.160.56 rac-server-07
172.20.160.57 rac-server-08

# Public Virtual IP (VIP) addresses
172.20.160.60 rac-server-vip-01
172.20.160.61 rac-server-vip-02
172.20.160.62 rac-server-vip-03
172.20.160.63 rac-server-vip-04
172.20.160.64 rac-server-vip-05
172.20.160.65 rac-server-vip-06
172.20.160.66 rac-server-vip-07
172.20.160.67 rac-server-vip-08

# RAC Interconnect - eth2
172.20.161.111         rac-server-01-priv
172.20.161.112         rac-server-02-priv
172.20.161.113         rac-server-03-priv
172.20.161.114         rac-server-04-priv
172.20.161.115         rac-server-05-priv
172.20.161.116         rac-server-06-priv
172.20.161.117         rac-server-07-priv
172.20.161.118         rac-server-08-priv
```

- Add a single client access name (SCAN) for the cluster that resolves to three IP addresses on the DNS server. SCAN is an Oracle RAC feature that provides a single name for clients to access Oracle databases running in a cluster. The SCAN IPs must not be added to the `/etc/hosts` file on the RAC nodes and must be resolved by DNS. The reverse lookup should also be enabled for the SCAN, the output of which is shown in the following example.

```
nslookup rac
Server:         10.61.186.19
Address:        10.61.186.19#53

Name:   rac.ice.rtp.netapp.com
Address: 172.20.160.101
Name:   rac.ice.rtp.netapp.com
Address: 172.20.160.102
Name:   rac.ice.rtp.netapp.com
Address: 172.20.160.100

[root@rac-server-01 oracle]# nslookup  172.20.160.101
Server:         10.61.186.19
Address:        10.61.186.19#53

101.160.20.172.in-addr.arpa     name = rac.ice.rtp.netapp.com.

[root@rac-server-01 oracle]# nslookup  172.20.160.102
Server:         10.61.186.19
Address:        10.61.186.19#53

102.160.20.172.in-addr.arpa     name = rac.ice.rtp.netapp.com.

[root@rac-server-01 oracle]# nslookup  172.20.160.100
Server:         10.61.186.19
Address:        10.61.186.19#53

100.160.20.172.in-addr.arpa     name = rac.ice.rtp.netapp.com.
```


1. Edit the /etc/nsswitch.conf file on each of the RAC nodes to change the search order of the name resolution for the hosts.

```
hosts: dns files nis
service nscd restart
```

RAC Node Time Synchronization

For this solution, the Network Time Protocol (NTP) server was disabled, and Oracle Clusterware Cluster Time Synchronization Service (CTSS) was used instead. To disable the NTP server, complete the following commands on all participating RAC nodes. If the requirement is to synchronize with an external time source, you must use NTPD, which makes CTSSD run in observer mode.

```
[root@rac-server-01 service ntpd stop
Shutting down ntpd:                                        [  OK  ]
[root@rac-server-01 mv /etc/ntp.conf /etc/ntp.conf.bkp
chkconfig ntpd off
```

Recommended Kernel Parameters

The following kernel parameters were used as part of this setup. All of the RAC nodes are required to be updated with the following settings. These parameters are modified by editing the file `/etc/sysctl.conf`. Customers implementing the Oracle RAC solution are advised to check the Oracle documentation for any updates on these kernel settings.

```
kernel.msgmnb = 65536
kernel.msgmax = 65536
kernel.shmmax = 68719476736
kernel.shmall = 4294967296
kernel.shmmni = 4096
kernel.sem = 8192 48000 8192 8192
fs.file-max = 6815744
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 4194304
net.core.rmem_max = 16777216
net.core.wmem_default = 262144
net.core.wmem_max = 16777216
net.ipv4.ipfrag_high_thresh = 524288
net.ipv4.ipfrag_low_thresh = 393216
net.ipv4.tcp_rmem = 4096 524288 16777216
net.ipv4.tcp_wmem = 4096 524288 16777216
net.ipv4.tcp_timestamps = 0
net.ipv4.tcp_sack = 0
net.ipv4.tcp_window_scaling = 1
net.core.optmem_max = 524287
net.core.netdev_max_backlog = 2500
net.ipv4.tcp_mem = 16384 16384 16384
fs.aio-max-nr = 1048576
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 0
vm.min_free_kbytes=262144
vm.swappiness=100
```

1. To make the changes permanent, run the following command on all of the RAC nodes:

```
[root@rac-server-01 /sbin/sysctl –p
```

User Limits (ulimit)

Oracle recommends setting the user limits to an optimal number; the affects application performance. For this solution, the following user limits were applied for both the grid and oracle users. The settings were applied on all RAC nodes. Refer to the Oracle documentation for any updates to these settings.

```
grid soft nproc 2047
grid hard nproc 16384
```



```
grid soft nofile 1024
grid hard nofile 65536
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
```

Home Directory Path

Before starting the Oracle installation, home directories must be created so that the installer GUI can be updated with the appropriate directory paths. For this installation, we used the following path names and the associated commands on all RAC nodes to set up the directory paths.

```
For the Oracle Grid
mkdir -p /u01/11.2.0/grid
chown -R grid:oinstall /u01/11.2.0/grid
chmod -R 775 /u01/11.2.0/grid
mkdir -p /u01/app/oraInventory
chown -R grid:oinstall /u01/app/oraInventory
chmod -R 775 /u01/app/oraInventory

For the Oracle Base
mkdir -p /u01/app/oracle
chown -R oracle:oinstall /u01/app/oracle
chmod -R 775 /u01/app/oracle
```

Installing Requisite Packages

NetApp recommends running the `runcluvfy.sh` script, which is part of the Oracle installation package. The script verifies that all the required packages exist for each RAC node. The script can be run from the first RAC node and validates all the other nodes that are part of the cluster. The script can be executed as follows and is available in the installer directory of the grid install software:

```
./runcluvfy.sh stage -pre crsinst -n rac-server-01,rac-server-02,rac-server-03,rac-server-
04,rac-server-05,rac-server-06,rac-server-07,rac-server-08 -fixup -verbose -asm -asmdev
dev/mapper/crs1,/dev/mapper/crs2,/dev/mapper/crs3
```

Oracle Software Installation

This section provides the details around installing the Oracle 12c RAC environment used in this solution.

Before you start the grid infrastructure installation, you must finish the section “Volume Creation for Oracle RAC Grid” so that all the CRS disks are available before the grid installation.

Oracle Grid Infrastructure Installation

To install the Oracle grid infrastructure, complete the following steps:

1. Run the installation application as a grid user.
2. Select Install and Configure Grid Infrastructure for a Cluster and click Next.

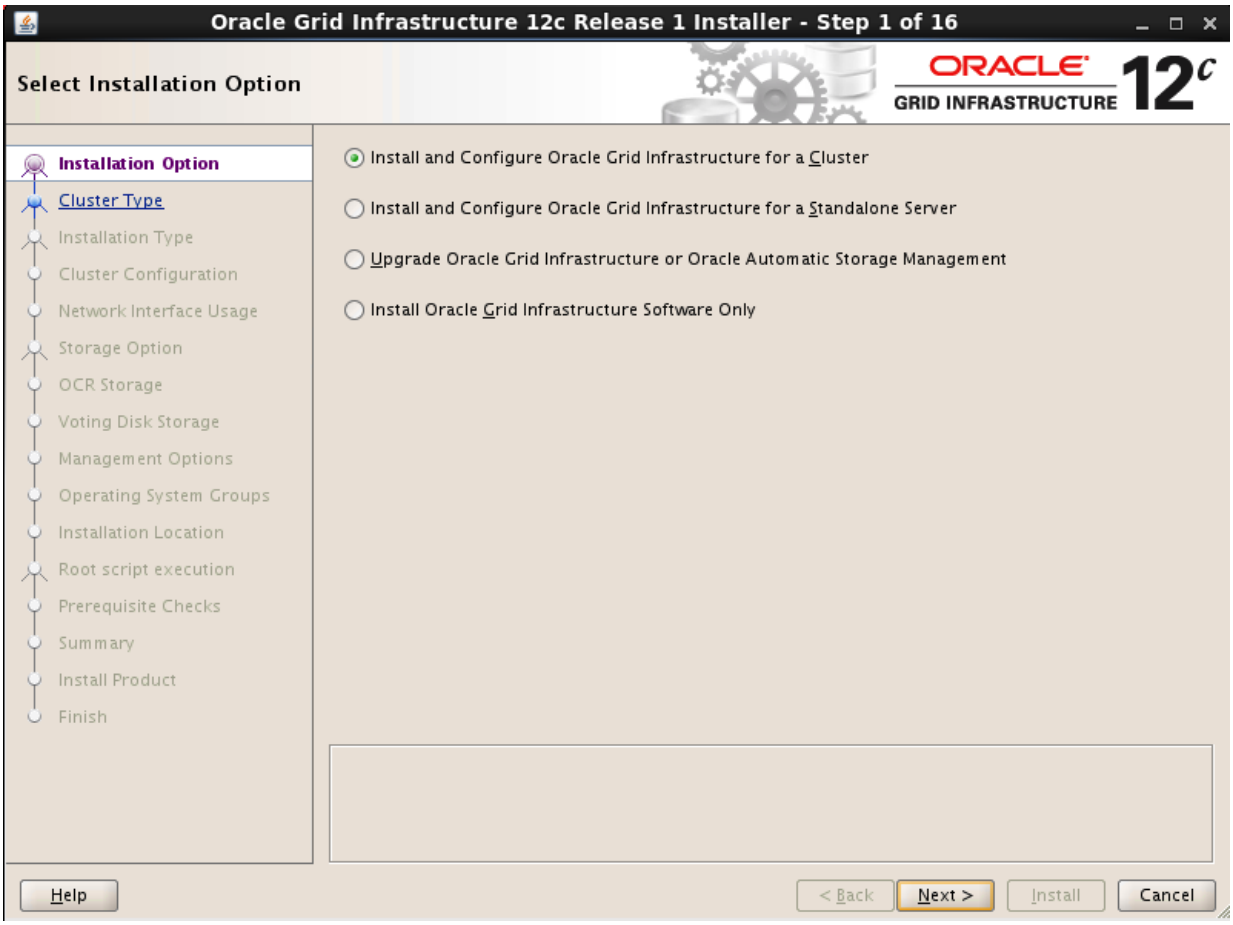

3. Select Configure a Standard cluster and click Next.

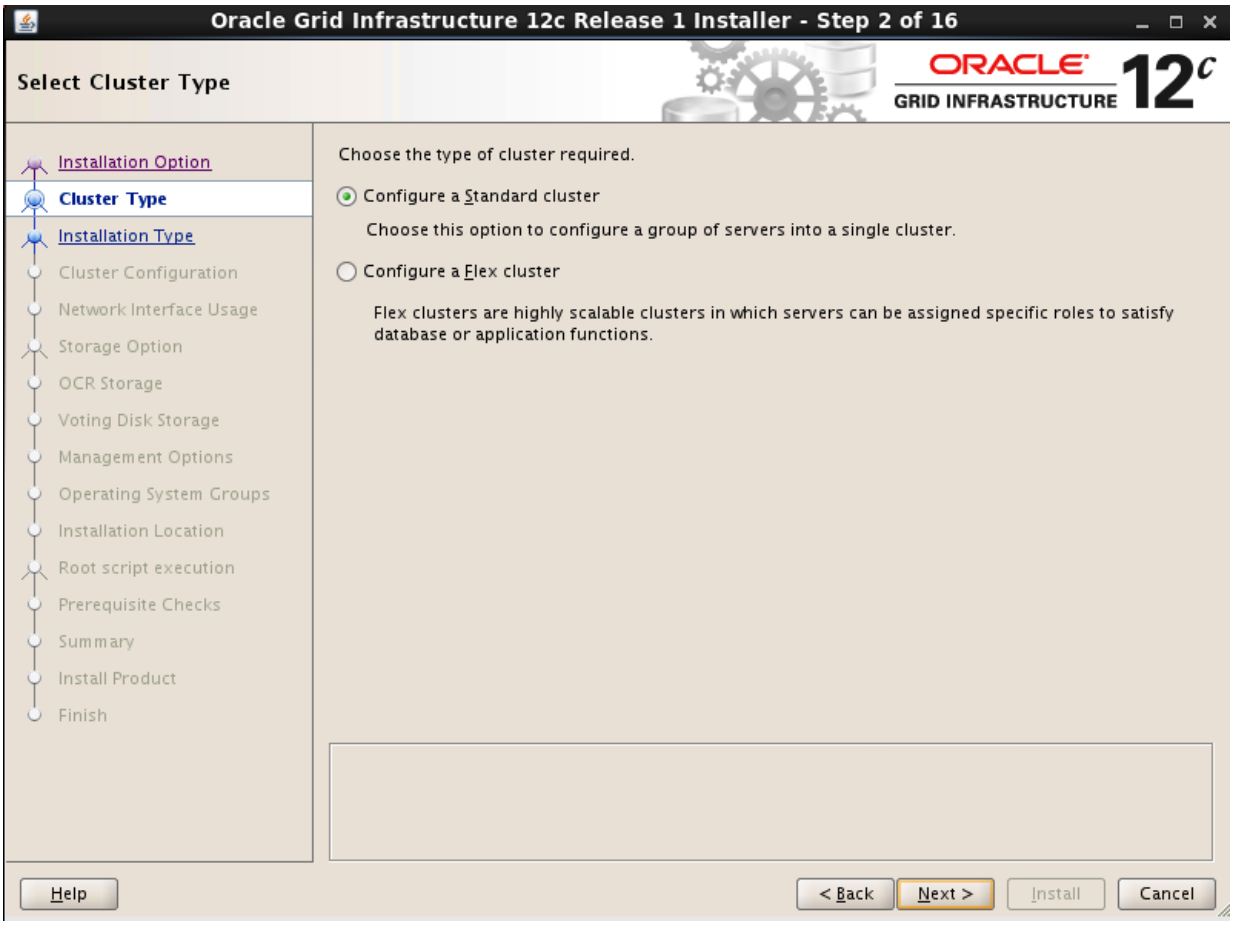

4. Select Advanced Installation and click Next.

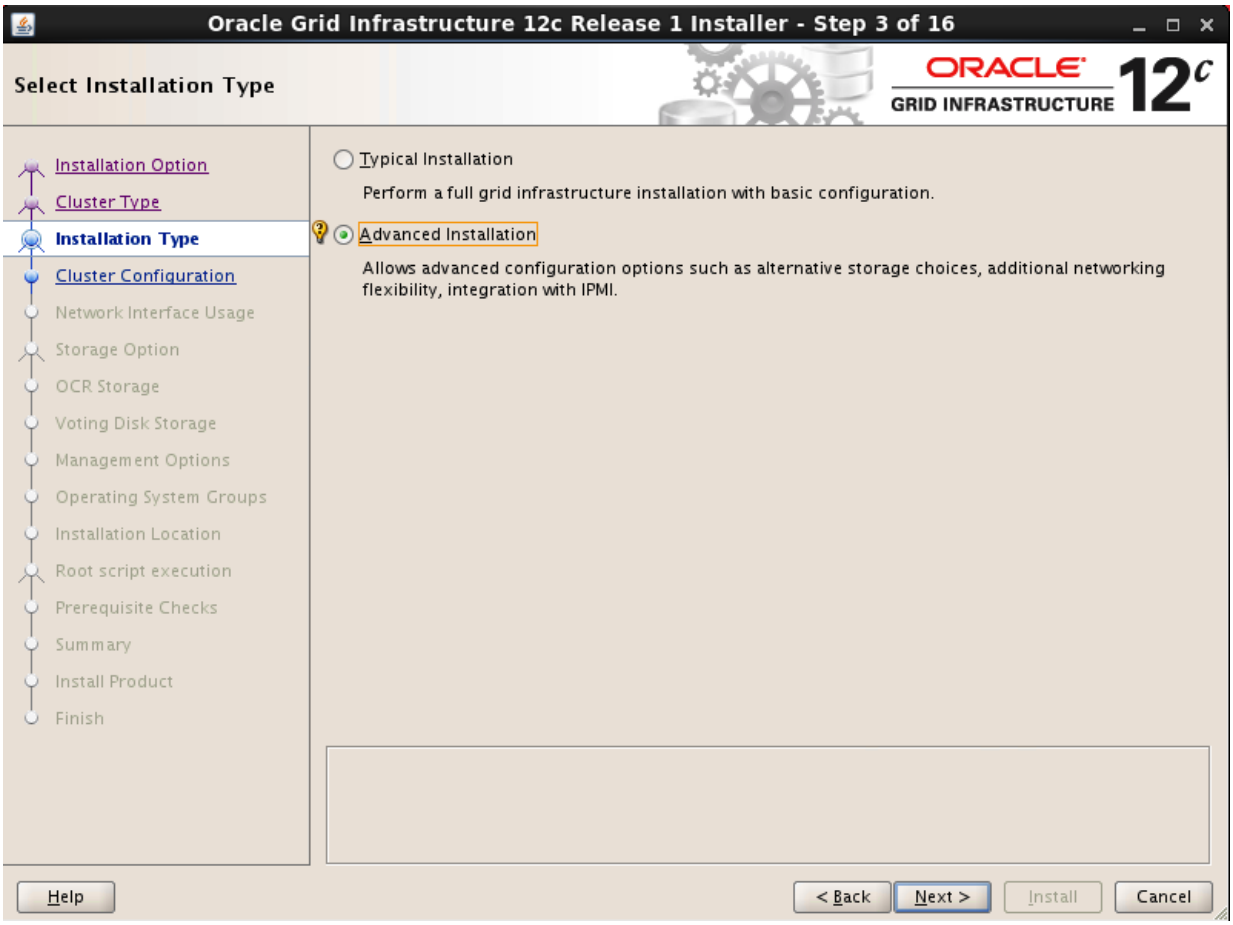

5. Select the preferred Product Language and click Next.

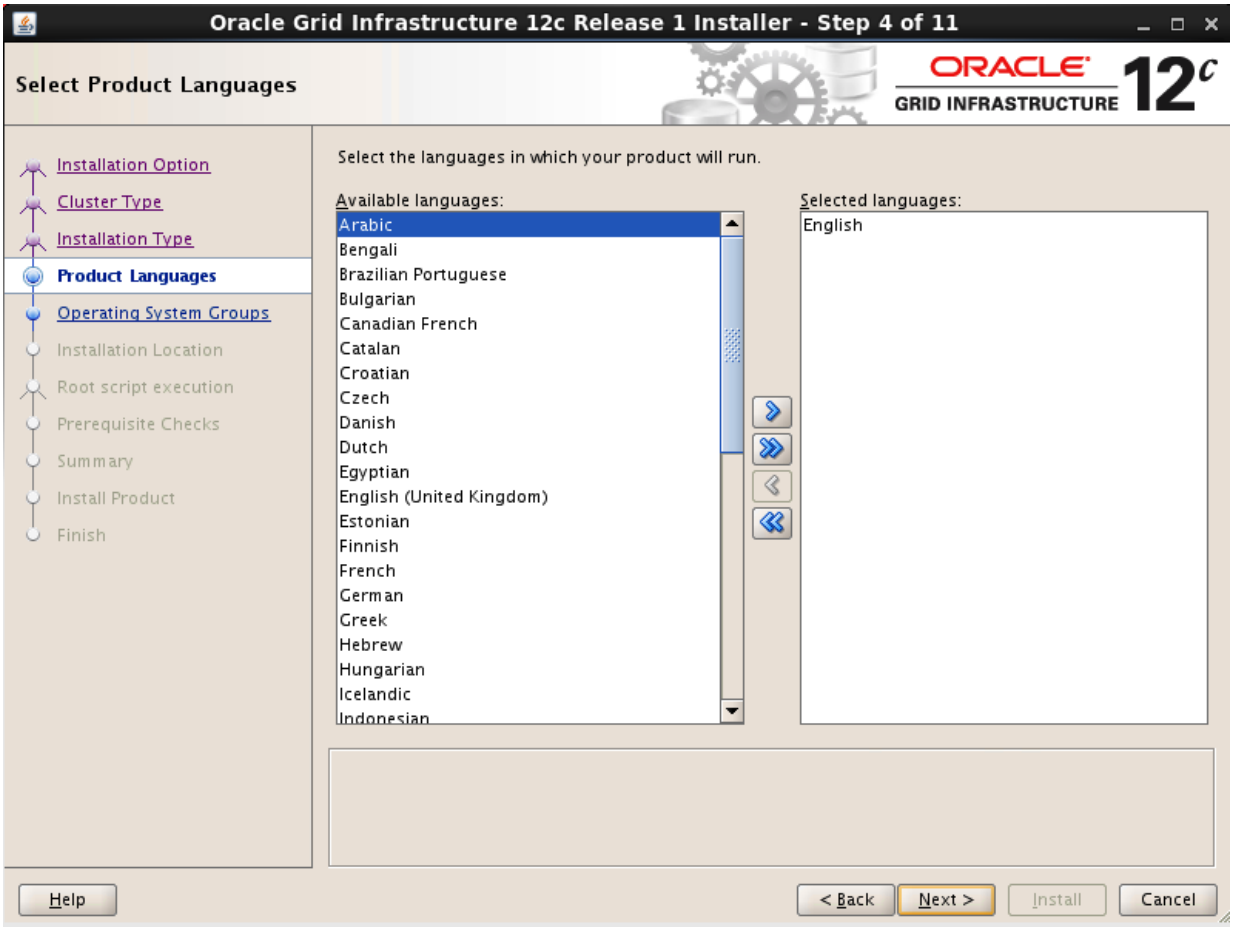

6. Provide the cluster name, SCAN name, and SCAN port name. GNS was disabled in this example.
7. Click Next.

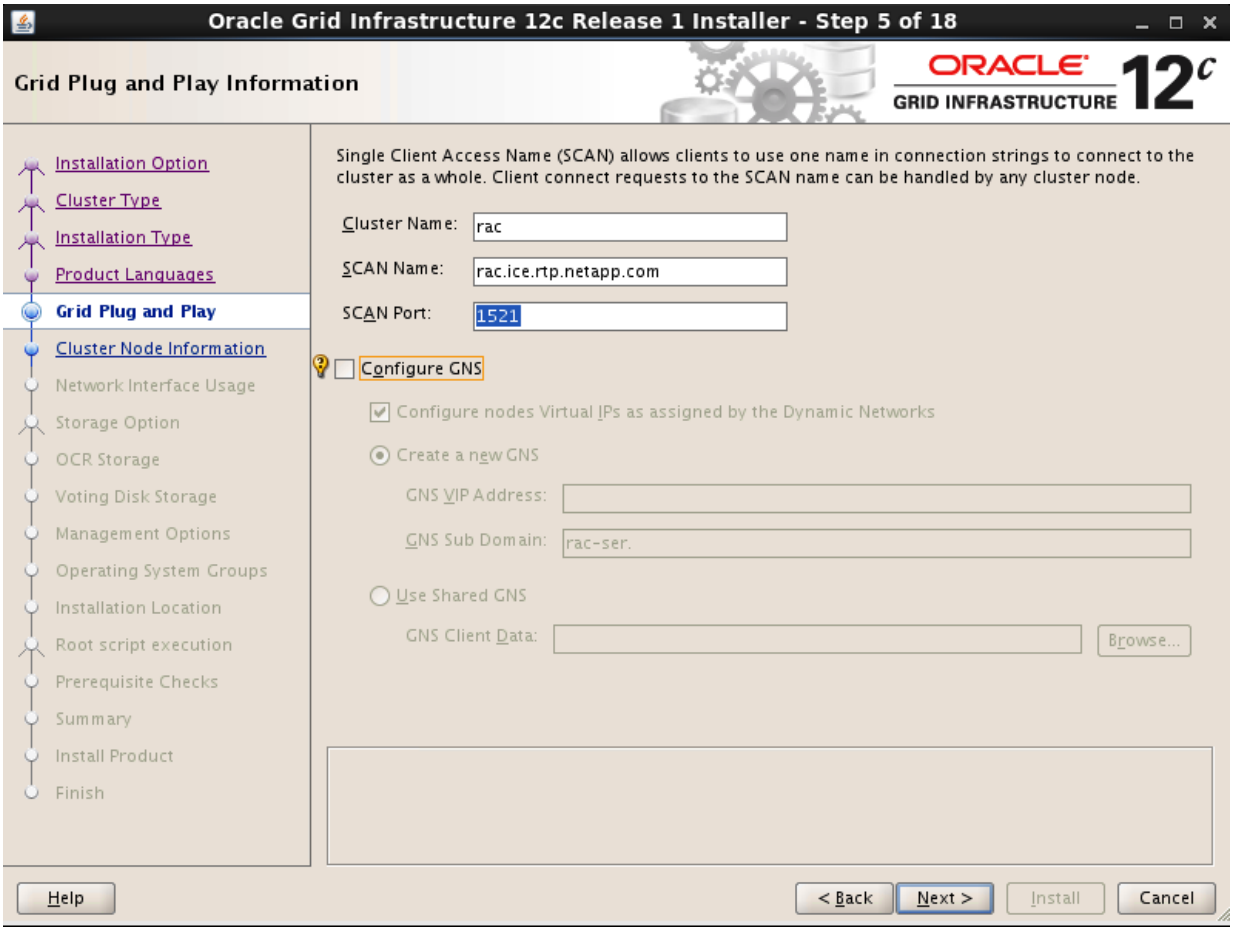

8. Provide the public and virtual host name for all the RAC nodes.
9. Click Add to add all the RAC nodes.
10. Click Next.

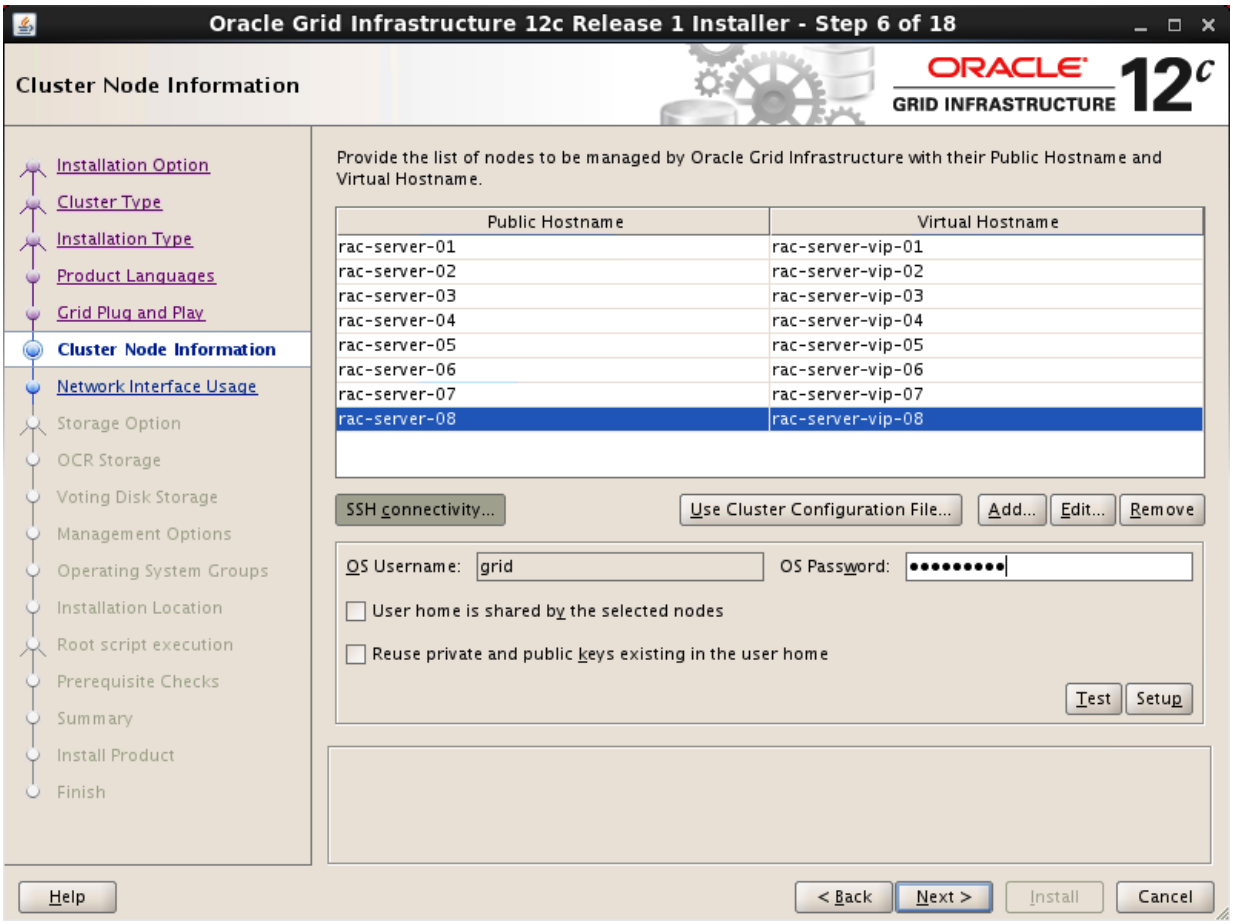

11. Make sure the network interfaces are displayed as required. Any anomaly is likely due to incorrect entries in the `/etc/hosts` file on the RAC nodes. The `/etc/hosts` file must be the same on all the RAC nodes.
12. Click Next.

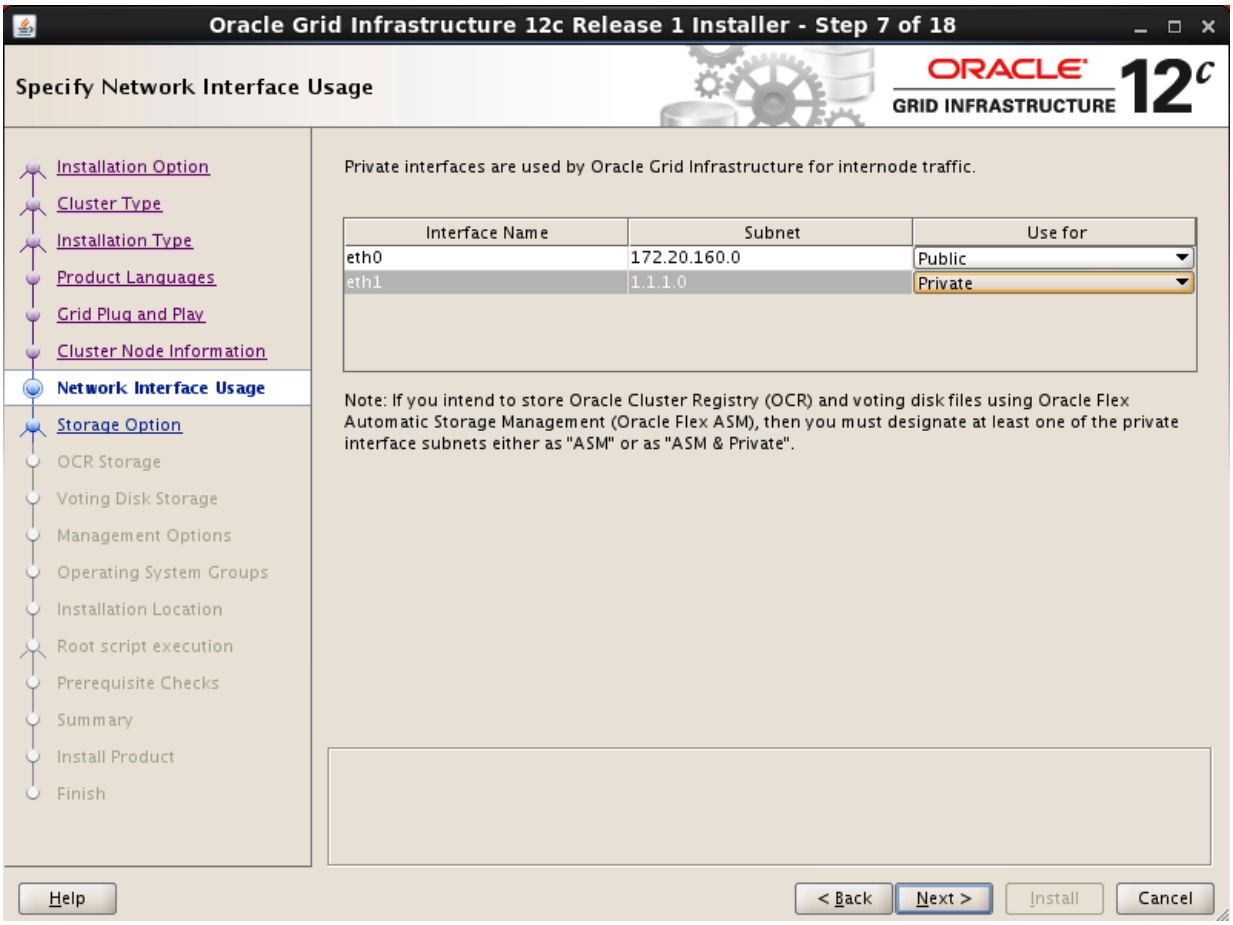

13. Select Use Standard ASM for storage and click Next.

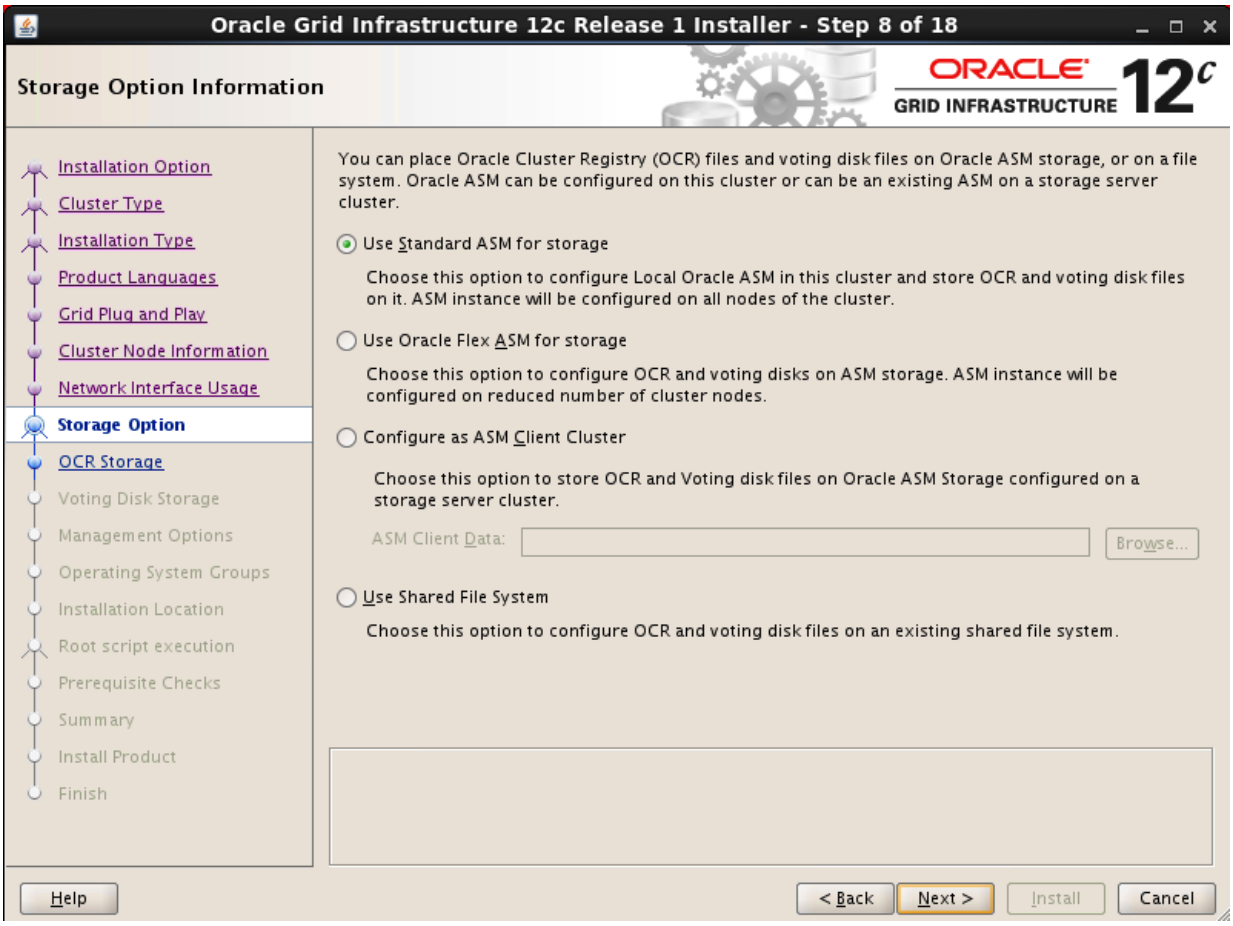

14. Configure the ASM disk group for the CRS disk. For this setup, external redundancy was selected for the Oracle ASM diskgroup CRS, and three disks were used.
15. Click Next.

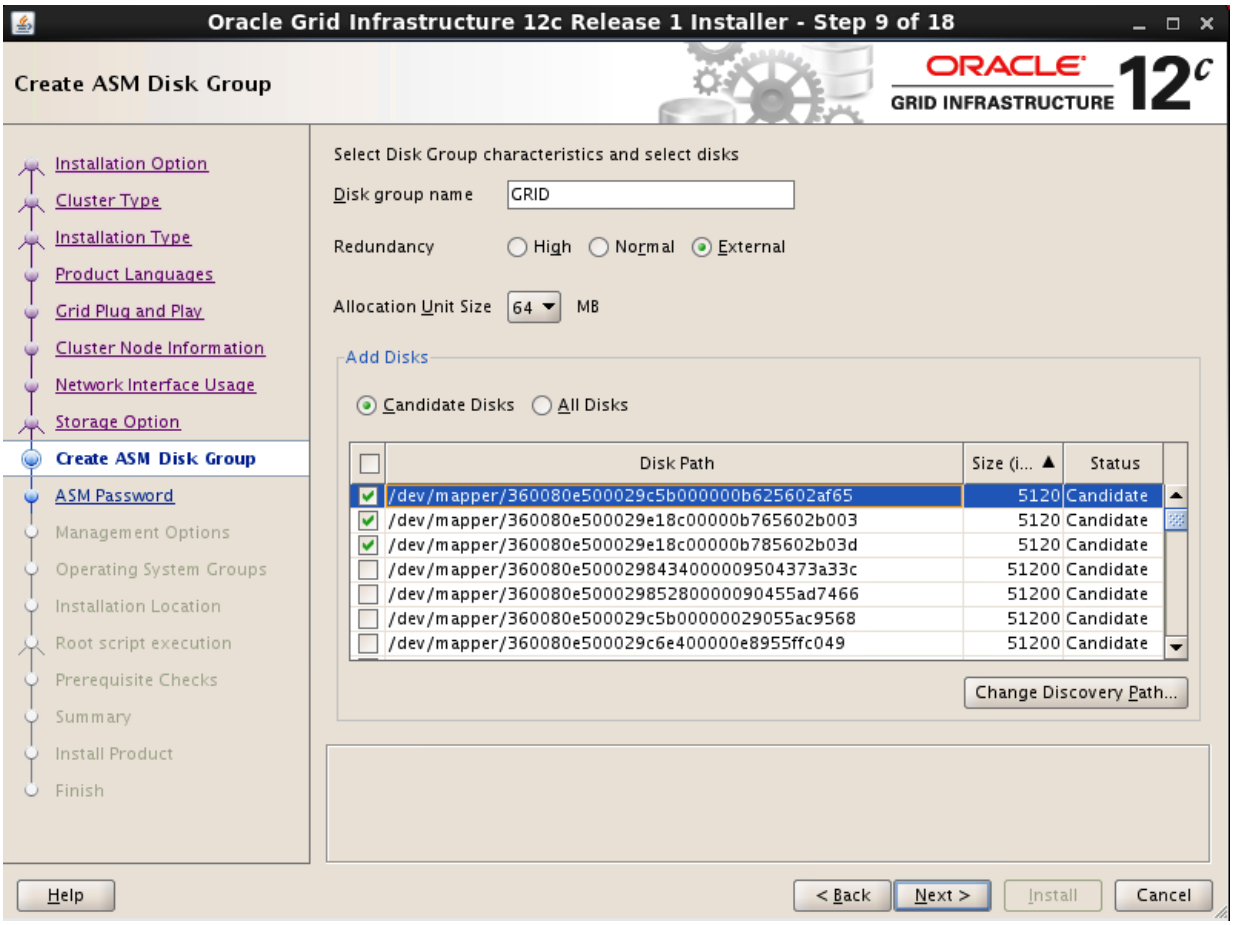

16. The same Oracle ASM passwords were used for all the accounts. In a production scenario, the passwords can be different.
17. Click Next.

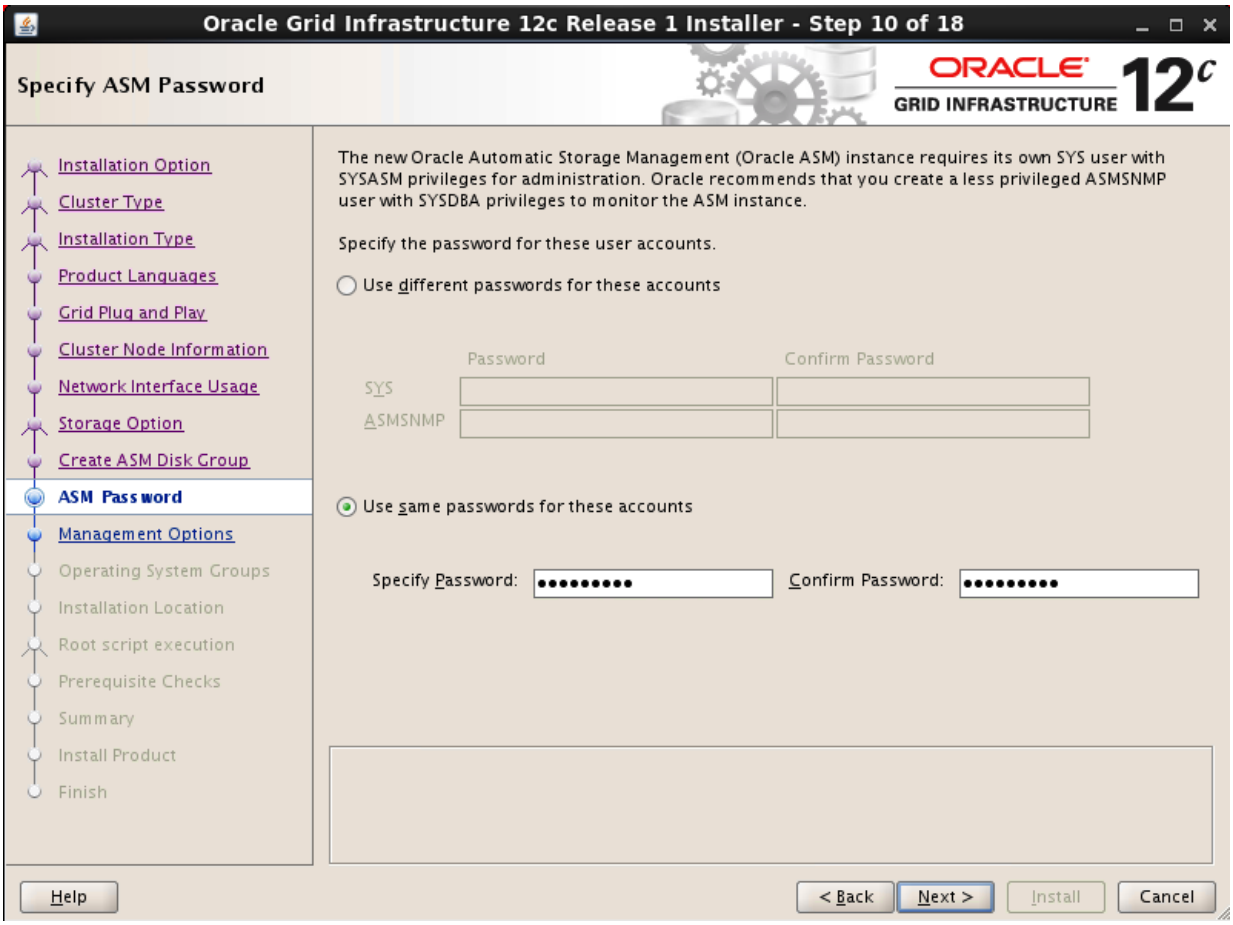

18. For this setup, the Intelligent Platform Management Interface (IPMI) option was disabled.
19. Click Next.

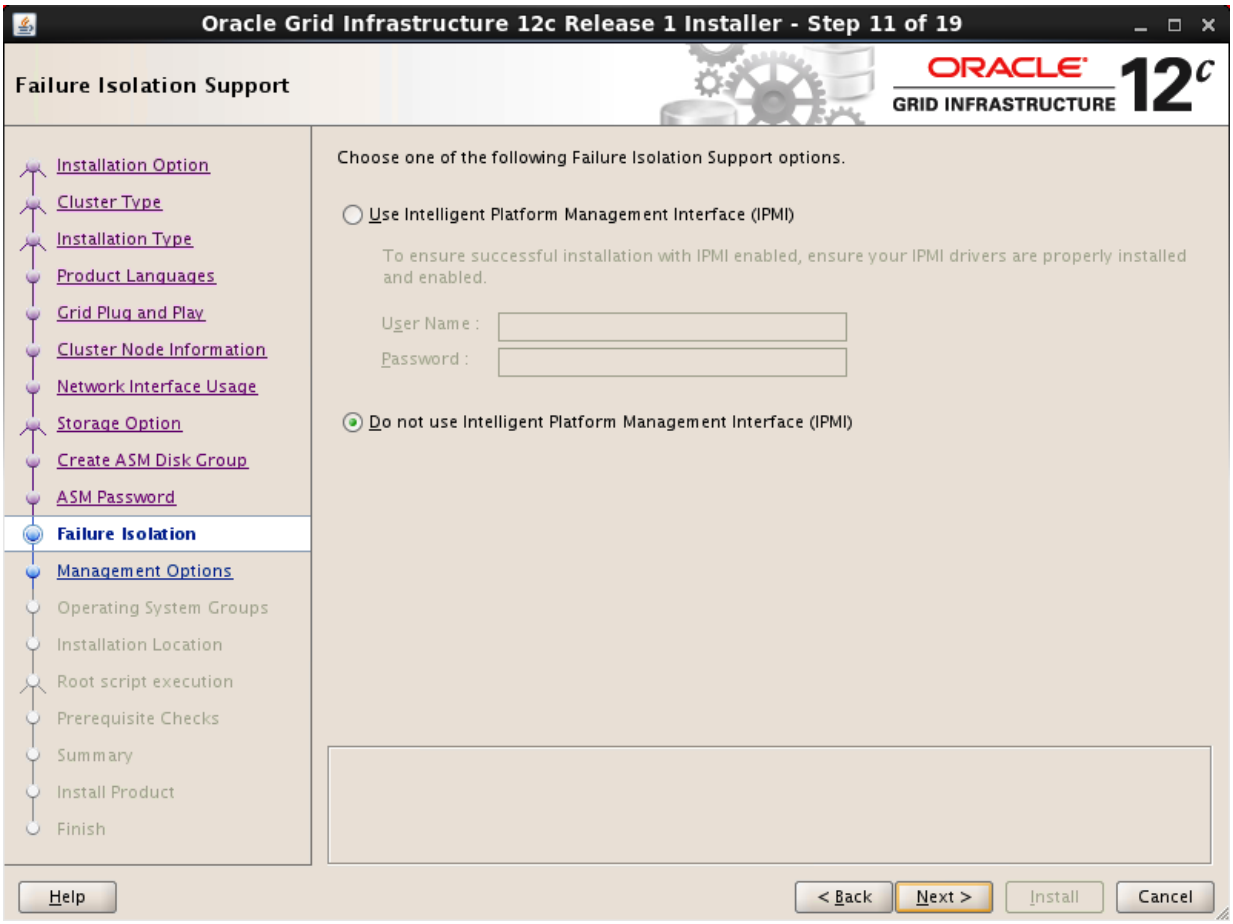

20. Make sure that Register with EM Cloud Control is unchecked and click Next.

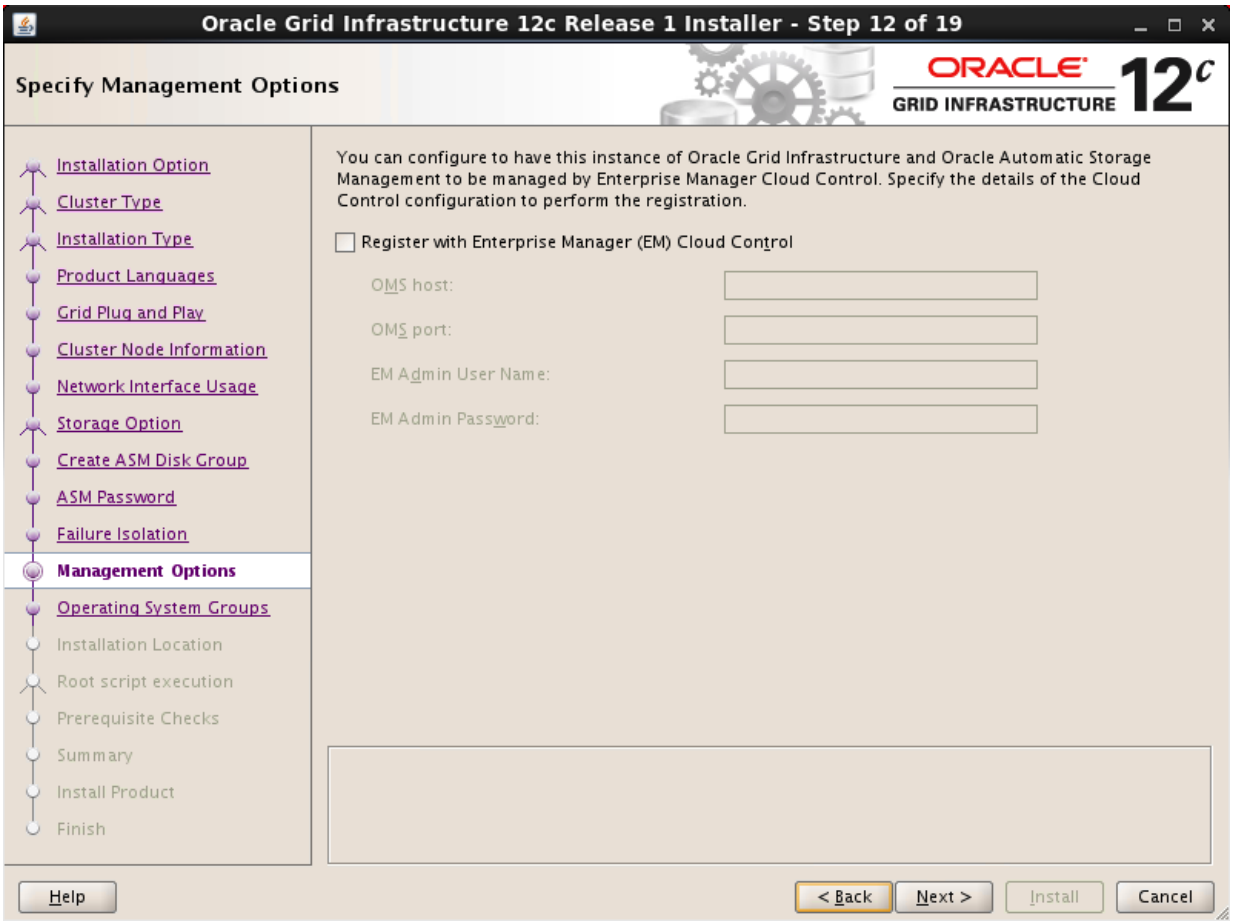

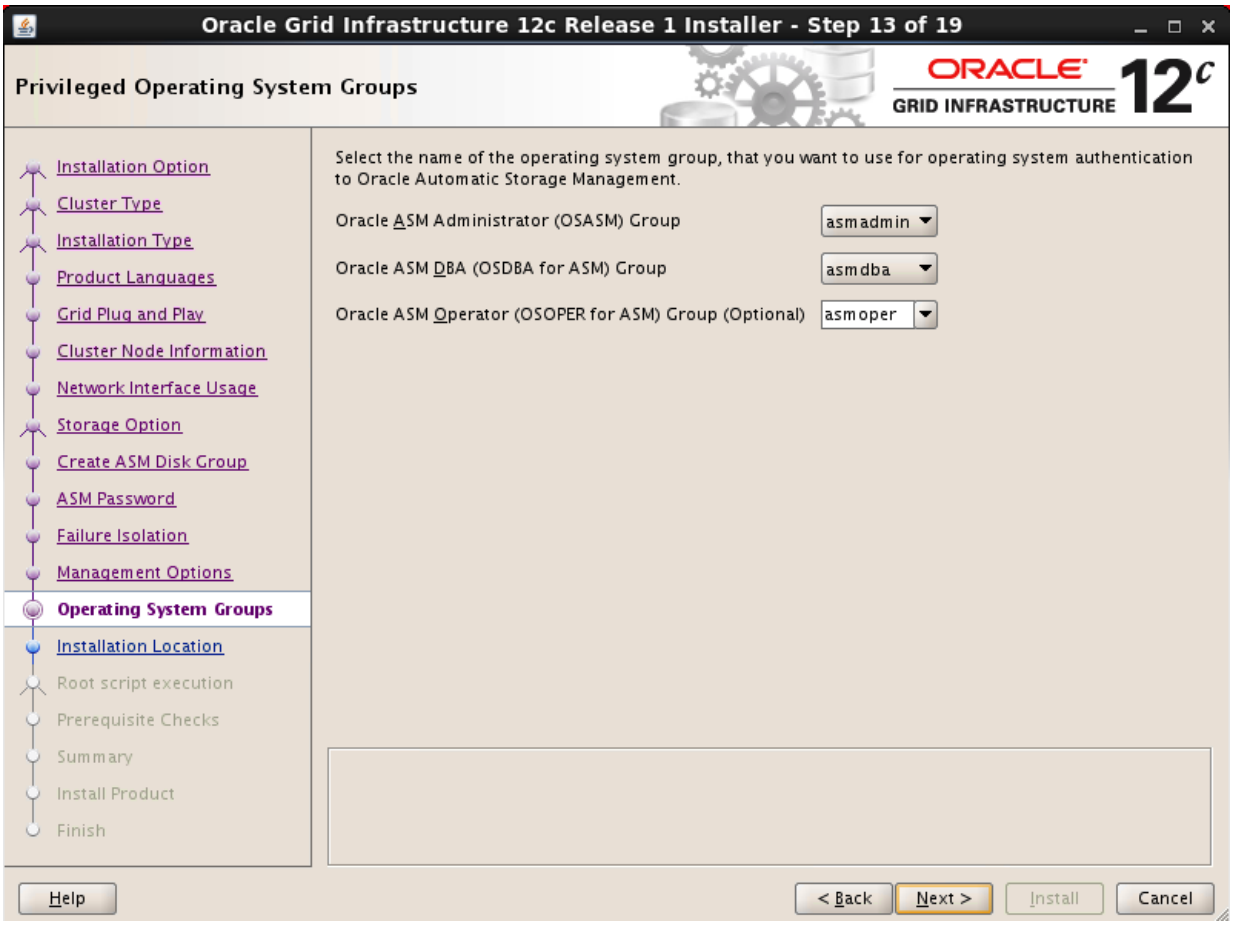

21. Provide the grid base software location directories. For this installation, we used /oracle/base for the grid base and /oracle/app/product/12c/grid for the grid software location.
22. Click Next.

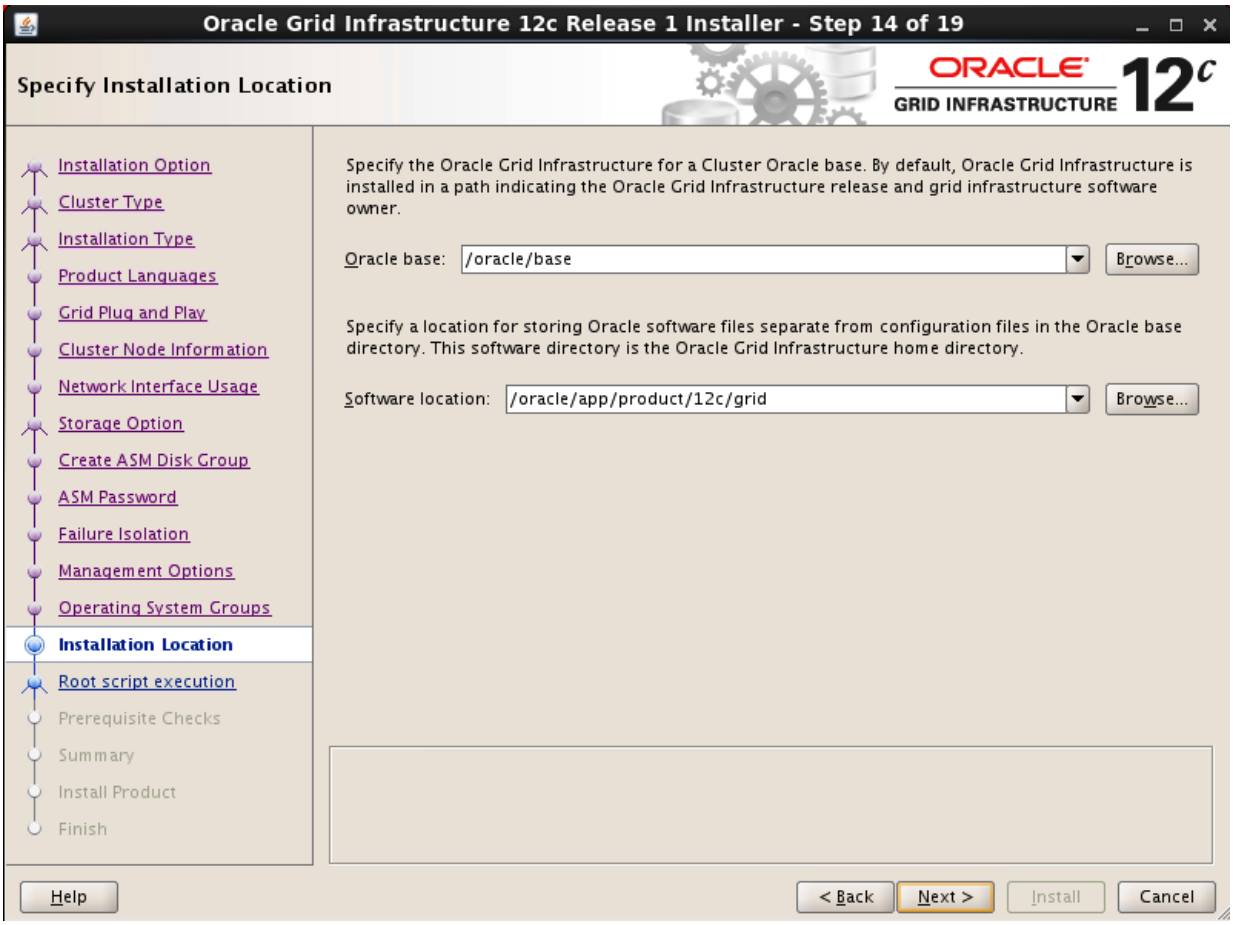

23. Provide the Inventory Directory.
24. Click Next.

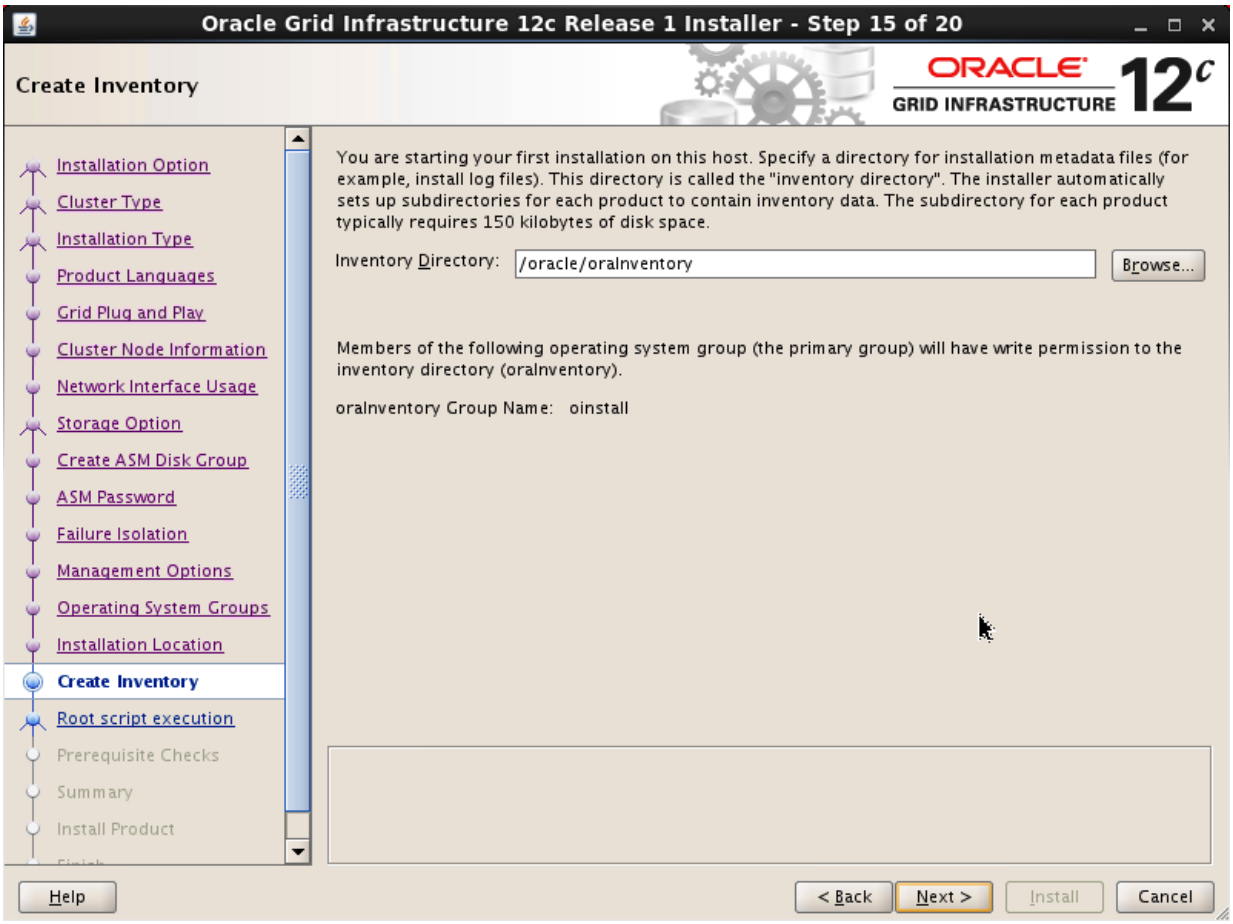

25. Click Next.

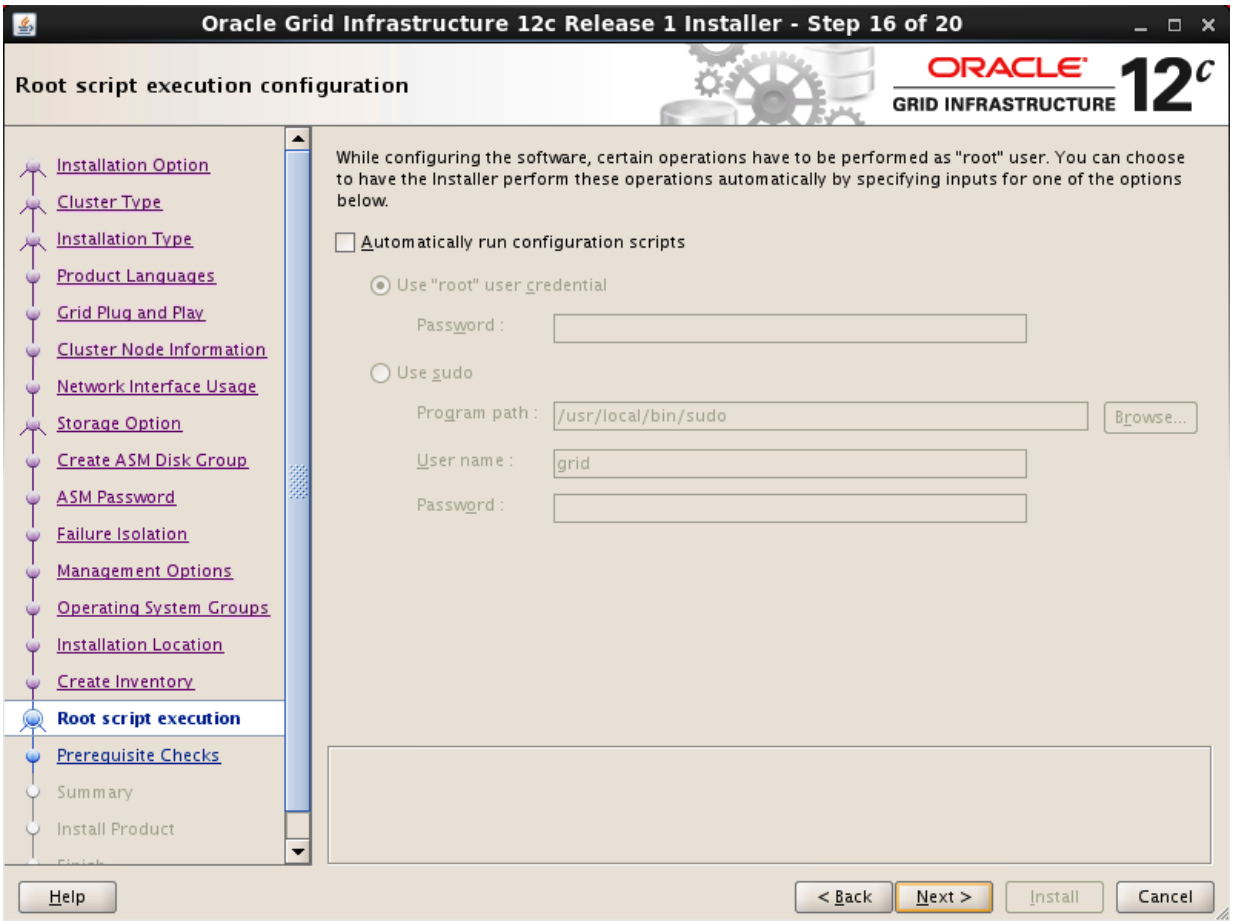

26. The next page performs the prerequisite verification. If the runcluvfy.sh script was run successfully in the section “Installing Requisite Packages,” the prerequisite verification should not generate any errors.
27. Click Next.

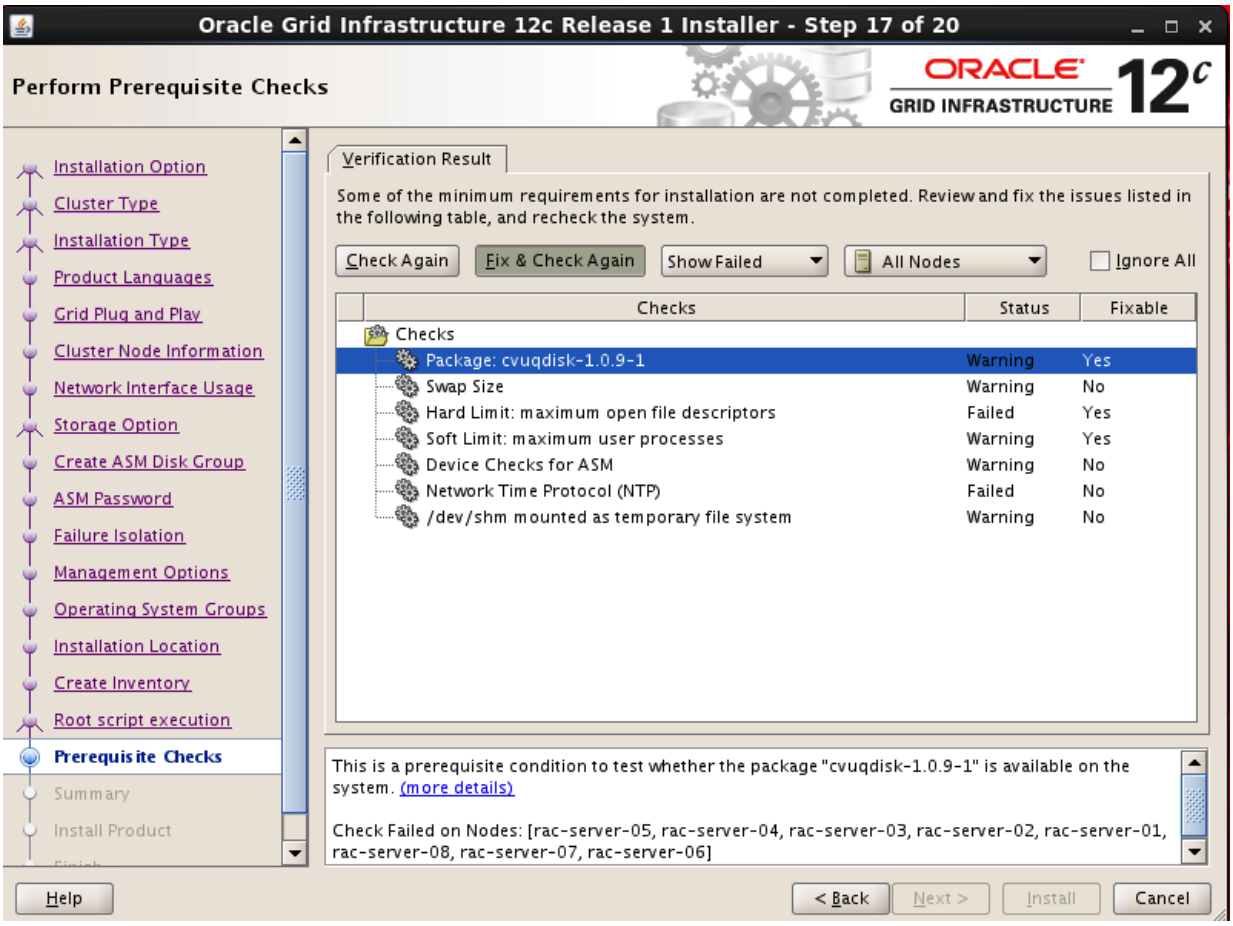

28. The installation prompts you to run the following two scripts as a root user. Run the scripts on all of the RAC nodes, one at a time. Make sure the script is successful on all of the nodes.

```
/u01/app/oraInventory/orainstRoot.sh
/u01/app/11.2.0/grid/root.sh.
```

29. Click Install.

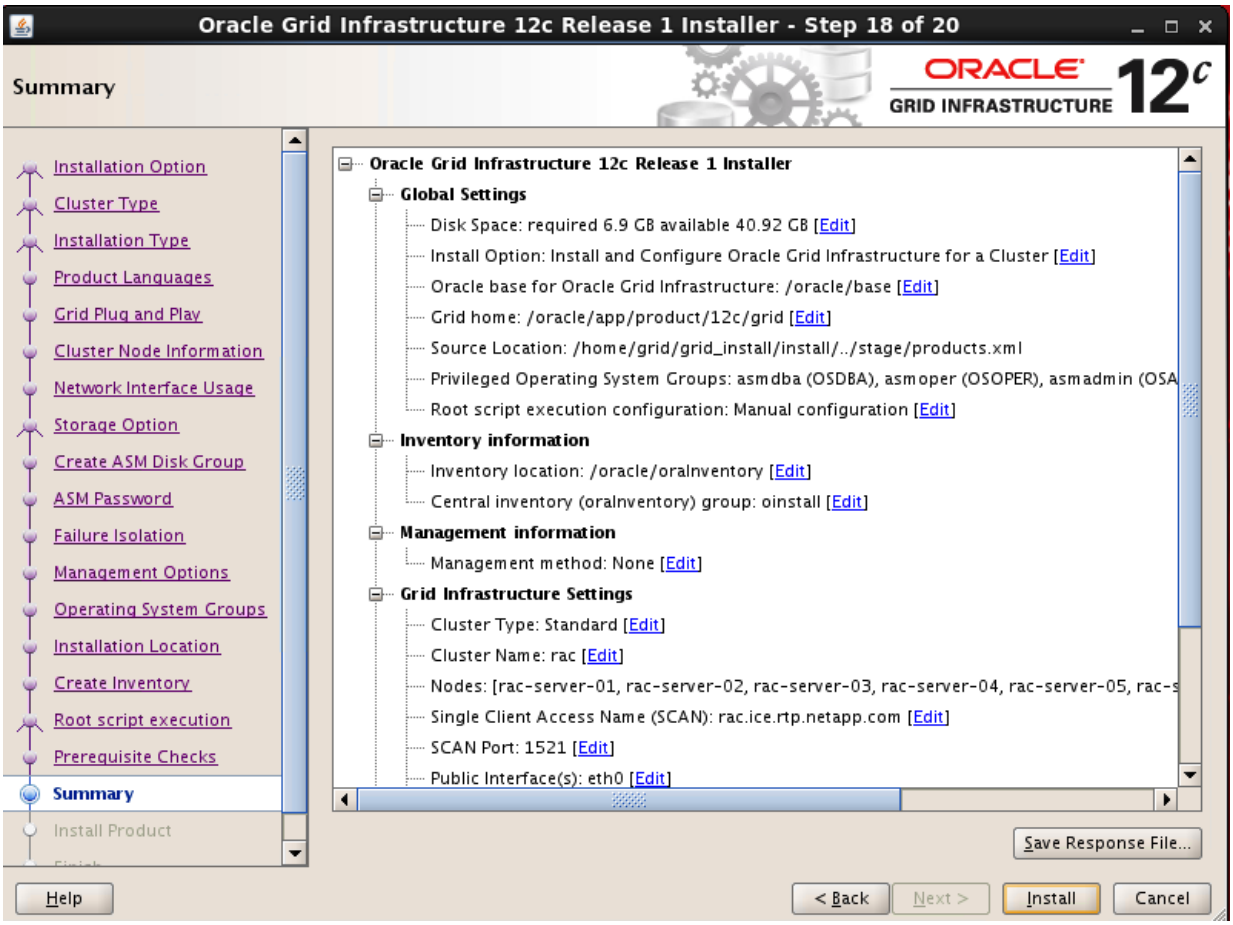

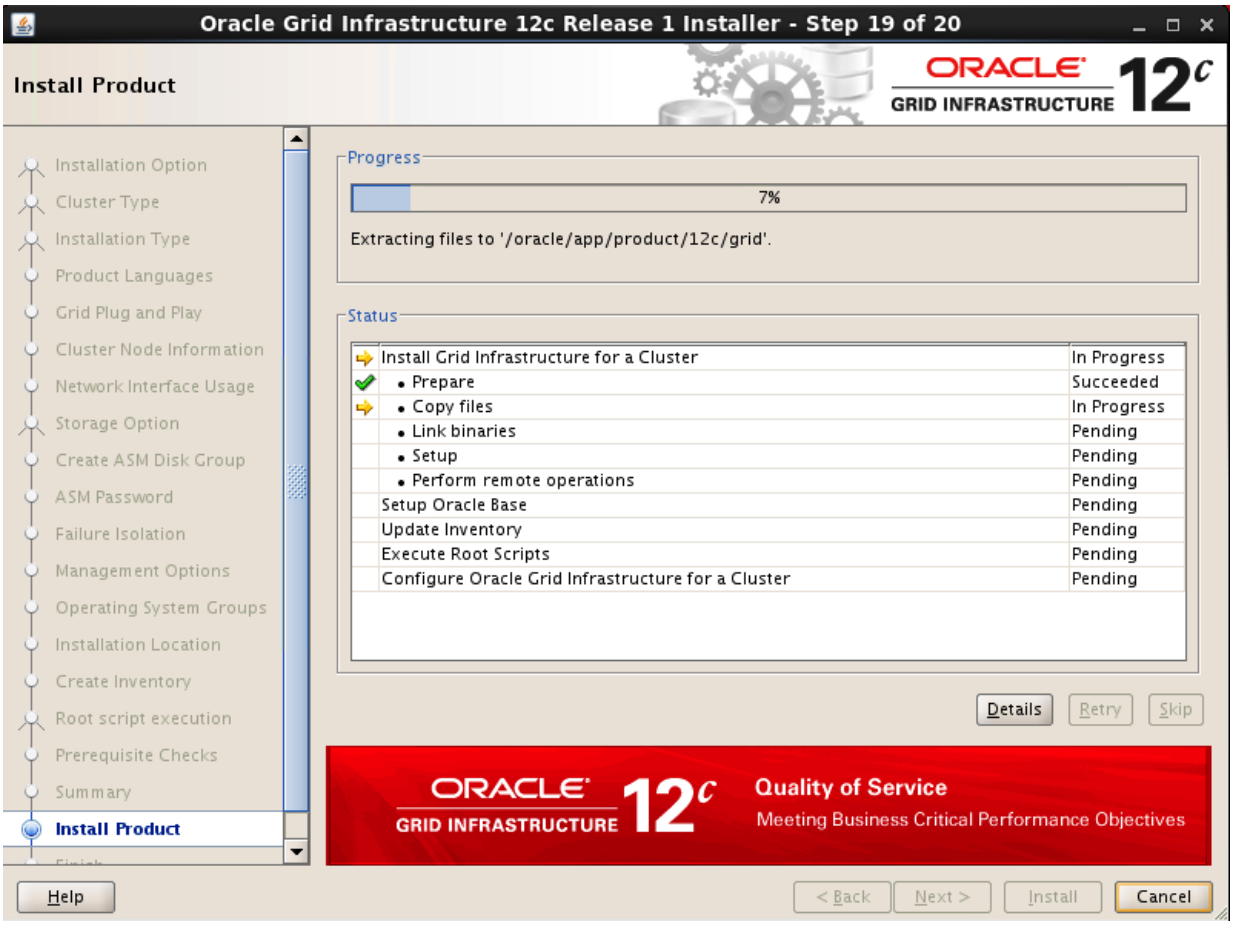

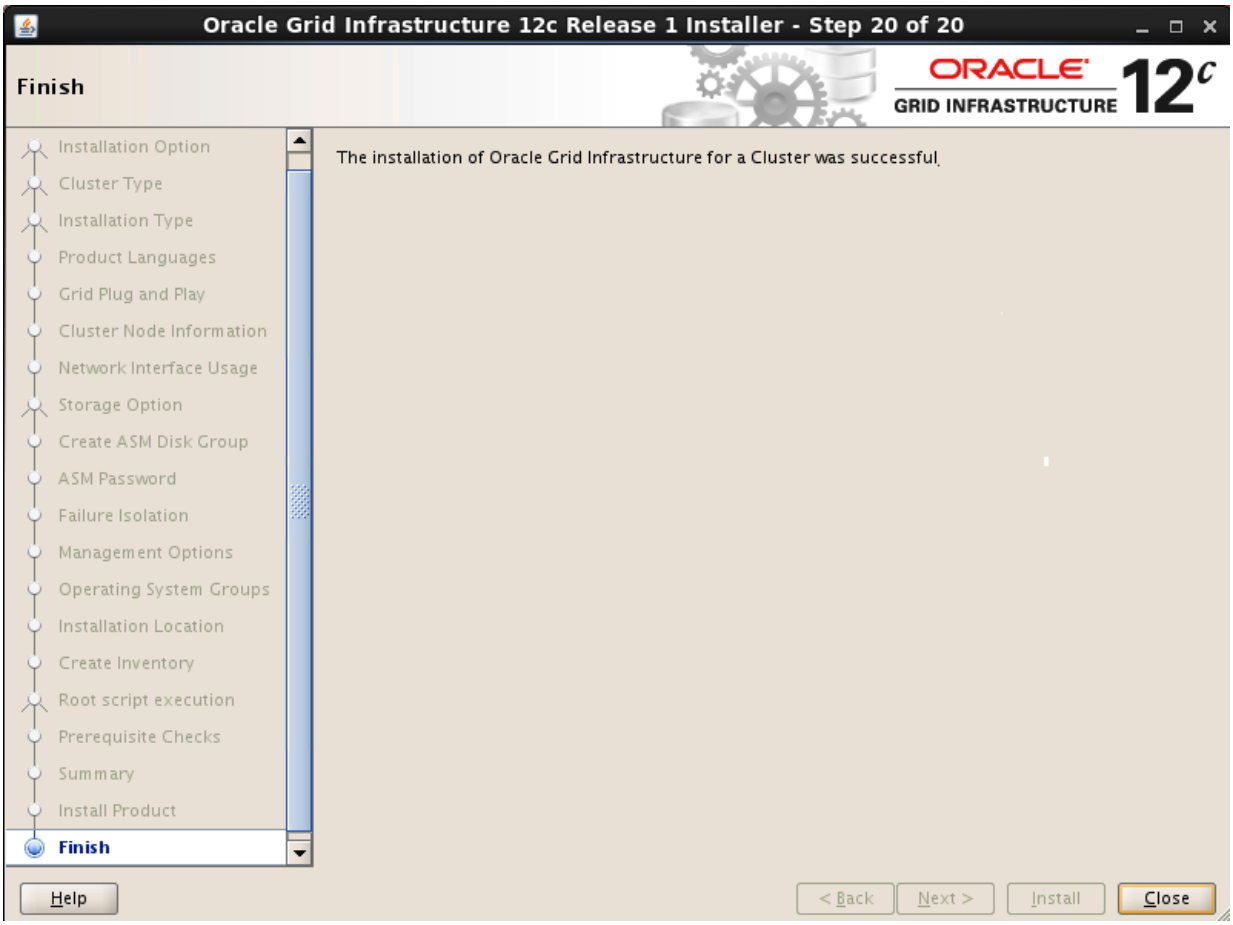

Oracle Database Software Installation

Launch the installer GUI to install the database software.

1. On the Configure Security Updates page, click Next.
2. On the next page, check the option to get the latest updates from Oracle, or skip the software updates, which is what was chosen for this installation.
3. Click Next.

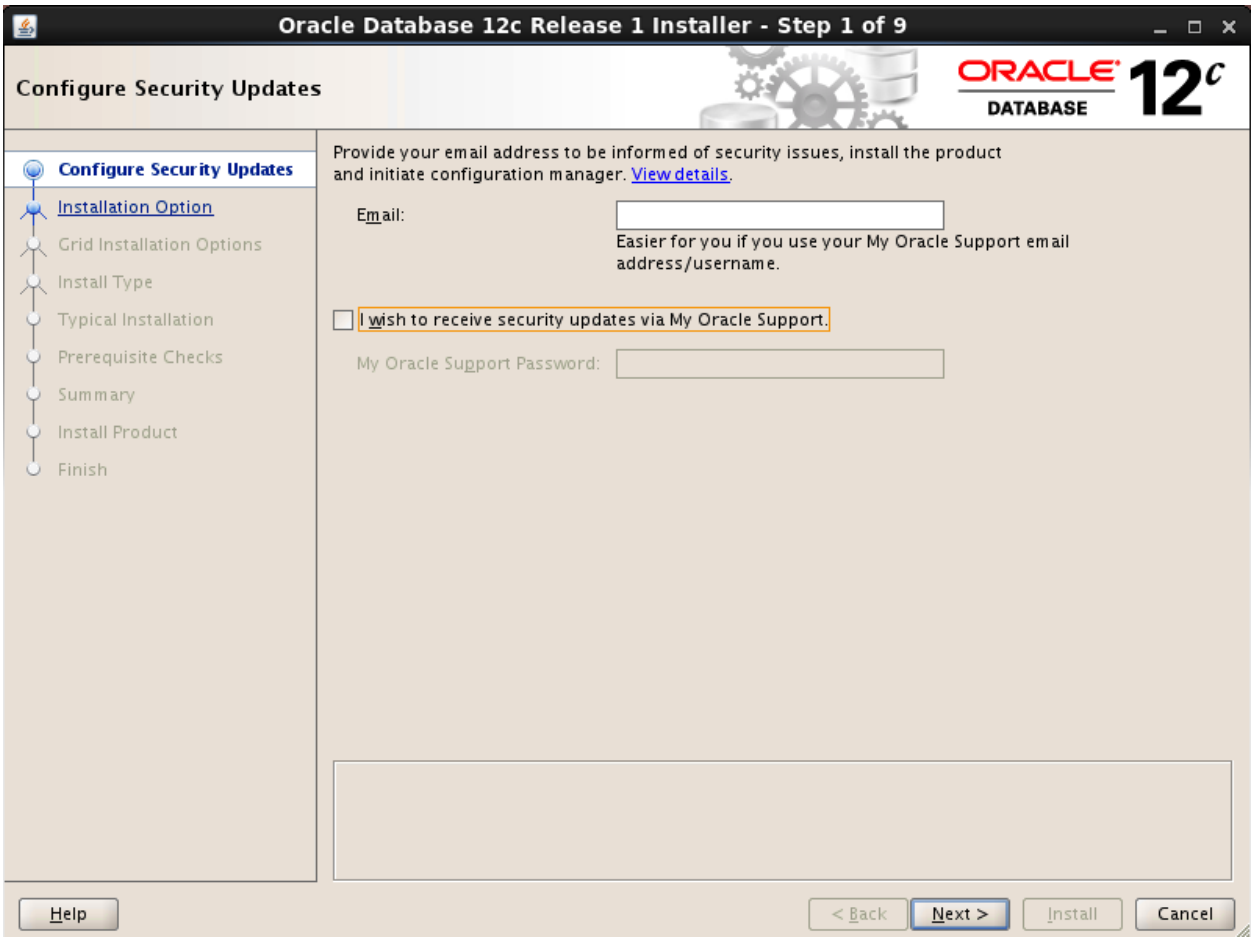

4. Select Install Database Software only.
5. Click Next.

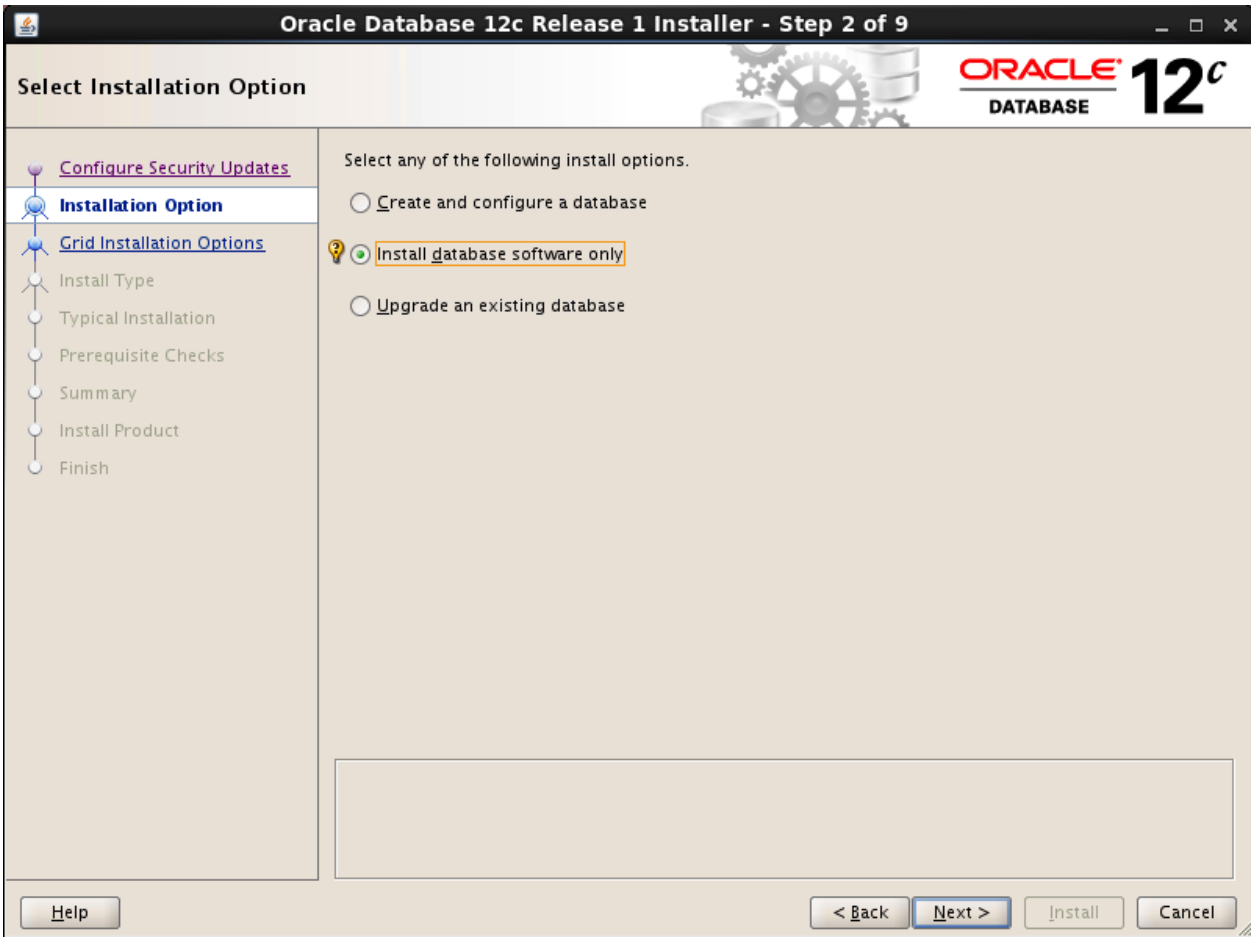

6. Select Oracle Real Application Clusters Database Installation.
7. Click Next.

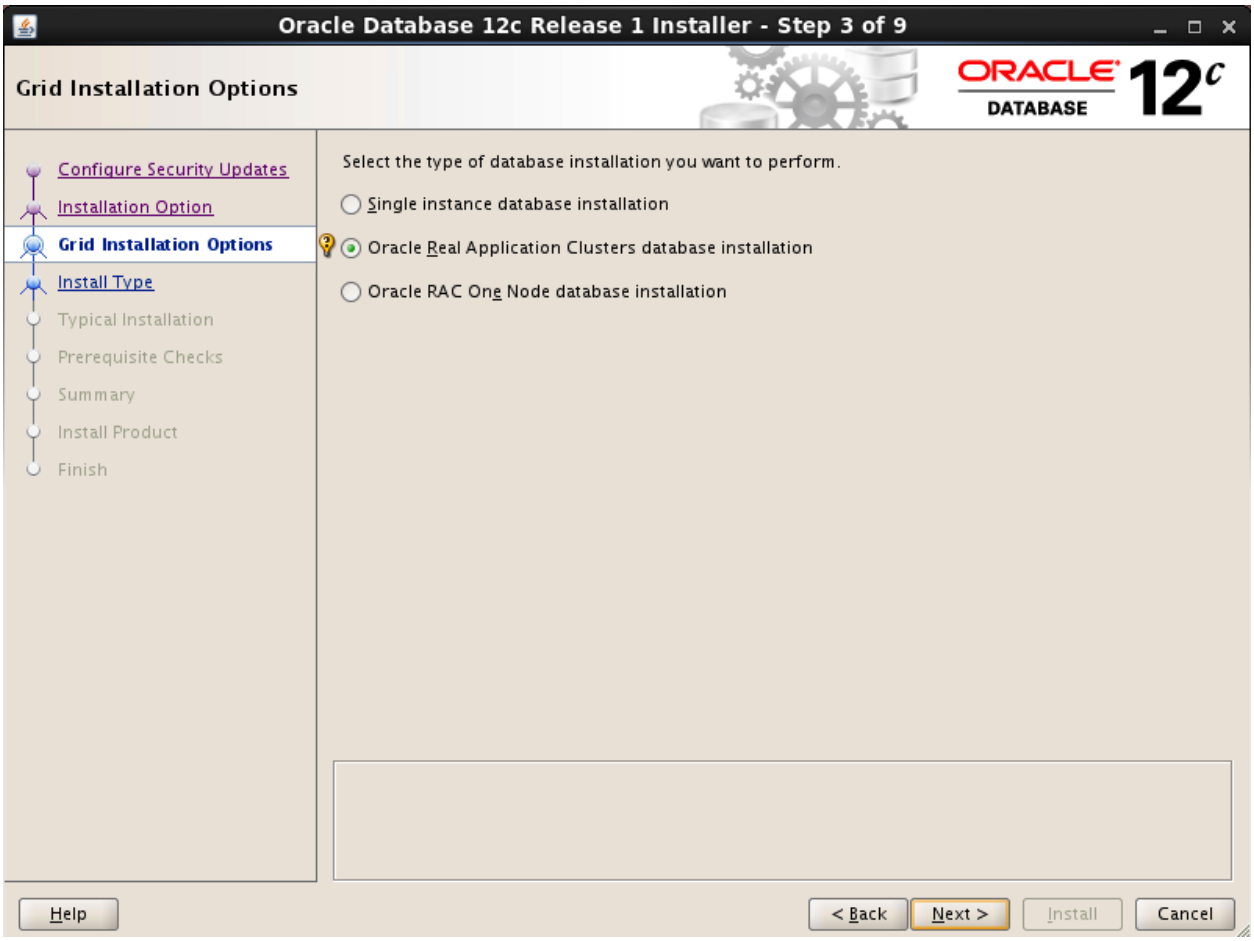

8. Select all of the RAC nodes.
9. Click SSH Connectivity.

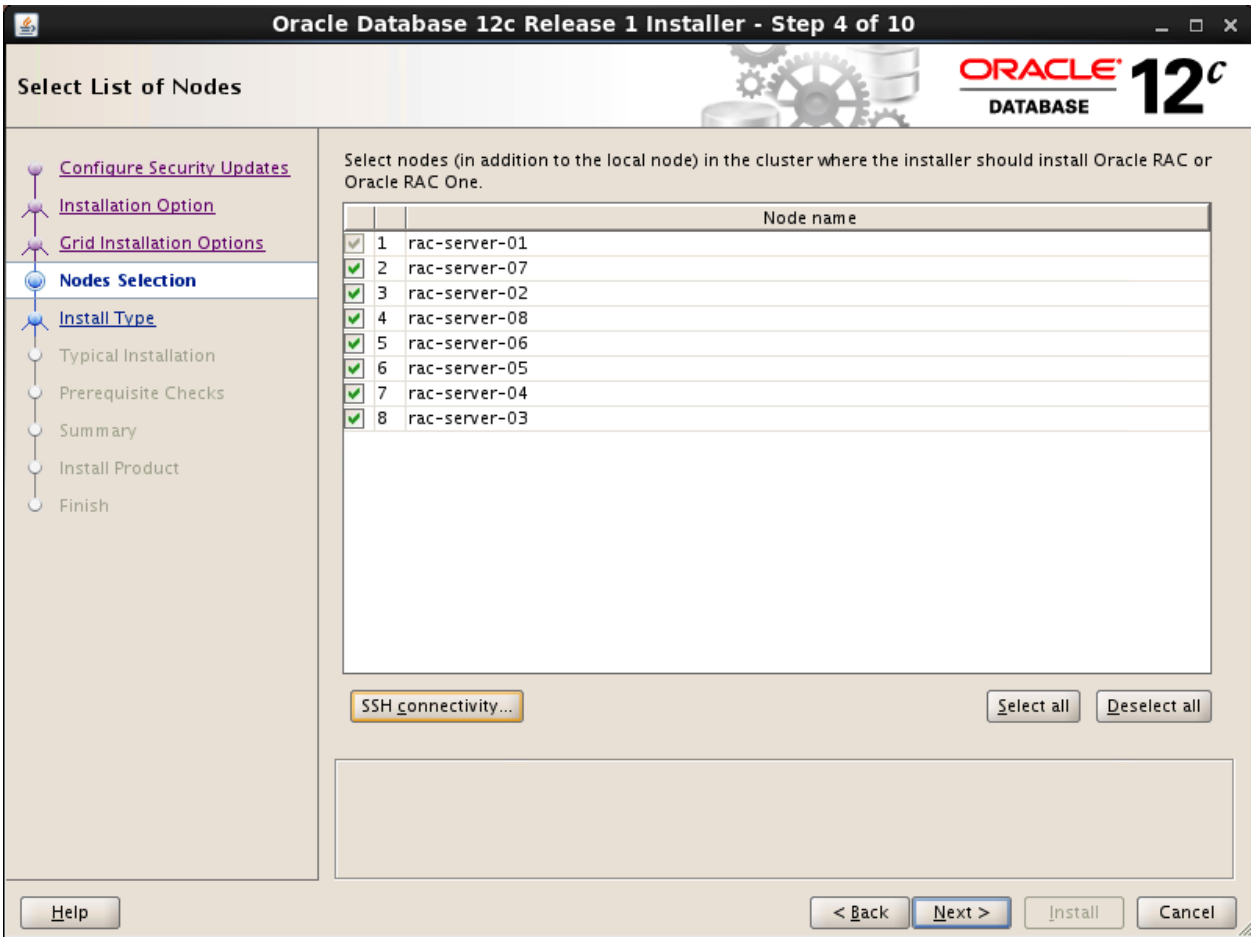

10. Enter the OS password.
11. Click Next.

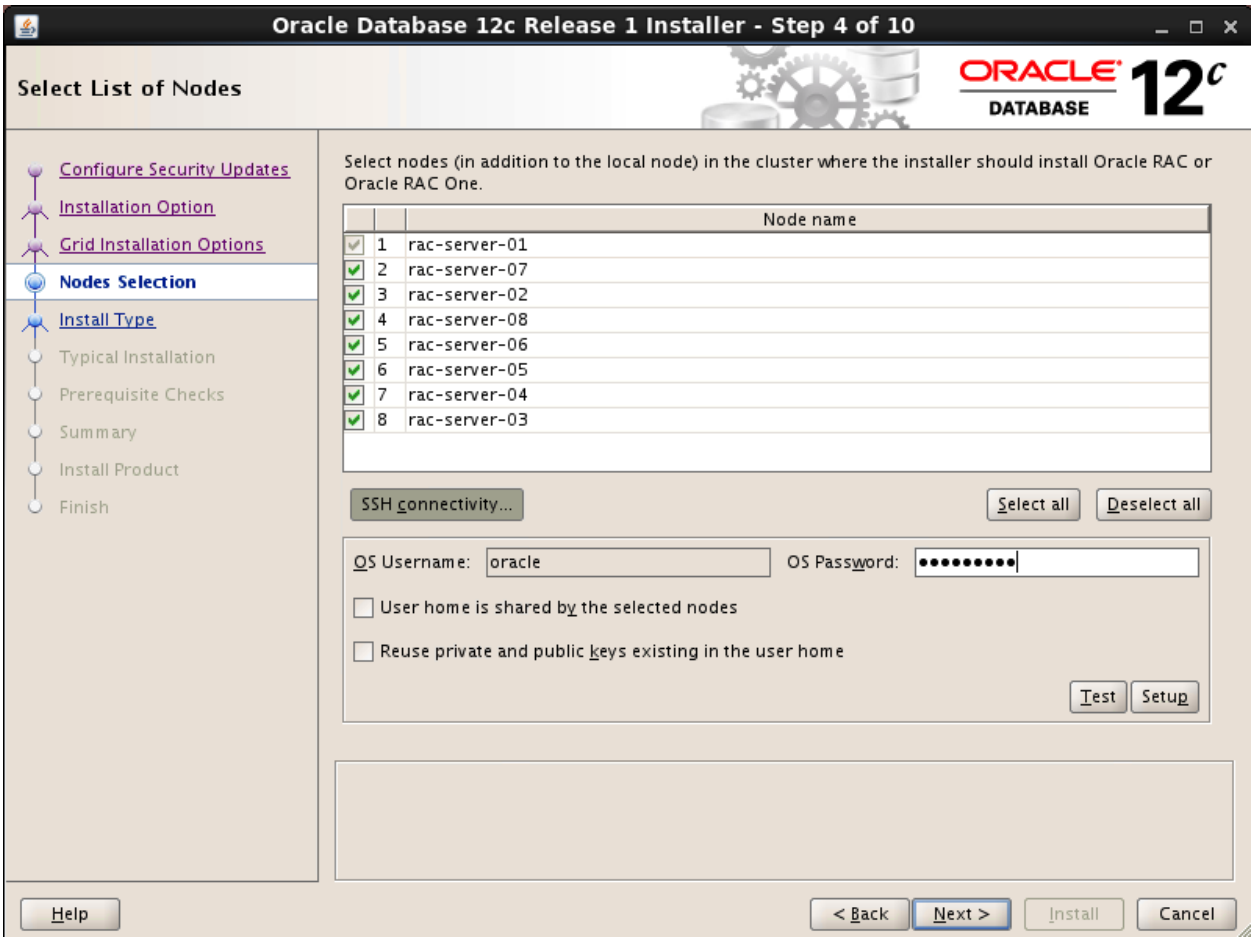

12. Select the preferred language.Click Next.

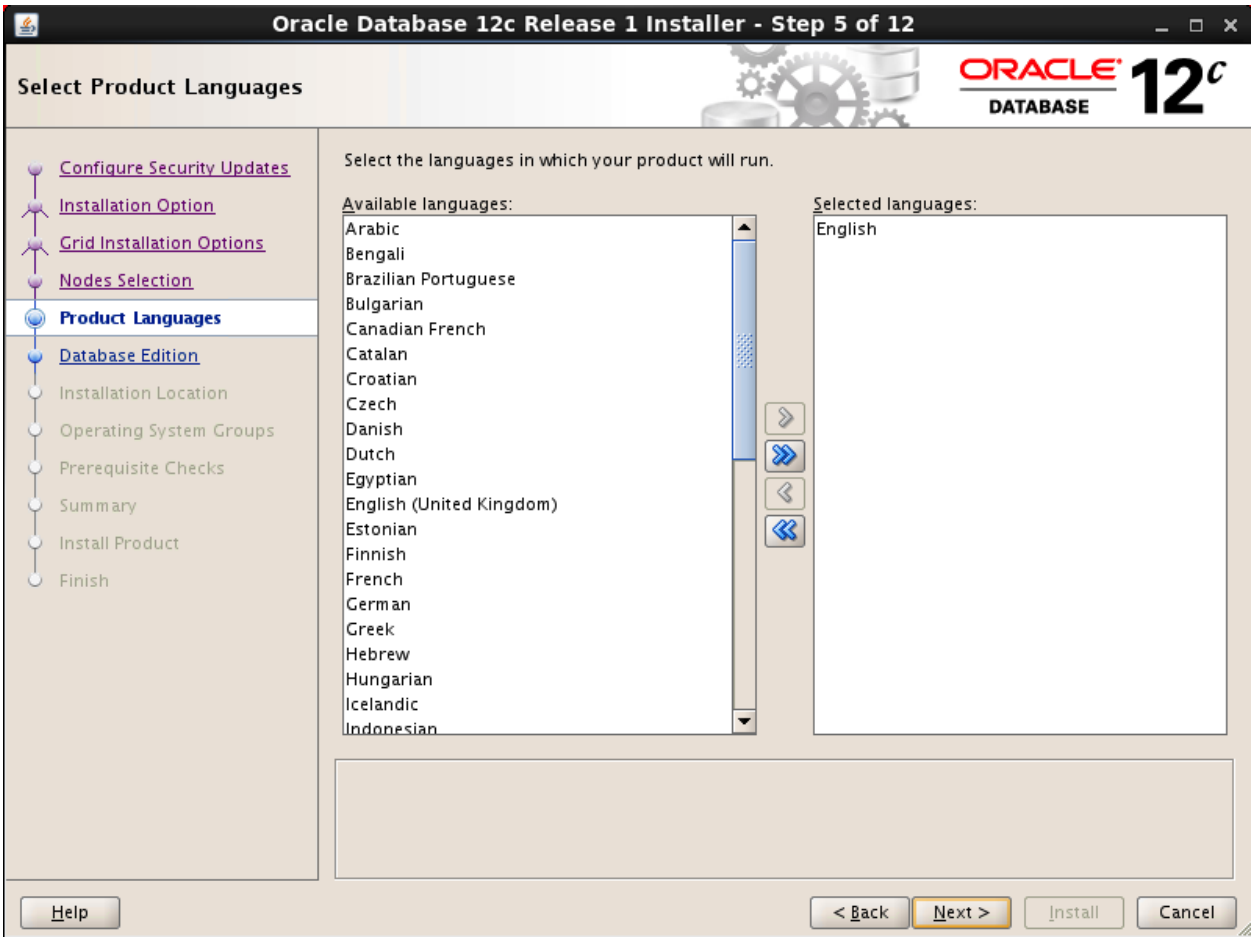

13. Select Enterprise Edition.
14. Click Next.

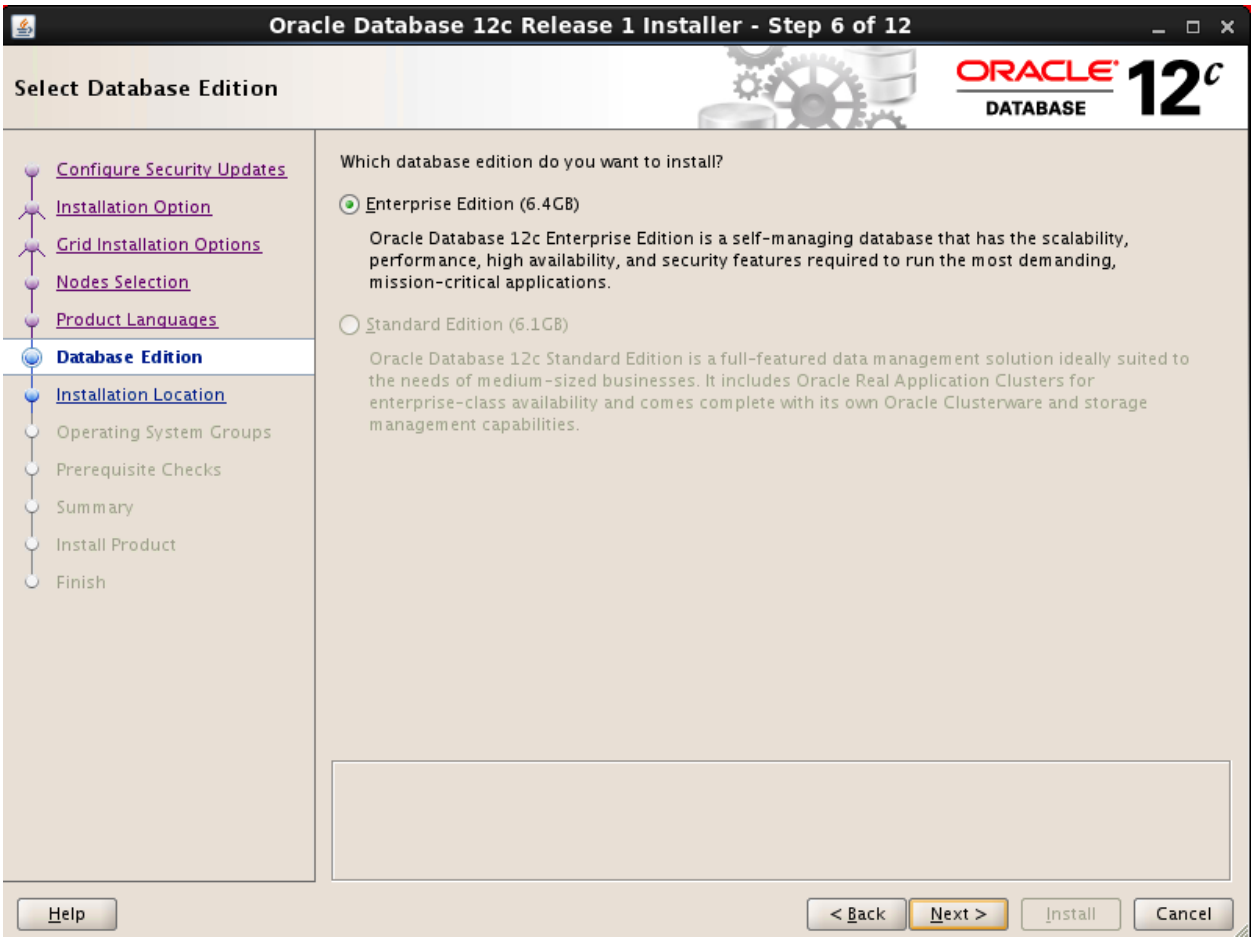

15. Provide the Oracle base and Oracle database software installation directory locations. For this installation, we used `/oracle/app` for Oracle base and `/oracle/app/product/12c/db` as the Oracle database software installation directory.
16. Click Next.

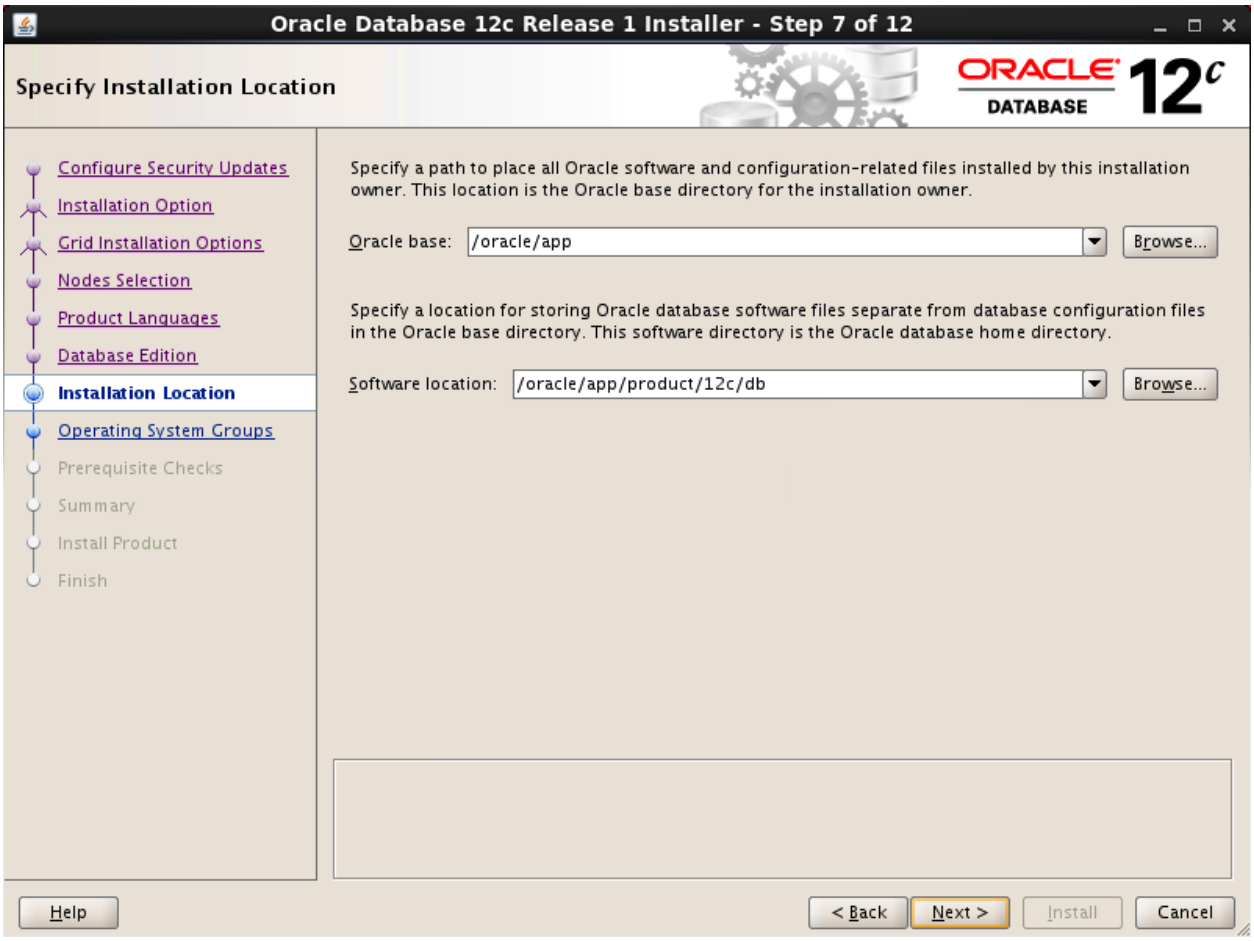

17. Provide the OSDBA and OSOPER group name in the next screen. For this setup, the oinstall group was selected.
18. Click Next.

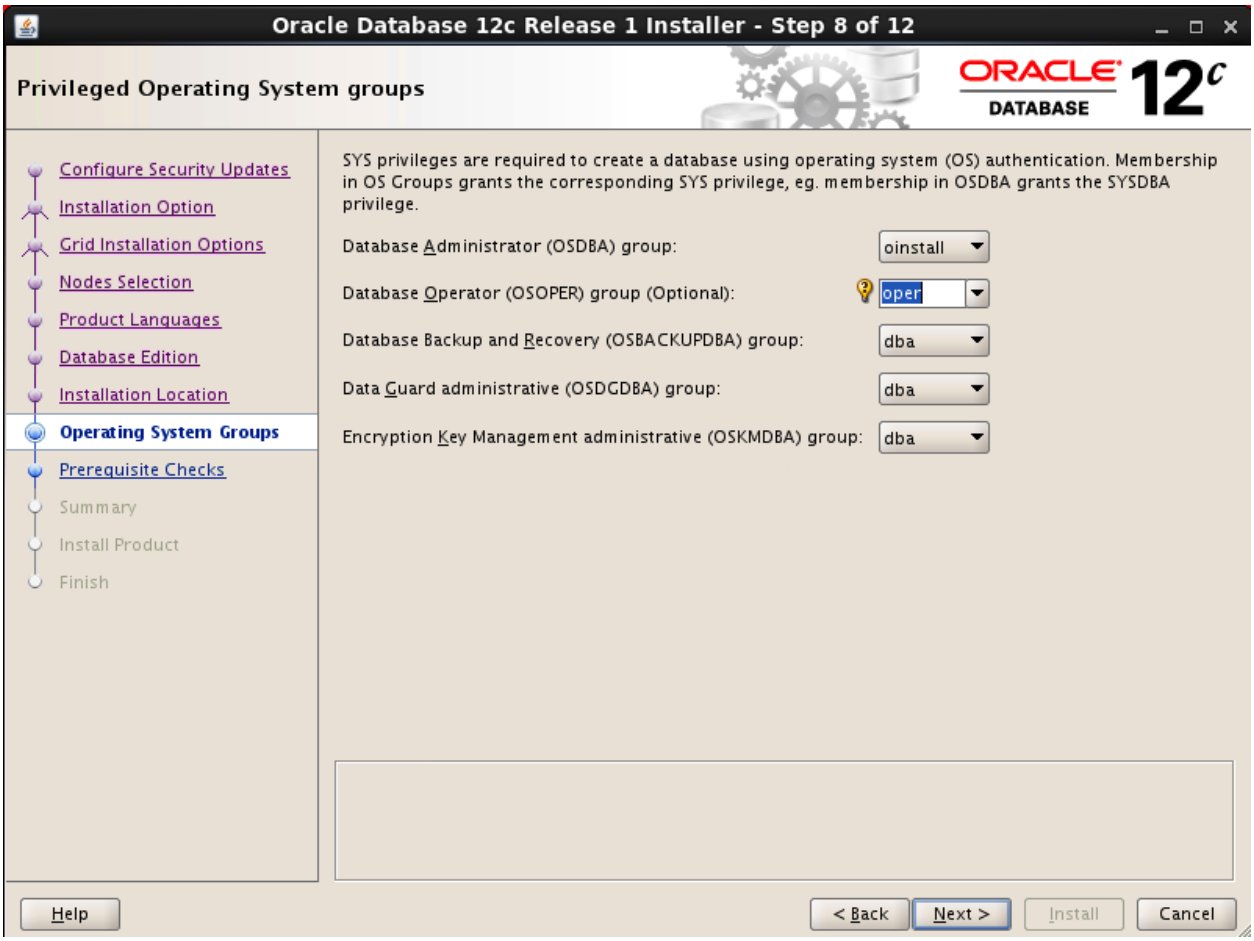

19. On the Perform Prerequisite Checks page, make sure the installer doesn't give any errors. If errors are detected, resolve them before continuing.
20. Click Next.

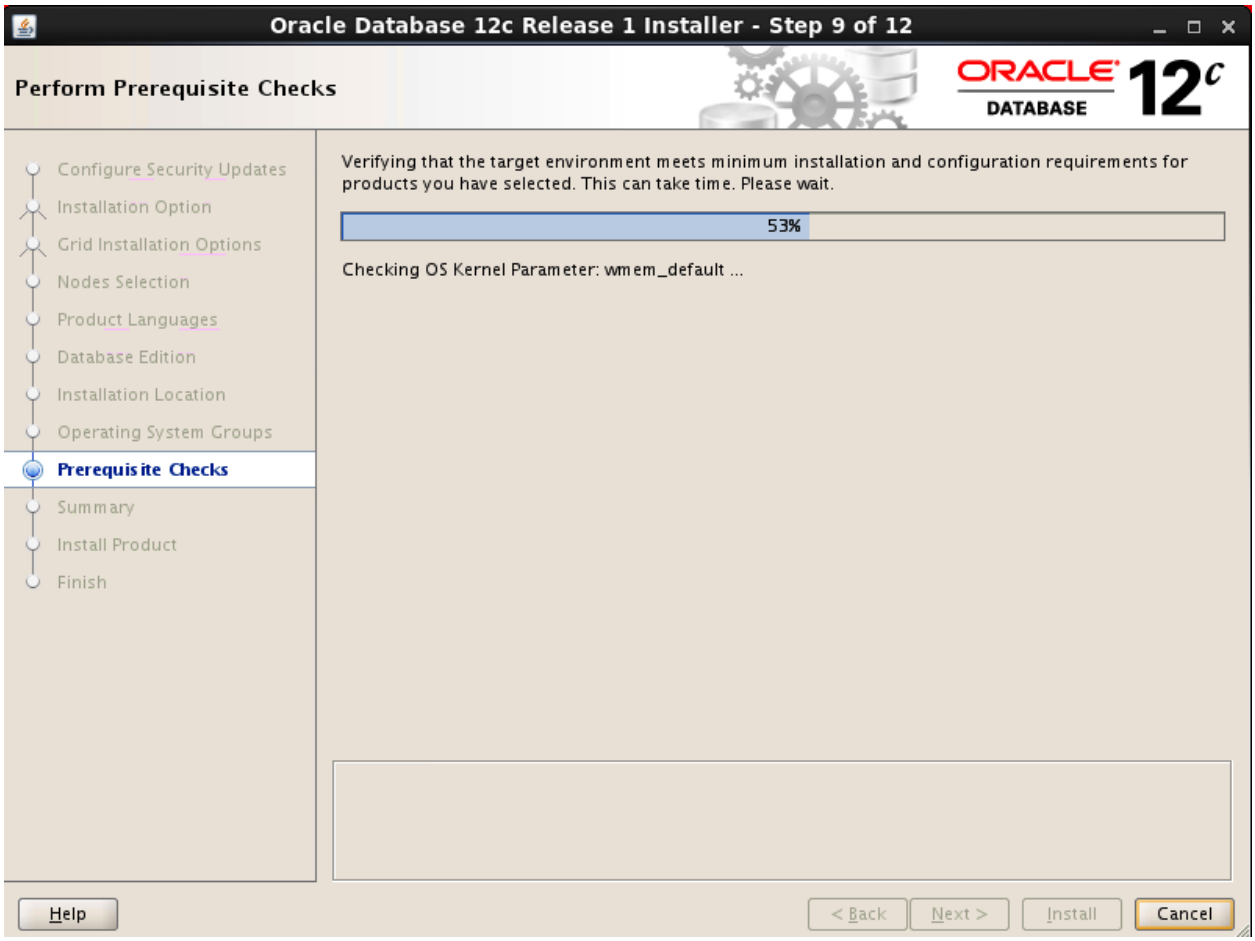

21. Click Finish on the summary page to finish the installation.

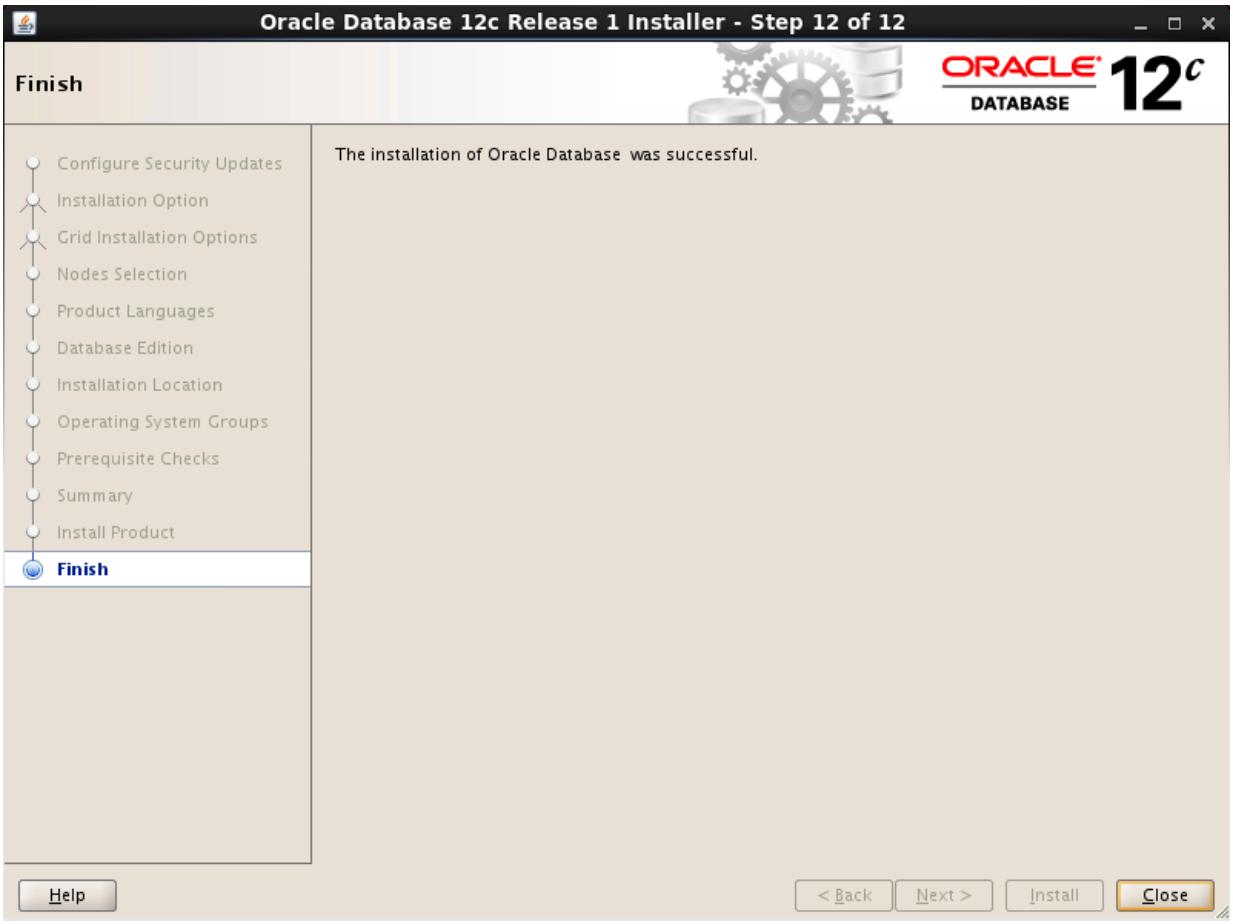

Oracle ASM Configuration

After provisioning the storage and successfully installing the Oracle RAC environment, create and configure the ASM disks on the RAC nodes. Launch the ASM GUI from any node as a grid user by running the command `amsca`, which opens the GUI.

1. Provide the ASM disk group name for DATA; in this case, it is ORADATA.
2. Select the appropriate drives. For this solution, 32 drives were selected from the list as per Table 24.
3. For Redundancy, select External.
4. Click Show Advanced Options.
5. Under Disk Group Attributes, set Allocation Unit (AU) Size to 64MB and click OK.
6. Click OK on the confirmation message.

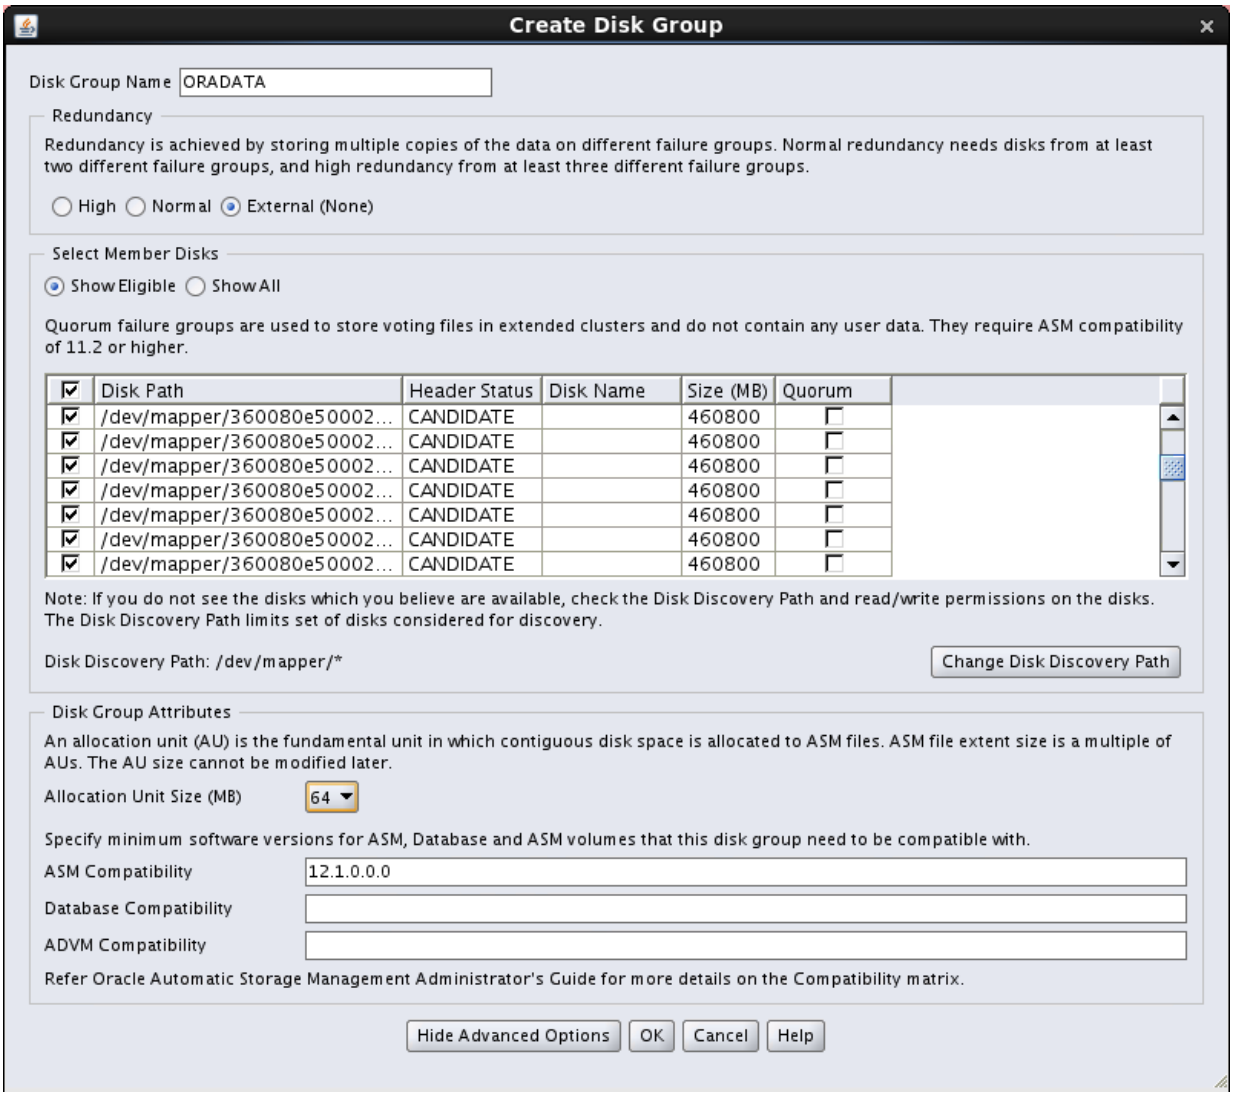

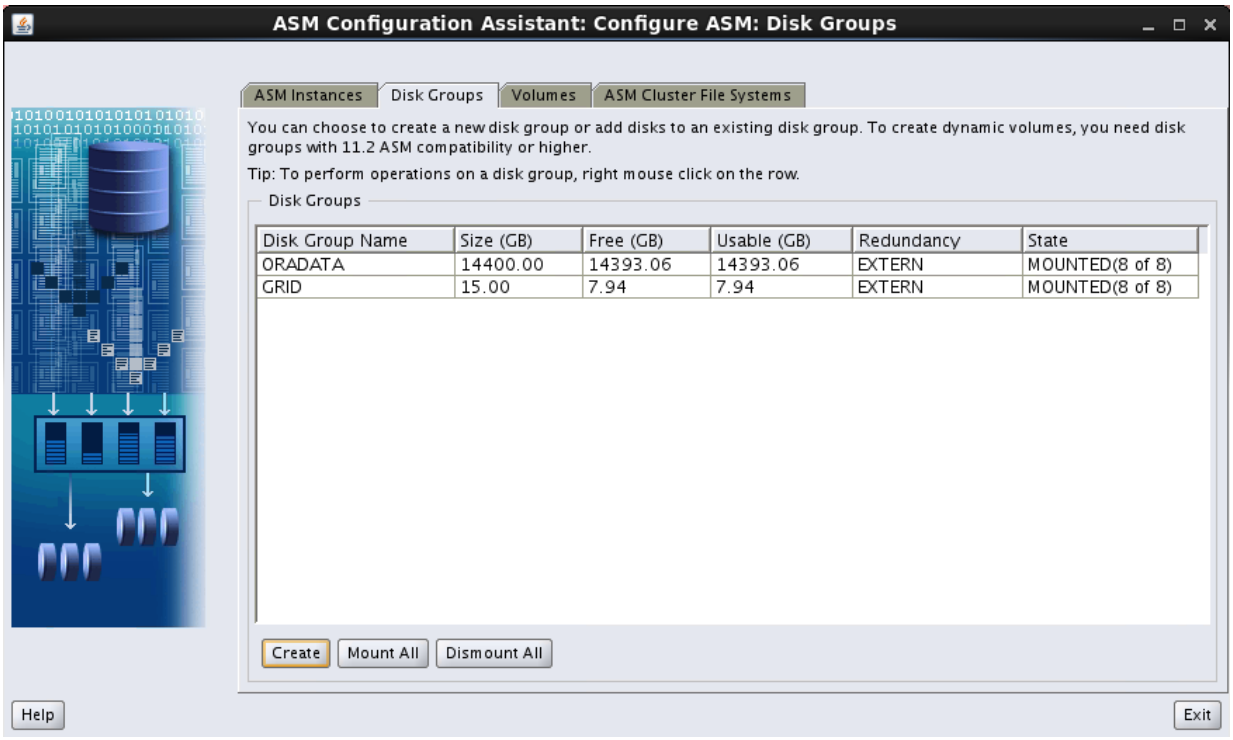

7. Repeat step 1 through step 6 for the ORALOG disk group. For the ORALOG disk group, eight ASM disks were used. The ASM AU size was set at 64M for ORALOG as well.

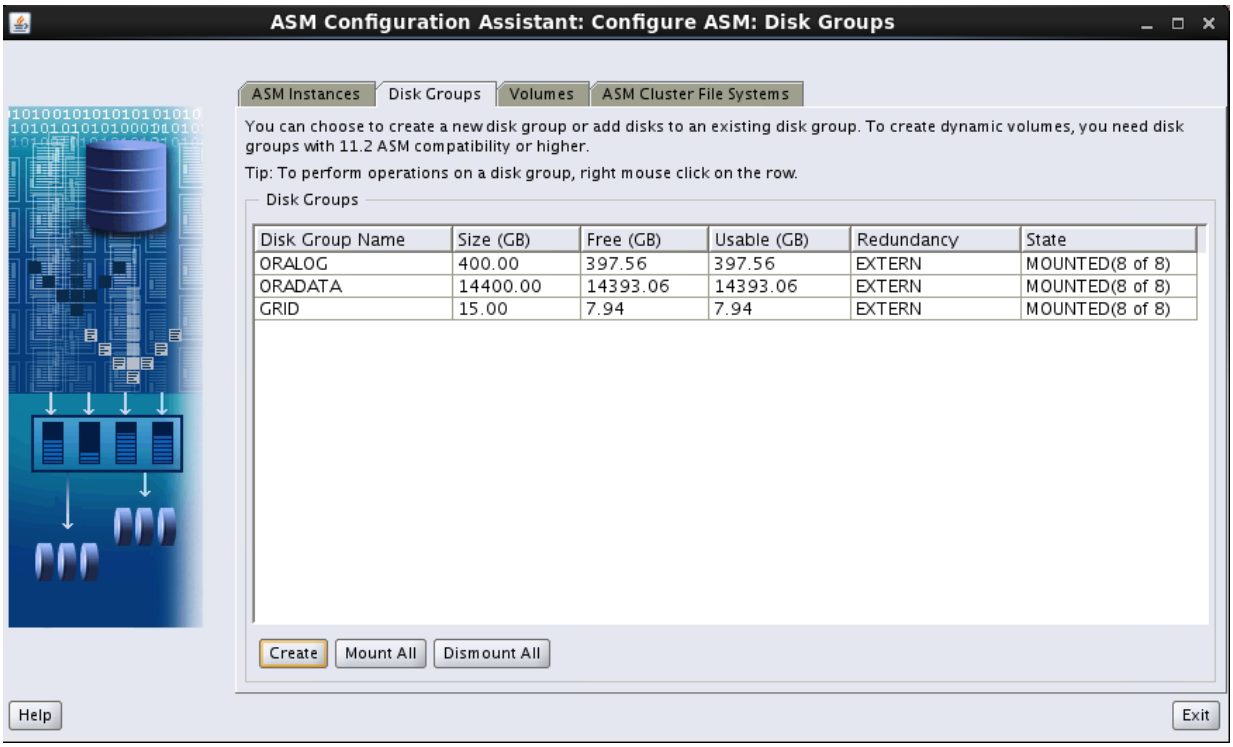

Oracle Database Creation

After the ASM disk groups are available, the Oracle RAC database can be created using them. In order to create the Oracle RAC database, run the Database Configuration Assistant (DBCA) utility as an Oracle user, which launches the DBCA GUI. After the DBCA GUI is successfully launched, complete the following steps to create the Oracle RAC database:

1. Select Oracle Real Application Clusters (RAC) Database as the database type and click Next.
2. Select Create Database.
3. Select Next.

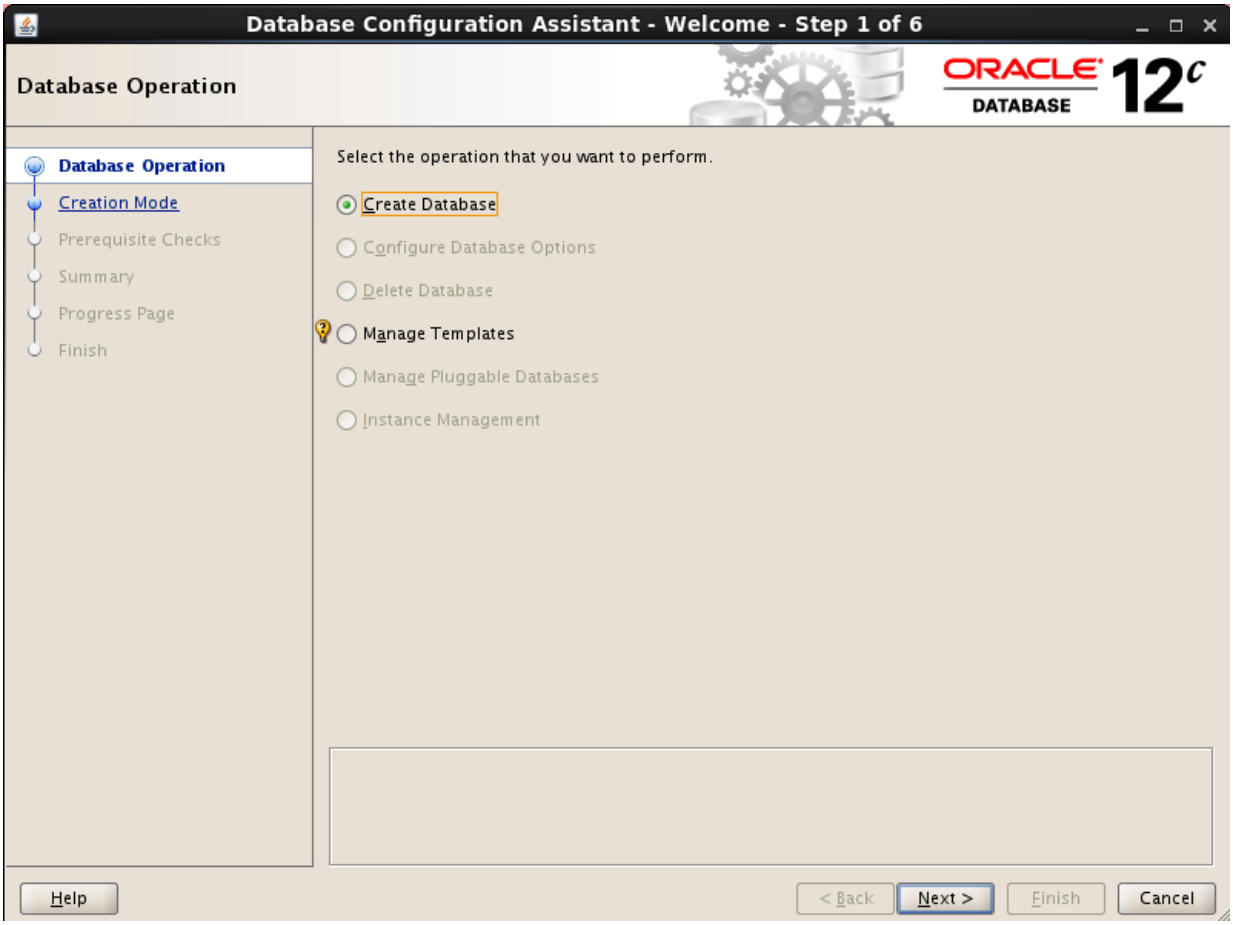

4. Select Advanced Mode and click Next.

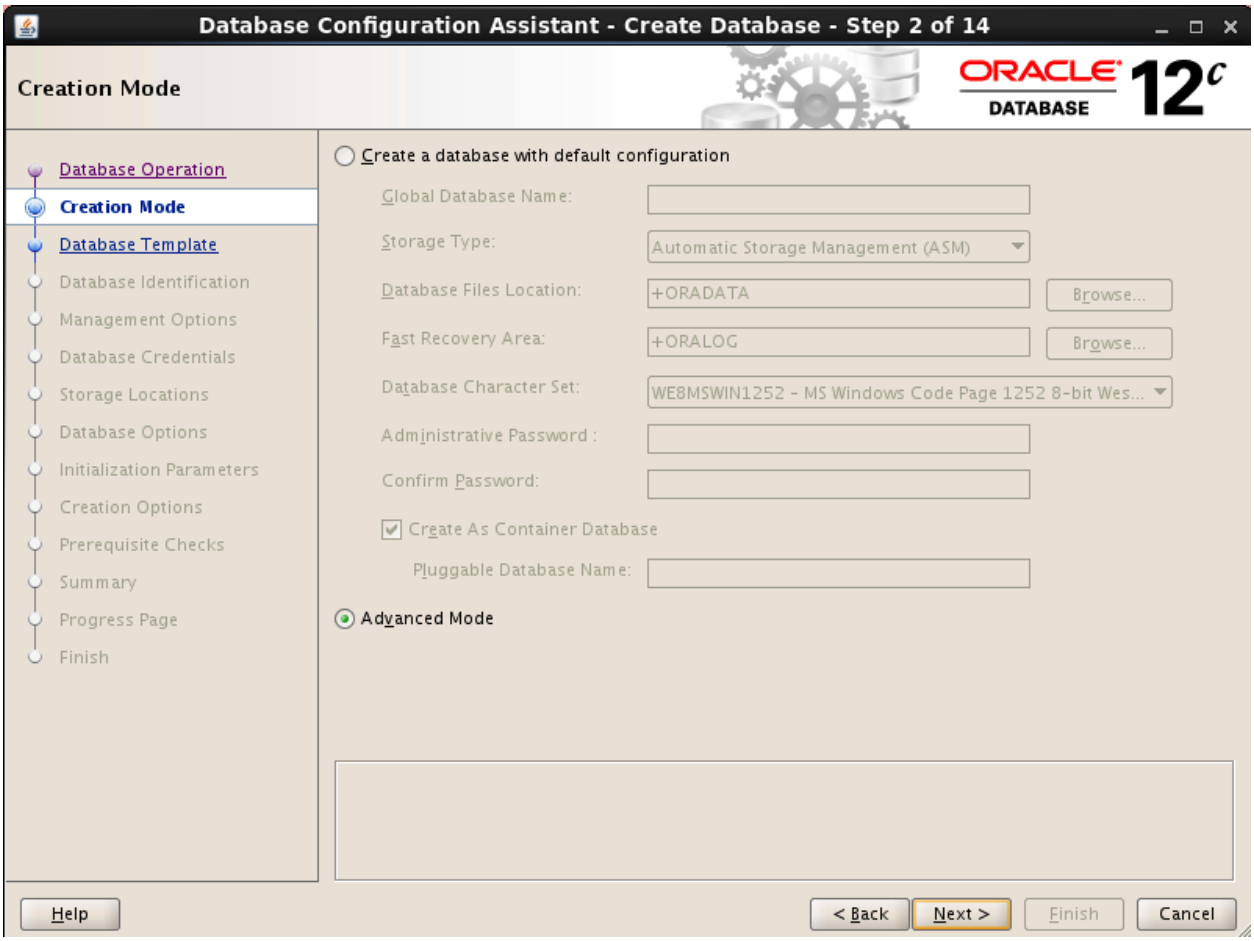

5. Select General Purpose or Transaction Processing
6. Select Admin Managed as the configuration type and click Next.

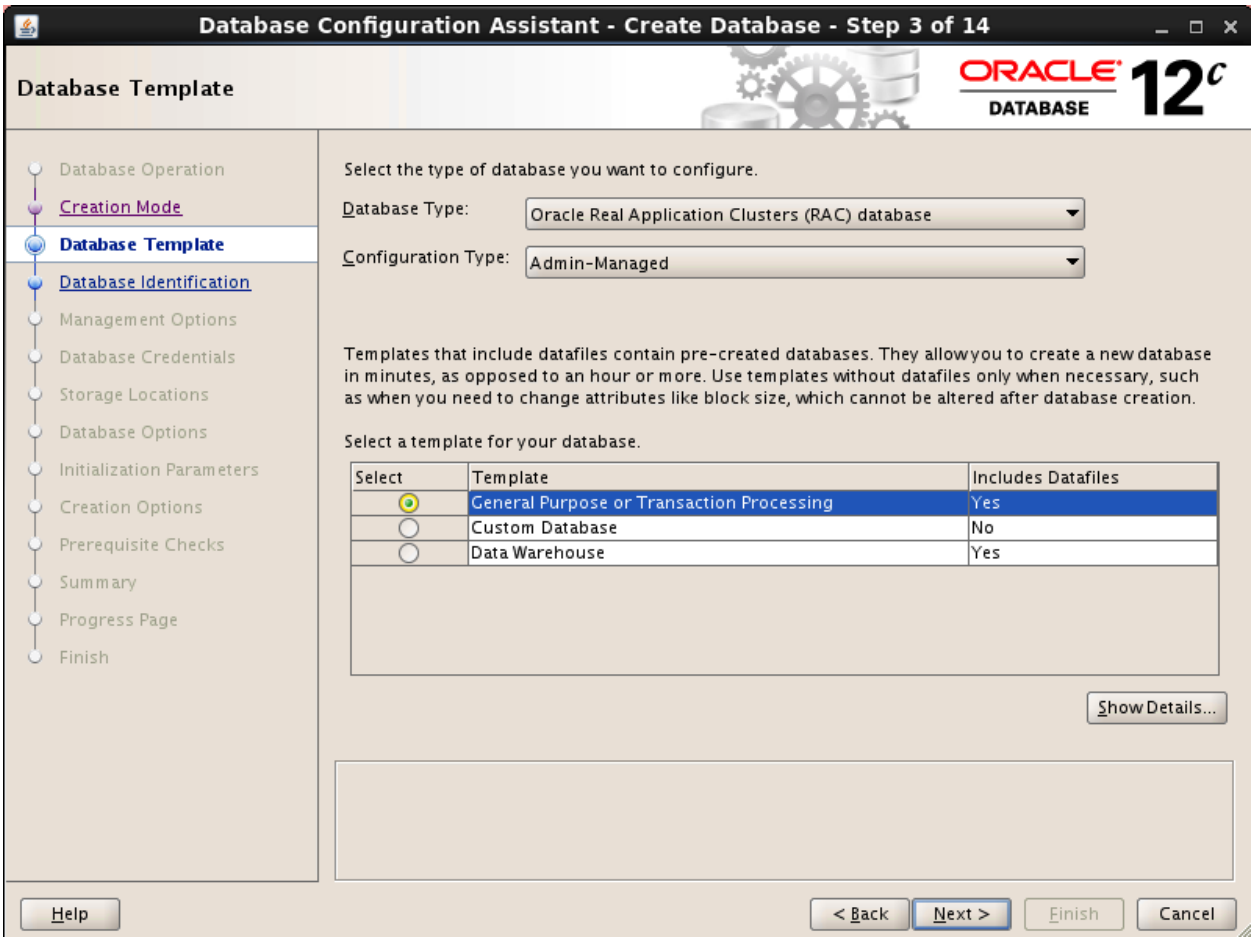

7. Provide a Global Database Name; in this case, LWNDB was used.
8. Click Next.

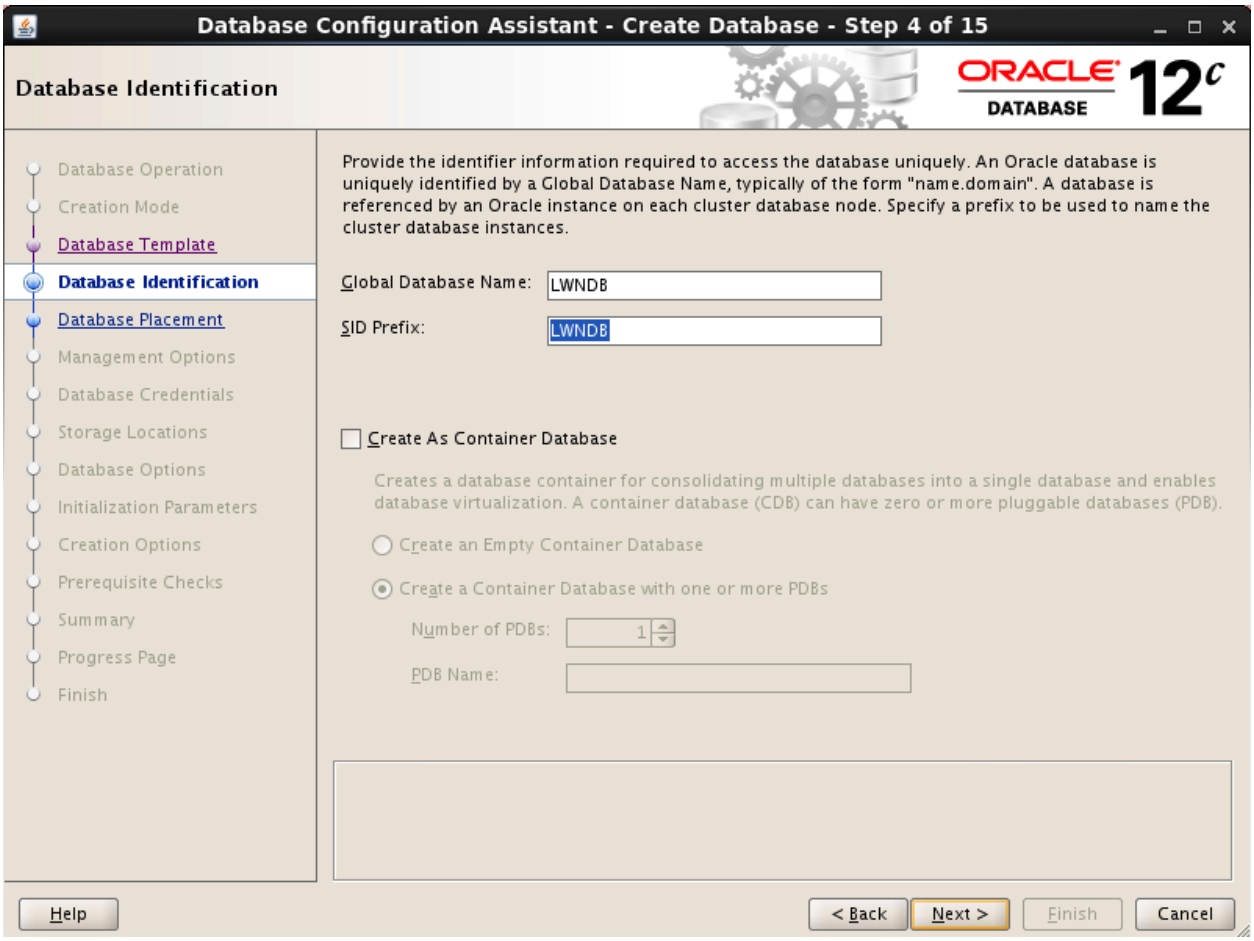

9. Click the “>>” button to add all eight RAC nodes and click Next.

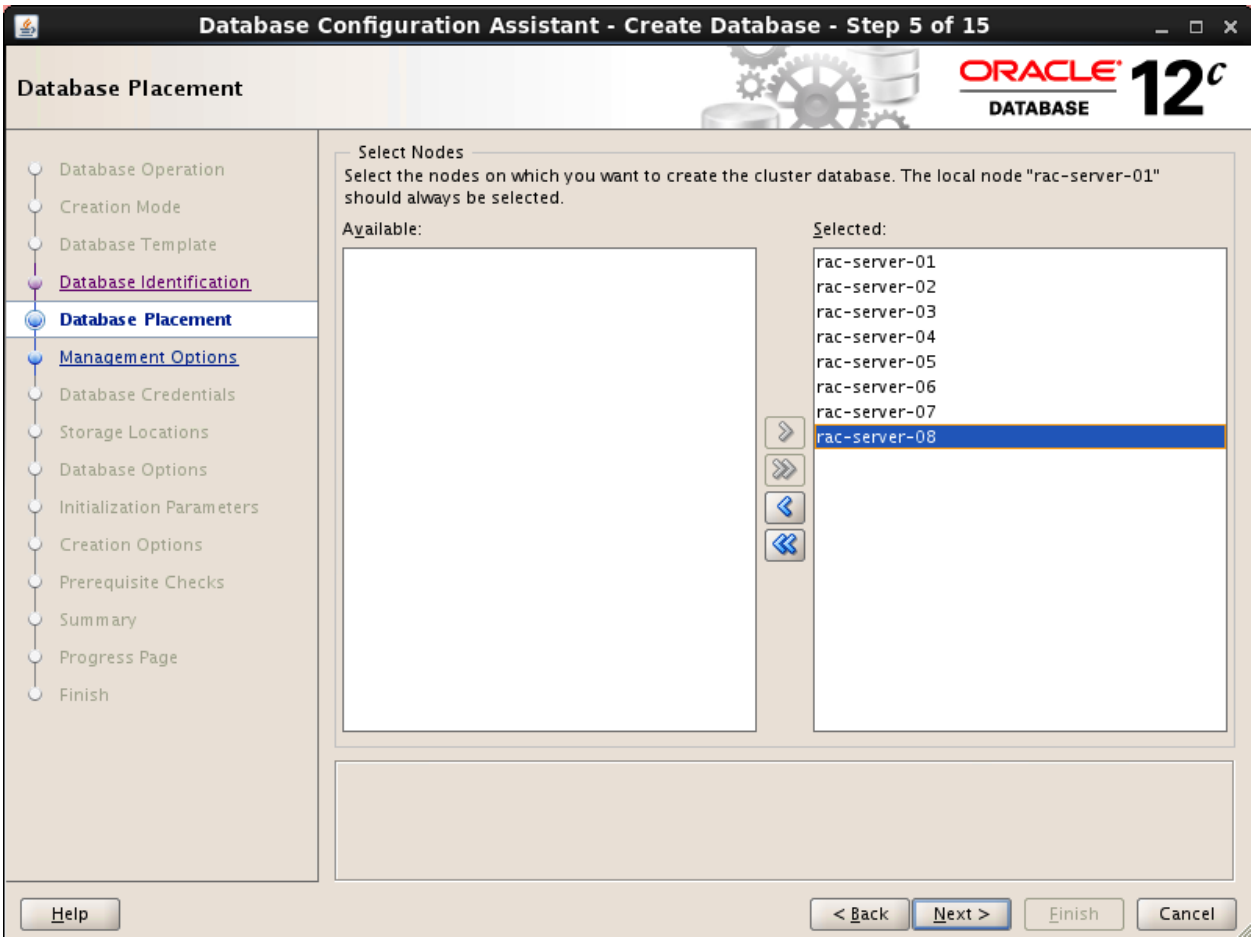

10. Unselect Configure Enterprise Manager. Although this was not selected for this configuration, this could change depending upon the user need.

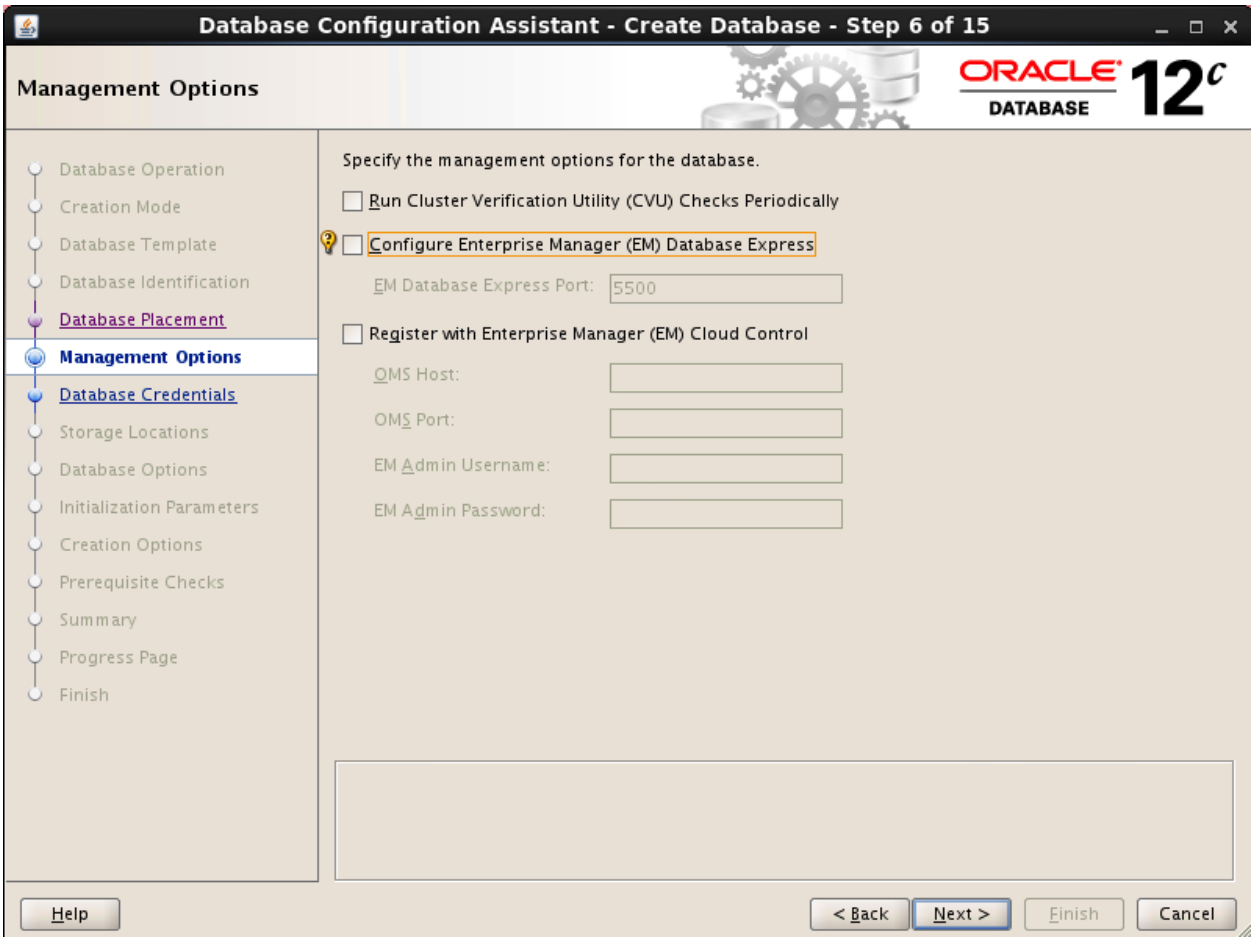

11. Enter and confirm the administrative passwords. Although the same password was used for all accounts for this configuration, this could be different depending on the customer requirement.

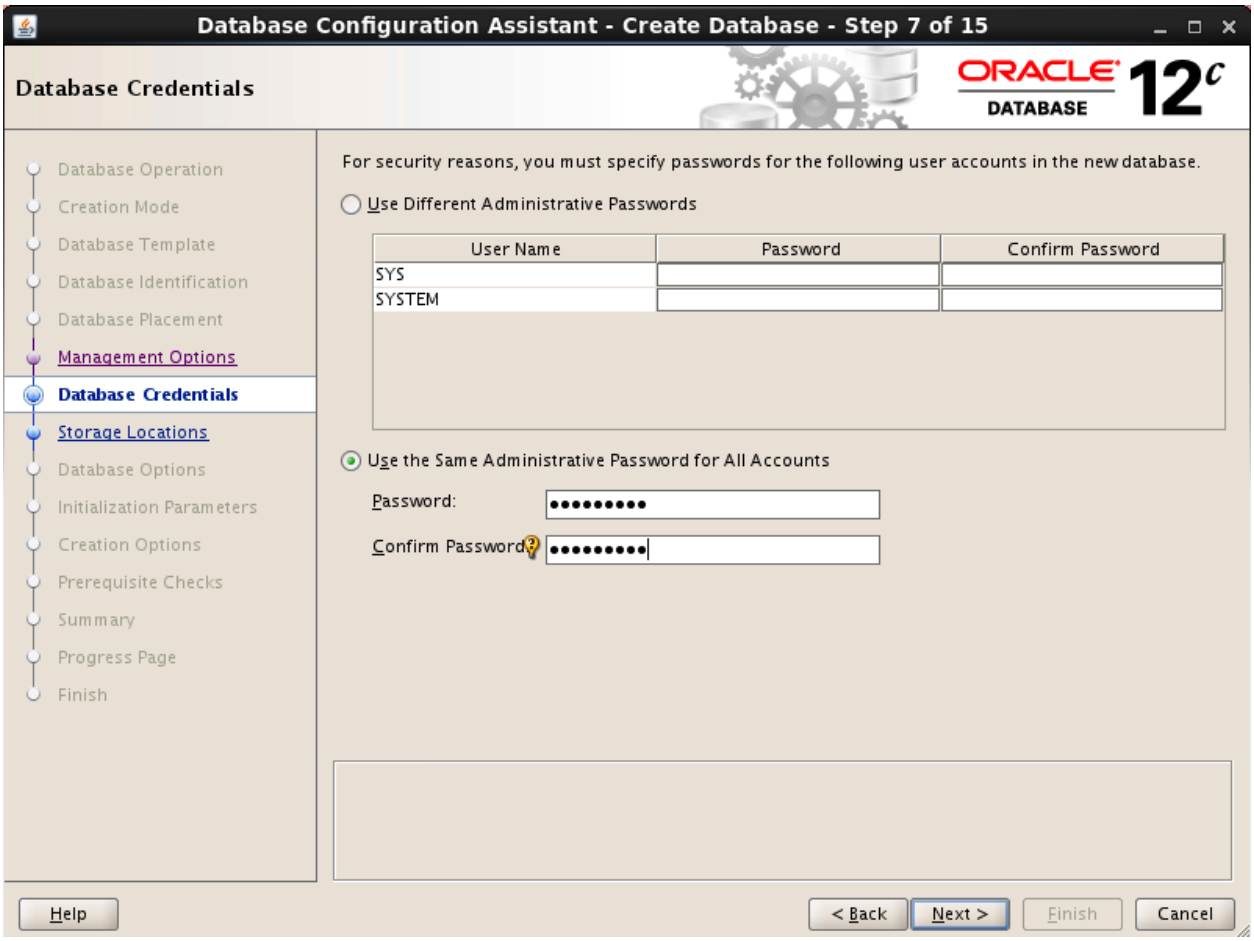

12. Select Use Common Location for All Database Files.
13. Use +ORADATA ASM disk group for the database files location.
14. Click Next.

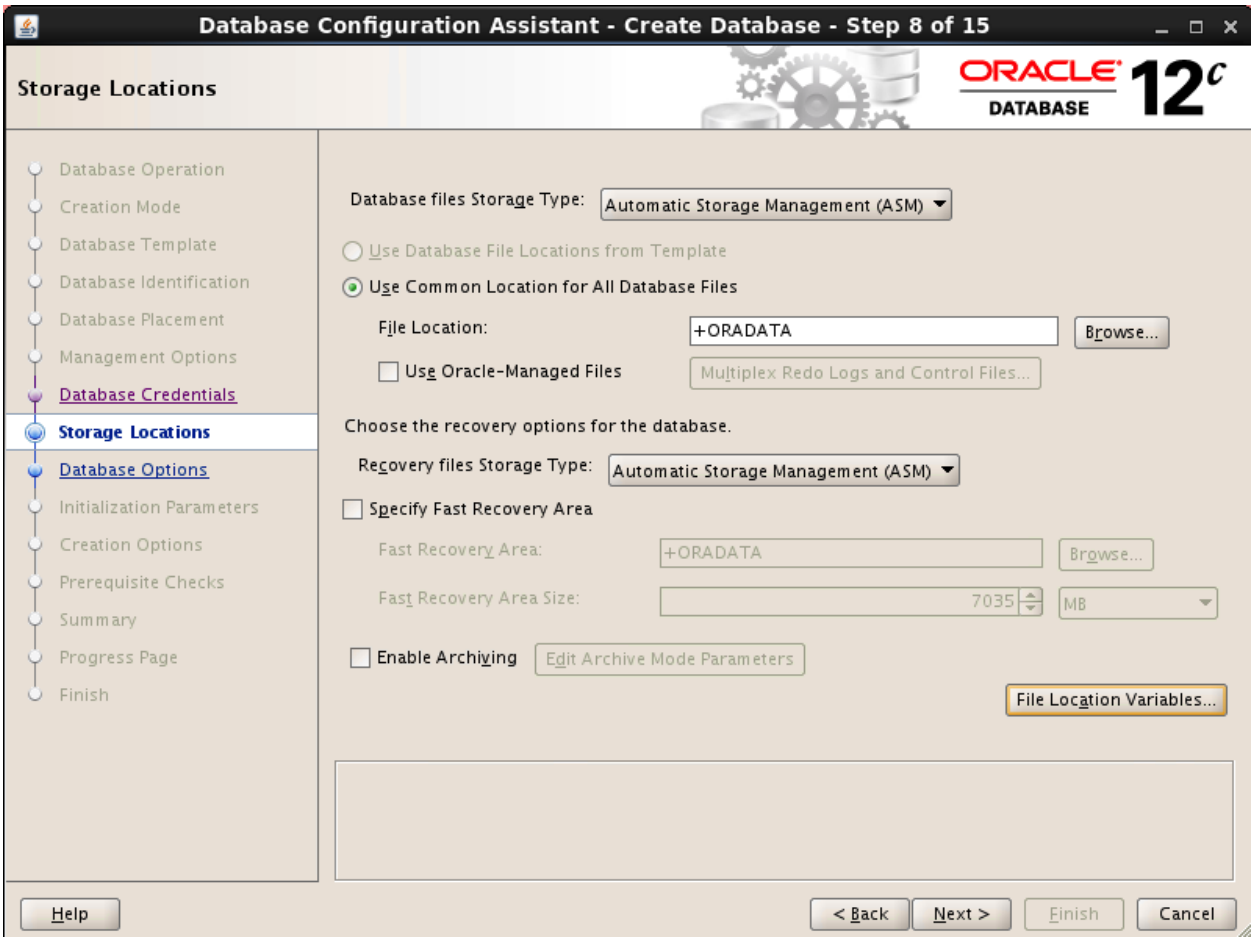

Specify Fast Recovery Area was not selected for this configuration.

Note: Flash recovery area (FRA) is a centralized storage area that is used to keep the entirety of the database backup and recovery files. The flash recovery area is managed by Oracle Managed Files (OMF).

15. Click Next.

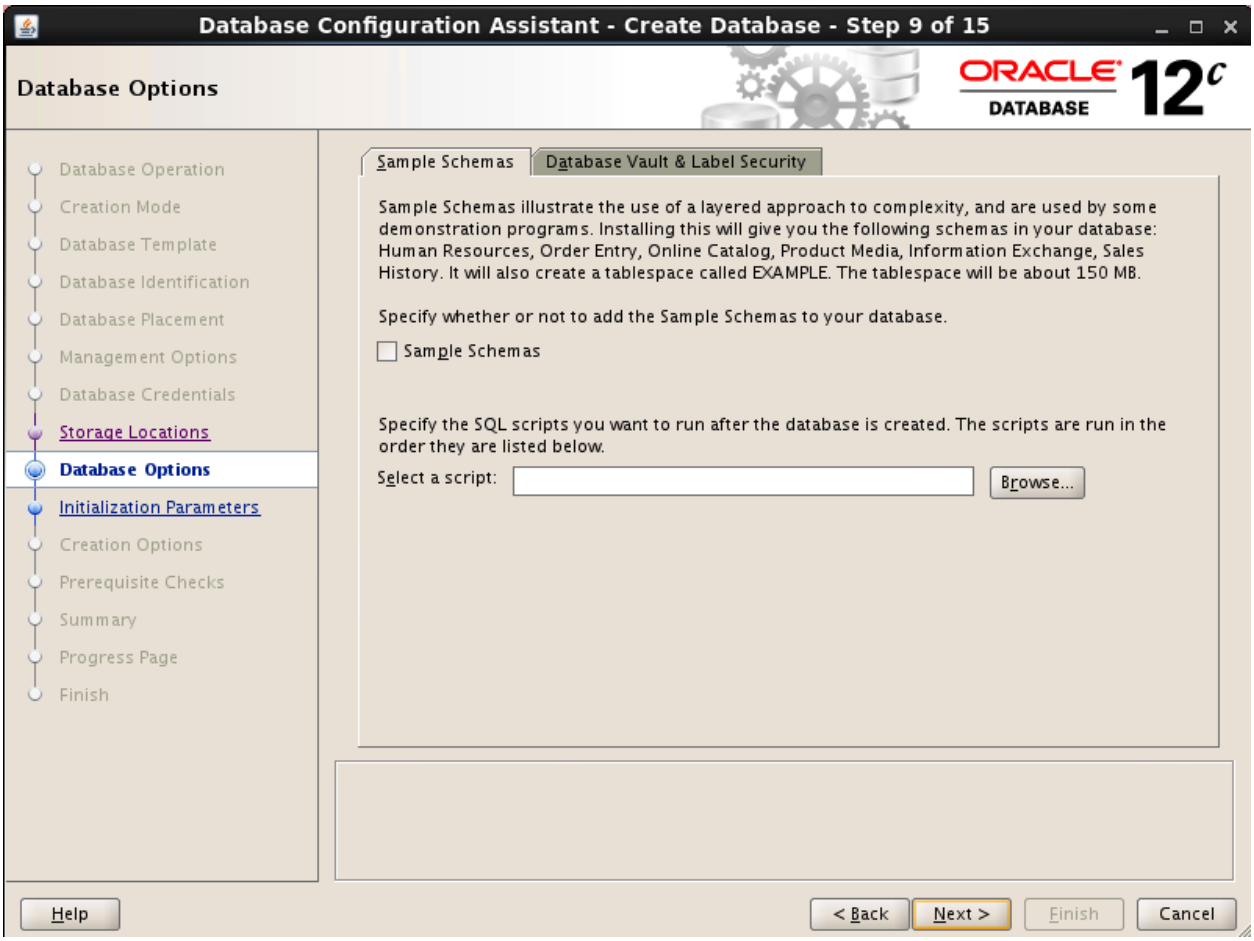

16. For Initialization Parameters, keep the default and click Next. The sample `init.ora` file used in this solution is provided in Figure 5 in the appendix.

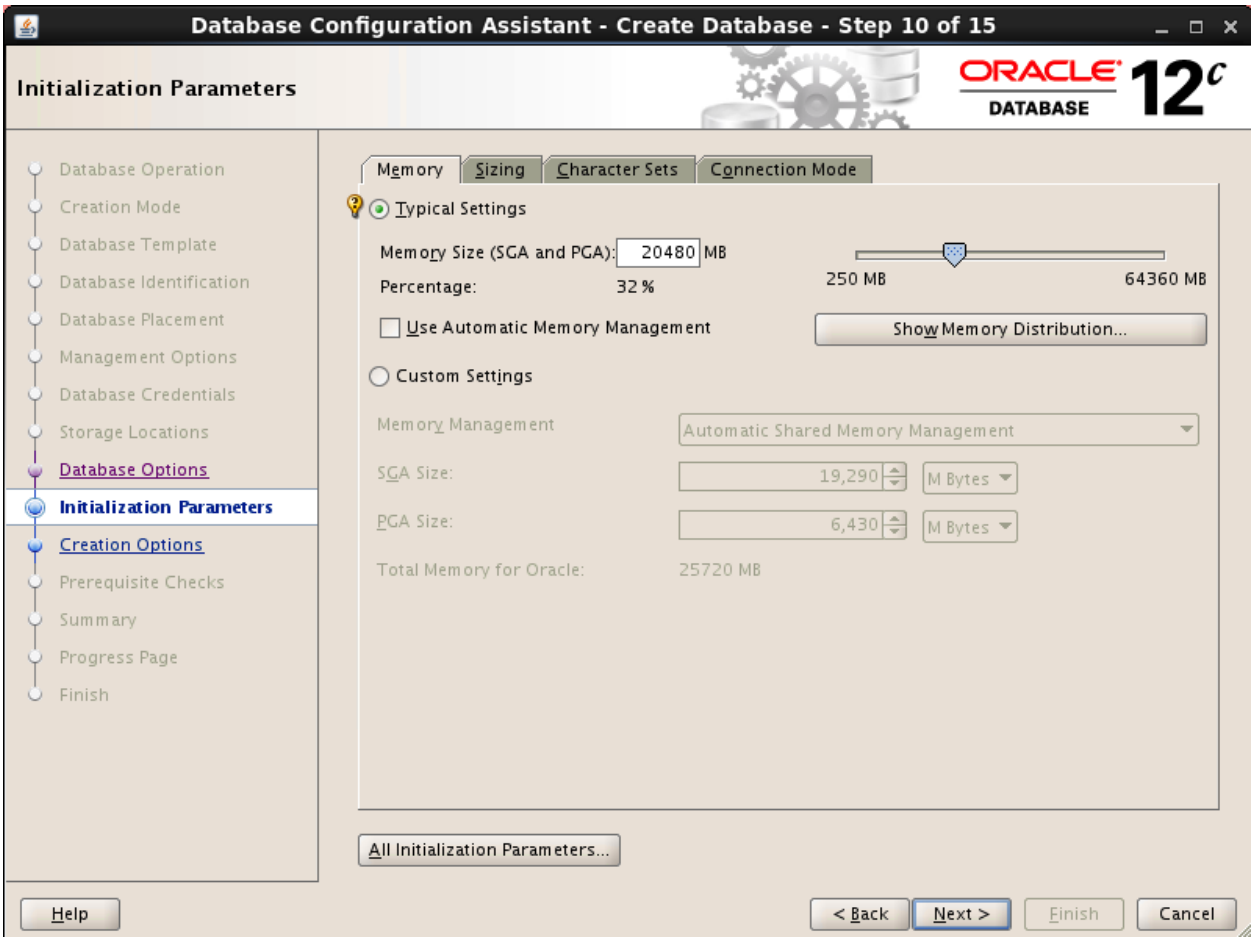

17. Select Create Database and click Next.

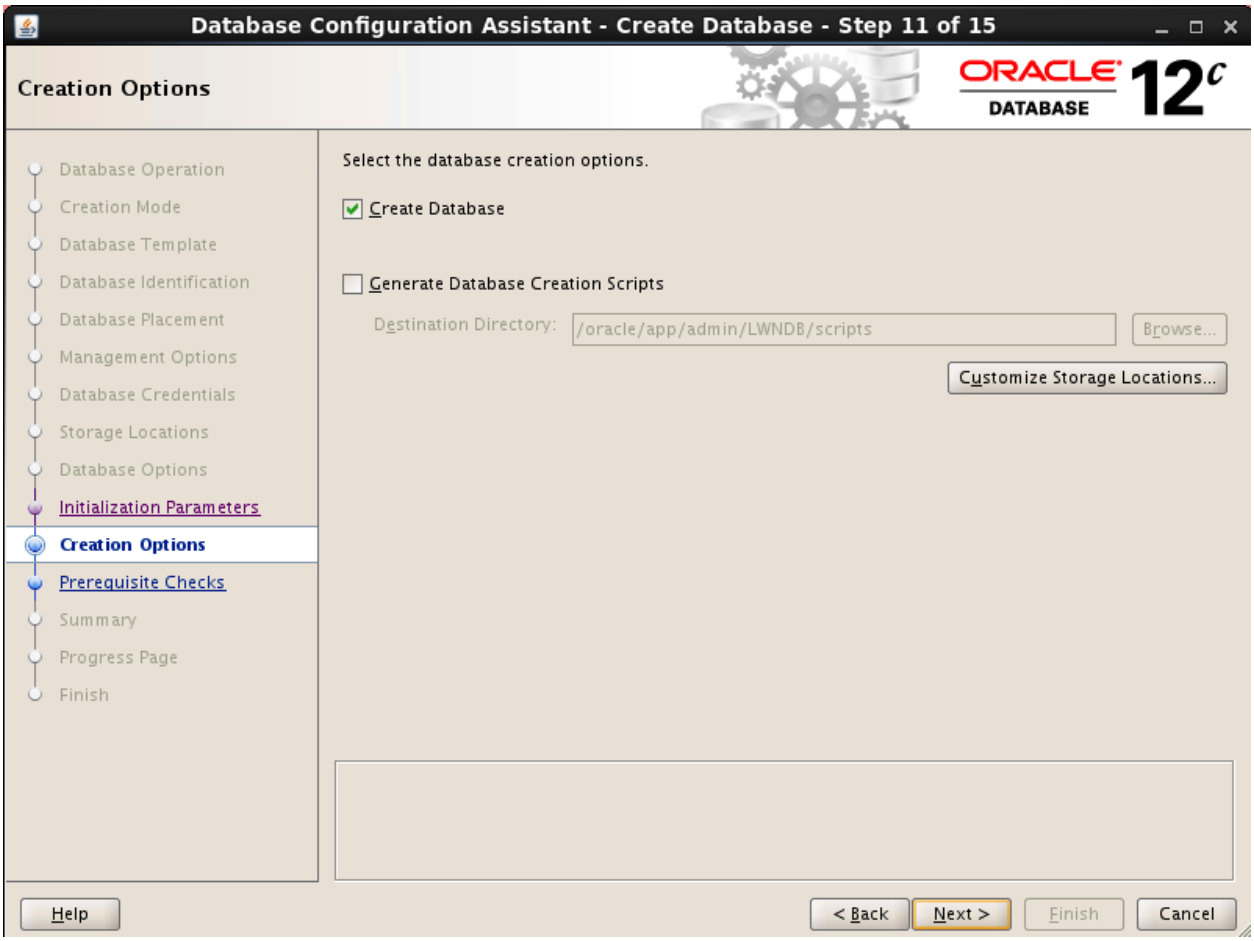

18. In the navigation pane, expand Redo Log Groups.
19. Change the file directory for each Redo Log Group (1–16) to the +ORALOG ASM disk group and click OK.

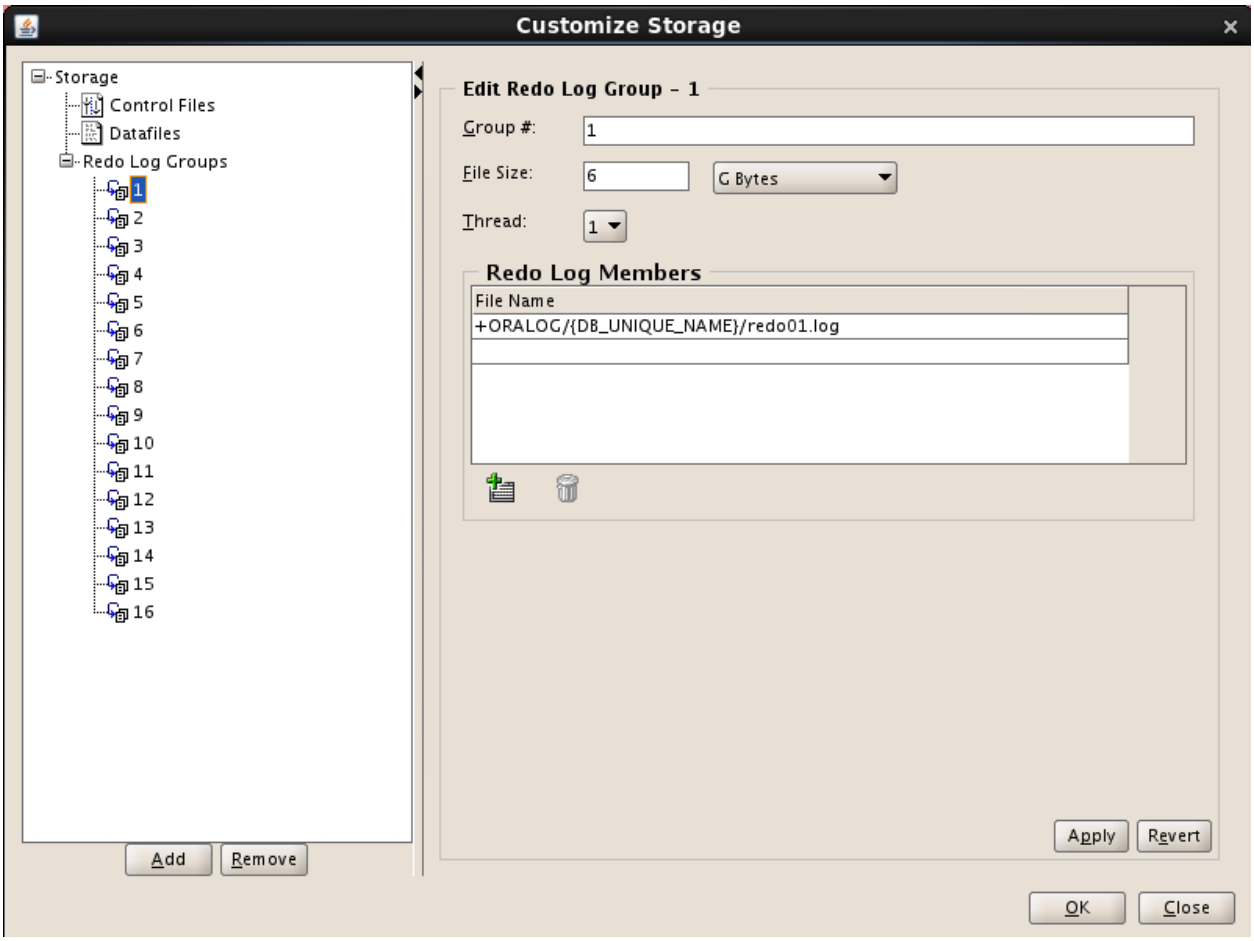

20. After the database is created, the Configuration Assistant performs prerequisite checks.

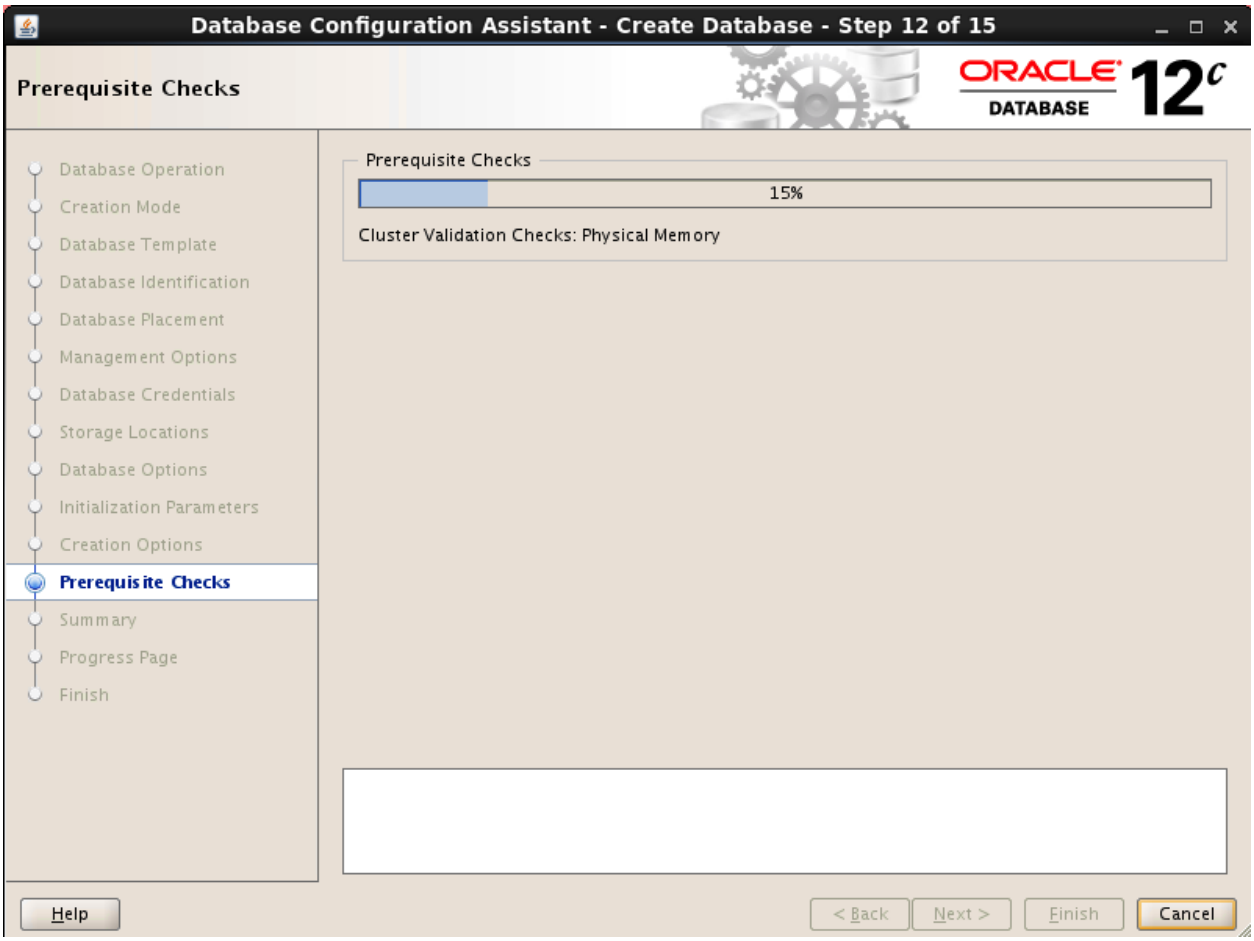

21. Resolve any issues that are found and click Next.

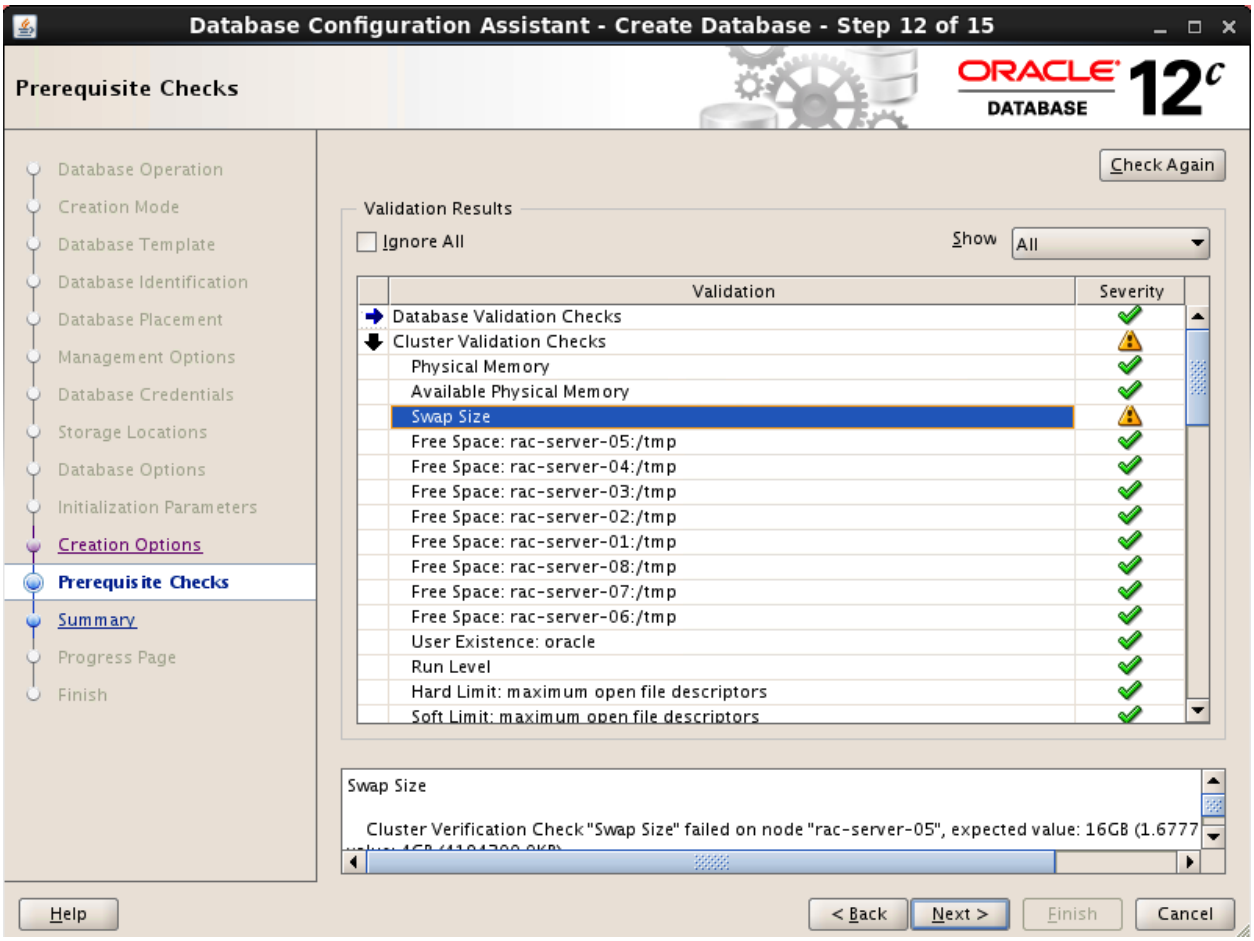

22. Click Finish to complete.

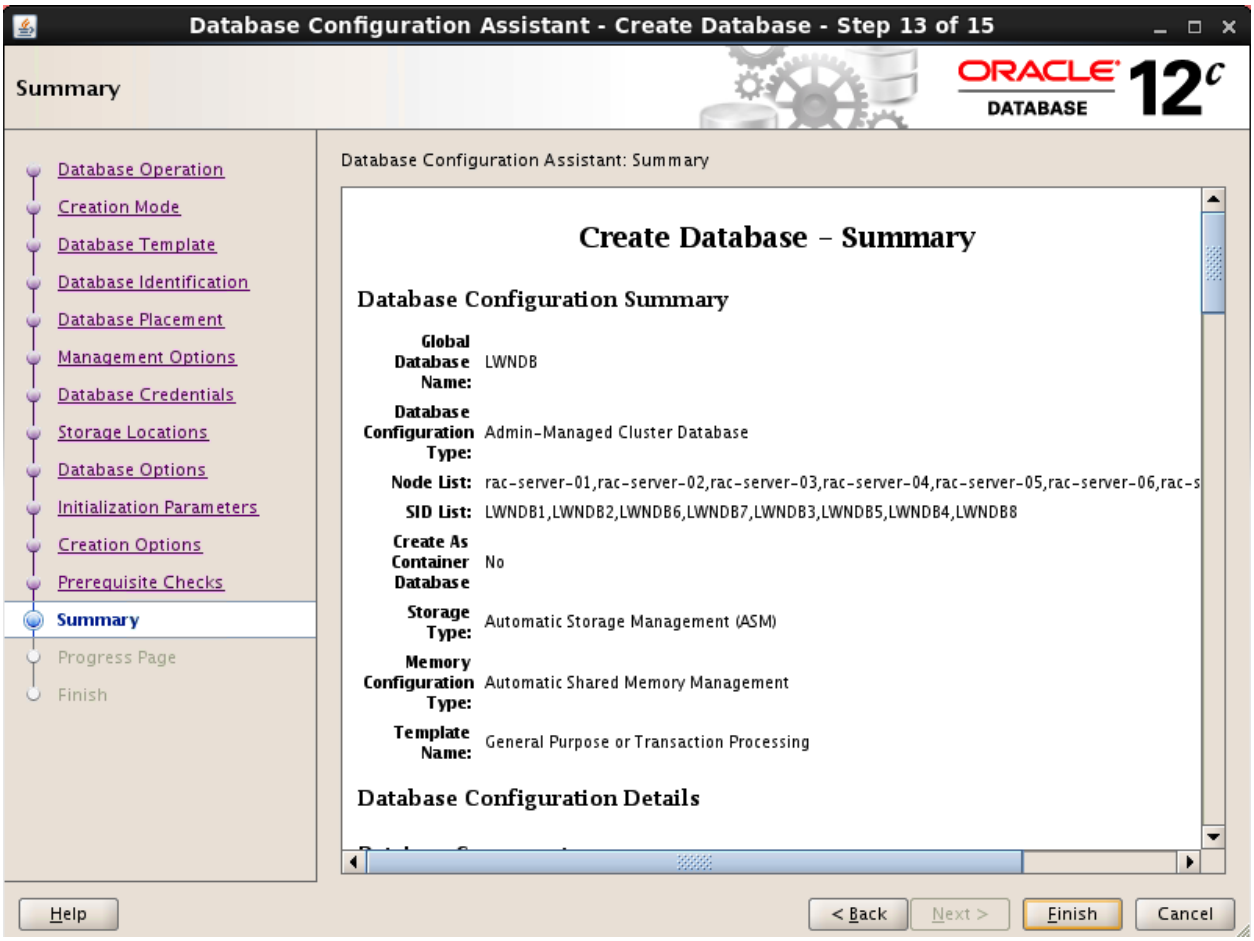

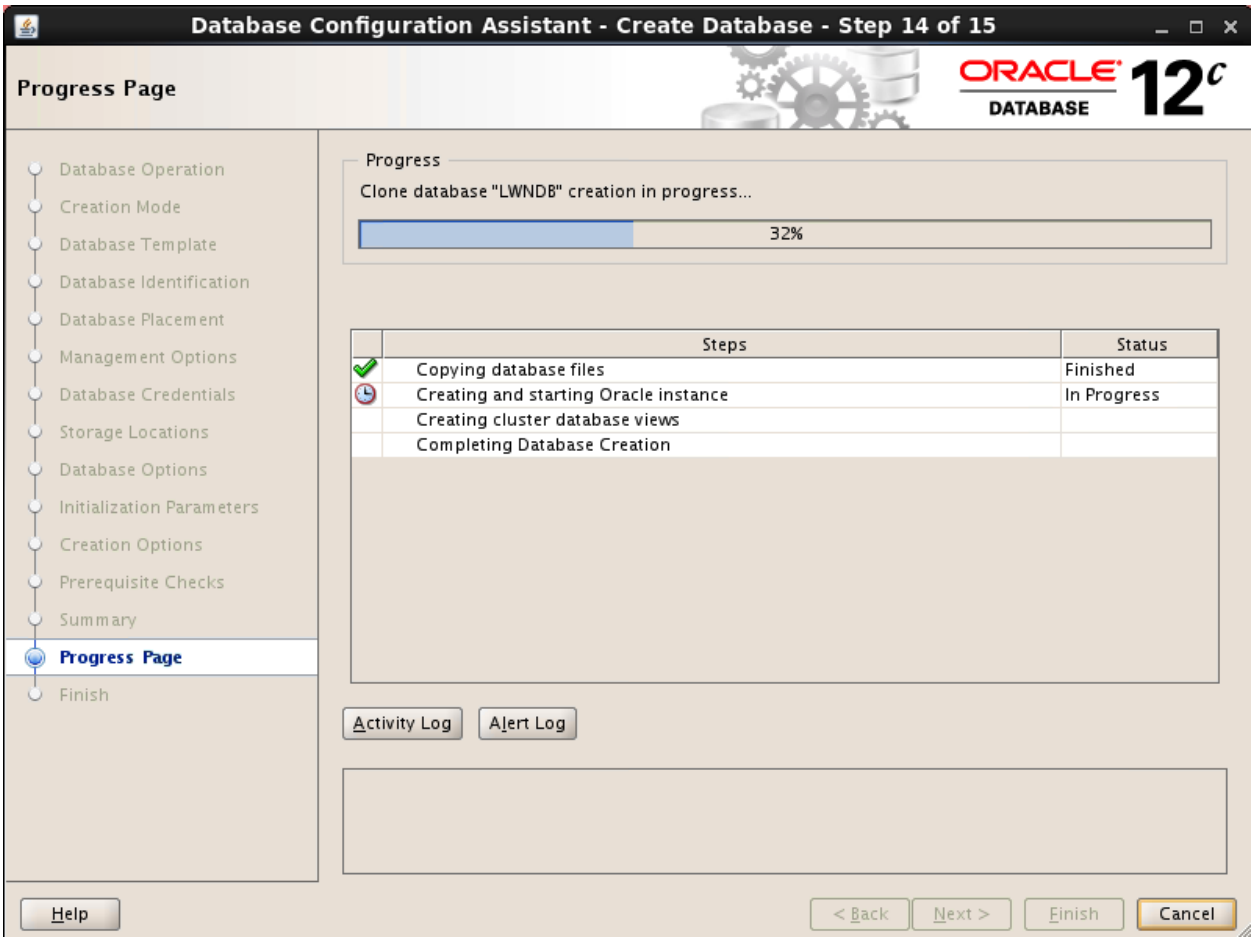

23. Close the database configuration assistant.

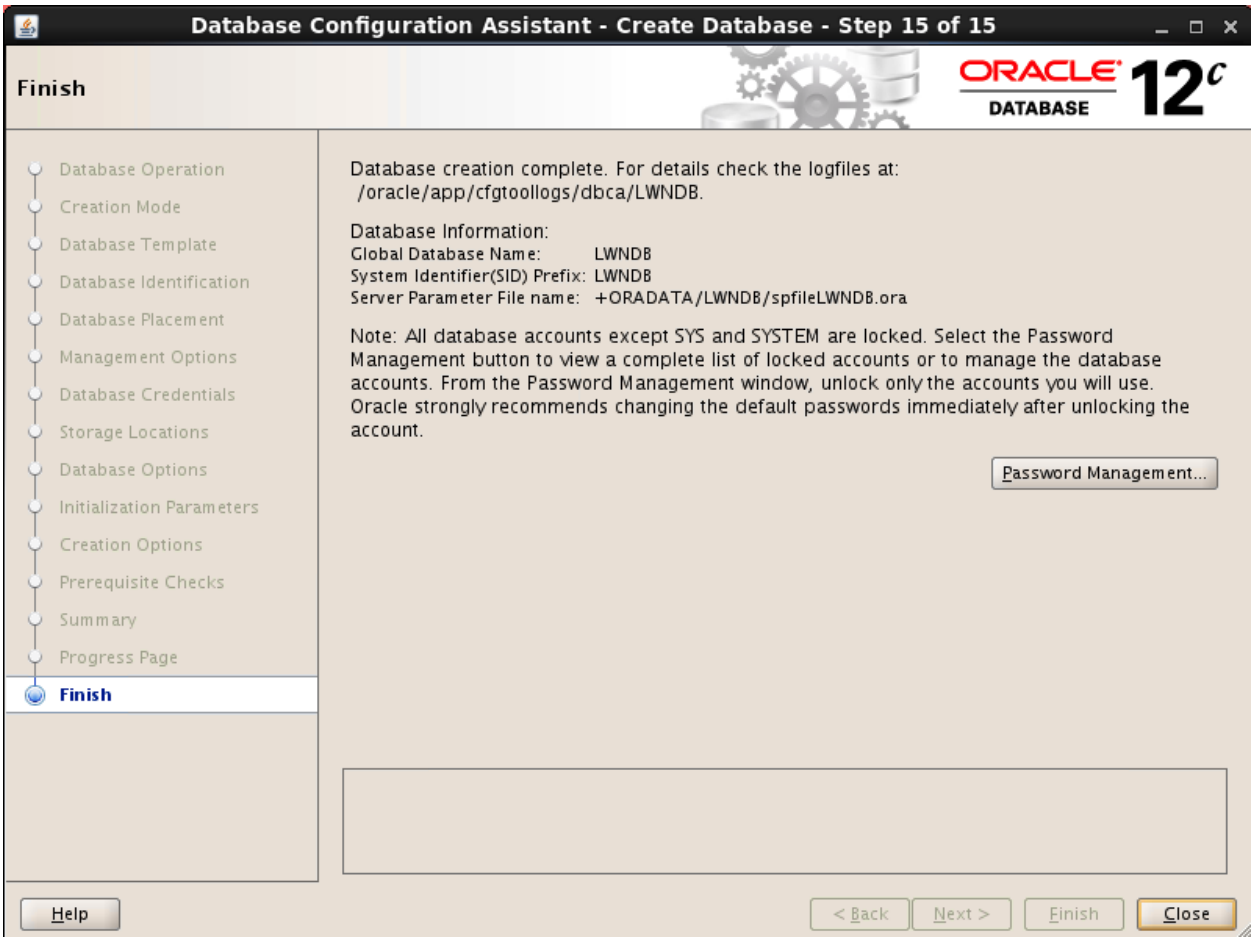

5 Conclusion

The FlexPod Select for High-Performance Oracle RAC solution is designed for applications that are looking for extreme performance and reliability. The architecture is also highly scalable, which allows the customer to select the number of Cisco UCS server blades and EF-Series storage units required for the workload. The NetApp EF560 all-flash storage array is a true enterprise-class storage array designed for applications that expect high IOPS with microsecond-level latencies. The fully redundant enterprise-class high-availability feature on the storage controllers provides maximum availability and delivers the required performance for mission-critical applications for which superior performance and low latency are imperative.

Appendix

Identifying Fibre Channel over Ethernet (FCoE) WWPNs in Cisco UCS

Use the following steps to identify the WWPNs for each server in the Cisco UCS environment.

1. In Cisco UCS Manager, click the Server tab in the navigation pane.
2. Select Service Profiles > root > Sub-organizations > [organization-name].
3. For each server, select the server name.

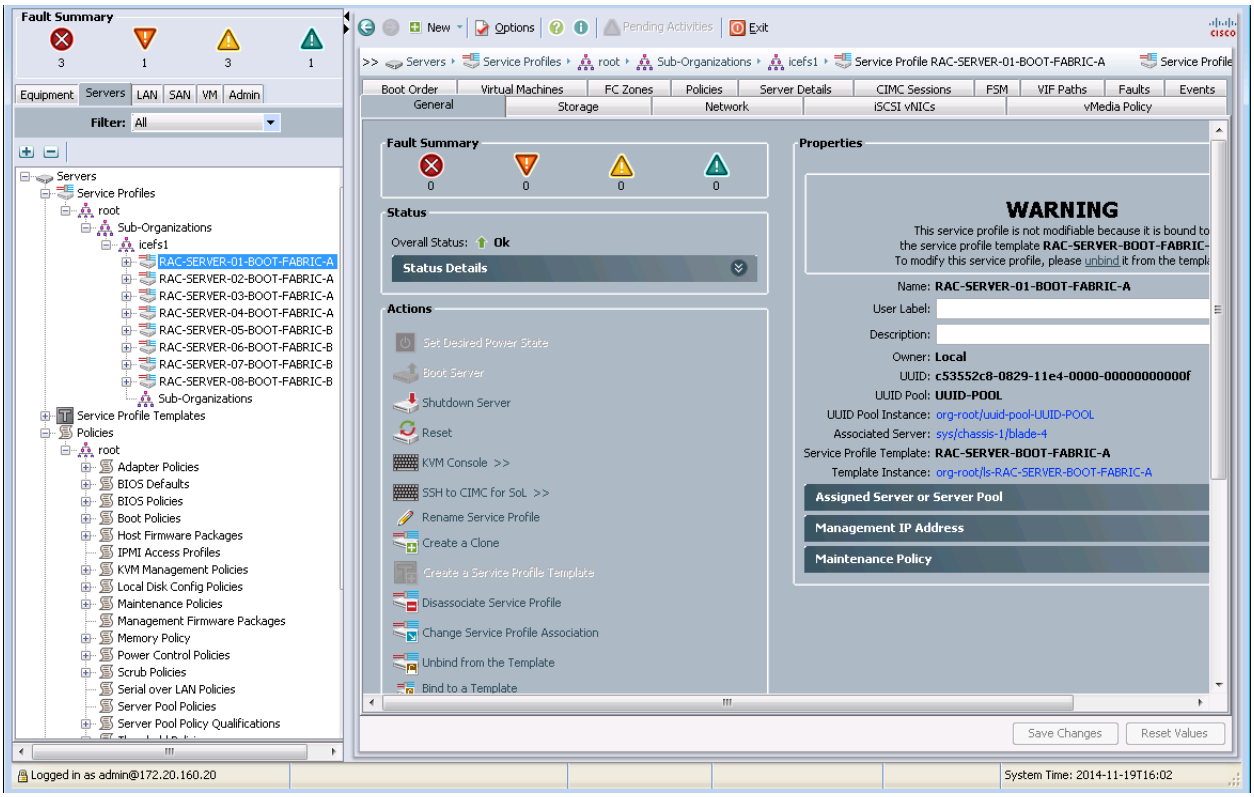

4. In the management pane, select the Storage tab. Expand the WWPN column until the full WWPN is visible. Each server has one WWPN for each fabric. The fabric to which a WWPN corresponds can be identified by the Fabric ID column.

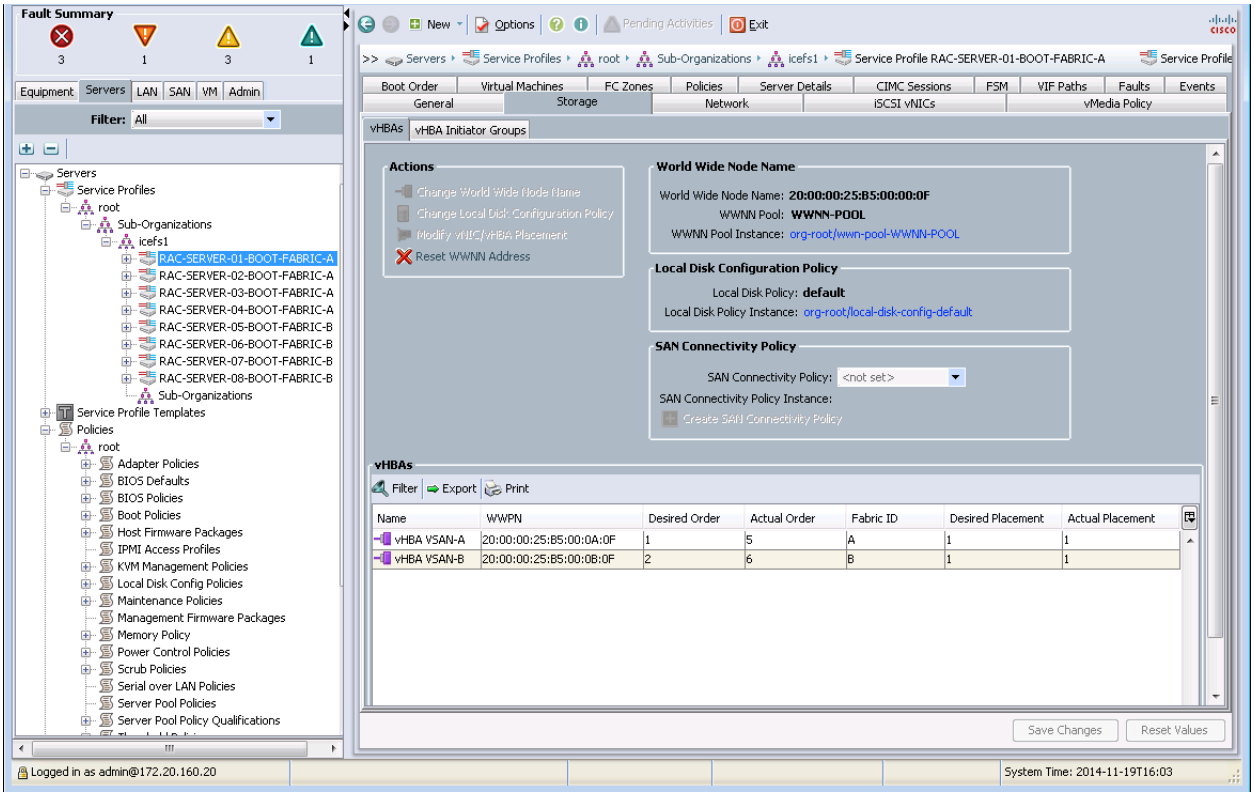

Identifying Fibre Channel WWPNs in SANtricity 11.10

The information gathered using the procedure to identify FC WWPNs is useful during the configuration of zoning within the Cisco Nexus switches and during configuration of the SAN boot path during Cisco UCS service profile configuration.

Use the following steps to identify the WWPN of each port on an EF560 array using SANtricity 11.10:

1. In the SANtricity Enterprise Management Console, right-click the name of the array and select Manage Storage Array.

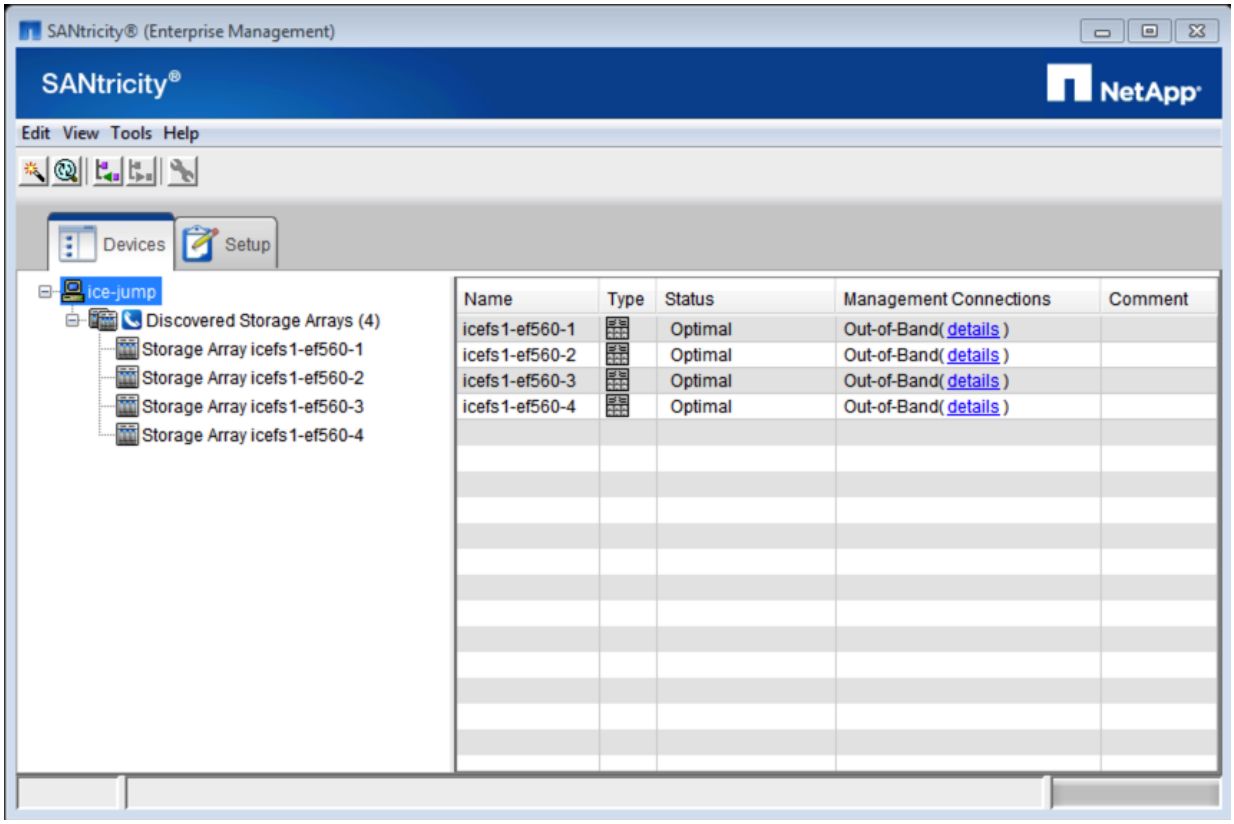

Note: This opens a new Array Management window.

2. From the Array Management window menu bar, select Monitor > Reports > Storage Array Profile.

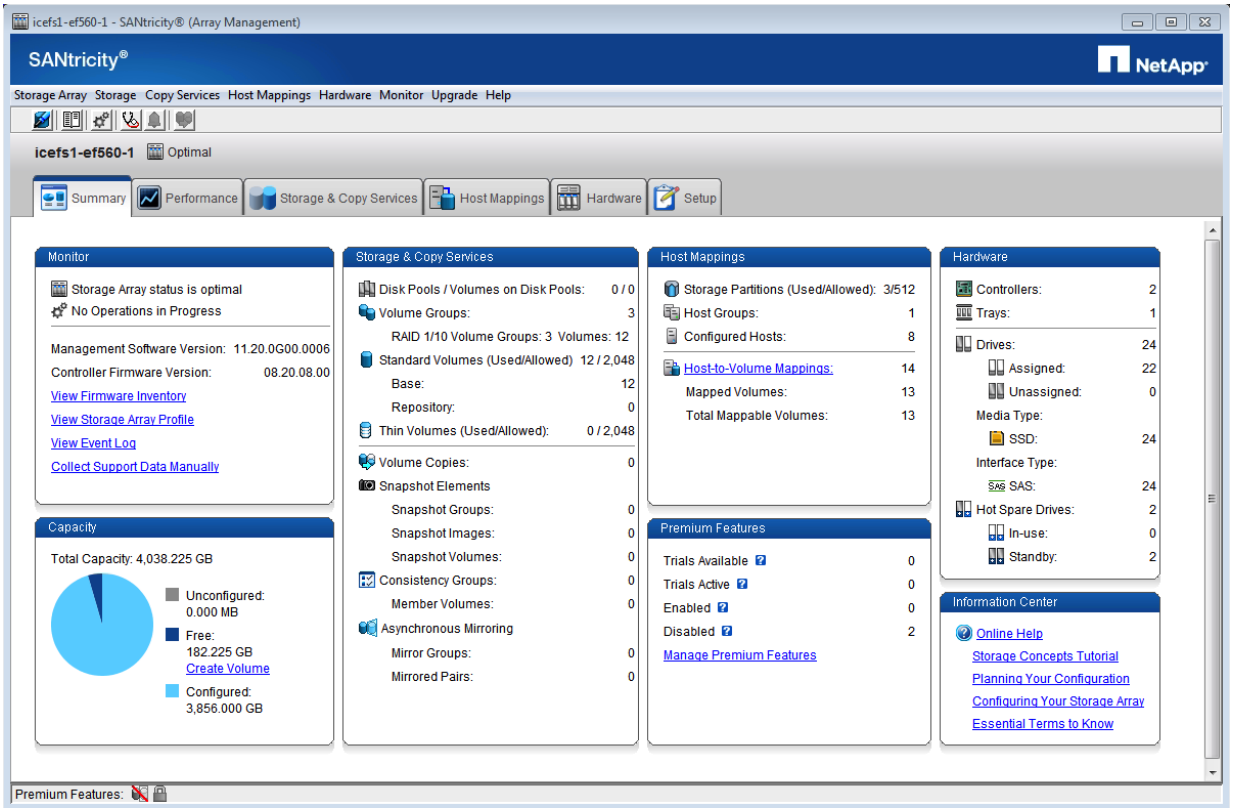

Note: This opens a new Storage Array Profile window.

3. From the Storage Array Profile window, select the Hardware tab.

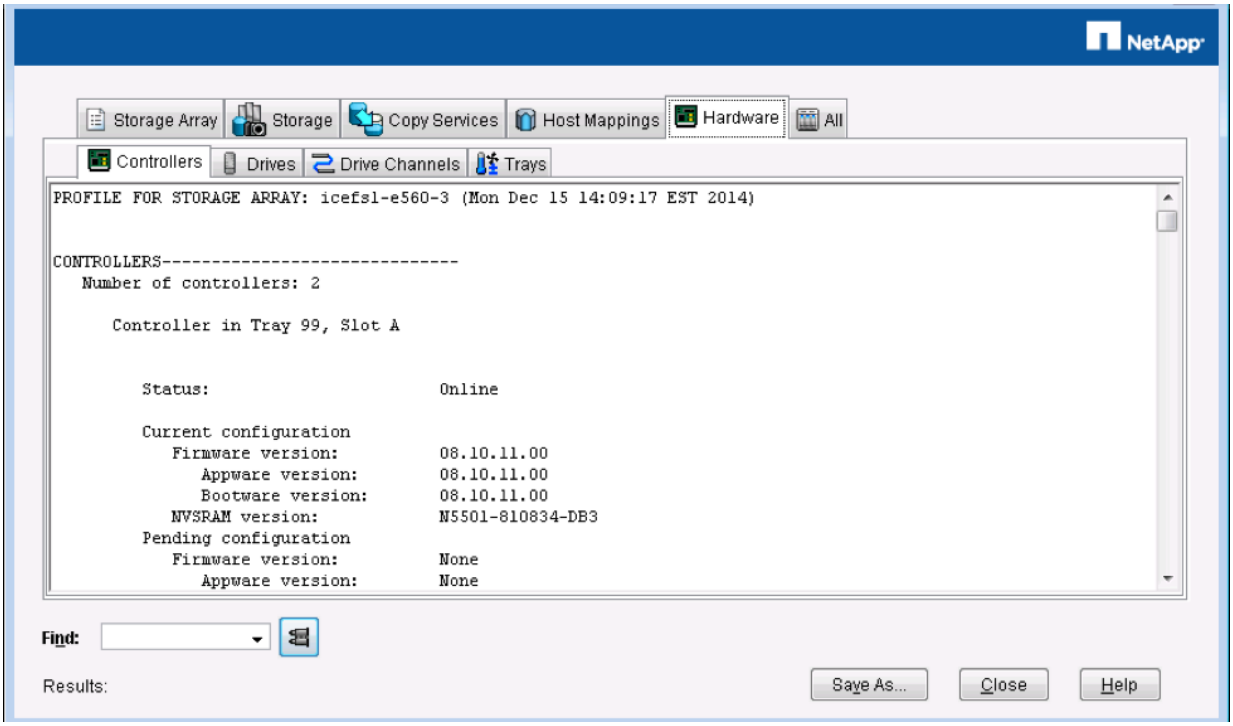

4. In the Find text box at the bottom left side of the window, enter `world-wide port`. The interface finds 8 results.

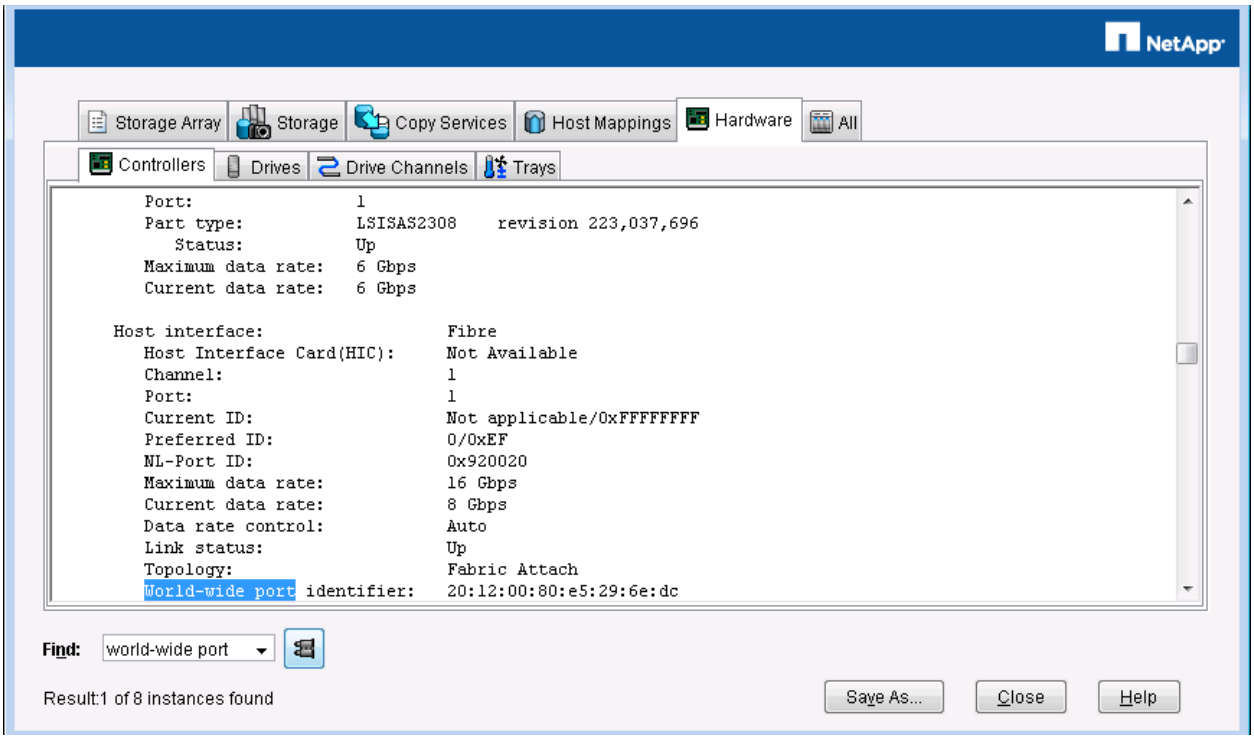

5. Note the port number and the worldwide port identifier for each port. Click the Next button (binoculars icon) to navigate to the next port. The first set of ports corresponds to controller 1, and the second set of ports corresponds to controller 2.
6. Repeat step 1 to step 5 for each array in the environment.

Storage Layout for Eight-Node RAC

Table 24 describes the consolidated storage layout for an eight-node Oracle RAC environment.

Table 24) Storage layout for eight-node RAC.

Storage Array	Type	Volume Group RAID 10	Number of Physical Disks	Volume/ LUN Name	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name/Host Name
EF560-1	Redo logs and boot LUNs	A1LOGVG	2	RAC1BOOT	78	256GB	2	RAC-A01
				RAC2BOOT	78			RAC-A02
				LOG1	50			RAC
				LOG2	50			
	Data	A1VG1	10	DATALUN1	450	3.6TB		
				DATALUN2	450			
				DATALUN3	450			
				DATALUN4	450			

Storage Array	Type	Volume Group RAID 10	Number of Physical Disks	Volume/ LUN Name	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name/Host Name
		A1VG2	10	DATALUN5	450			
				DATALUN6	450			
				DATALUN7	450			
				DATALUN8	450			
EF560-2	Redo logs and boot LUNs	A2LOGVG	2	LOG3	50	271GB	2	RAC
				LOG4	50			
				RAC3BOOT	78			RAC-A03
				RAC4BOOT	78			RAC-A04
				GRID1	5			RAC
				GRID2	5			
				GRID3	5			
	Data	A2VG1	10	DATALUN9	450	3.6TB		RAC
				DATALUN10	450			
				DATALUN11	450			
				DATALUN12	450			
		A2VG2	10	DATALUN13	450			
				DATALUN14	450			
				DATALUN15	450			
				DATALUN16	450			
EF560-3	Redo logs and boot LUNs	A3LOGVG	2	LOG5	50	256GB	2	RAC
				LOG6	50			
				RAC5BOOT	78			RAC-B01
				RAC6BOOT	78			RAC-B02
	Data	A3VG1	10	DATALUN17	450	3.6TB		RAC
				DATALUN18	450			
				DATALUN19	450			
				DATALUN20	450			
		A3VG2	10	DATALUN21	450			
				DATALUN22	450			

Storage Array	Type	Volume Group RAID 10	Number of Physical Disks	Volume/ LUN Name	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name/Host Name
				DATALUN23	450			
				DATALUN24	450			
EF560-4	Redo logs and boot LUNs	A4LOGVG	2	LOG7	50	256G	2	RAC
				LOG8	50			
				RAC7BOOT	78			RAC-B03
				RAC8BOOT	78			RAC-B04
	Data	A4VG1	10	DATALUN25	450	3.6TB		RAC
				DATALUN26	450			
				DATALUN27	450			
				DATALUN28	450			
		A4VG2	10	DATALUN29	450			
				DATALUN30	450			
				DATALUN31	450			
				DATALUN32	450			

Sample init.ora File

The sample `init.ora` file in Figure 5 was used for this solution.

Note: The SGA size in this setup was kept to the lowest in order to get the maximum hits on the storage arrays. For a real production environment, the databases must be provided with appropriate SGA.

Figure 5) Sample init.ora file for this solution.

```
cat init.ora
*.audit_file_dest='/oracle/app/admin/LWNDB/adump'
*.audit_trail='db'
*.cluster_database=true
*.compatible='12.1.0.2.0'
*.control_files='+ORADATA/LWNDB/control01.ctl','+ORADATA/LWNDB/control02.ctl'
*.db_block_size=8192
*.db_create_file_dest='+ORADATA'
*.db_create_online_log_dest_1='+ORALOG'
*.db_domain=''
*.db_name='LWNDB'
*.diagnostic_dest='/oracle/app'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=LWNDBXDB)'
LWNDB3.instance_number=3
LWNDB6.instance_number=6
LWNDB7.instance_number=7
LWNDB5.instance_number=5
LWNDB1.instance_number=1
LWNDB4.instance_number=4
LWNDB2.instance_number=2
LWNDB8.instance_number=8
LWNDB2.thread=2
LWNDB6.thread=6
LWNDB8.thread=8
```



```
LWNDB7.thread=7
LWNDB3.thread=3
LWNDB4.thread=4
LWNDB5.thread=5
LWNDB1.thread=1
LWNDB8.undo_tablespace='UNDOTBS8'
LWNDB5.undo_tablespace='UNDOTBS6'
LWNDB1.undo_tablespace='UNDOTBS1'
LWNDB2.undo_tablespace='UNDOTBS2'
LWNDB6.undo_tablespace='UNDOTBS3'
LWNDB7.undo_tablespace='UNDOTBS4'
LWNDB3.undo_tablespace='UNDOTBS5'
LWNDB4.undo_tablespace='UNDOTBS7'
*.db_writer_processes=20
*.filesystemio_options='setall'
*.db_files=2000
*.open_cursors=300
*.pga_aggregate_target=2G
*.processes=3000
*.remote_login_passwordfile='exclusive'
*.shared_pool_size=4G
*.db_cache_size=3G
*.log_buffer=134217728
*.pga_aggregate_target=1G
```

References

This report references the following documents and resources:

- Cisco UCS Virtual Interface Card Drivers for Linux Installation Guide
- Cisco UCS 2.2(5) Hardware and Software Interoperability Matrix
- Cisco UCS fnic Tunables
- Oracle Database Quick Installation Guide 12c Release 1 (12.1) for Linux x86
- NetApp EF560 Flash Array Installation Guide
- TR-4305: NetApp Extreme Performance Solution for Oracle Database

Version History

Version	Date	Document Version History
Version 1.0	January 2015	Engineering content creation
Version 2.0	December 2015	Updated test configuration data

NetApp Contact Information

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: 1 408 822 6000
Fax: 1 408 822 4501
Support telephone: 1 888 463 8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.