



NetApp Verified Architecture

Oracle 19c RAC databases on FlexPod Datacenter with Cisco UCS and NetApp AFF A800 over FC

Design and deployment guide

Allen Cao, NetApp

June 2021 | NVA-1155 | Version 1.0

Abstract

This design and deployment guide for Oracle 19c RAC databases on FlexPod® Datacenter with Cisco UCS and NetApp® AFF A800 over FC provides details of the solution design as well as step-by-step deployment processes for hosting Oracle RAC databases on most recent FlexPod Datacenter infrastructure with the Oracle Linux 8.2 operating system and a Red Hat compatible kernel.

In partnership with



TABLE OF CONTENTS

Oracle 19c RAC databases on FlexPod Datacenter with Cisco UCS and NetApp AFF A800 over FC	3
Customer value	3
Use cases	4
Target audience.....	4
Architecture	5
Solution technology	5
Architectural diagram.....	5
Hardware requirements	6
Software requirements	7
Design considerations	7
Network design.....	7
Compute design	11
Storage design	12
Oracle 19c database deployment considerations	12
Best practices	16
Deploying Oracle 19c solution	17
Manual deployment	17
Automated deployment (optional).....	145
Ansible controller setup	145
Deployment summary.....	147
Validation results	148
FIO benchmark validation.....	148
SLOB benchmark validation	151
Infrastructure high availability and resiliency validation	153
Conclusion	156
Where to find additional information	157

LIST OF TABLES

Table 1) Hardware requirements	6
Table 2) Software requirements.	7
Table 3) UCS 6454 ports.....	8
Table 4) Nexus switch port connections.	9

Table 5) MDS 9132 FC switch port connections.	9
Table 6) VLAN and VSAN details.	10
Table 7) vPC configuration.	10
Table 8) vPC FC traffic aggregation.	11
Table 9) Connection ports for an AFF A800 HA controller pair.	12
Table 10) Oracle 19c RAC cluster storage layout.	14
Table 11) Port connections between MDS, FI, and AFF A800.	58
Table 12) Port connections between MDS, FI, and AFF A800, continued.	58
Table 13) ONTAP node configuration details.	70
Table 14) Cluster details.	72

LIST OF FIGURES

Figure 1) Recommended infrastructure architecture.	6
Figure 2) UCS 6454 Fabric Interconnect.	8
Figure 3) Nexus 9336C-FX2 switch.	9
Figure 4) MDS 9132 FC switch.	9
Figure 5) AFF A800 ports and available expansion slots.	12
Figure 6) Solution deployment in the CDB/PDB model in an RAC configuration.	14
Figure 7) SnapCenter architecture for Oracle data protection and management.	16

Oracle 19c RAC databases on FlexPod Datacenter with Cisco UCS and NetApp AFF A800 over FC

Oracle Database 19c is the final, long-term-support release of the Oracle Database 12c family of products, which includes Oracle Database 18c. Long term support means that Oracle Database 19c comes with four years of premium support and a minimum of three years of extended support.

This solution provides the details of design considerations and step-by-step deployment guidance to help customers and/or partners who are interested in deploying Oracle 19c RAC databases on the latest release of FlexPod Datacenter converged infrastructure over the FC protocol and the latest supported Oracle Linux version 8.2.

This solution also showcases how Oracle 19c can be deployed using a root container database that hosts several pluggable databases (the CDB/PDB model) for resource sharing on bare-metal infrastructure and database workload consolidation to reduce the total cost of ownership.

This solution demonstrates that, by using the NetApp SnapCenter UI, you can dramatically speed up Oracle database backup, recovery, clone, and refresh. SnapCenter provides a quick database backup, restore, and recovery solution that helps to ease complicated application and database migration to Oracle 19c. By using NetApp SnapCenter, you can increase DBA productivity while improving database performance.

Customer value

The current Oracle installed base is over 400,000 database customers across many versions of Oracle DB, largely between 9i and 18c, on different infrastructure platforms. As such, there is a large installed base that has been running on now outdated hardware architecture and Oracle versions that represents a number of challenges for customers:

- **Database sprawl.** Many databases running on different servers are scattered between different locations, which creates management difficulty, inefficient resource utilization, and security concerns of shadow IT.
- **Performance and availability.** The current Oracle version and infrastructure no longer provide the level of performance and availability required.
- **Migration urgency.** The need to migrate to latest version of Oracle 19c before time runs out on unsupported versions of Oracle.
- **Deployment and scale out.** The need for deployment guidance on the latest version of Oracle 19c and database deployment scale up and out.

FlexPod Datacenter is a market-leading and proven converged infrastructure that answers these customer challenges. It provides an ideal platform for customers to consolidate their Oracle database workload for efficient resource sharing and management.

FC is the top data protocol for Oracle workloads. The best practices documented here provide you with confidence that Cisco and NetApp have verified that this FC solution provides the best performance and availability possible with the latest hardware infrastructure.

Customers do not typically have the time and idle resources to research which Cisco servers, NetApp storage products, and Oracle Linux setup are best for their IT requirements, nor do they have the time to test and determine best practices. This solution provides you with a known solution or that greatly reduces the time required to migrate, deploy, and scale out their Oracle databases.

Traditionally, Oracle DBAs use Oracle Recovery Manager (RMAN) to perform database backup and recovery. RMAN is a reliable tool for Oracle DBAs that is built into the Oracle database software stack. However, backup can take a long time, especially for large databases. It can also take a long time to restore the database in a recovery scenario, which might take a heavy toll on the recovery time objective

(RTO) or recovery point objective (RPO). A significant amount of precious storage space must be reserved for storing the backup of Oracle data and log files.

Compared to RMAN, NetApp SnapCenter is a very useful tool for DBAs. It allows almost instantaneous backup, recovery, and cloning of production Oracle databases in a storage-efficient way. It can provide real storage cost savings and make the daily activities of the Oracle DBA easier. Fast NetApp Snapshot copies remediate throttle issues with storage bandwidth in backing up databases and improving database performance. Fast backup and restore can also dramatically improve RTO and RPO. Beyond data protection, customer can leverage SnapCenter tool for database migration to Oracle 19c in migration test dry runs and go-live checkpointing and fallback planning.

Oracle DBAs generally spend large amounts of time running database backup and data refresh activities. With NetApp SnapCenter, Oracle DBAs can be relieved from daily database administration tasks and turn their efforts instead toward business application development, support, user experience improvement, and/or other revenue generating activities.

Use cases

The goal of this solution is to establish bare-metal FlexPod with AFF and FC SAN as a leading comprehensive infrastructure solution to run Oracle 19c RAC. The solution demonstrates the ability to scale IOPS while maintaining a near constant low latency with FC SAN for Oracle RAC on FlexPod.

The solution showcases a new version of NetApp ONTAP® 9.7 capabilities, and NetApp SnapCenter 4.4 as a viable tool for managing Oracle 19c databases running on FlexPod.

Where applicable, the solution demonstrates how automation can be used to improve productivity and time to value with FlexPod Oracle deployment and management.

This solution applies to the following use cases:

- Replace current aging servers and storage hosting Oracle DB and RAC. Customers cite performance, availability, and/or a desire to migrate from a prior version to Oracle Database 19c.
- New customers and installations of Oracle RAC 19c.
- Current Oracle 18c (or prior version) single-instance database customers looking to improve availability and performance of Oracle DB by implementing RAC.
- Consolidate several servers and Oracle DBs onto fewer server and container database (CDB) on FlexPod with 19c RAC, thus reducing both operating and licensing costs.

Target audience

The target audience for the solution includes the following groups:

- Sales engineers
- Field consultants
- Database administrators
- Oracle database architects
- IT managers
- Any customers or partners who want to deploy an Oracle 19c RAC database solution on FlexPod converged infrastructure with NetApp ONTAP® and the Cisco UCS platform.

Architecture

Solution technology

In an era of data-driven business and applications, data is the lifeblood of any enterprise. Data powers essentially every operation in the modern enterprise, from keeping the supply chain operating efficiently to managing customer relationships.

Built on innovative technology from NetApp and Cisco, the FlexPod converged infrastructure platform meets and exceeds the challenges of simplifying deployments for best-in-class data center infrastructure.

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and nonvirtualized solutions. Composed of pre-validated storage, networking, and server technologies, FlexPod is designed to increase IT responsiveness to organizational needs and reduce the cost of computing with maximum uptime and minimal risk. Simplifying the delivery of data center platforms gives enterprises an advantage in delivering new services and applications.

FlexPod provides the following differentiators:

- A single platform built from unified compute, fabric, and storage technologies, allowing you to scale to large-scale data centers without architectural changes.
- Flexible design with a broad range of reference architectures and validated designs.
- Elimination of costly, disruptive downtime through Cisco UCS and NetApp ONTAP.
- A pre-validated platform to minimize business disruption and improve IT agility and reduce deployment time from months to weeks.
- Cisco Validated Designs (CVDs) and NetApp Validated Architectures (NVAs) covering a variety of use cases.
- A proven, market-leading converged infrastructure platform.

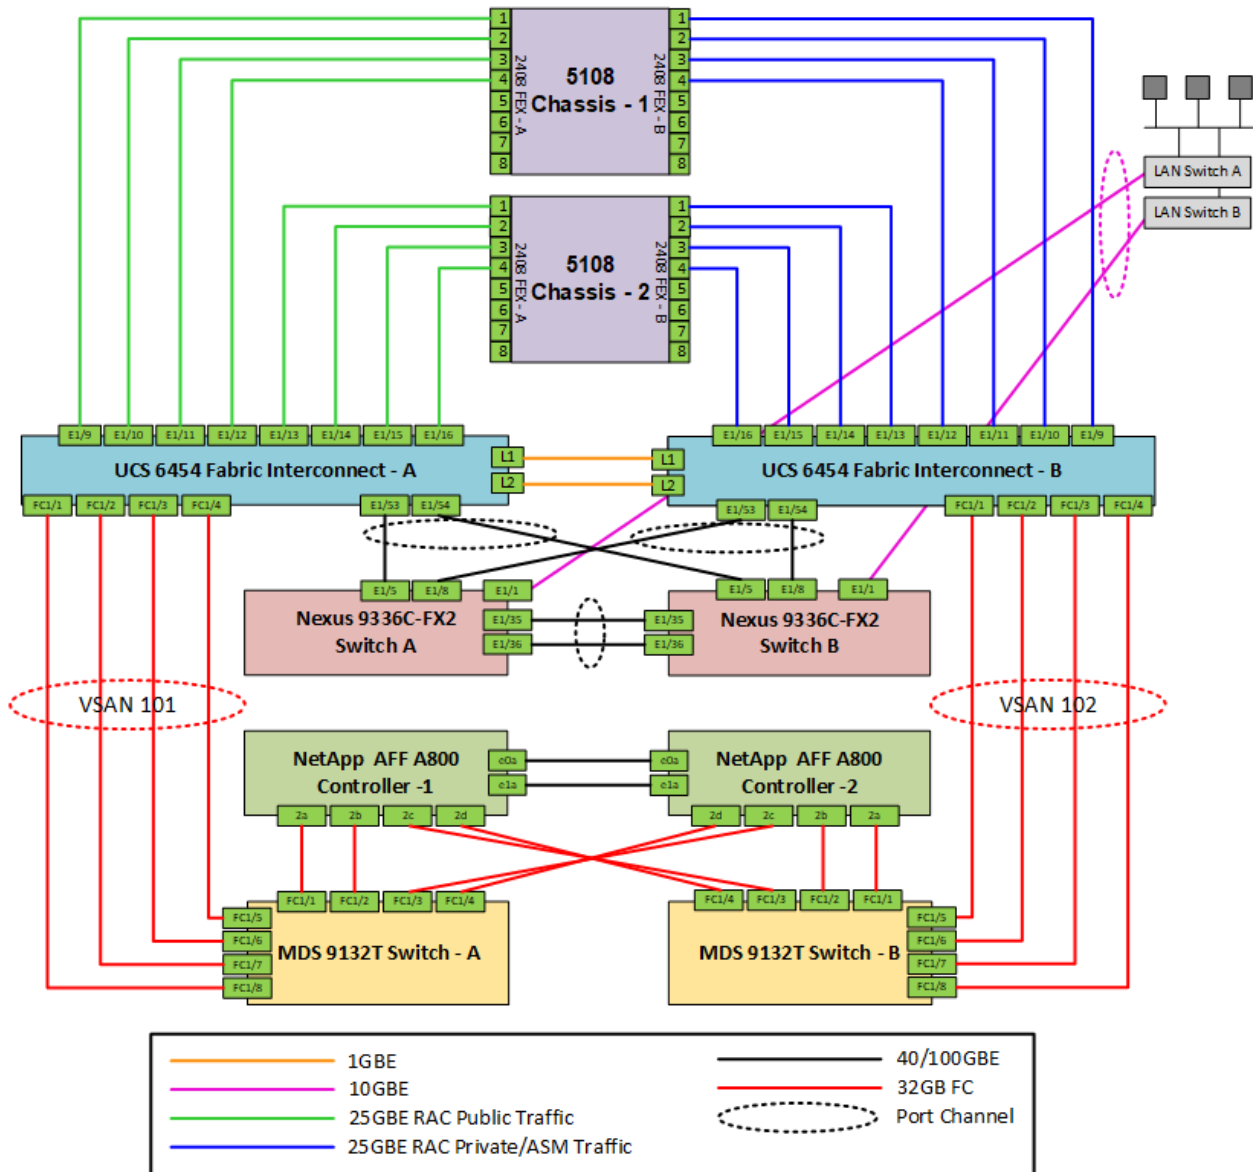
Cisco and NetApp have carefully validated and verified the FlexPod solution architecture and its many use cases while creating a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model.

Cisco and NetApp work closely with Oracle to support the demanding transactional and response-time-sensitive Oracle databases required by today's enterprise businesses.

Architectural diagram

Figure 1 shows the recommended infrastructure solution architecture for Oracle 19c on FlexPod Datacenter.

Figure 1) Recommended infrastructure architecture.



Hardware requirements

Table 1 lists the hardware components that are required to implement the solution. The hardware components that are used in any implementation of the solution might vary based on customer requirements.

Table 1) Hardware requirements.

Hardware	Quantity
Cisco UCS 5108 blade server chassis	2
Cisco UCS B200 M5 blade server with VIC 1440 2x Intel® Xeon® CPUs RAM, supported minimum balanced GB (192GB)	8
Cisco UCS IOM 2408	4

Hardware	Quantity
Cisco UCS 6454-16UP fabric interconnect	2
Cisco Nexus 9336C-FX2 switch in NX-OS standalone mode	2
Cisco MDS 9132T 32Gbps 32-port FC switch	2
Dual NetApp AFF A800 dual controllers with 24 1.92TB NVME SSD drives	2

Software requirements

Table 2 lists the software components that are required to implement the solution. The software components that are used in any implementation of the solution might vary based on customer requirements.

Table 2) Software requirements.

Software	Version or other information
Oracle Linux Oracle Linux	8.2 and RHCK kernel
Oracle 19c Enterprise Edition	19.3
Oracle 19c RU patch	19.8
NetApp ONTAP	9.7
Cisco UCS Manager	4.1(2b)
NetApp SnapCenter	4.4
Linux FIO	3.19
SLOB	2.5.2.4
Ansible	2.10.3

Design considerations

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. FlexPod components are connected and configured according to best practices of both Cisco and NetApp to provide the ideal platform for running a variety of enterprise workloads with confidence. This solution provides an end-to-end architecture with the Cisco Unified Computing System, Oracle, and NetApp technologies and demonstrates the FlexPod configuration benefits for running Oracle RAC databases 19c workloads with high availability and server redundancy.

The reference FlexPod architecture covered in this document is built on the NetApp All Flash AFF A800s for storage, Cisco B200 M5 blade servers for compute, Cisco Nexus 9336-FX2 switches, Cisco MDS 9132T 32G multilayer fabric switches and Cisco 6454 fabric interconnects for system management in a single package. The design is flexible enough that networking, computing, and storage can fit in one data center rack or can be deployed according to a customer's data center design.

The reference architecture reinforces the wire-once strategy, because, as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect. Additional blades can be added to UCS chassis to expand compute if needed. The UCS Manager server service profile can be easily applied to new blades to enable fast provision of new blades to predefined configuration.


Network design

This FlexPod for Oracle 19c solution networking design uses three pairs of Cisco switches for ethernet and FC communication with redundancy.

UCS 6454 Fabric Interconnect

The UCS 6454 Fabric Interconnect is the management and communication backbone for Cisco UCS B-Series blade servers and 5108 Series blade server chassis. The 6454 FI is a 10/25/40/100 Gigabit Ethernet, FCoE and FC switch offering up to 3.82Tbps throughput and up to 54 ports. The switch has 16 unified ports (port numbers 1-16) that can support 10/25Gbps SFP28 Ethernet ports or 8/16/32Gbps FC ports, 28 10/25Gbps Ethernet SFP28 ports (port numbers 17-44), 4 1/10/25Gbps Ethernet SFP28 ports (port numbers 45-48), and 6 40/100Gbps Ethernet QSFP28 uplink ports (port numbers 49-54). All Ethernet ports can support FCoE. Figure 2 and Table 3 outline UCS 6454 port assignment and the ports used for cabling with the validated solution.

Figure 2) UCS 6454 Fabric Interconnect.



At the front, the UCS 6454 Fabric Interconnect also has one network management port, one console port for setting up the initial configuration, and one USB port for saving or loading configurations. The fabric interconnect also includes L1/L2 ports in the front for connecting two fabric interconnects for high availability.

Table 3) UCS 6454 ports.

Connection	Port	Location
Uplink ports to Nexus switches	Port 53, 54 - 40/100 GBE	Rear
Connections to MDS 9132T	Port 1, 2, 3, 4 – 32G FC	Rear
Connections to UCS blades	Port 9, 10, 11, 12, 13, 14, 15, 16 – 25 GBE	Rear
HA links	Port L1, L2 – 1 GBE	front

Note: Refer to the architecture diagram for a detailed cabling setup.

Nexus 9336C-FX2 switch

The Cisco Nexus 9336C-FX2 switch provides user connections to applications hosted inside the FlexPod infrastructure. Nexus 9336C-FX2 (next figure) is a 1RU switch that supports 7.2Tbps of bandwidth and over 2.4bps. The switch can be configured to work as 1/10/25/40/50/100Gbps offering flexible options in a compact form factor with 36 ports. Breakout is supported on all ports. Table 4 shows design port connections for Nexus switches.

Figure 3) Nexus 9336C-FX2 switch.



Table 4) Nexus switch port connections.

Connection	Port	Location
FI uplink	Port 5, 8 – 40 GBE	Front
LAN connection	Port 1, breakout to 4/10GBE	Front
Peer link	Port 35, 36 – 40 GBE	Front

MDS 9132T FC switch

MDS 9132T 32Gbps 32-port FC switch provides high-speed FC connectivity from the UCS server rack to the SAN backend via the fabric interconnect. Figure 4 and Table 5 show the ports used for this solution design.

Figure 4) MDS 9132 FC switch.

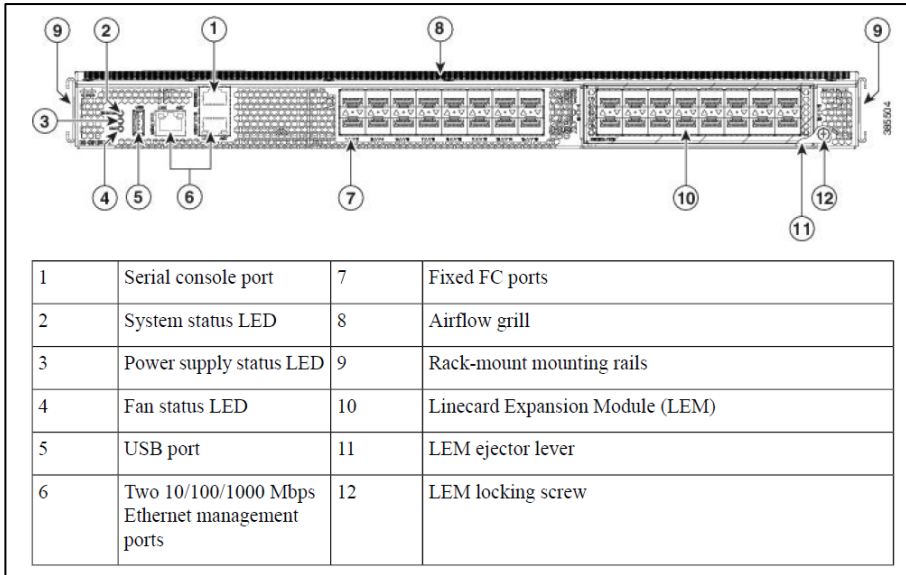


Table 5) MDS 9132 FC switch port connections.

Connection	Port	Location
Uplink to AFF A800 controller	Port 1, 2, 3, 4 – 32 GB FC	Front
Link to UCS blades via FI	Port 5, 6, 7, 8 – 32 GB FC	Front

VLAN and VSAN

To manage the traffic flow, we have set up a VLAN and a VSAN to segregate the infrastructure management, Oracle public and private traffic, and VSAN traffic to the AFF A800 storage into separate A and B fabric as listed in Table 6.

Table 6) VLAN and VSAN details.

Name	ID	Description
Default VLAN	1	Native VLAN
Management VLAN	184	Infrastructure management
Public VLAN	180	VLAN for Oracle public traffic
Private VLAN	3357	VLAN for Oracle private/asm traffic
VSAN – A	101	SAN communication for fabric interconnect A
VSAN – B	102	SAN communication for fabric interconnect B

Port channel

A virtual port channel or vPC provides HA, faster convergence in the event of a failure, and greater throughput with LACP. With vPC configuration, the Nexus switch pair are viewed by the connected devices as single connection domain and automatically failover to any surviving switch in the event of a switch failure.

Cisco Nexus vPC configurations with the vPC domains and corresponding vPC names and IDs for this Oracle 19c solution are shown in Table 7.

Table 7) vPC configuration.

vPC name	vPC ID	Connections	Ports
		Allowed VLAN	
Peer-link	po10	NX-ORA-01 to NX-ORA-02	NX-ORA-01 E1/35 to NX-ORA-02 E1/35
		1,180,184,3357	NX-ORA-01 E1/36 to NX-ORA-02 E1/36
FI-ORA-01	po15	FI-ORA-01 to NX-ORA-01/02	FI-ORA-01 E1/53 to NX-ORA-01 E1/5
		1,180,184,3357	FI-ORA-01 E1/54 to NX-ORA-02 E1/5
FI-ORA-02	po16	FI-ORA-02 to NX-ORA-01/02	FI-ORA-02 E1/53 to NX-ORA-01 E1/8
		1,180,184,3357	FI-ORA-02 E1/54 to NX-ORA-02 E1/8
Uplink-cie-c5672-g0845 LAN Switch	po100	g0845 to NX-ORA-01/02	g0845 E1/48 to NX-ORA-01 E1/1/1
		180,184	g0845 E1/47 to NX-ORA-02 E1/1/1

Port channels on the MDS switch are also configured to aggregate Oracle FC traffic from the fabric interconnect into VSAN 101 and VSAN 102 on MDS switch A and B. This is shown in Table 8.

Table 8) vPC FC traffic aggregation.

vPC Name	vPC ID	Connections	Ports
		Allowed VSAN	
UCS-FI-A	po101	FI-ORA-01 to MDS-ORA-01	FI-ORA-01 FC1/1, FC1/2, FC1/3, FC1/4 to MDS-ORA-01 FC1/5, FC1/6, FC1/7, FC1/8
		101	
UCS-FI-B	po102	FI-ORA-02 to MDS-ORA-02	FI-ORA-02 FC1/1, FC1/2, FC1/3, FC1/4 to MDS-ORA-02 FC1/5, FC1/6, FC1/7, FC1/8
		102	

Compute design

The compute design in this Oracle 19c solution includes two UCS 5108 chassis with each hosting four B200 M5 blade servers for an eight-node Oracle RAC cluster setup.

UCS 5180 chassis

The UCS 5108 chassis is a 6RU chassis that can accommodate up to eight half-width or four full-width blade form factor servers for easy and standardized deployment. Mixed generations of blades can be deployed in the same chassis. The chassis includes four onboard PDUs for redundant power supply.

Using two chassis for Oracle 19c solution provides redundancy in the event of failure of one chassis for mission critical Oracle RAC cluster deployment.

2408 fabric extender

The 5108 chassis has two I/O bays for fabric extenders or in-chassis fabric interconnects. For Oracle deployment, NetApp recommends to deploying a 2408 fabric extender for high throughput and expansion capability down the road.

The Oracle 19c solution deployment uses four ports per fabric extender, but 2408 I/O module provides eight external 25G GBE/FCoE ports or an aggregate of 200Gps total throughput per fabric extender.

B200 M5 blade server

The Cisco UCS B200 M5 blade server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for enterprise applications such as Oracle databases.

For this Oracle 19c solution, the M5 B200 blade servers are deployed with Intel Xeon Gold 5218 CPU with 16 cores and a 2.3 GHZ clock speed. Each blade server is configured with a minimum of 192G RAM, four HBAs, and two ethernet connections with Cisco VIC 1440.

CPU options


The processing capabilities of CPUs have increased much faster than the processing demands of most database workloads. Sometimes databases are limited by CPU work, but this is generally a result of the processing limits of a single core and not a limitation of the CPU. The result is an increasing number of idle cores on database servers that must still be licensed for Oracle Database software. This underutilization of CPU resources is a waste of capital expenditure, not only in terms of licensing costs, but also in terms of the cost of the server itself, heat output, and so on. Cisco UCS servers come with [different CPU options](#) in terms of higher clock-speed that can assist with database workloads. Therefore, customers have the option of using higher clock-speed CPUs with fewer cores to keep Oracle licensing costs down.

Storage design

The solution is powered by the top-of-the-line ONTAP all-flash storage array AFF A800 from NetApp. The AFF A800 is designed for workloads that demand the most performance such as Oracle database. With ONTAP 9.7 and FC, an HA pair of A800 controllers delivers over one million IOPS at under 500 microseconds latency.

For the Oracle 19c solution, the AFF A800 controllers are loaded with 24 1.92TB internal all flash SSD drives and set up as a cluster-less HA pair. Figure 5 shows the ports on a single AFF A800 controller and available expansion slots.

Figure 5) AFF A800 ports and available expansion slots.



In addition to management ports such as the serial console port, e0M, and BMC ports, e0a and e0b are onboard HA ports designed for storage cluster high availability connectivity. There are a total of five expansion slots for adding ethernet or FC ports on each AFF A800 controller. In this solution, we added an FC card with four FC ports (2a, 2b, 2c, and 2d) in slot 1 for FC connectivity to MDS. Table 9 shows the connection ports on an AFF A800 HA controller pair.

Table 9) Connection ports for an AFF A800 HA controller pair.

Connection	Port	Location
Terminal switch term1g2	Serial console port	Rear
Management switch g0246	e0M	Rear
Management switch g0246	BMC	Rear
FlexPod-A800-01 to FlexPod-A800-02 HA	e0a to e0a, e0b to e0b – 100GBE	Rear
FlexPod-A800-01 to MDS-ORA-01 FC	2a, 2b – 32G FC	Rear
FlexPod-A800-01 to MDS-ORA-02 FC	2c, 2d – 32G FC	Rear
FlexPod-A800-02 to MDS-ORA-01 FC	2c, 2d – 32G FC	Rear
FlexPod-A800-02 to MDS-ORA-02 FC	2a, 2b – 32G FC	Rear

Oracle 19c database deployment considerations

The goal of the solution is to demonstrate Oracle 19c database deployment in a CDB/PDB model in an Oracle RAC configuration. The benefits of the CDB/PDB deployment model allow enterprises to consolidate their database and application workload to reduce Oracle licensing and operational cost on a bare-metal infrastructure.

See the following considerations for Oracle 19c database and application deployment on FlexPod.

The type of database workload

Determining the type of Oracle database workload is an important first step. OLTP and DSS are common Oracle database workload types. Some workloads might include both. OLTP and DSS workloads have very different characteristics, so they should be grouped and deployed into different Oracle 19c database containers and ideally in different RAC clusters.

In addition to workload grouping, storage configuration should be set differently based on the workload type. Storage I/O performance generally plays an important role in Oracle database application performance. OLTP workload I/O characteristics typically include single-block random I/O, and random I/O latency is a primary performance indicator. On the other hand, DSS workload I/O is characterized by multiblock sequential I/O, and I/O throughput is a key performance indicator. Therefore, storage evaluation and configuration optimization should be considered accordingly.

A NetApp AFF A800 storage array is designed to support both OLTP, DSS, or mixed workloads.

Sizing Oracle RAC cluster

About 80% of Oracle RAC clusters deployed in production are two-node because of simplicity and HA capability, while achieving the highest cache fusion efficiency. As the number of RAC nodes increases, the cache fusion efficiency tends to go down, and contention can become an issue. If, however, the application code is written to be RAC aware, then contention is mitigated at the application level. A larger RAC cluster is less likely to affect application performance in this scenario.

Besides contention, the other factors to consider include the number of applications, the transaction volume, the size of the database and of course the availability of hardware resources in terms of compute cycles and memory on a cluster node.

It makes sense to set up multiple two-node RAC clusters to support many smaller applications with smaller databases. You might also need to build a larger cluster to support a single application with larger data sets and a large number of concurrent users that require more compute for the required transaction volumes.

In summary, the sizing of an Oracle RAC cluster depends on the workload requirements that the RAC cluster is built to support.

Database and application consolidation

For customers who struggle with database sprawl, FlexPod offers an excellent platform for consolidation. The future of Oracle database deployment is the CDB/PDB model as signified by Oracle phasing out standard single instances from release 20. CDB/PDB provides an option for resources sharing and database consolidation without relying on virtualization, which adds on overhead as well as additional virtualization licensing costs. You can, in theory, consolidate as many as 252 PDBs in a container with the enterprise edition.

The CDB/PDB model is an Oracle licensing option. The good news is that from Oracle 19c, Oracle allows up to three PDBs in a CDB container free of multitenant license cost. This could be good enough for many organizations to consolidate multiple applications into fewer container databases, with each hosting three PDBs in a container. Therefore, you can achieve a degree of resource sharing and application consolidation while enjoying free license and complying with Oracle multitenancy requirement going into Oracle version 21.

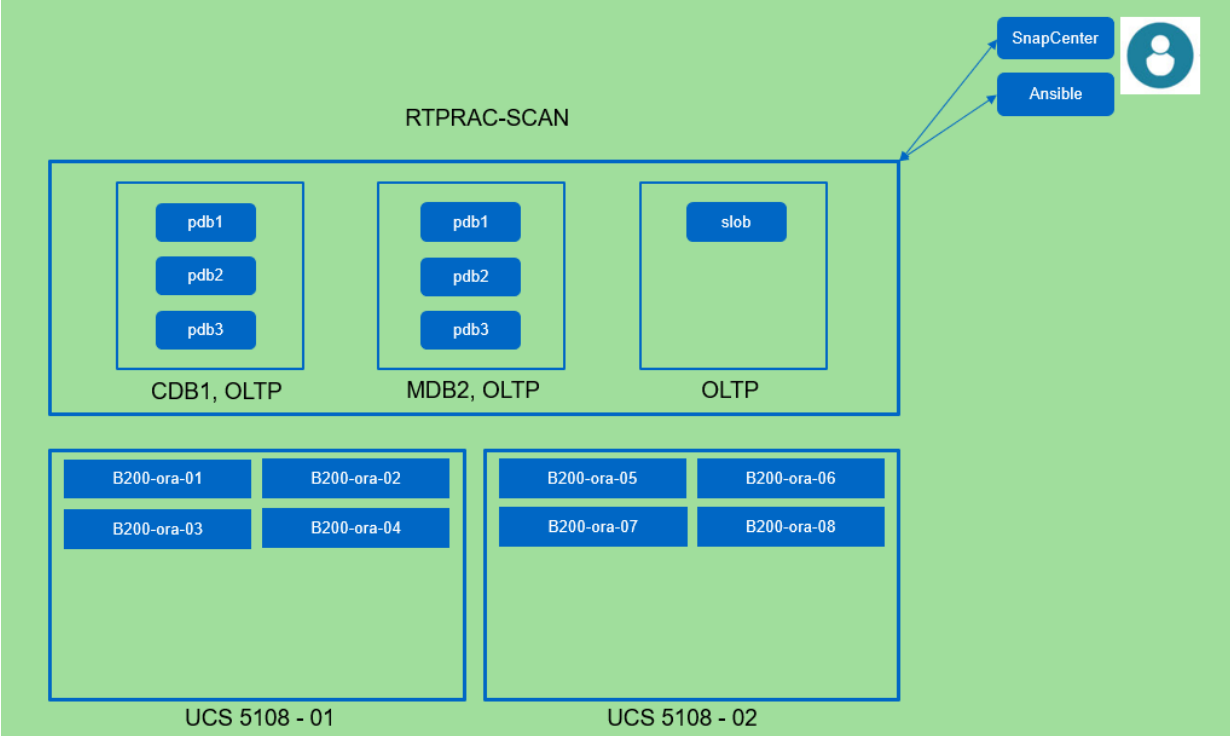
For further applications consolidation within a PDB, multiple applications can be hosted within a subset of a pluggable database in isolated tablespaces inside a pluggable database.

In this solution, we have built single RAC cluster with eight RAC cluster nodes. Figure 6 demonstrates solution deployment in the CDB/PDB model in an RAC configuration.

The goal of the solution deployment demonstration is not to cover all possible deployment scenarios but rather to provide a starting point that customer can use to customize to their own systems.

To that purpose, we have created two container databases or CDBs and within each CDB that host up to three pluggable databases or PDBs. We also created a SLOB database for testing a SLOB benchmark workload.

Figure 6) Solution deployment in the CDB/PDB model in an RAC configuration.



Storage layout

In this eight-node 19c RAC cluster, we created 60 LUNs to build three Oracle ASM disk groups to maximize storage performance for a mission critical application. Table 10 shows the details of the storage layout for the Oracle 19c RAC cluster.

Table 10) Oracle 19c RAC cluster storage layout.

Node1 - FlexPod-A800-01-02-01					Node2 - FlexPod-A800-01-02-02							Total, GB
Aggr - FlexPod_A800_01_02_01_NVME_SSD_1 RAID_DP, 16.28 TB					Aggr - FlexPod_A800_01_02_02_NVME_SSD_1 RAID_DP, 16.28 TB							
ASM	Vol	Size, GB	LUN	Size, GB	Note	Vol	Size, GB	LUN	Size, GB	Note	SVM	
DATA+	ora_data_01	2000	ora_data_01_1	400	Data file	ora_data_02	2000	ora_data_02_1	400	Data file	ora19c_svm	19200
			ora_data_01_2	400	Data file			ora_data_02_2	400	Data file		
			ora_data_01_3	400	Data file			ora_data_02_3	400	Data file		
			ora_data_01_4	400	Data file			ora_data_02_4	400	Data file		
	ora_data_03	2000	ora_data_03_1	400	Data file	ora_data_04	2000	ora_data_04_1	400	Data file		
			ora_data_03_2	400	Data file			ora_data_04_2	400	Data file		
			ora_data_03_3	400	Data file			ora_data_04_3	400	Data file		
			ora_data_03_4	400	Data file			ora_data_04_4	400	Data file		

	ora_data_05	2000	ora_data_05_1	400	Data file	ora_data_06	2000	ora_data_06_1	400	Data file							
			ora_data_05_2	400	Data file			ora_data_06_2	400	Data file							
			ora_data_05_3	400	Data file			ora_data_06_3	400	Data file							
			ora_data_05_4	400	Data file			ora_data_06_4	400	Data file							
	ora_data_07	2000	ora_data_07_1	400	Data file	ora_data_08	2000	ora_data_08_1	400	Data file							
			ora_data_07_2	400	Data file			ora_data_08_2	400	Data file							
			ora_data_07_3	400	Data file			ora_data_08_3	400	Data file							
			ora_data_07_4	400	Data file			ora_data_08_4	400	Data file							
	ora_data_09	2000	ora_data_09_1	400	Data file	ora_data_10	2000	ora_data_10_1	400	Data file							
			ora_data_09_2	400	Data file			ora_data_10_2	400	Data file							
			ora_data_09_3	400	Data file			ora_data_10_3	400	Data file							
			ora_data_09_4	400	Data file			ora_data_10_4	400	Data file							
	ora_data_11	2000	ora_data_11_1	400	Data file	ora_data_12	2000	ora_data_12_1	400	Data file							
			ora_data_11_2	400	Data file			ora_data_12_2	400	Data file							
			ora_data_11_3	400	Data file			ora_data_12_3	400	Data file							
			ora_data_11_4	400	Data file			ora_data_12_4	400	Data file							
	REDO+	ora_redo_01	1000	ora_redo_01_1	100	Redo	ora_redo_02	1000	ora_redo_02_1	100			Redo, recovery		800		
																Redo, recovery	
		ora_redo_03	1000	ora_redo_03_1	100	Redo	ora_redo_04	1000	ora_redo_04_1	100			Redo, recovery				200
VOTE+	ora_crs_01	200	ora_crs_01_1	50	OCR, vote	ora_crs_02	200	ora_crs_02_1	50	OCR, vote		1200					
			ora_crs_01_2	50	OCR, vote			ora_crs_02_2	50	OCR, vote							
Binary	ora_bin_01	800	ora_bin_01_1	150	Binary	vol_bin_02	800	ora_bin_02_1	150	Oracle binary							
			ora_bin_01_2	150	Binary			ora_bin_02_2	150	Oracle binary							
			ora_bin_01_3	150	Binary			ora_bin_02_3	150	Oracle binary							
			ora_bin_01_4	150	Binary			ora_bin_02_4	150	Oracle binary							
Boot	ora_boot_01	1000	ora_b200_01_1	200	OS	ora_boot_02	1000	ora_b200_02_1	200	OS	infra_svm	1600					
			ora_b200_01_2	200	OS			ora_b200_02_2	200	OS							
			ora_b200_01_3	200	OS			ora_b200_02_3	200	OS							
			ora_b200_01_4	200	OS			ora_b200_02_4	200	OS							
Total, TB		16	11.5			16		11.5			23						

Note: This storage layout is optimized to support an application with a large data set and high concurrent application users. See the best practices for tips for configuring the database storage layout.

Oracle data protection and management

Protecting enterprise data is paramount for any organization. Among the many data protection options, NetApp is a pioneer and proponent for using storage-level snapshots for database protection.

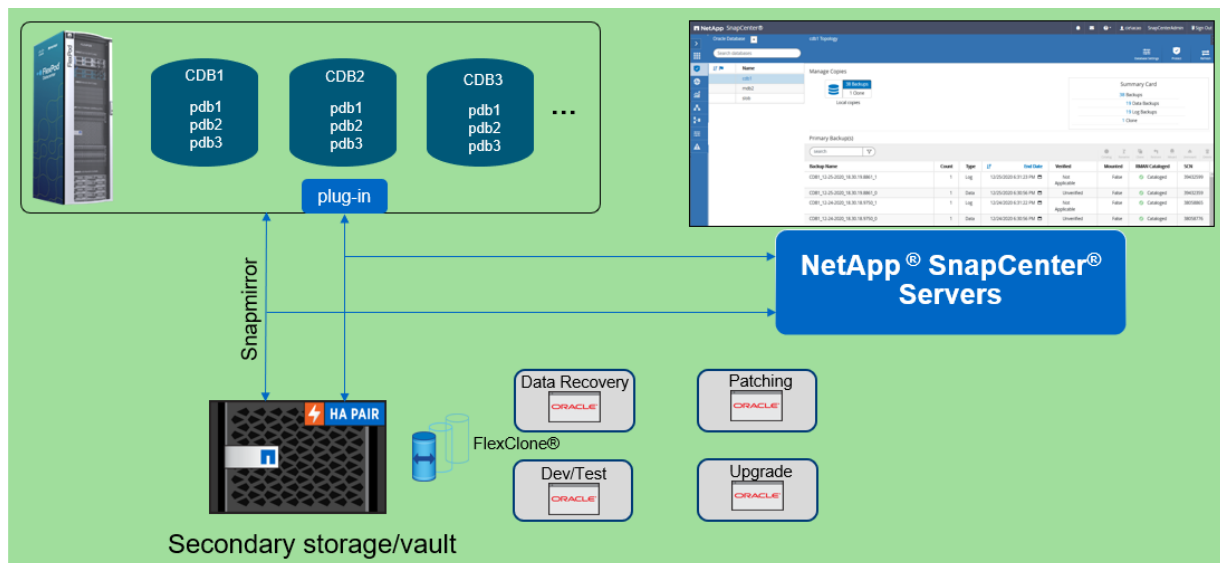
NetApp SnapCenter software is a UI tool used to deliver storage-level data protection functions. It is a proven, simple, centralized, scalable platform that provides application-consistent protection for databases running on ONTAP systems anywhere in the hybrid cloud.

SnapCenter uses NetApp Snapshot, SnapRestore®, FlexClone®, SnapMirror®, and SnapVault® technologies to provide the following benefits:

- Fast, space-efficient, application-consistent, disk-based backups
- Rapid, granular restore and application-consistent recovery
- Quick, space-efficient cloning

In release 4.4, SnapCenter incorporated many new features tailored for protecting Oracle databases deployed with the CDB/PDB model. We demonstrate many of those data protection capability in the section “SnapCenter deployment.” Figure 7 illustrates the SnapCenter architecture for Oracle data protection and management at a high level.

Figure 7) SnapCenter architecture for Oracle data protection and management.



Best practices

- Considerations for configuring FC SAN:
 - In most cases, NetApp recommends using between 8 and 16 volumes to maximize performance. This is assuming that these are the only volumes on a given controller. If other volumes are present, then you can generally consider using up to eight additional LUNs to maximize performance.
 - LUNs that are related to each other and have similar performance and management requirements can be hosted by a single volume for the benefits of reduced administrative complexity, being in a common consistency group, and storage efficiency.
 - NetApp recommends using more smaller LUNs versus using fewer bigger ones, ideally, between 8 and 16 LUNs.
 - NPIV is required for FC LIFs to operate correctly. Before creating FC LIFs, make sure that any fabrics attached to an ONTAP system have NPIV enabled.

- When creating FC LIFs for the first time for an existing storage virtual machine (SVM), make sure that the FC for that SVM has been created and is turned on by using the `fc show` command.
- Host integration:
 - Installation of the NetApp Host Utilities Kit sets timeout and other operating system–specific values to their recommended defaults and includes utilities for examining LUNs provided by NetApp storage.
- Consideration for ASM disk groups:
 - In a multitenant CDB/PDB deployment, set up separate ASM disk groups for different container databases with dedicated volumes. As such, ONTAP QoS can be implemented at the CDB level if needed. SnapCenter backup and recovery can also be isolated to the CDB-volumes level.

Deploying Oracle 19c solution

In this section, detailed step-by-step procedures are demonstrated to meet design objectives as outlined in previous section with primarily manual steps. Selective manual steps can also be accomplished through automation playbook in the optional section “Automated Deployment.”

Manual deployment

Nexus switch deployment

Initial configuration

To set up the initial configuration for the Cisco Nexus A switch on NX-ORA-01, complete the following steps:

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

```

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: NX-ORA-01
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: 10.61.184.188
Mgmt0 IPv4 netmask: 255.255.255.0
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: 10.61.184.1
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: 10.61.185.177
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [noshut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

Review the configuration summary before enabling the configuration.

```

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

```

Complete identical steps for Nexus switch B on NX-ORA-02.

Configure global settings for Cisco Nexus A and Cisco Nexus B

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

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

```
[admin@mgmt oracle]$ ssh 10.61.184.188
NX-ORA-01# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
NX-ORA-01(config)# feature interface-vlan
NX-ORA-01(config)# feature udld
NX-ORA-01(config)# feature hsrp
NX-ORA-01(config)# feature lacp
NX-ORA-01(config)# feature vpc
NX-ORA-01(config)# spanning-tree port type network default
NX-ORA-01(config)# spanning-tree port type edge bpduguard default
NX-ORA-01(config)# port-channel load-balance src-dst l4port
NX-ORA-01(config)# policy-map type network-qos jumbo
NX-ORA-01(config-pmap-nqos)# class type network-qos class-default
NX-ORA-01(config-pmap-nqos-c)# mtu 9216
NX-ORA-01(config-pmap-nqos-c)# exit
NX-ORA-01(config-pmap-nqos)# exit
NX-ORA-01(config)# system qos
NX-ORA-01(config-sys-qos)# service-policy type network-qos jumbo
NX-ORA-01(config-sys-qos)# exit
NX-ORA-01(config)# vrf context management
NX-ORA-01(config-vrf)# ip route 0.0.0.0/0 10.61.184.1
NX-ORA-01(config-vrf)# exit
NX-ORA-01(config)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```

Complete same procedure on Switch B.

Configure VLANs for Cisco Nexus A and Cisco Nexus B switches

To create the necessary VLANs, complete the follow steps on both Nexus switches:

1. Log in as admin user into the Nexus Switch A. Create VLAN 180 for Public Network Traffic and VLAN 3357 for Private Network Traffic, and VLAN 184 for management traffic.

```
NX-ORA-01# configure terminal
NX-ORA-01(config)# vlan 180
NX-ORA-01(config-vlan)# name ora_pub
NX-ORA-01(config-vlan)# no shutdown
NX-ORA-01(config-vlan)# exit
NX-ORA-01(config)# vlan 3357
NX-ORA-01(config-vlan)# name ora_pri
NX-ORA-01(config-vlan)# no shutdown
NX-ORA-01(config-vlan)# exit
NX-ORA-01(config)# vlan 184
NX-ORA-01(config-vlan)# name mgmt
NX-ORA-01(config-vlan)# no shutdown
NX-ORA-01(config-vlan)# exit
NX-ORA-01(config)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```

2. Log in as admin user into the Nexus Switch B to complete the same VLAN setup.

Create virtual port channel for data network

In the Cisco Nexus 9336C-FX2 switch topology, a single virtual port channel (vPC) feature is enabled to provide HA, faster convergence in the event of a failure, and greater throughput. Cisco Nexus vPC

configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers is shown in the section “Network Design.”

We created three virtual port channels on the Nexus switches. vPC ID 10 is defined as peer link communication between the two Nexus switches in Fabric A and B. vPC IDs 15 and 16 are for public and private network traffic from the Cisco UCS fabric interconnects. To setup the vPCs, complete the steps in the following sections.

1. Create vPC peer link between the two Nexus switches.
2. Log in as the admin user for Nexus Switches A and B and create the peer-link port channels between devices on both Nexus switches:

On Switch A:

```
NX-ORA-01# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
NX-ORA-01(config)# vpc domain 1
NX-ORA-01(config-vpc-domain)# peer-keepalive destination 10.61.184.189 source 10.61.184.188
NX-ORA-01(config-vpc-domain)# role priority 10
NX-ORA-01(config-vpc-domain)# peer-switch
NX-ORA-01(config-vpc-domain)# delay restore 150
NX-ORA-01(config-vpc-domain)# peer-gateway
NX-ORA-01(config-vpc-domain)# ip arp synchronize
NX-ORA-01(config-vpc-domain)# auto-recovery
NX-ORA-01(config-vpc-domain)# system-priority 1
NX-ORA-01(config-vpc-domain)# interface port-channel10
NX-ORA-01(config-if)# description vPC peer-link
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# spanning-tree port type network
NX-ORA-01(config-if)# vpc peer-link
NX-ORA-01(config-if)# interface Ethernet1/35
NX-ORA-01(config-if)# description Peer link to NX-ORA-02-eth1/35
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# channel-group 10 mode active
NX-ORA-01(config-if)# no shutdown
NX-ORA-01(config-if)# interface Ethernet1/36
NX-ORA-01(config-if)# description Peer link to NX-ORA-02-eth1/36
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# channel-group 10 mode active
NX-ORA-01(config-if)# no shutdown
NX-ORA-01(config-if)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```

Note: Reverse the keepalive destination IPs for Switch B (NX-ORA-02) accordingly and set role priority 20.

3. Create the vPC for HA connections between the Nexus switches and the fabric interconnect.

```
NX-ORA-01# configure
Enter configuration commands, one per line. End with CNTL/Z.
NX-ORA-01(config-if)# interface port-channel15
NX-ORA-01(config-if)# description Port-Channel FI-ORA-1
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# spanning-tree port type edge trunk
Edge port type (portfast) should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when edge port type (portfast) is enabled, can cause temporary bridging loops.
Use with CAUTION

NX-ORA-01(config-if)# mtu 9216
NX-ORA-01(config-if)# vpc 15
NX-ORA-01(config-if)# no shutdown
```



```

NX-ORA-01(config-if)# interface port-channel16
NX-ORA-01(config-if)# description Port-Channel FI-ORA-2
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# spanning-tree port type edge trunk
NX-ORA-01(config-if)# mtu 9216
NX-ORA-01(config-if)# vpc 16
NX-ORA-01(config-if)# no shutdown

NX-ORA-01(config)# interface Ethernet1/5
NX-ORA-01(config-if)# description FI-ORA-1 e1/53
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# spanning-tree port type edge trunk
NX-ORA-01(config-if)# mtu 9216
NX-ORA-01(config-if)# channel-group 15 mode active
NX-ORA-01(config-if)# no shutdown

NX-ORA-01(config-if)# interface Ethernet1/8
NX-ORA-01(config-if)# description FI-ORA-2 e1/53
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-01(config-if)# spanning-tree port type edge trunk
NX-ORA-01(config-if)# mtu 9216
NX-ORA-01(config-if)# channel-group 16 mode active
NX-ORA-01(config-if)# no shutdown

NX-ORA-01(config-if)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.

```

On Nexus Switch B:

```

NX-ORA-02# configure
Enter configuration commands, one per line. End with CNTL/Z.
NX-ORA-02(config-if)# interface port-channel15
NX-ORA-02(config-if)# description Port-Channel FI-ORA-1
NX-ORA-02(config-if)# switchport mode trunk
NX-ORA-02(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-02(config-if)# spanning-tree port type edge trunk
Edge port type (portfast) should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when edge port type (portfast) is enabled, can cause temporary bridging loops.
Use with CAUTION

NX-ORA-02(config-if)# mtu 9216
NX-ORA-02(config-if)# vpc 15
NX-ORA-02(config-if)# no shutdown

NX-ORA-02(config-if)# interface port-channel16
NX-ORA-02(config-if)# description Port-Channel FI-ORA-2
NX-ORA-02(config-if)# switchport mode trunk
NX-ORA-02(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-02(config-if)# spanning-tree port type edge trunk
NX-ORA-02(config-if)# mtu 9216
NX-ORA-02(config-if)# vpc 16
NX-ORA-02(config-if)# no shutdown

NX-ORA-02(config)# interface Ethernet1/5
NX-ORA-02(config-if)# description FI-ORA-1 e1/54
NX-ORA-02(config-if)# switchport mode trunk
NX-ORA-02(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-02(config-if)# spanning-tree port type edge trunk
NX-ORA-02(config-if)# mtu 9216
NX-ORA-02(config-if)# channel-group 15 mode active
NX-ORA-02(config-if)# no shutdown

NX-ORA-02(config-if)# interface Ethernet1/8
NX-ORA-02(config-if)# description FI-ORA-2 e1/54

```

```

NX-ORA-02(config-if)# switchport mode trunk
NX-ORA-02(config-if)# switchport trunk allowed vlan 1, 180, 184, 3357
NX-ORA-02(config-if)# spanning-tree port type edge trunk
NX-ORA-02(config-if)# mtu 9216
NX-ORA-02(config-if)# channel-group 16 mode active

NX-ORA-02(config-if)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.

```

Note: The spanning-tree port type edge trunk command triggers the following warning, which can be ignored:

```

Edge port type (portfast) should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface when edge port type (portfast) is enabled, can cause temporary bridging loops.
Use with CAUTION

```

4. Configure connectivity to the LAN switch. You need to first break out a QSFP28 port on Nexus switch to setup 10G connections to local LAN switch for Oracle environment access.

- a. Port breakout.

```

NX-ORA-01# configure
Enter configuration commands, one per line. End with CNTL/Z.
NX-ORA-01(config-if)# interface breakout module 1 port 1 map 10g-4x
NX-ORA-01(config-if)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.

```

- b. Configure port channel 100 for LAN access.

```

NX-ORA-01(config-if)# interface port-channel100
NX-ORA-01(config-if)# description Uplink to uplink cie-c5672-g0845
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk native vlan 2
NX-ORA-01(config-if)# switchport trunk allowed vlan 180, 184
NX-ORA-01(config-if)# spanning-tree port type normal
NX-ORA-01(config-if)# mtu 9216
NX-ORA-01(config-if)# vpc 100
NX-ORA-01(config-if)# no shutdown
NX-ORA-01(config-if)# interface Ethernet1/1/1
NX-ORA-01(config-if)# description connect to uplink cie-c5672-g0845 eth1/48
NX-ORA-01(config-if)# switchport mode trunk
NX-ORA-01(config-if)# switchport trunk native vlan 2
NX-ORA-01(config-if)# switchport trunk allowed vlan 180,184
NX-ORA-01(config-if)# mtu 9216
NX-ORA-01(config-if)# channel-group 100 mode active
NX-ORA-01(config-if)# copy run start
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.

```

Complete the same procedure to create port-channel 100 for connection to g0845 port eth1/47 on Nexus Switch B.

5. Validate the vPC configuration.
6. After vPC configuration, verify the configuration on both switches to make sure all vPCs are in SU status.

```

NX-ORA-01# show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        s - Suspended     r - Module-removed
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lACP mode (member)
        M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
10    Po10(SU)    Eth       LACP      Eth1/35(P)  Eth1/36(P)
15    Po15(SU)    Eth       LACP      Eth1/5(P)
16    Po16(SU)    Eth       LACP      Eth1/8(P)
100   Po100(SU)    Eth       LACP      Eth1/1/1(P)
NX-ORA-01#

```

```

NX-ORA-02# show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        s - Suspended     r - Module-removed
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lACP mode (member)
        M - Not in use. Min-links not met
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
10    Po10(SU)    Eth       LACP      Eth1/35(P)  Eth1/36(P)
15    Po15(SU)    Eth       LACP      Eth1/5(P)
16    Po16(SU)    Eth       LACP      Eth1/8(P)
100   Po100(SU)    Eth       LACP      Eth1/1/1(P)
NX-ORA-02#

```

Cisco UCS setup

Initial UCS switch setup using console

To configure the Cisco UCS for use in a FlexPod environment in UCSM managed mode, follow these steps.

1. Connect to the console port on the first Cisco UCS fabric interconnect.

```

Enter the configuration method. (console/gui) ? console
Enter the management mode. (ucsm/intersight)? ucsm
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
You have chosen to setup a new Fabric interconnect in "ucsm" managed mode. Continue? (y/n): y
Enforce strong password? (y/n) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
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: <ucs-cluster-name>
Physical Switch Mgmt0 IP address : <ucsa-mgmt-ip>
Physical Switch Mgmt0 IPv4 netmask : <ucsa-mgmt-mask>
IPv4 address of the default gateway : <ucsa-mgmt-gateway>
Cluster IPv4 address : <ucs-cluster-ip>
Configure the DNS Server IP address? (yes/no) [n]: y
DNS IP address : <dns-server-1-ip>
Configure the default domain name? (yes/no) [n]: y
Default domain name : <ad-dns-domain-name>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

2. Wait for the login prompt for UCS Fabric Interconnect A before proceeding to the next section.
3. Connect to the console port on the second Cisco UCS 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: <password>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucsa-mgmt-mask>
Cluster IPv4 address : <ucs-cluster-ip>
Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address
Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>
Local fabric interconnect model(UCS-FI-6454)
Peer fabric interconnect is compatible with the local fabric interconnect. Continuing with the
installer...
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

4. Wait for the login prompt for UCS Fabric Interconnect B before proceeding to the next section.

Log in to Cisco UCS Manager

1. Open a web browser and navigate to the Cisco UCS fabric interconnect cluster address.
2. Click the Launch UCS Manager link to launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the username and enter the administrative password.
5. Click Login to log into Cisco UCS Manager.

Note: You might need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to open. When logging into UCS manager for the first time, you might be prompted to enable anonymous reporting, which send anonymous data to Cisco for improving future products. If you choose yes, enter the IP address of your SMTP Server. Click OK.

Upgrade Cisco UCS Manager Software to Version 4.1(2b)

This document assumes the use of Cisco UCS 4.1(2b). To upgrade the Cisco UCS Manager software and the Cisco UCS fabric interconnect software to version 4.1(2b), see the [Cisco UCS Manager Install and Upgrade Guides](#).

Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in Cisco UCS Manager. Configuring Call Home accelerates the resolution of support cases. To configure Call Home, complete the following steps:

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

Set fabric interconnects to FC End Host mode

To set the fabric interconnects to the FC End Host mode, complete the following steps:

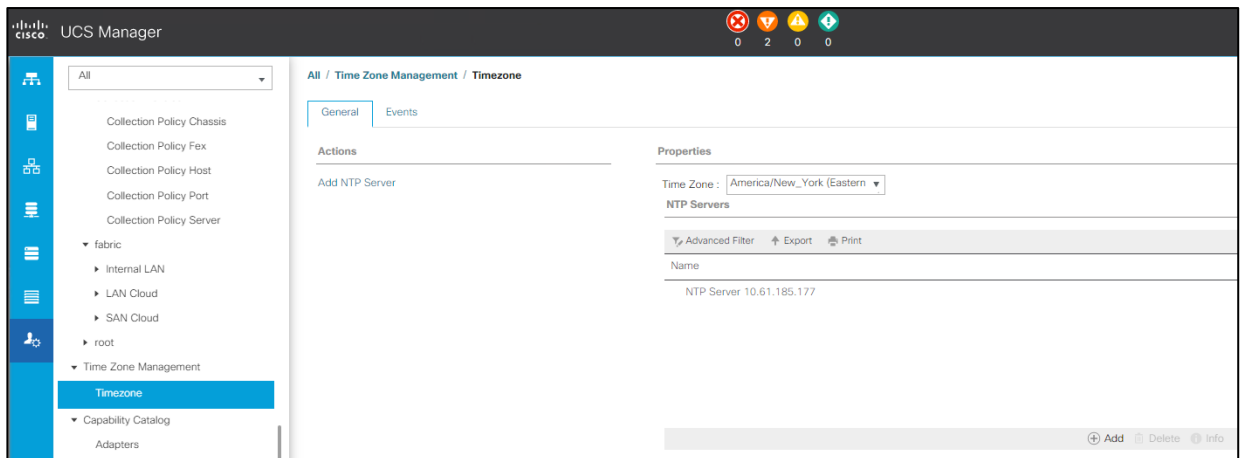
1. On the Equipment tab, expand the fabric interconnects node and click Fabric Interconnect A.
2. On the General tab in the Actions pane, click Set FC End Host mode.
3. Follow the dialogs to complete the change.

Note: Both fabric interconnects automatically reboot sequentially when you confirm you want to operate in this mode.

Synchronize UCS to NTP

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

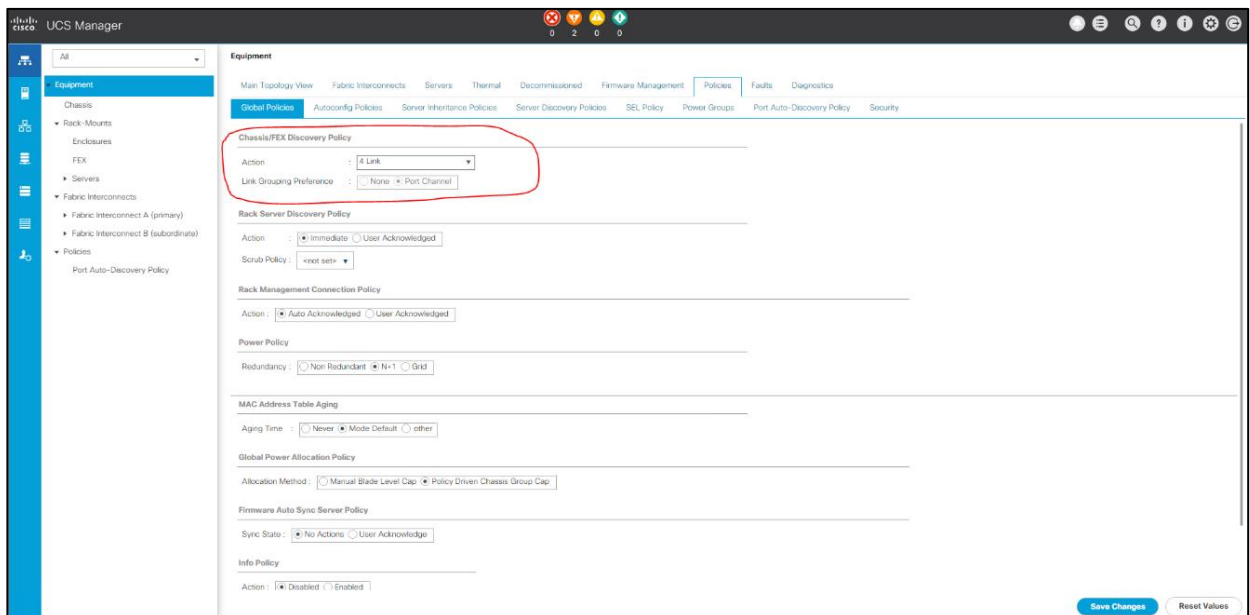
1. In Cisco UCS Manager in the navigation pane, click the Admin tab.
2. Select All > Time zone Management.
3. In the Properties pane, select the appropriate time zone in the Time zone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.
6. Enter the NTP server IP address and click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



Configure chassis discovery policy

To configure global policies, complete the following steps:

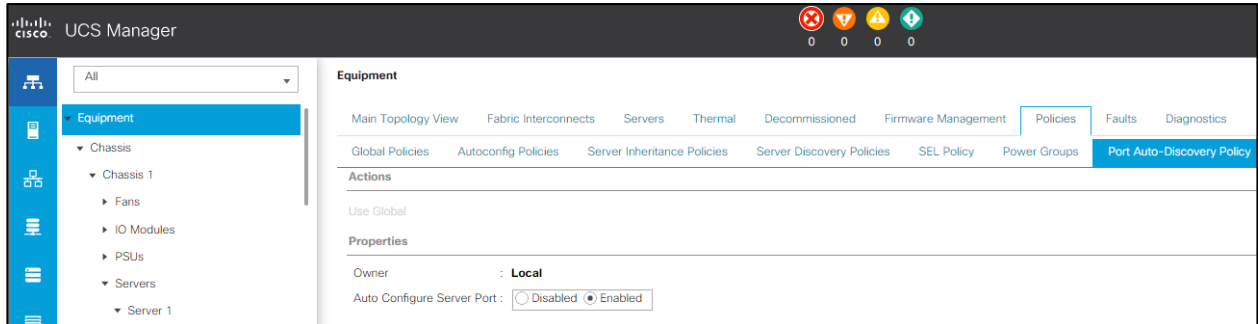
1. Go to Equipment > Policies (right pane) > Global Policies > Chassis/FEX Discovery Policies.
2. Select Action as 4 Links from the drop-down list and set Link Grouping to Port Channel as shown below. Click Save Changes and then click OK.



Enable port auto-discovery policy

Setting the port auto-discovery policy enables automatic discovery of Cisco UCS B-Series chassis server ports. To modify the port auto-discovery policy, complete the following steps:

1. In Cisco UCS Manager, click Equipment, choose All > Equipment in the Navigation Pane, and choose the Policies tab.
2. Under Port Auto-Discovery Policy, set Auto Configure Server Port to Enabled.

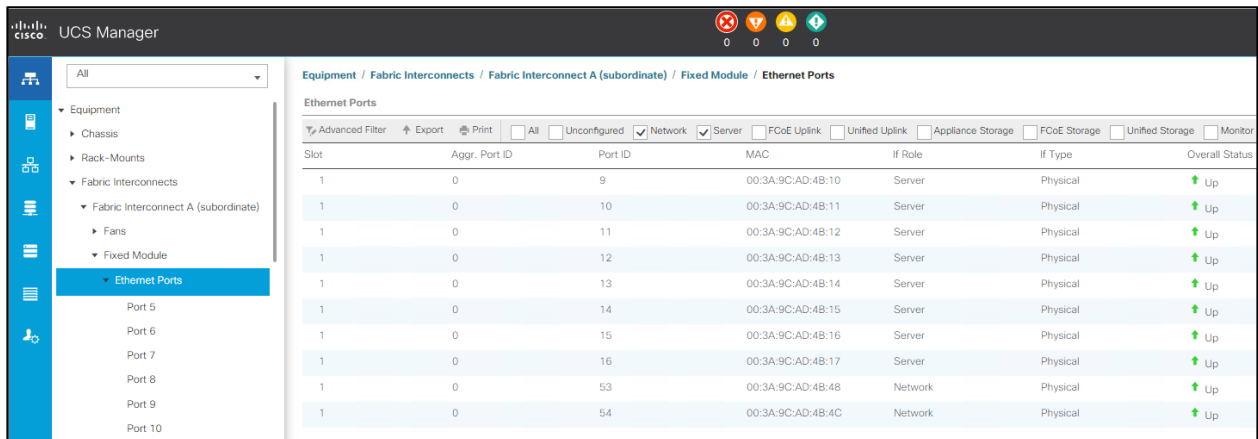


3. Click Save Changes and then OK.

Enable server and uplink ports

Enable ports 9 through 16 as server ports for connections to the UCS 5108 chassis. Enable port 53 and port 54 as the uplink ports to the Nexus switches. Complete the following steps:

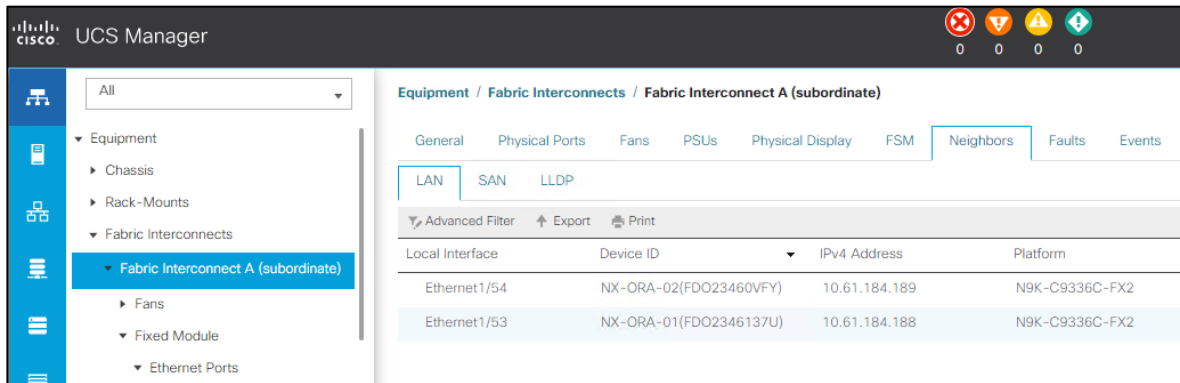
1. In Cisco UCS Manager, click Equipment.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand and choose Ethernet Ports.
4. Right click on the ethernet port and set port 9 through 16 as server ports and ports 53 and 54 as uplink ports.
5. Verify that all ports connected to the UCS chassis are configured as server ports and have a status of up. Ports 53 and 54 are configured as uplink network ports.



Enable info policy for neighbor discovery

Enabling the info policy enables fabric interconnect neighbor information to be displayed. To modify the info policy, follow these steps:

1. In Cisco UCS Manager, click Equipment, choose All > Equipment in the Navigation Pane, and choose the Policies tab on the right.
2. Under Global Policies, scroll down to Info Policy and choose Enabled for Action.
3. Click Save Changes and then click OK.
4. Under Equipment, choose Fabric Interconnect A or B. On the right, choose the Neighbors tab. Nexus switches are discovered with CDP.



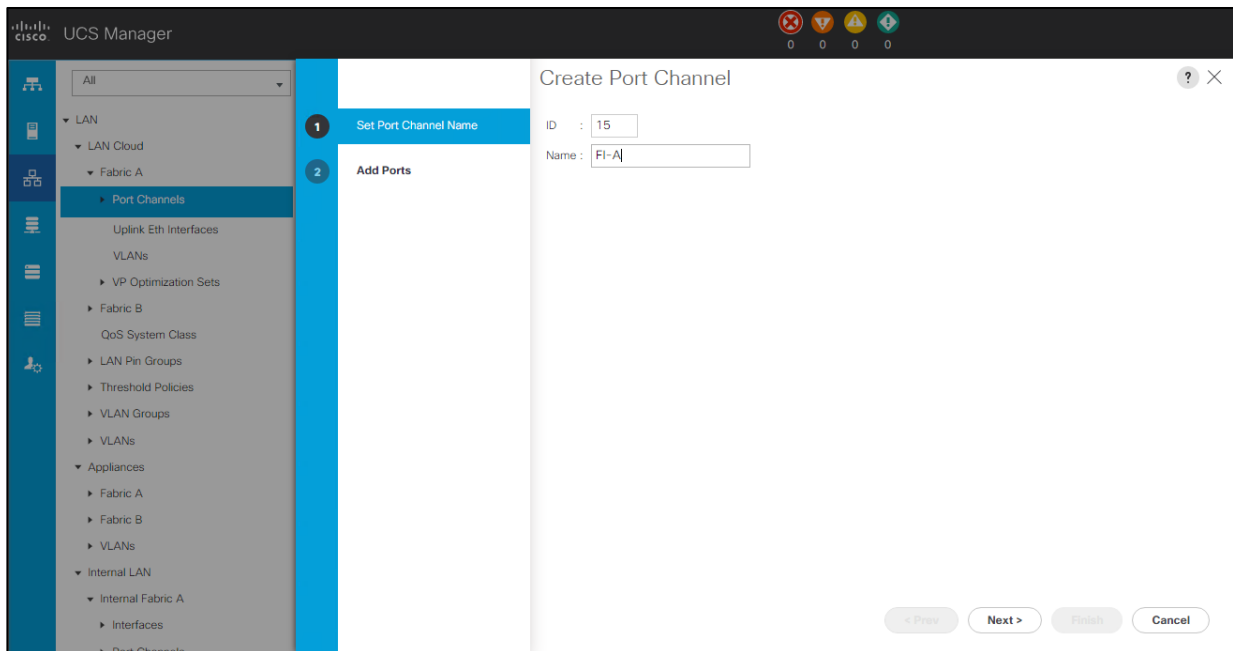
Acknowledge chassis and servers

1. After configuring server ports, acknowledge both of the chassis.
2. Go to Equipment > Chassis > Chassis 1 > General > Actions > select Acknowledge Chassis. Similarly, acknowledge Chassis 2.
3. After acknowledging both chassis, re-acknowledge all servers placed in the chassis.
4. Go to Equipment > Chassis 1 > Servers > Server 1 > General > Actions > select Server Maintenance > select the option Re-acknowledge and click OK. Similarly, repeat the process to re-acknowledge all eight servers.

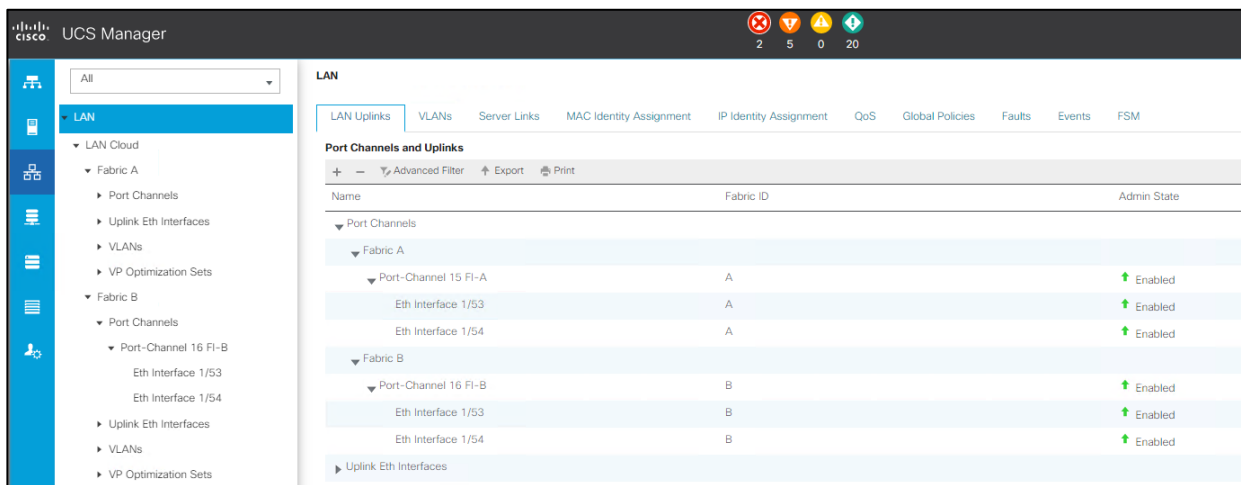
Note: Re-acknowledging causes a reboot of blades.

Create uplink port channels to Cisco Nexus switches

1. Create uplink port channels to Cisco Nexus switches. In this procedure, two port channels were created: one from Fabric A to both Cisco Nexus switches and one from Fabric B to both Cisco Nexus switches. To configure the necessary port channels in the Cisco UCS environment, complete the following steps:
 - a. In Cisco UCS Manager, click the LAN tab in the navigation pane.
 - b. Under LAN > LAN Cloud, expand the node Fabric A tree:
 - c. Right-click Port Channels.
 - d. Select Create Port Channel.
 - e. Enter 15 as the unique ID of the port channel.
 - f. Enter FI-A as the name of the port channel.



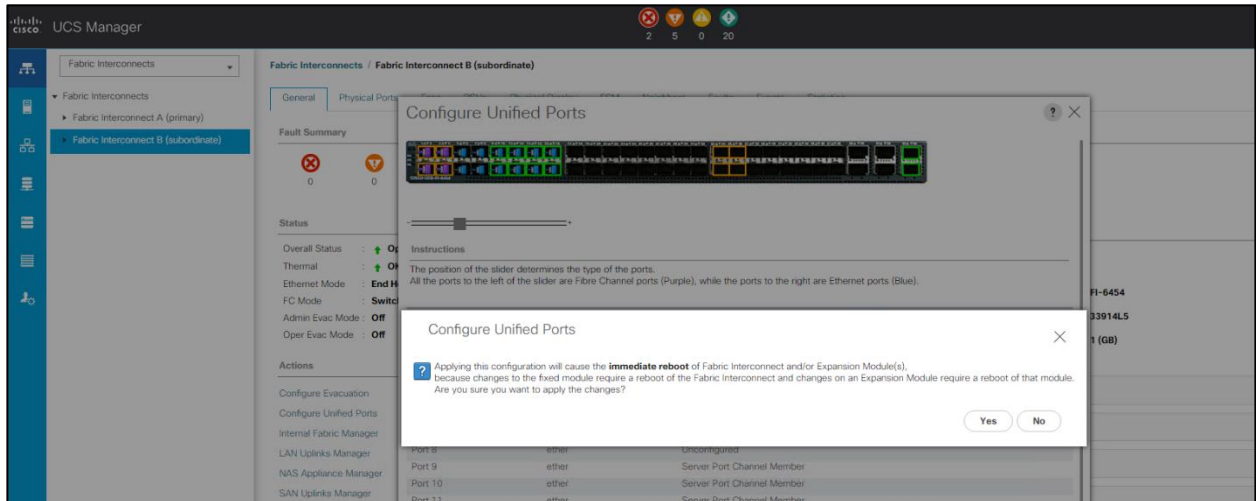
- g. Click Next.
 - h. Select Ethernet ports 53 and 54 for the port channel.
 - i. Click >> to add the ports to the port channel.
 - j. Click Finish to create the port channel and then click OK.
2. Complete same procedure on Switch B. Enter ID 16 and the name FI-B for the port-channel.
 3. Review the port-channels and port configuration summary in the LAN Uplink tab.



Configure FC SAN uplink ports

1. Configure FC SAN uplink ports. To enable the FC ports, follow these steps for FI-6454 A and B:
 - a. In Cisco UCS Manager, click Equipment on the left.
 - b. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
 - c. Select Configure Unified Ports.

- d. Click Yes on the pop-up window warning that changes to the fixed module require a reboot of the fabric interconnect and changes to the expansion module require a reboot of that module.
- e. Within the Configured Fixed Ports pop-up window, move the gray slider bar from the left to the right to select four ports to be set as FC uplinks.
- f. Click OK, then Yes, then OK to continue.

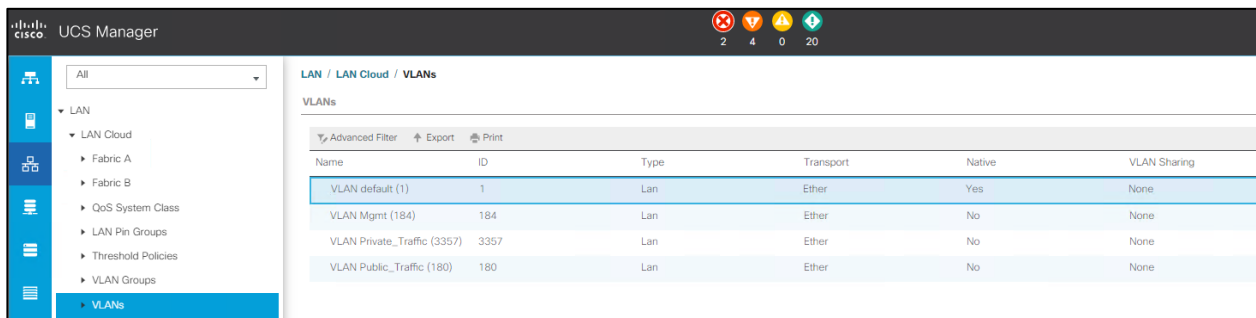


2. Wait for both fabric interconnects to reboot.
3. Log back into Cisco UCS Manager.

Configure VLAN

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

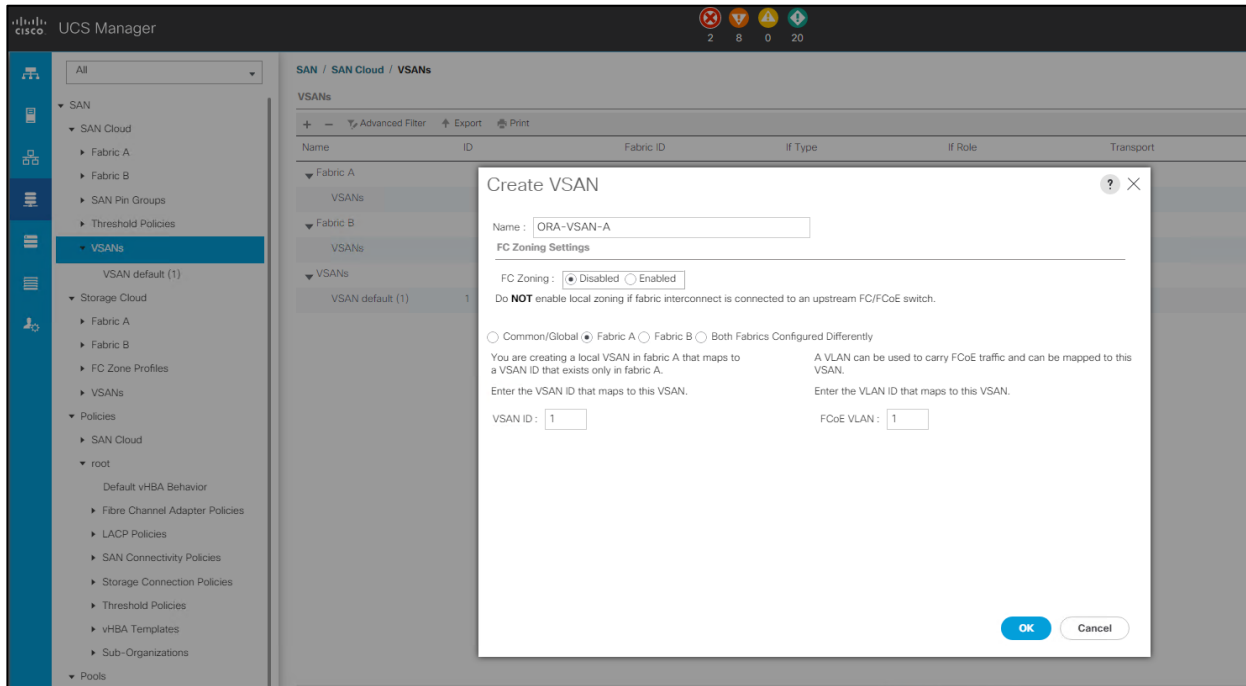
1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter Public_Traffic as the name of the VLAN to be used for public network traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 180 as the ID of the VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.



Configure VSAN

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > SAN Cloud.
3. Under VSANs, right-click VSANs.
4. Select Create VSANs.
5. Enter ORA-VSAN-A as the name of the VSAN.
6. Leave FC Zoning set at Disabled.



Note: Select Fabric A for the scope of the VSAN.

7. Enter 101 as the ID of the VSAN.

Enter a unique VSAN ID that matches the configuration in the MDS switch for Fabric A. NetApp recommends using the same ID for both parameters and to use something other than 1.

1. Click OK and then click OK again.
2. Repeat these steps to create the VSANs necessary for this solution. VSAN 101 and 102 are configured as shown below:

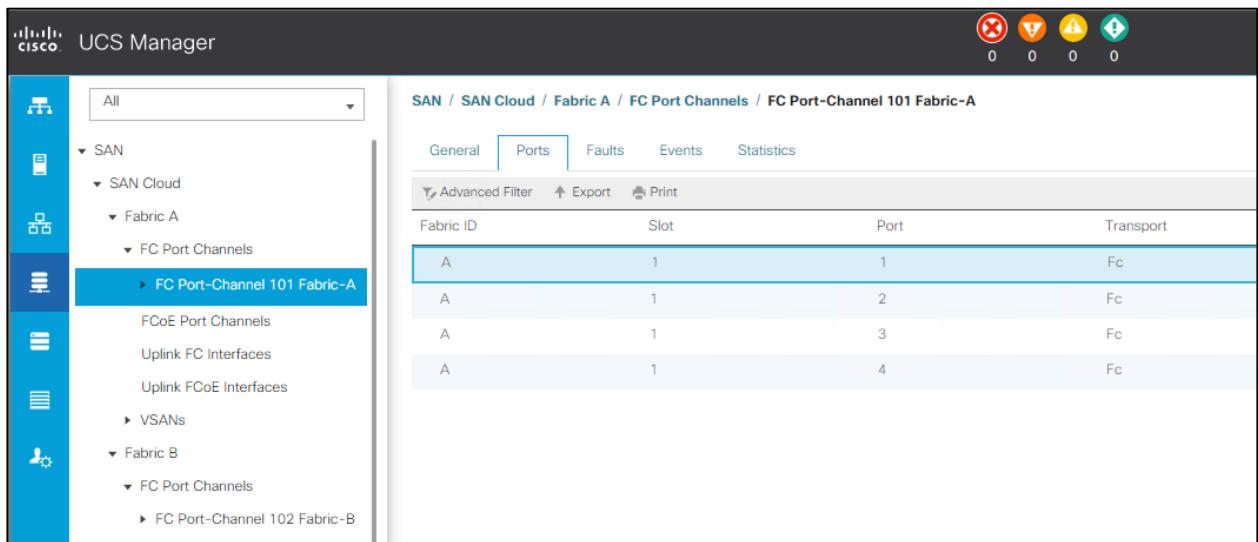
The screenshot shows the Cisco UCS Manager interface with the 'VSANs' table. The table has the following columns: Name, ID, Fabric ID, If Type, If Role, Transport, FCoE VLAN ID, and Operational Status. The table contains the following data:

Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operational Status
VSAN ORA-VSAN-A (101)	101	A	Virtual	Network	Fc	101	OK
VSAN ORA-VSAN-B (102)	102	B	Virtual	Network	Fc	102	OK
VSAN default (1)	1	Dual	Virtual	Network	Fc	4048	OK

Create FC uplink port channels (FCP)

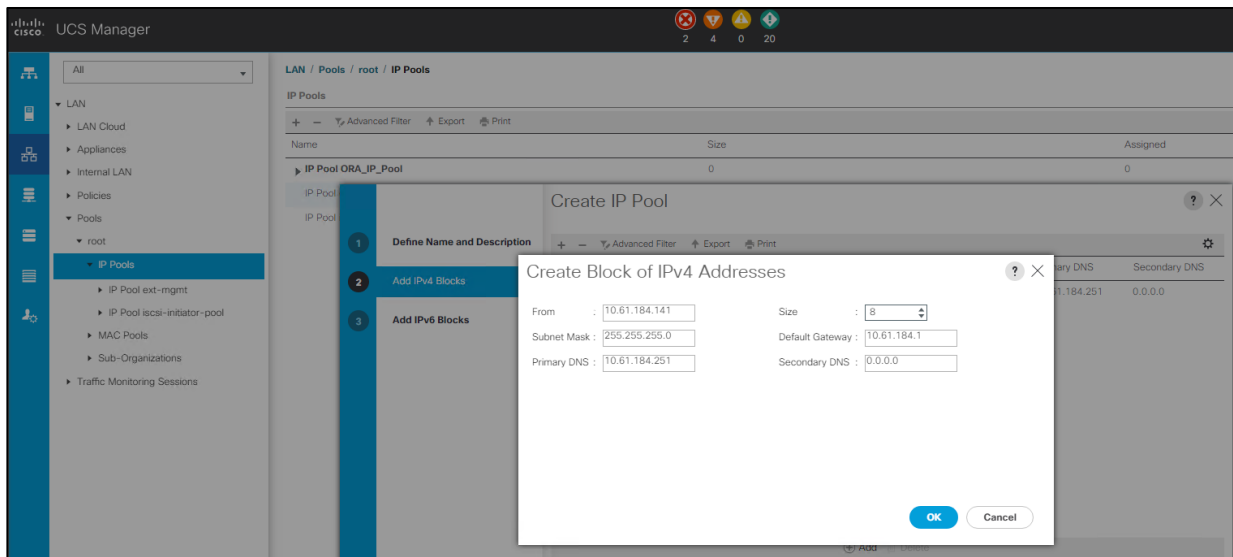
To create the FC uplink port channels and assign the appropriate VSANs to them for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Choose SAN > SAN Cloud.
3. Expand Fabric A and choose the FC port channels.
4. Right-click FC Port Channels and choose Create FC Port Channel.
5. Set a unique ID for the port channel and provide a unique name for the port channel.
6. Click Next.
7. Choose the appropriate Port Channel Admin Speed.
8. Choose the four ports connected to Cisco MDS 9132T A and use >> to add them to the port channel.
9. Click Finish to complete creating the port channel.
10. Under FC Port-Channels, choose the newly created port channel.
11. From the drop-down list, choose ORA-VSAN-A.
12. Click Save Changes to assign the VSAN.

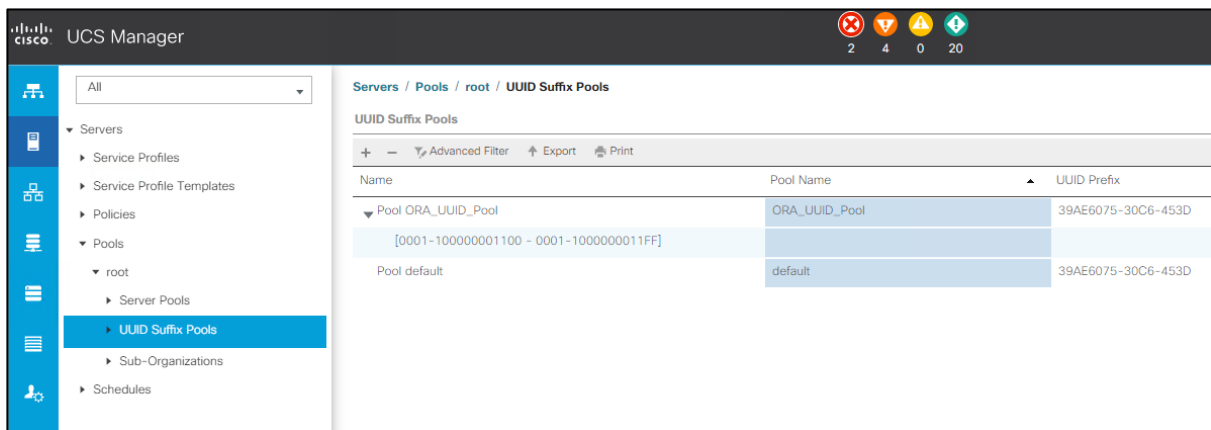


Configure IP, UUID, Server, MAC, WWNN and WWPN pools

1. IP pool creation. An IP address pool on the out-of-band management network must be created to facilitate KVM access to each compute node in the Cisco UCS domain. To create a block of IP addresses for server KVM access in the Cisco UCS environment, follow these steps:
 - a. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
 - b. Select Pools > root > IP Pools > Click Create IP Pool.
 - c. We have named the IP Pool as ORA-IP-Pool for this solution.
 - d. Select option Sequential to assign IPs in sequential order, and then click Next.
 - e. Click Add IPv4 Block.
 - f. Enter the starting IP address of the block, the number of IP addresses required, and the subnet and gateway information as shown below.

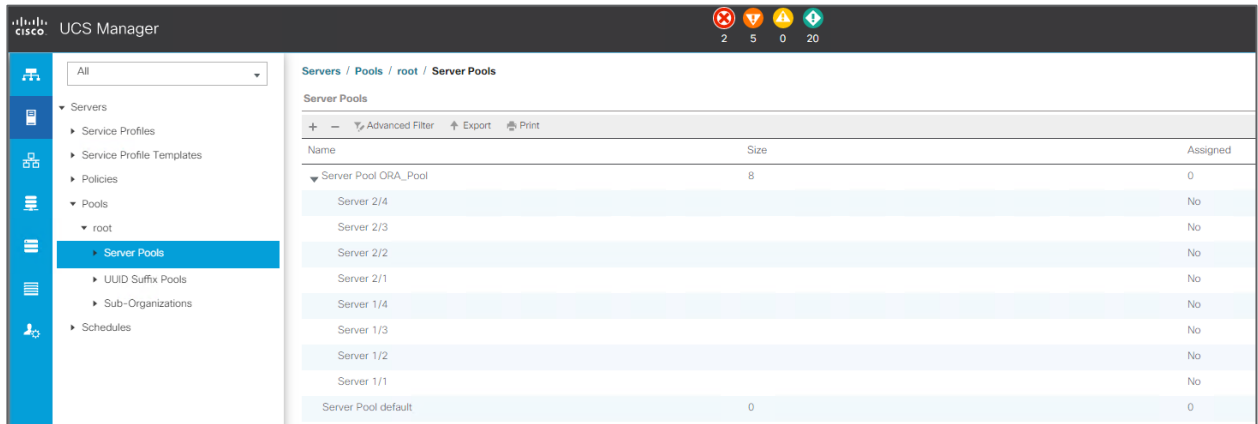


- g. Click Next, and then click Finish to create the IP block.
2. UUID Suffix Pool Creation. To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, follow these steps:
 - a. In Cisco UCS Manager, click the Servers tab in the navigation pane.
 - b. Select Pools > Root.
 - c. Right-click UUID Suffix Pools, and then select Create UUID Suffix Pool.
 - d. Enter ORA-UUID-Pool as the name of the UUID name.
 - e. Optional: Enter a description for the UUID pool.
 - f. Keep the prefix at the derived option, and select Sequential as the Assignment Order. Then click Next.
 - g. Click Add to add a block of UUIDs.
 - h. Create a starting point UUID appropriate for your environment.

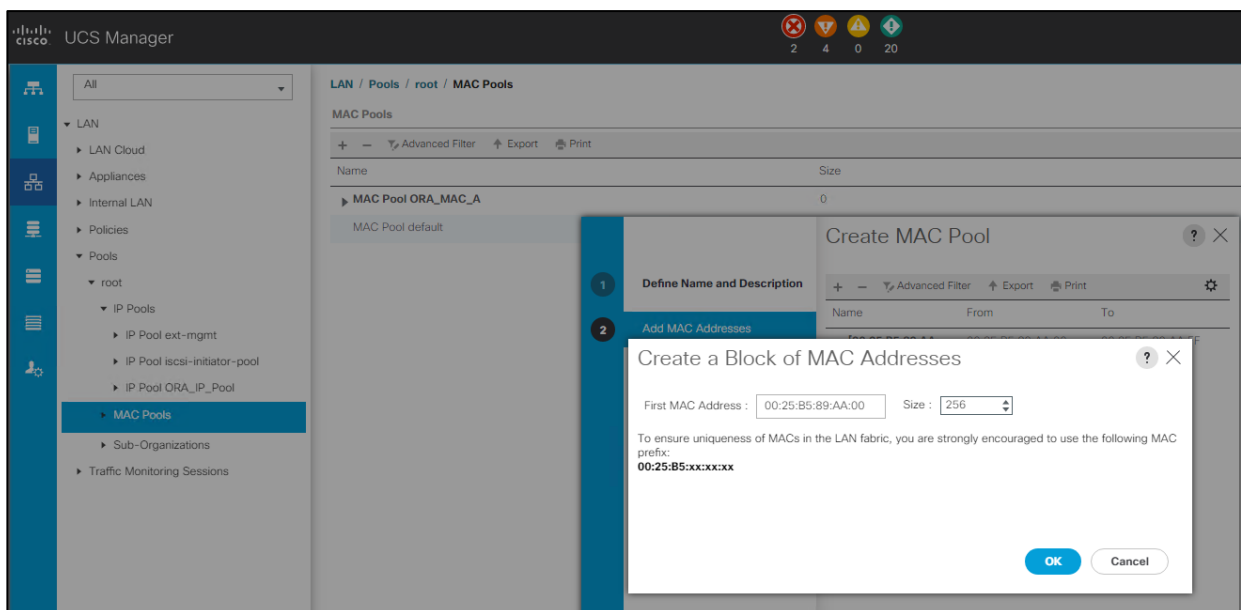


- i. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
3. Server pool creation. To configure the necessary server pool for the Cisco UCS environment, complete the following steps:

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

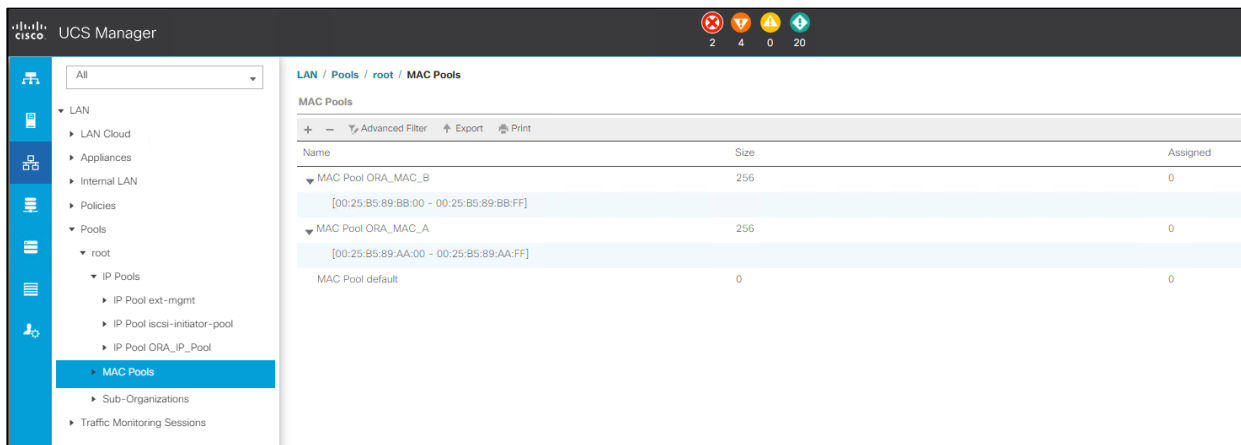


- MAC pool creation. To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:
 - In Cisco UCS Manager, click the LAN tab in the navigation pane.
 - Select Pools > Root > right-click MAC Pools under the root organization.
 - Select Create MAC Pool to create the MAC address pool
 - Enter ORA-MAC-A as the name for MAC pool.
 - Enter the seed MAC address and provide the number of MAC addresses to be provisioned.

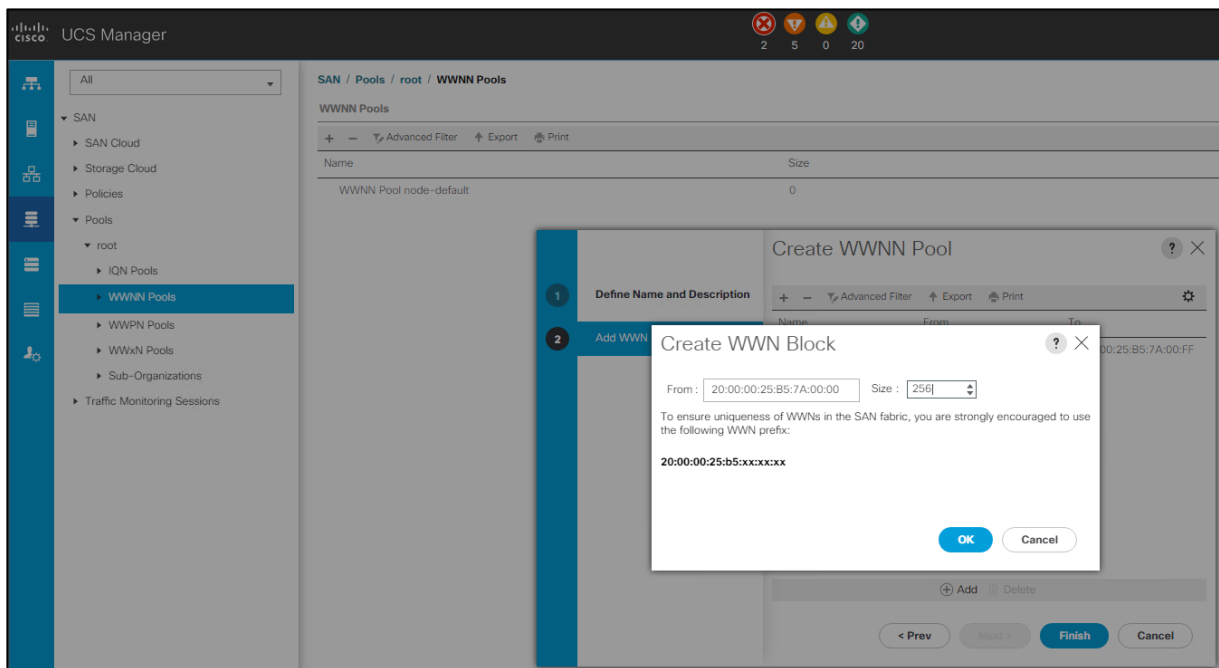


- Click OK and then click Finish.

- g. In the confirmation message, click OK.
- h. Create MAC Pool B and assign unique MAC addresses as shown below.

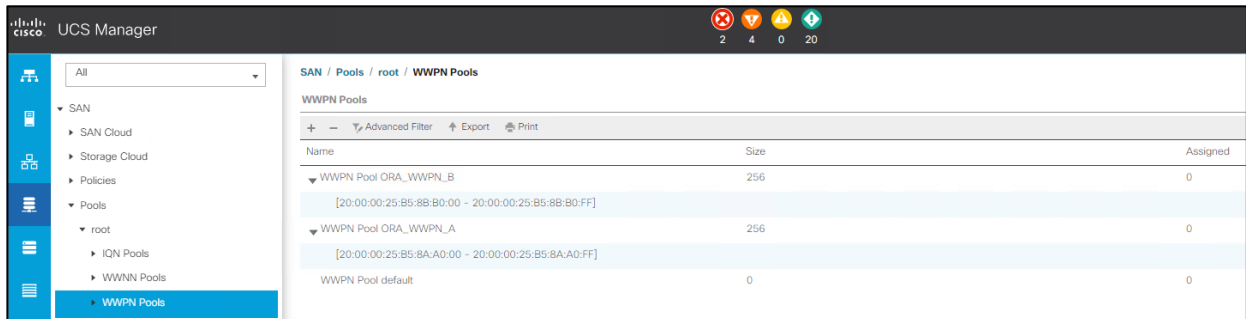


5. WWNN pool creation. To configure the necessary WWNN pools for the Cisco UCS environment, complete the following steps:
 - a. In Cisco UCS Manager, click the SAN tab in the navigation pane.
 - b. Select Pools > Root > WWNN Pools > right click WWNN Pools > Select Create WWNN Pool.
 - c. Assign the name and the Assignment Order as sequential as shown below.



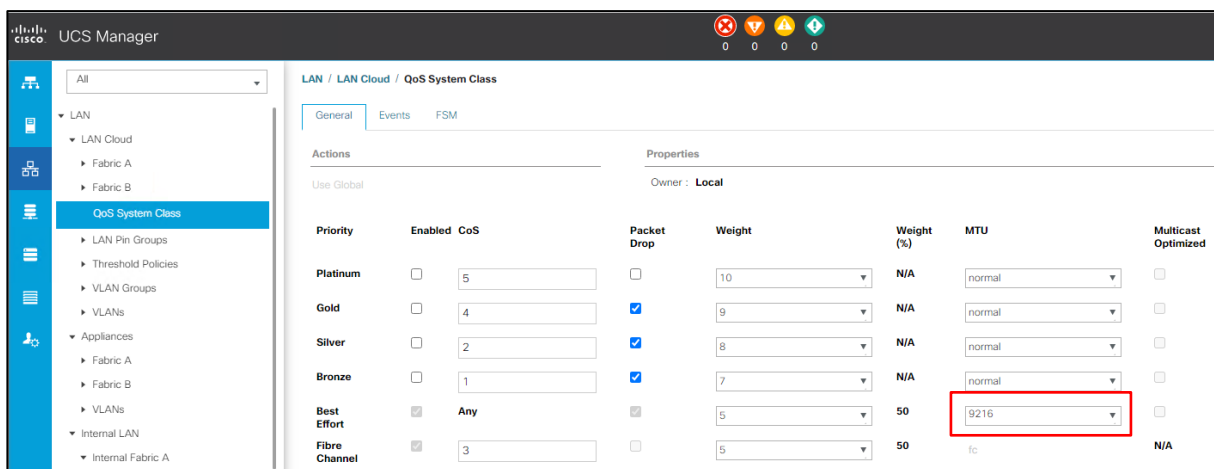
- d. Click OK, and then click Finish.
6. WWPN pool creation. To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:
 - a. In Cisco UCS Manager, click the SAN tab in the navigation pane.
 - b. Select Pools > Root > WWPN Pools > right-click WWPN Pools > Select Create WWPN Pool.
 - c. Assign the name as ORA-WWPN-A and the Assignment Order as sequential.

- d. Click Next, and then click Add to add a block of ports.
- e. Enter Block for WWN and size.
- f. Click OK and then Finish.
- g. Repeat the same for ORA_WWPN_B, and two WWPN pools are created as showed below:



Note: When there are multiple adjacent UCS domains, it is important that each set of blocks have different values for the WWNN, WWPN, and MAC.

7. Set jumbo frames in both of the Cisco fabric interconnects. To configure jumbo frames and enable quality of service in the Cisco UCS fabric, follow these steps:
 - a. In Cisco UCS Manager, click the LAN tab in the navigation pane.
 - b. Select LAN > LAN Cloud > QoS System Class.
 - c. In the right pane, click the General tab.
 - d. In the Best Effort row, enter 9216 in the box under the MTU column.

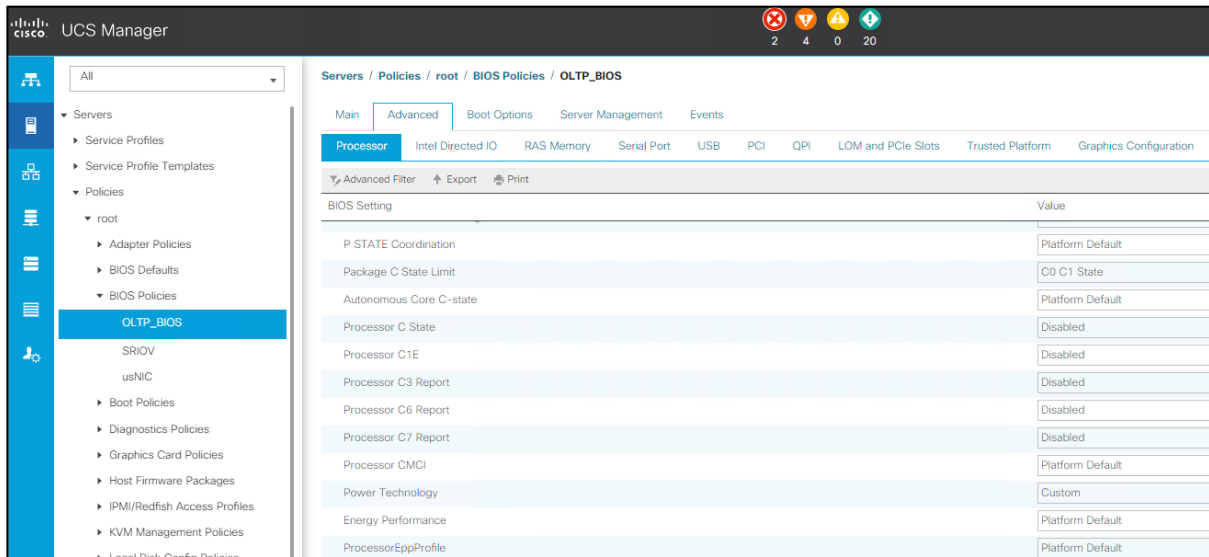


Configure server BIOS policy

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > Root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter OLTP_BIOS as the BIOS policy name.
6. Select and click the newly created BIOS policy.

7. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab.
8. Set the following within the Processor tab:

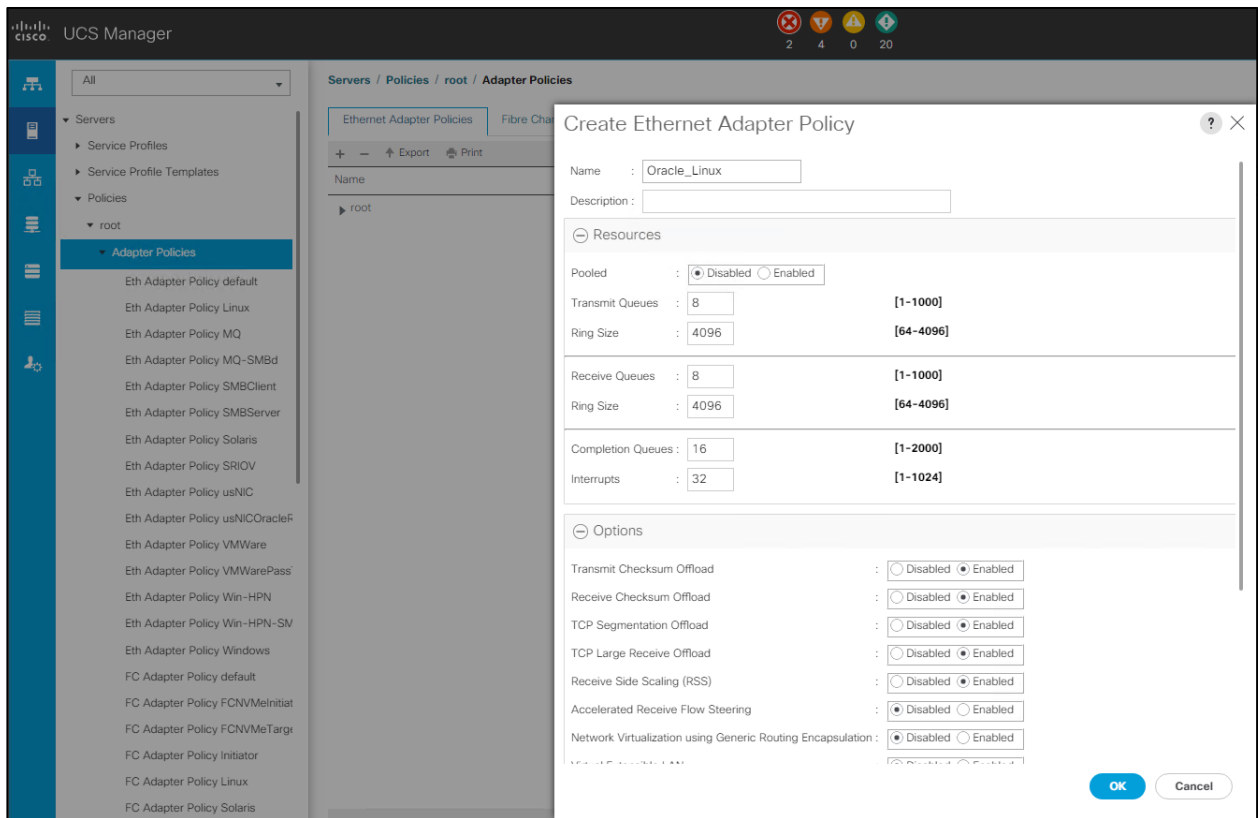


9. Scroll down to the bottom of the page and set Workload Configuration to IO Sensitive.
10. Click Save Changes, and then click OK.

Create adapter policy

To create an adapter policy for the Cisco UCS environment, complete the following steps:

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



6. Click OK.

Update default Maintenance Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > Root > Maintenance Policies > Default.
3. Change the Reboot Policy to User Ack.
4. Click Save Changes.
5. Click OK to accept the changes.

Configure vNIC and vHBA template

For this solution, we created two vNIC templates for the public network and private network traffic.

1. Create the public and private vNIC template.
 - a. In Cisco UCS Manager, click the LAN tab in the navigation pane.
 - b. Select Policies > Root > vNIC Templates > right-click to vNIC Template, and select Create vNIC Template.
 - c. Enter ORA-vNIC-A as the vNIC template name and keep Fabric A selected.
 - d. Select the Enable Failover checkbox for high availability of the vNIC.
 - e. Select the template type as Updating Template.
 - f. Under VLANs, select the checkboxes Default and Public_Traffic, and set Native-VLAN as the Public_Traffic.

- g. Keep the MTU value 1500 for Public Network Traffic.
- h. In the MAC Pool list, select ORA-MAC-A.
- i. Click OK to create the vNIC template as shown below.

Create vNIC Template

Name : ORA_vNIC_A

Description :

Fabric ID : ☒ Fabric A ☐ Fabric B ☒ Enable Failover

Redundancy

Redundancy Type : ☒ No Redundancy ☐ Primary Template ☐ Secondary Template

Target

☒ Adapter ☐ VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

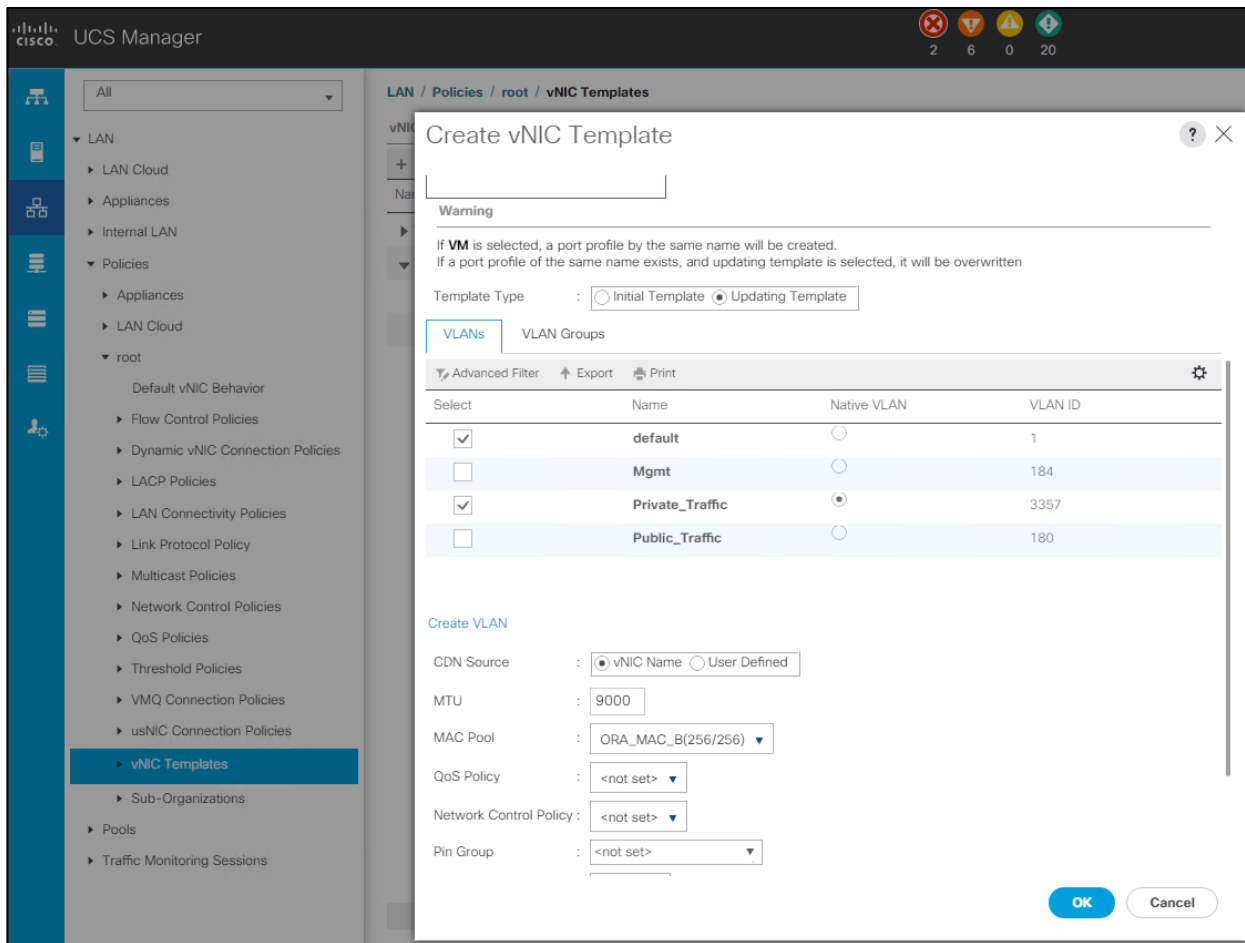
Template Type : ☐ Initial Template ☒ Updating Template

VLANs | VLAN Groups

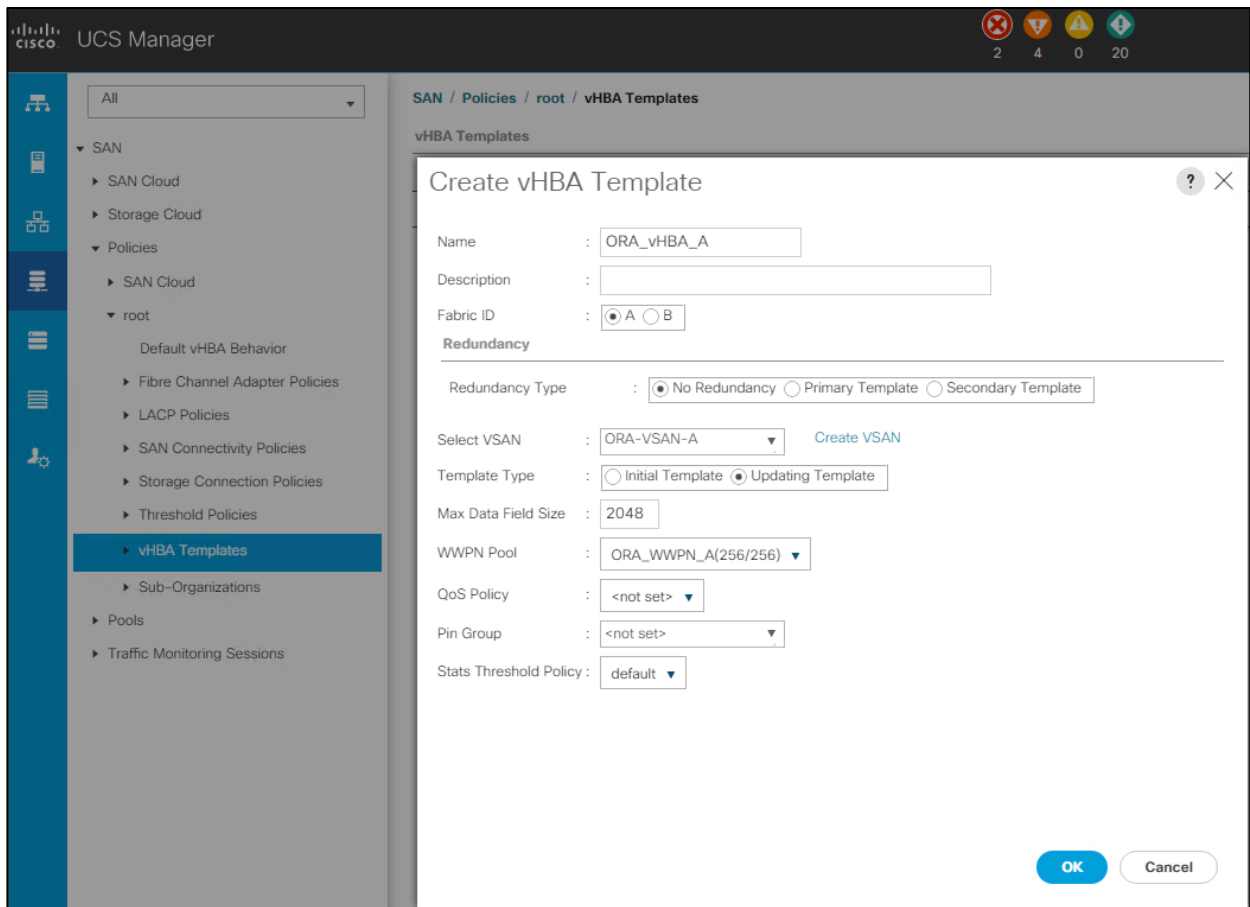
Select	Name	Native VLAN	VLAN ID
<input checked="" type="checkbox"/>	default	<input type="radio"/>	1
<input checked="" type="checkbox"/>	Mgmt	<input type="radio"/>	184
<input type="checkbox"/>	Private_Traffic	<input type="radio"/>	3357
<input checked="" type="checkbox"/>	Public_Traffic	<input checked="" type="radio"/>	180

OK Cancel

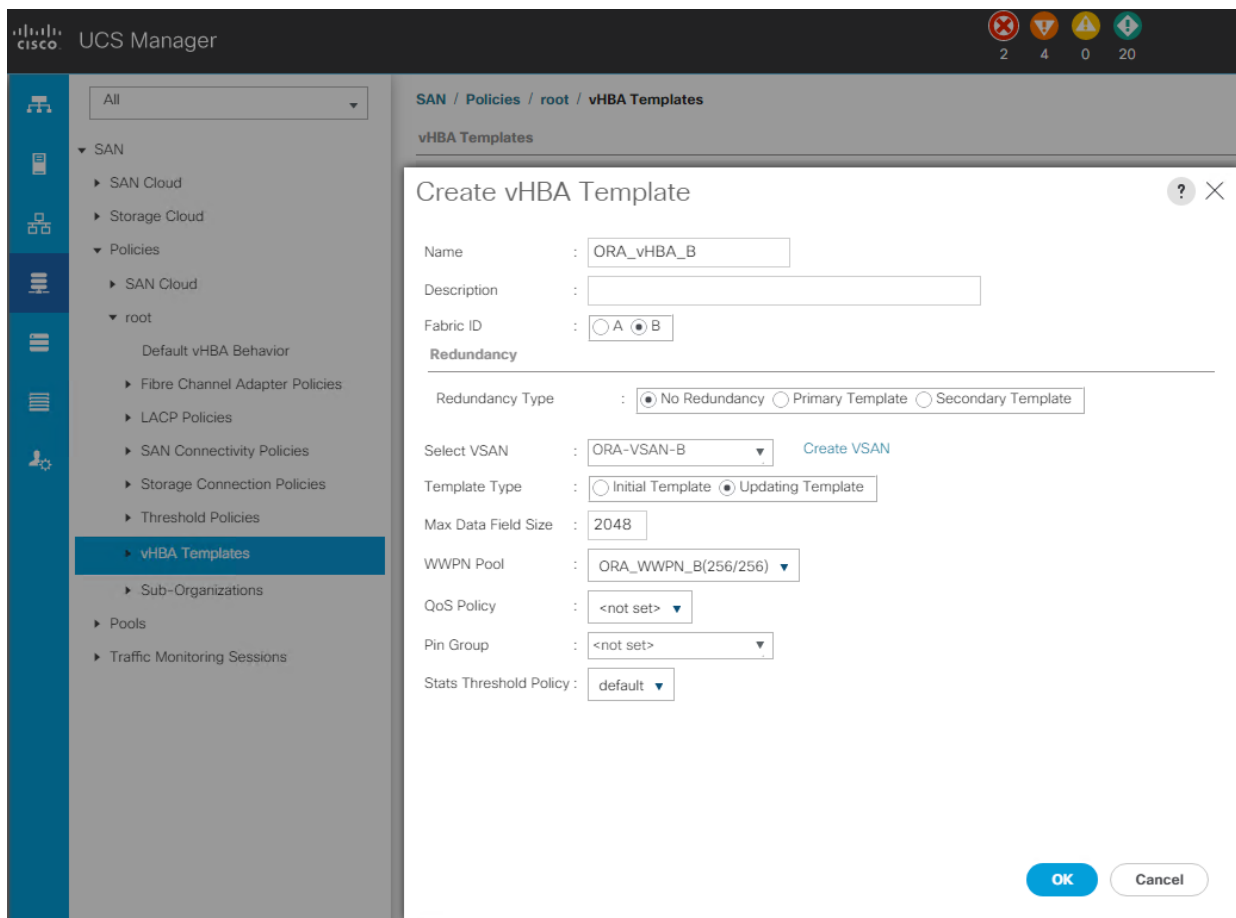
- j. Click OK to finish.
2. Create the private vNIC template.
 - a. Enter ORA-vNIC-B as the vNIC template name for Private Network Traffic.
 - b. Select Fabric B and Enable Failover for the Fabric ID options.
 - c. Select Updating Template for the template type.
 - d. Under VLANs, select the checkboxes Default and Private_Traffic, and set the Native-VLAN as Private_Traffic.
 - e. Set the MTU value to 9000 and the MAC Pool as ORA-MAC-B.
 - f. Click OK to create the vNIC template as shown below.



3. Create the storage vHBA template. To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:
 - a. In Cisco UCS Manager, click the SAN tab in the navigation pane.
 - b. Select Policies > Root > right-click vHBA Templates > Select Create vHBA Template to create vHBAs.
 - c. Enter the name ORA-vHBA-A and keep Fabric A selected.
 - d. Select VSAN as ORA-VSAN-A and set Updating Template as the template type.
 - e. Select ORA-WWPN_A for the WWPN Pool from the drop-down list as shown below.



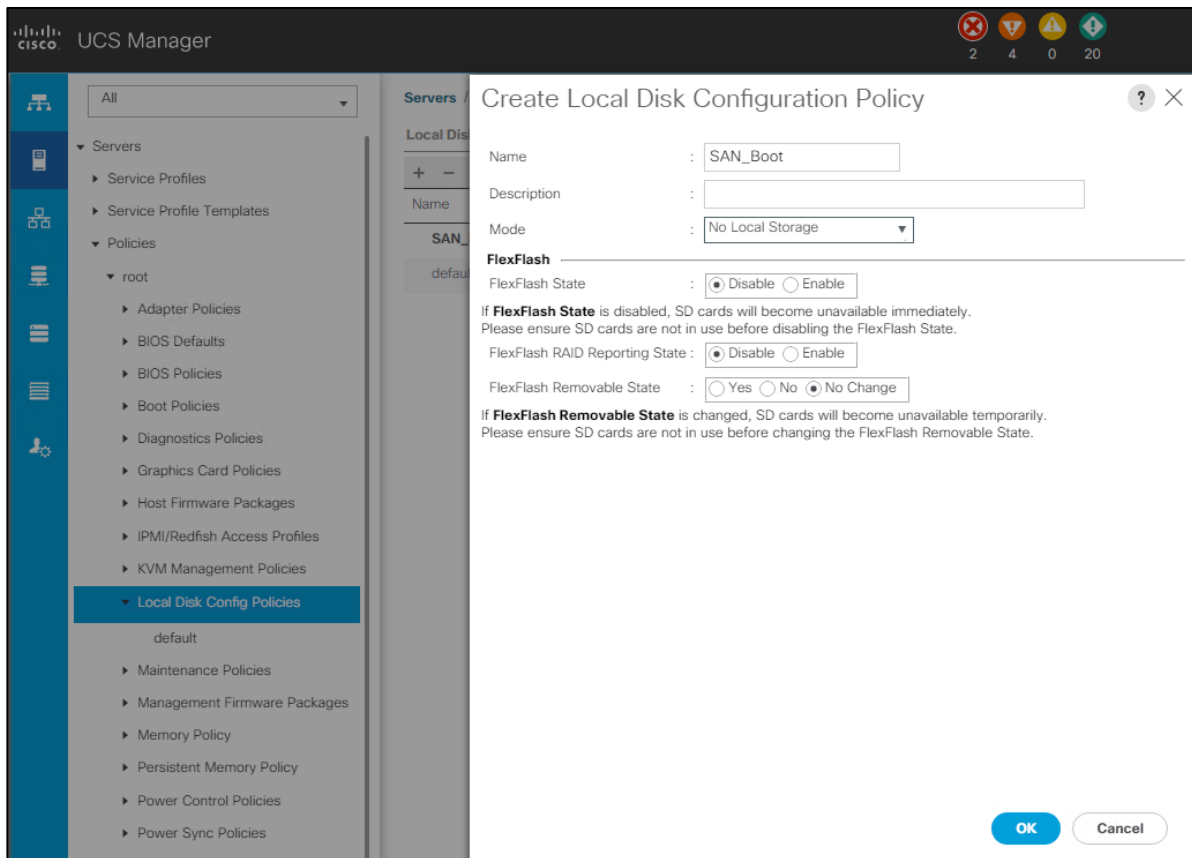
- f. Enter ORA-vHBA_B for the name, and select Fabric B. Select WWPN Pool for ORA-vHBA-B as ORA-WWPN-B as shown below.



Create server boot policy for SAN boot

All Oracle nodes were set to Boot from SAN for the CVD as part of the Service Profile template. The benefits of booting from SAN are numerous: disaster recovery, lower cooling, and power requirements for each server since a local drive is not required, and better performance. This process applies to a Cisco UCS environment in which the storage SAN ports are configured as detailed in the following sections.

1. Create a local disk configuration policy. A local disk configuration for the Cisco UCS is necessary if the servers in the environments have a local disk. To configure the local disk policy, complete the following steps:
 - a. Go to the tab Servers > Policies > Root > right-click Local Disk Configuration Policy > enter SAN-Boot as the local disk configuration policy name, and change the mode to No Local Storage.
 - b. Click OK to create the policy as shown below.



2. The SAN boot policy configures the SAN primary's primary-target to be network interface `infra-svm_data_fcp_lif_1_2a` on the NetApp storage cluster and the SAN primary's secondary-target to be network interface `infra-svm_data_fcp_lif_2_2c` on the NetApp storage cluster. Similarly, the SAN secondary's primary-target is network interface `infra-svm_data_fcp_lif_1_2a` on the NetApp storage cluster and the SAN secondary's secondary-target is network interface `srm_data_fcp_lif_1_2c` on the NetApp storage cluster. Log in to the NetApp storage controller and verify that the port information is correct. This information can be found in the NetApp storage UI under Network > Network Interfaces.

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

- a. Retrieve the interface LIF configuration on the AFF A800 controller for `infra_svm` as follows:

Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Type
ora19c_svm_data_fcp_lif_2_2c	✓	ora19c_svm		20:1a:d0:39:ea:20:bad3	FlexPod-A800-01-02-02	2c	FC	Data
infra-svm_data_fcp_lif_1_2a	✓	infra_svm		20:03:d0:39:ea:20:bad3	FlexPod-A800-01-02-01	2a	FC	Data
infra-svm_data_fcp_lif_1_2b	✓	infra_svm		20:05:d0:39:ea:20:bad3	FlexPod-A800-01-02-01	2b	FC	Data
infra-svm_data_fcp_lif_2_2a	✓	infra_svm		20:08:d0:39:ea:20:bad3	FlexPod-A800-01-02-02	2a	FC	Data
infra-svm_data_fcp_lif_2_2b	✓	infra_svm		20:09:d0:39:ea:20:bad3	FlexPod-A800-01-02-02	2b	FC	Data
infra-svm_data_fcp_lif_1_2c	✓	infra_svm		20:13:d0:39:ea:20:bad3	FlexPod-A800-01-02-01	2c	FC	Data
infra-svm_data_fcp_lif_1_2d	✓	infra_svm		20:14:d0:39:ea:20:bad3	FlexPod-A800-01-02-01	2d	FC	Data
infra-svm_data_fcp_lif_2_2c	✓	infra_svm		20:15:d0:39:ea:20:bad3	FlexPod-A800-01-02-02	2c	FC	Data
infra-svm_data_fcp_lif_2_2d	✓	infra_svm		20:16:d0:39:ea:20:bad3	FlexPod-A800-01-02-02	2d	FC	Data

- b. Go to UCS Manager, and then go to the tab Servers > Policies > Root > Boot Policies.

- c. Right-click and select Create Boot Policy. Enter SAN_Boot as the name of the boot policy as shown below:

UCS Manager

Create Boot Policy

Name : SAN_Boot

Description :

Reboot on Boot Order Change : ☐

Enforce vNIC/vHBA/iSCSI Name : ☒

Boot Mode : ☒ Legacy ☐ Uefi

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

Name	Order	vNIC/vHBA	Type	LUN Name	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
No data available									

Move Up Move Down Delete

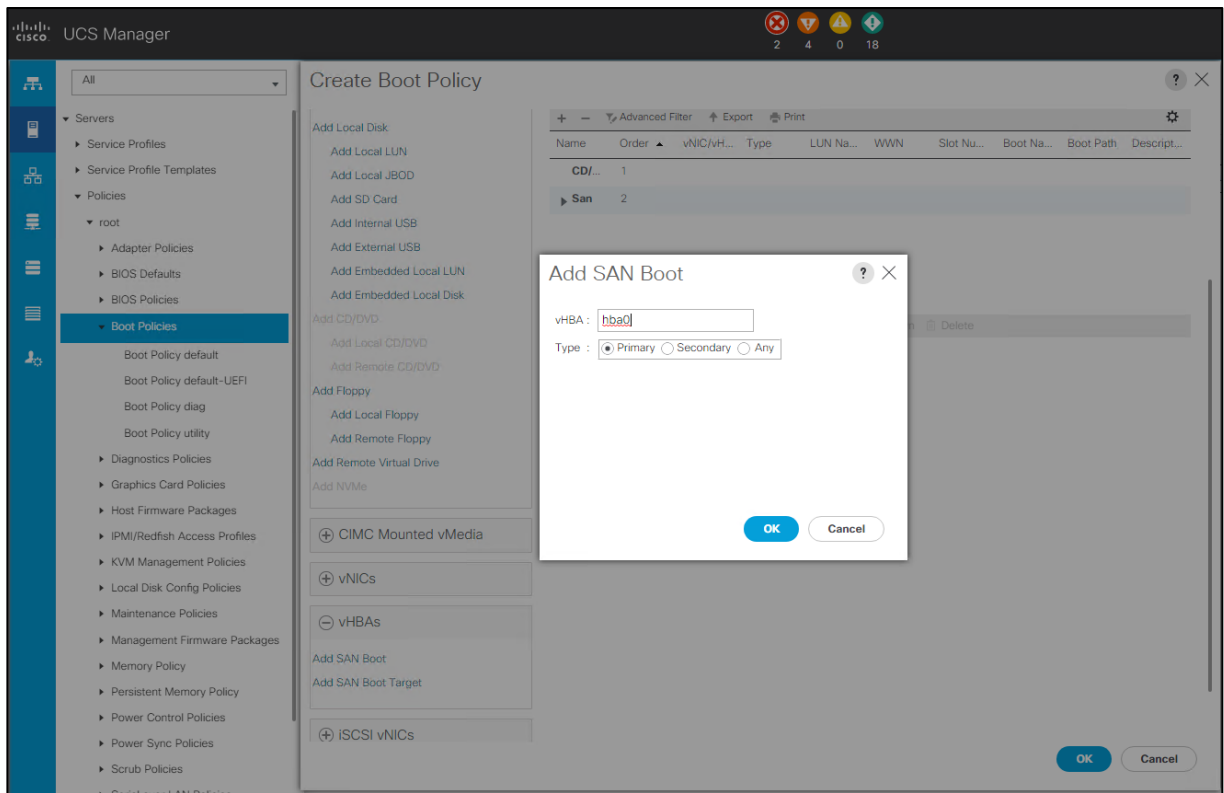
Set Last Boot Parameters

OK Cancel

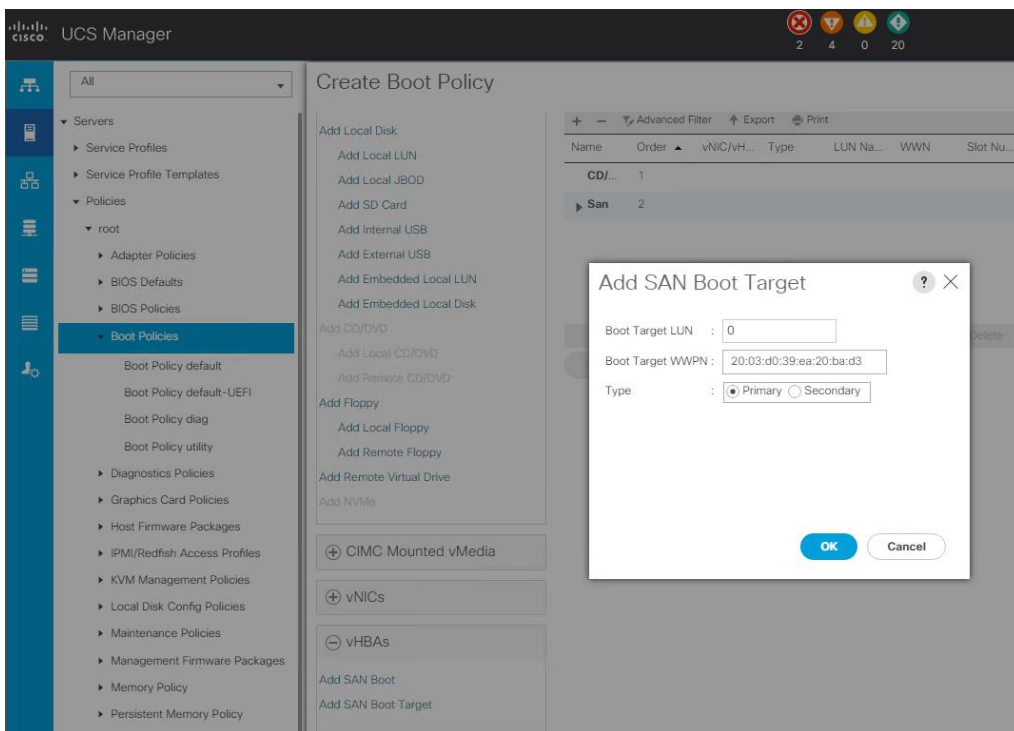
- d. Expand the Local Devices drop-down menu and Choose Add CD/DVD. Expand the vHBAs drop-down menu and select Add SAN Boot.

Note: The SAN boot paths and targets include primary and secondary options to maximize resiliency and the number of paths.

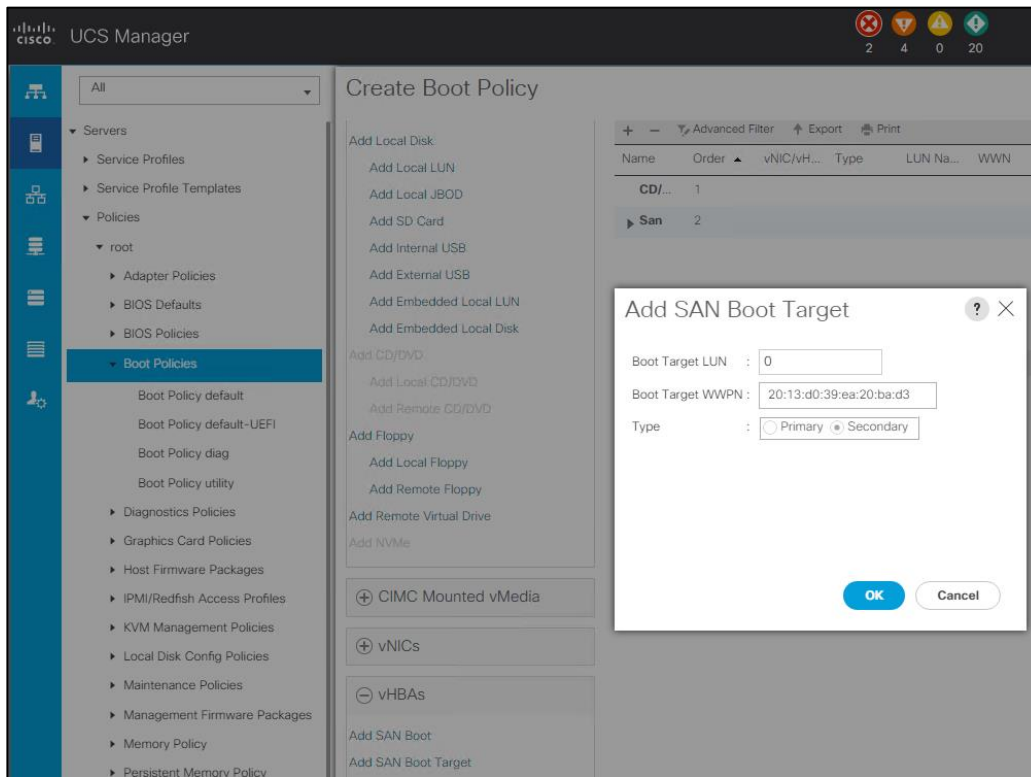
- e. In the Add SAN Boot dialog box, select the type as Primary and name vHBA as hba0. Click OK to add SAN Boot.



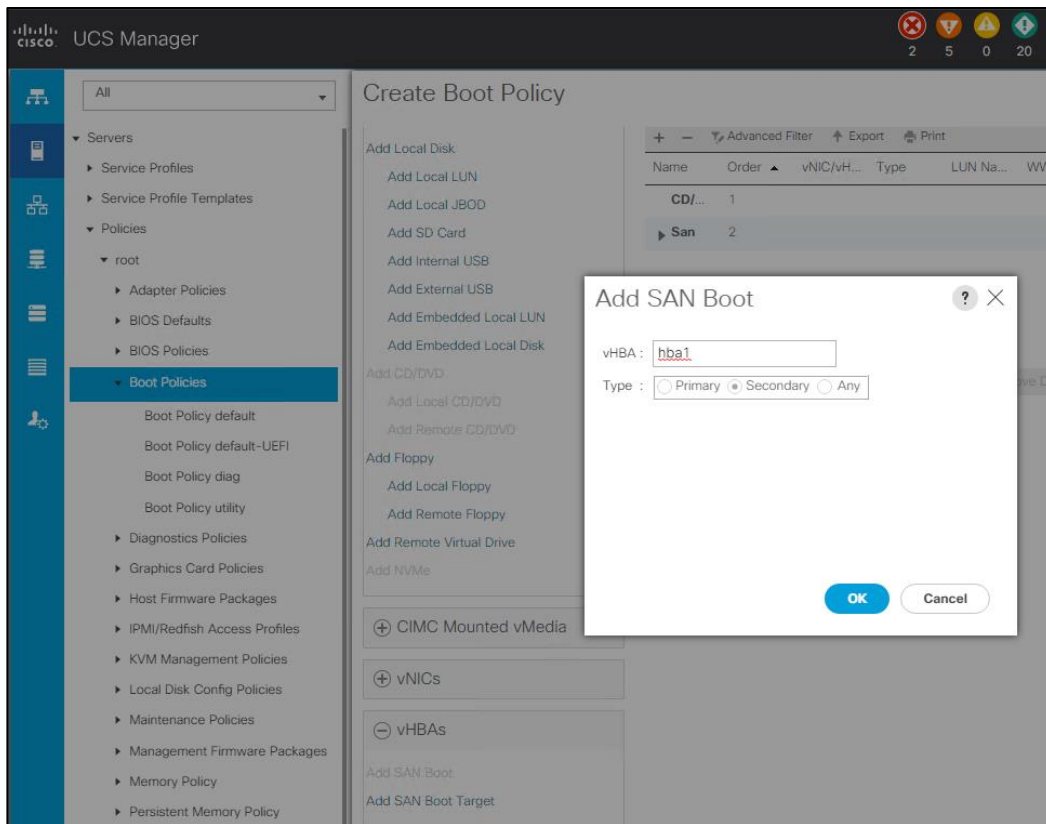
- f. Select Add SAN Boot Target to enter the WWPN address of the storage LIF. Keep 0 as the value for Boot Target LUN. Enter the WWPN of the NetApp storage cluster interface infra-svm_data_fcp_lif_1_2a, and add the SAN boot primary target.



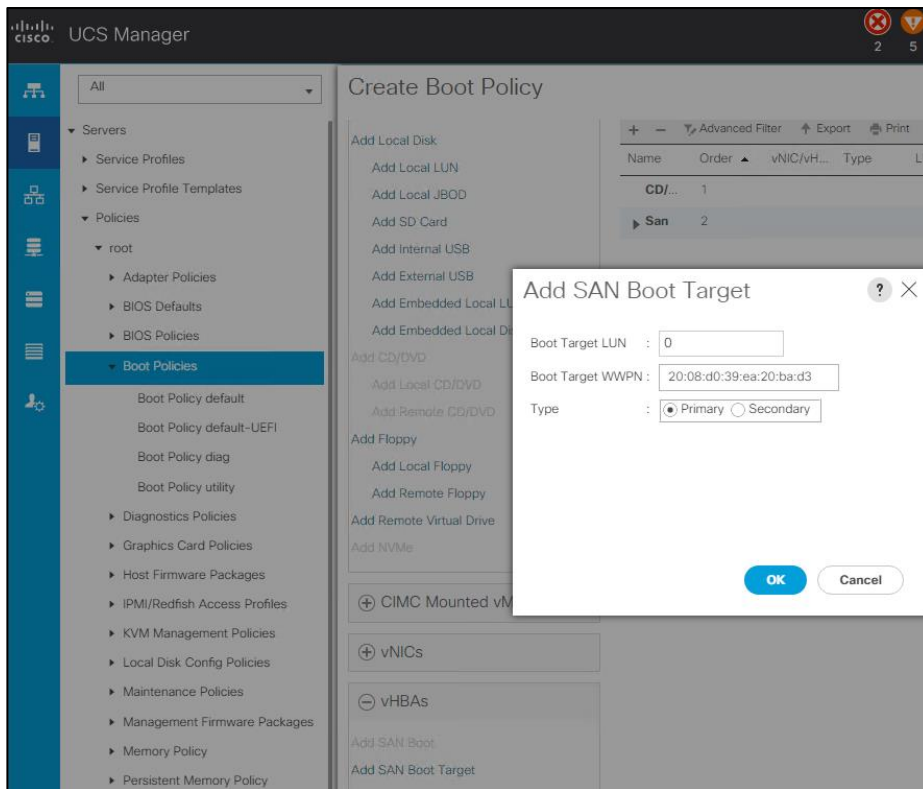
- g. Add the secondary SAN Boot target in same hba0. Enter the boot target LUN as 0 and the WWPN of the NetApp storage cluster AFF A800 interface infra-svm_data_fcp_lif_2_2c, and add the SAN boot secondary target.



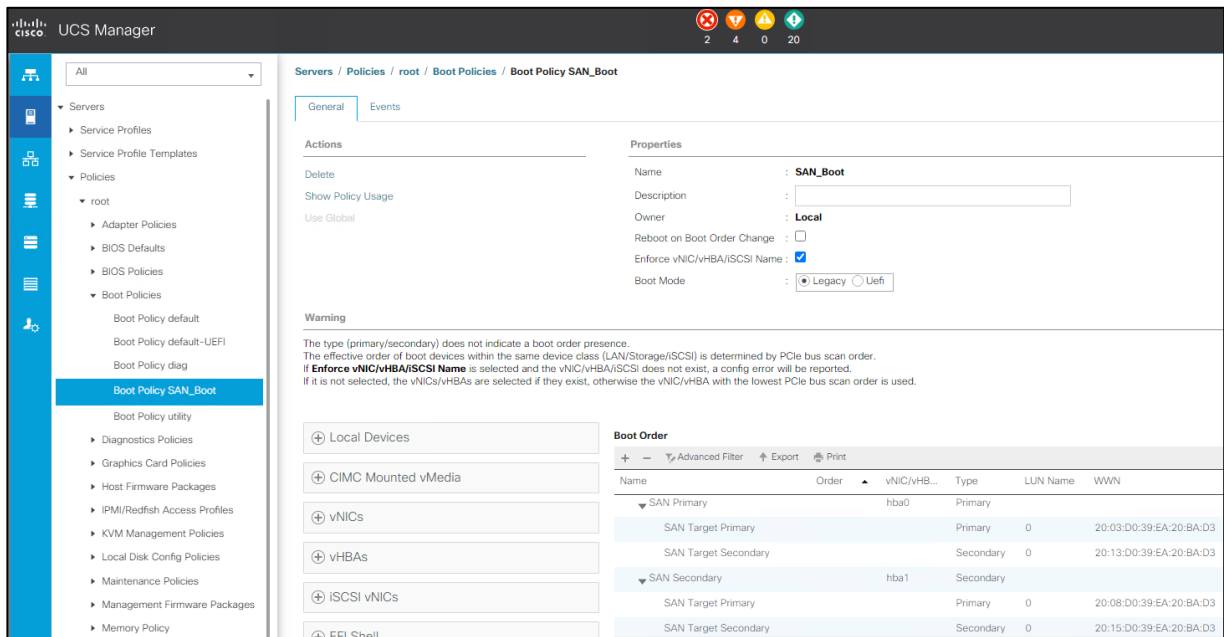
- h. From the vHBA drop-down list, choose Add SAN Boot. In the Add SAN Boot dialog box, enter hba1 in the vHBA field.



- i. Click OK to SAN boot, and then choose Add SAN Boot Target. Enter 0 for the Boot Target LUN. Enter the WWPN of the NetApp storage cluster AFF A800 interface infra-svm_data_fcp_lif_2_2a and add SAN Boot Primary Target.



- j. Add the secondary SAN Boot target into the same hba1, and enter boot target LUN as 0 and WWPN of NetApp Storage cluster AFF A800 interface infra-svm_data_fcp_lif_1_2c and address 20:15:d0:39:ea:20:ba:d3 to be added to that SAN Boot Secondary Target.
- k. Verify SAN Boot order in UCSM boot policy SAN_Boot:



For this solution, we created one Boot Policy as SAN_Boot. For all eight Oracle Database RAC nodes, we assign this boot policy to the service profiles as explained in the following section.

Configure and create a service profile template

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

Create one service profile template called ORA_FLEXPOD using the boot policy created earlier to use four LIF ports from NetApp storage for high availability in case of any FC links go down.

The following sections describe how to create ORA_FLEXPOD.

3. Create a service profile template.

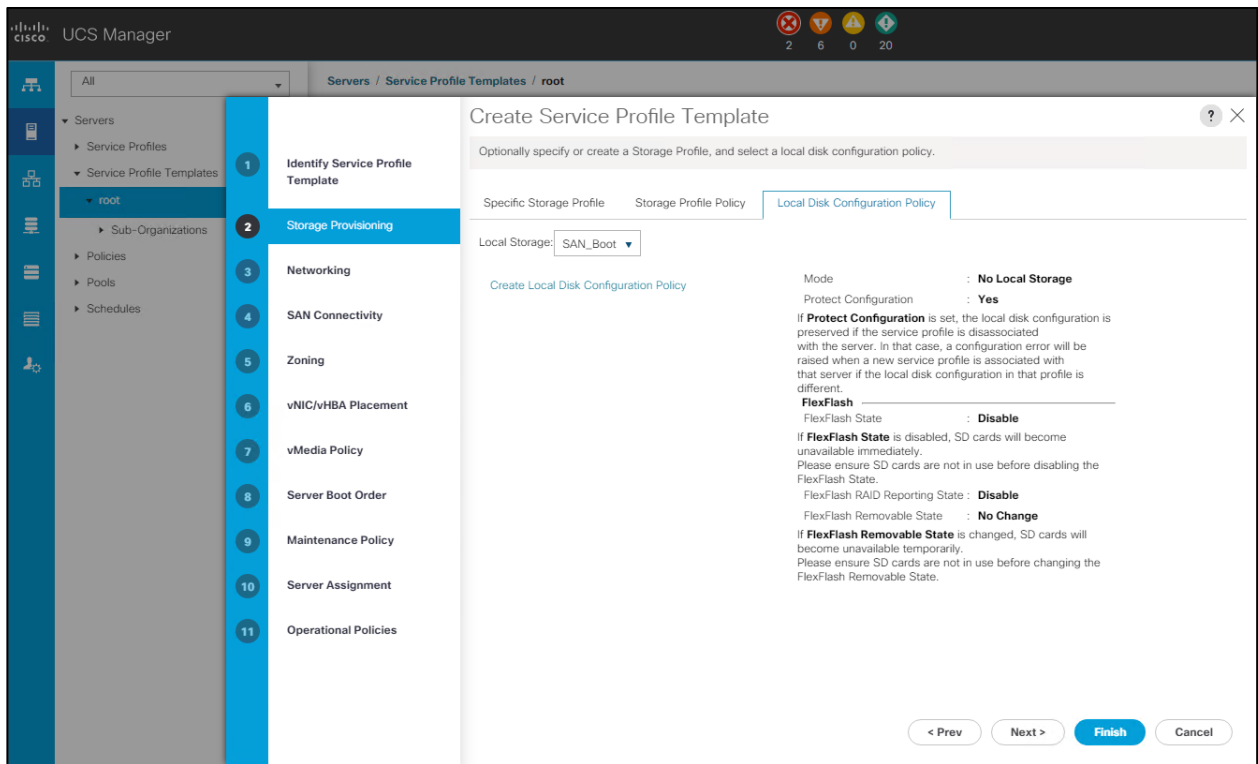
- a. In the Cisco UCS Manager, go to Servers > Service Profile Templates > root and right-click to “Create Service Profile Template” as shown below.

The screenshot shows the Cisco UCS Manager interface. The breadcrumb navigation is 'Servers / Service Profile Templates / root'. The main window is titled 'Create Service Profile Template'. It contains the following fields and options:

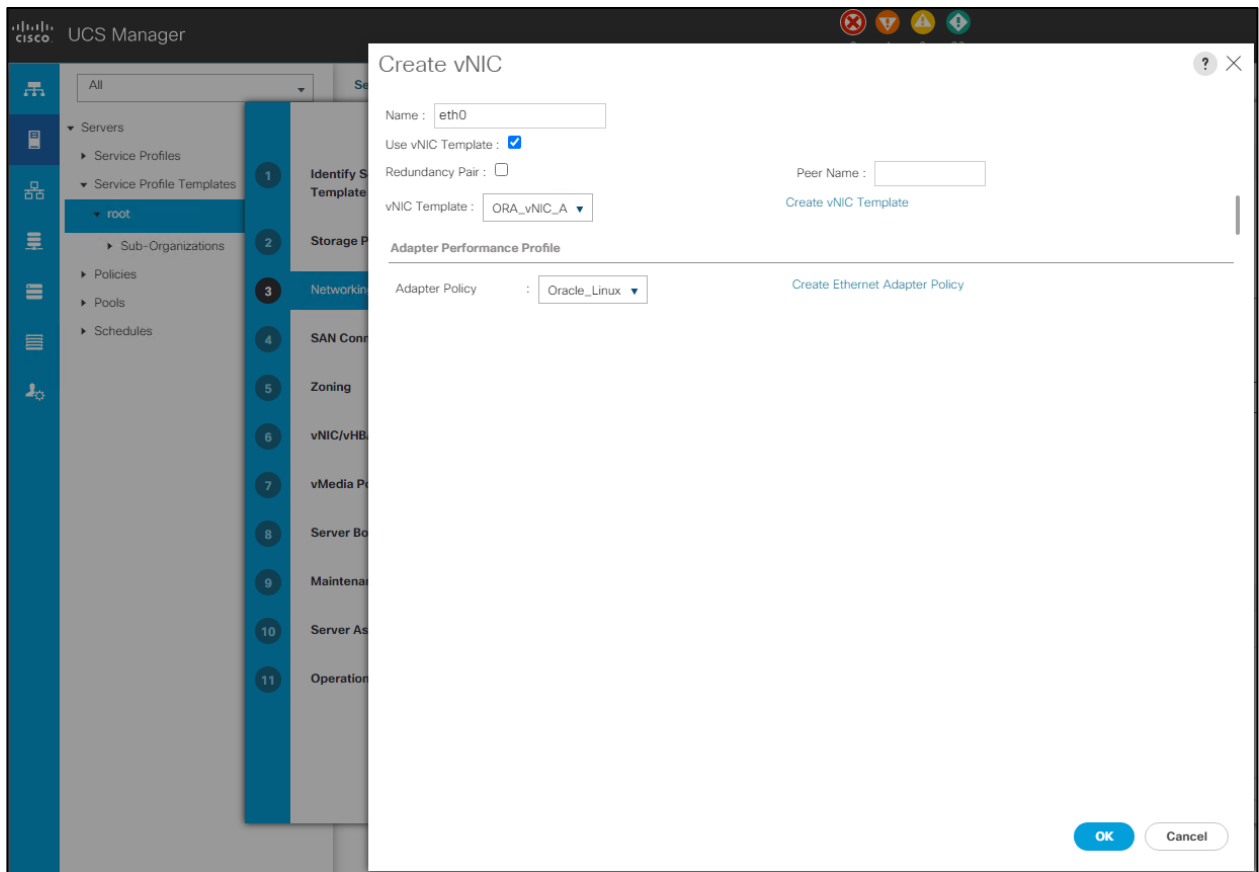
- Name :** ORA_FLEXPOD
- Where :** org-root
- Type :** ☐ Initial Template ☒ Updating Template
- UUID Assignment:** ORA_UUID_Pool(256/256)
- Description:** A text area for an optional description.

At the bottom, there are buttons for '< Prev', 'Next >', 'Finish', and 'Cancel'.

- b. Enter the Service Profile Template name, select the UUID pool that was created earlier, and click Next.
- c. Set the Local Disk Configuration Policy to SAN_Boot as no Local Storage.



- d. In the networking window, select Expert and click Add to create vNICs. Add one or more vNICs that the server should use to connect to the LAN.
- e. We created two vNICs in the create vNIC menu. We named the first vNIC eth0 and the second vNIC eth1.
- f. Select ORA-vNIC-A as the vNIC template as and ORA_Linux_ as the adapter policy, as was done earlier for the vNIC eth0.



- g. Select ORA-vNIC-B as the vNIC template and ORA_Linux as the adapter policy, as was done for the vNIC eth1. The vNICs eth0 and eth1 allow servers to connect to the LAN.

UCS Manager

2 4 0 20

All

Servers

Service Profiles

Service Profile Templates

root

Sub-Organizations

Policies

Pools

Schedules

Create vNIC

Name : eth1

Use vNIC Template : ☒

Redundancy Pair : ☐

Peer Name :

vNIC Template : ORA_vNIC_B

Create vNIC Template

Adapter Performance Profile

Adapter Policy : Oracle_Linux

Create Ethernet Adapter Policy

OK Cancel

UCS Manager

2 4 0 20

All

Servers / Service Profile Templates / root

Servers

Service Profiles

Service Profile Templates

root

Sub-Organizations

Policies

Pools

Schedules

- 1 Identify Service Profile Template
- 2 Storage Provisioning
- 3 Networking
- 4 SAN Connectivity
- 5 Zoning
- 6 vNIC/vHBA Placement
- 7 vMedia Policy
- 8 Server Boot Order
- 9 Maintenance Policy
- 10 Server Assignment
- 11 Operational Policies

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default)

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
vNIC eth1	Derived	derived	
vNIC eth0	Derived	derived	

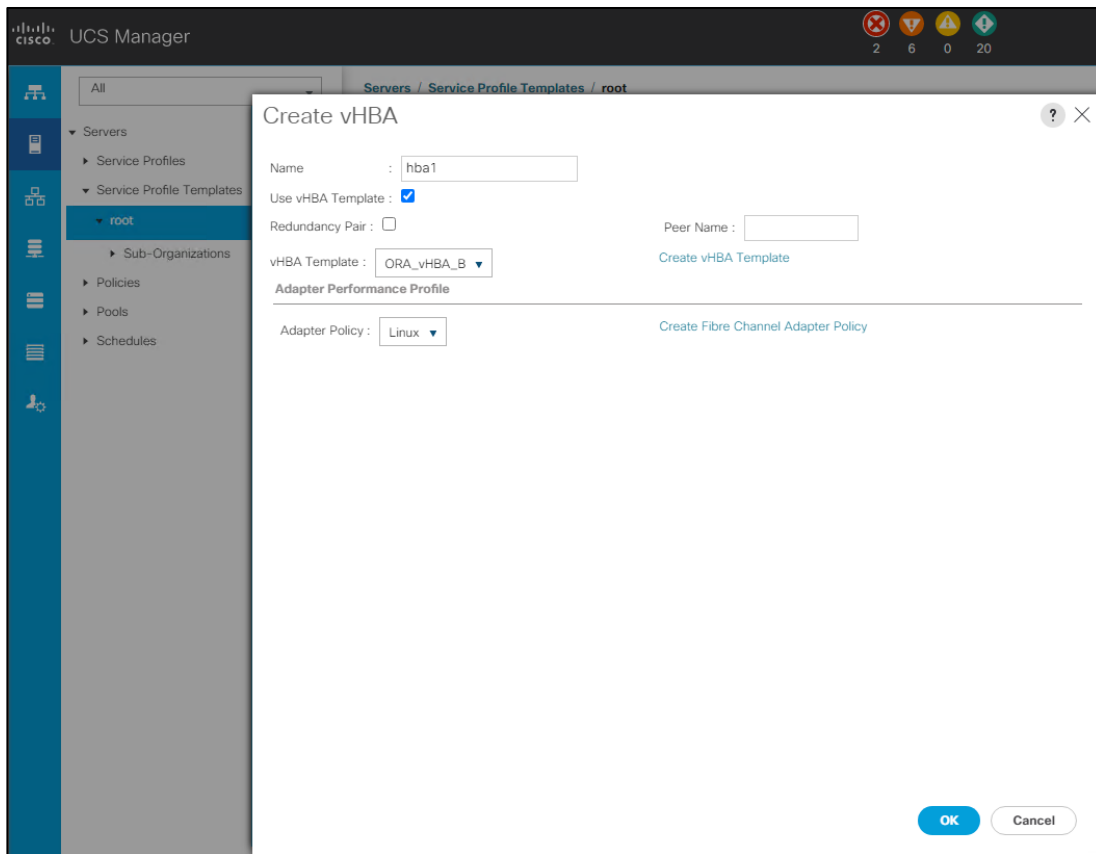
Delete Add Modify

+ iSCSI vNICs

< Prev Next > Finish Cancel

- h. Click Next. In the SAN Connectivity menu, select Expert to configure SAN connectivity. Select the WWNN (World Wide Node Name) pool created previously. Click Add to add vHBAs as shown below. The following four HBA were created:
- hba0 using vHBA template Oracle-HBA-A
 - hba1 using vHBA template Oracle-HBA-B
 - hba2 using vHBA template Oracle-HBA-A
 - hba3 using vHBA template Oracle-HBA-B

The screenshot displays the 'Create vHBA' window in the Cisco UCS Manager. The window has a title bar with the Cisco logo and 'UCS Manager' text. On the right side of the title bar, there are four status icons with counts: 2 (red), 5 (orange), 0 (yellow), and 20 (green). The main content area is divided into two sections. The top section contains the following fields: 'Name' with the value 'hba0', 'Use vHBA Template' with a checked checkbox, 'Redundancy Pair' with an unchecked checkbox, 'vHBA Template' with a dropdown menu showing 'ORA_vHBA_A', and 'Adapter Policy' with a dropdown menu showing 'Linux'. There are also two links: 'Create vHBA Template' and 'Create Fibre Channel Adapter Policy'. The bottom section contains 'OK' and 'Cancel' buttons. The background shows the UCS Manager interface with a sidebar on the left containing icons for Servers, Service Profiles, Service Profile Templates, root, Sub-Organizations, Policies, Pools, and Schedules. A vertical bar on the right side of the sidebar contains numbers 1 through 11.



UCS Manager

Servers / Service Profile Templates / root

Create vHBA

Name : hba2

Use vHBA Template : ☒

Redundancy Pair : ☐

Peer Name :

vHBA Template : ORA_vHBA_A

Adapter Performance Profile

Adapter Policy : Linux

[Create vHBA Template](#)

[Create Fibre Channel Adapter Policy](#)

OK Cancel

UCS Manager

2 5 0 20

All

- Servers
 - Service Profiles
 - Service Profile Templates
 - root
 - Sub-Organizations
 - Policies
 - Pools
 - Schedules

Create vHBA

Name : hba3

Use vHBA Template : ☒

Redundancy Pair : ☐

Peer Name :

vHBA Template : ORA_vHBA_B [Create vHBA Template](#)

Adapter Performance Profile

Adapter Policy : Linux [Create Fibre Channel Adapter Policy](#)

OK Cancel

UCS Manager

2 4 0 19

All

- Servers
 - Service Profiles
 - Service Profile Templates
 - root
 - Sub-Organizations
 - Policies
 - Pools
 - Schedules

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

☐ Simple ☒ Expert ☐ No vHBAs ☐ Use Connectivity Policy

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.

World Wide Node Name

WWNN Assignment: ORA_WWNN_Pool(256/256)

The WWNN will be assigned from the selected pool.
The available/total WWNNs are displayed after the pool name.

Name	WWPN
▶ vHBA hba3	Derived
▶ vHBA hba2	Derived
▶ vHBA hba1	Derived
▶ vHBA hba0	Derived

< Prev Next > Finish Cancel

- i. For this Oracle RAC configuration, the Cisco MDS 9132T is used for zoning. Skip zoning and go to the next step.
- j. In the vNIC/vHBA Placement menu, keep the option Let System Perform Placement.
- k. Do not configure any vMedia Policy. Click Next.
- l. For the Server Boot Order, select SAN_Boot as the boot policy created previously.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **SAN_Boot** [Create Boot Policy](#)

Name : **SAN_Boot**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Legacy**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

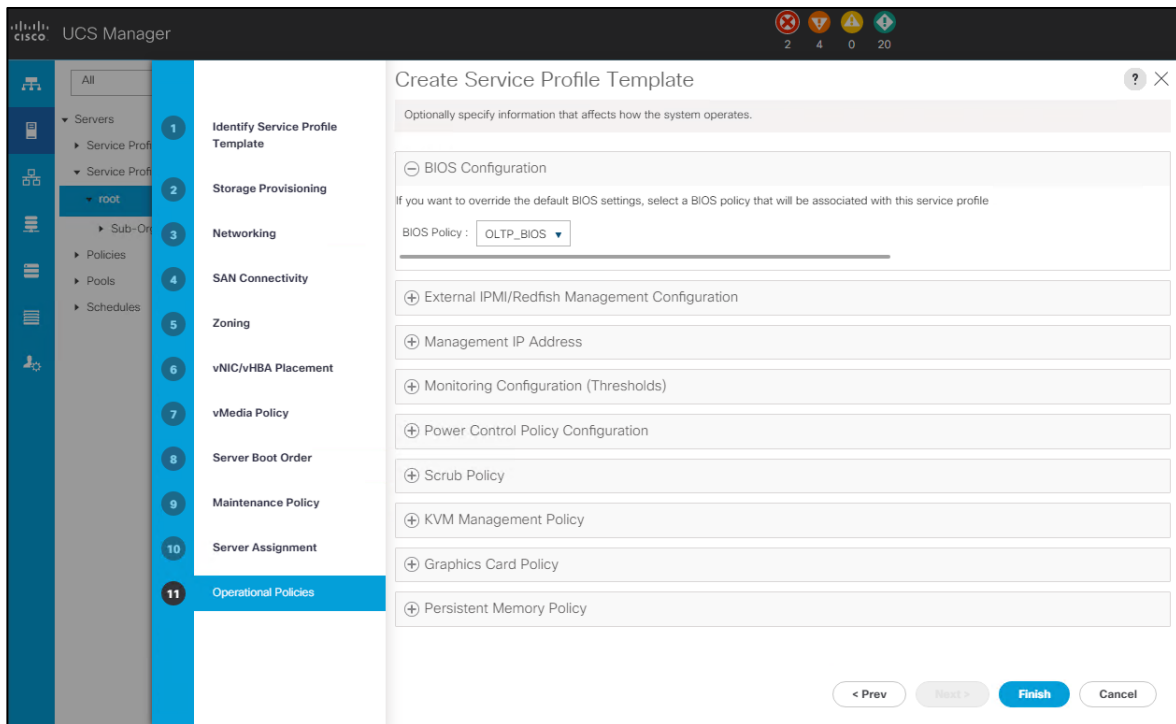
Boot Order

Name	O...	vNIC...	Type	LUN ...	WWN	Slot ...	Boot...	Boot ...	Desc...
CD/DVD	1								
▼ San	2								
▼ SAN Primary		hba0	Prim...						
SAN Target Primary			Prim...	0	20:03:D0:39:EA:20:BA:D3				
SAN Target Secon...			Sec...	0	20:13:D0:39:EA:20:BA:D3				
▼ SAN Secondary		hba1	Sec...						

[Create iSCSI vNIC](#) [Set iSCSI Boot Parameters](#) [Set iSCSI Boot Parameters](#)

< Prev Next > **Finish** Cancel

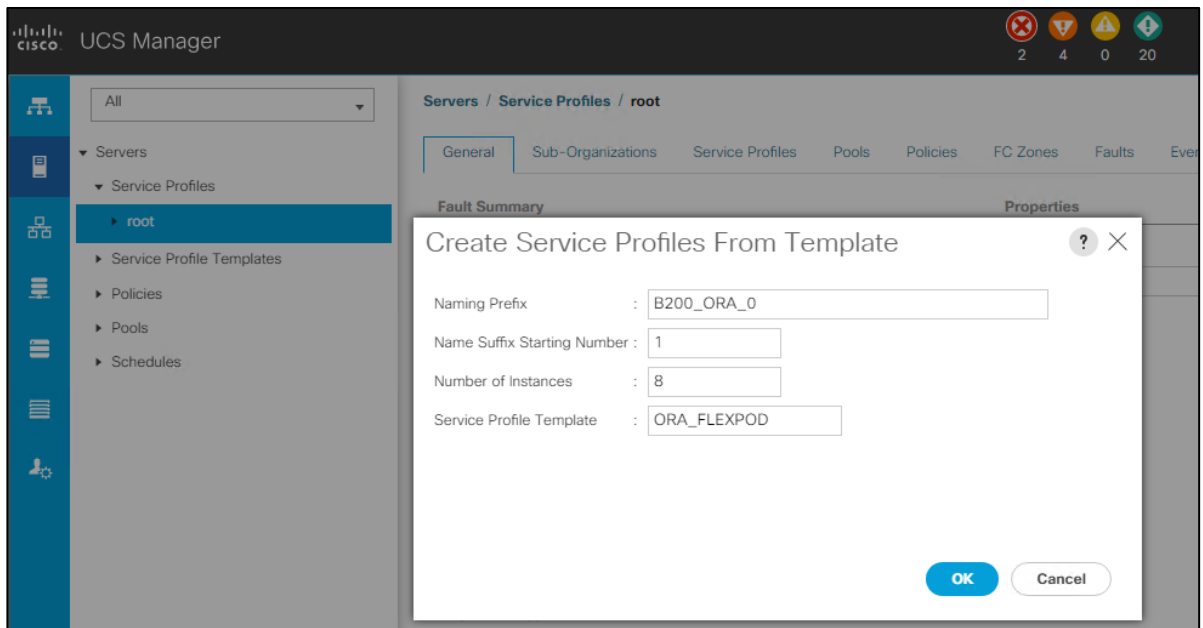
- m. The maintenance policy and server assignment options were left as default in the configuration.
- n. In the Operational Policies menu, select OLTP_BIOS for the BIOS policy created previously.



- o. The rest of the configuration is left as the default configuration. However, your configuration might vary from site-to-site depending on workloads, best practices, and policies.
- p. Click Finish to create the ORA_FLEXPOD service profile template. This service profile template is used to create eight service profiles for eight oracle RAC nodes (b200_ora_01, b200_ora_02, b200_ora_03, b200_ora_04, b200_ora_05, b200_ora_06, b200_ora_07, and b200_ora_08).

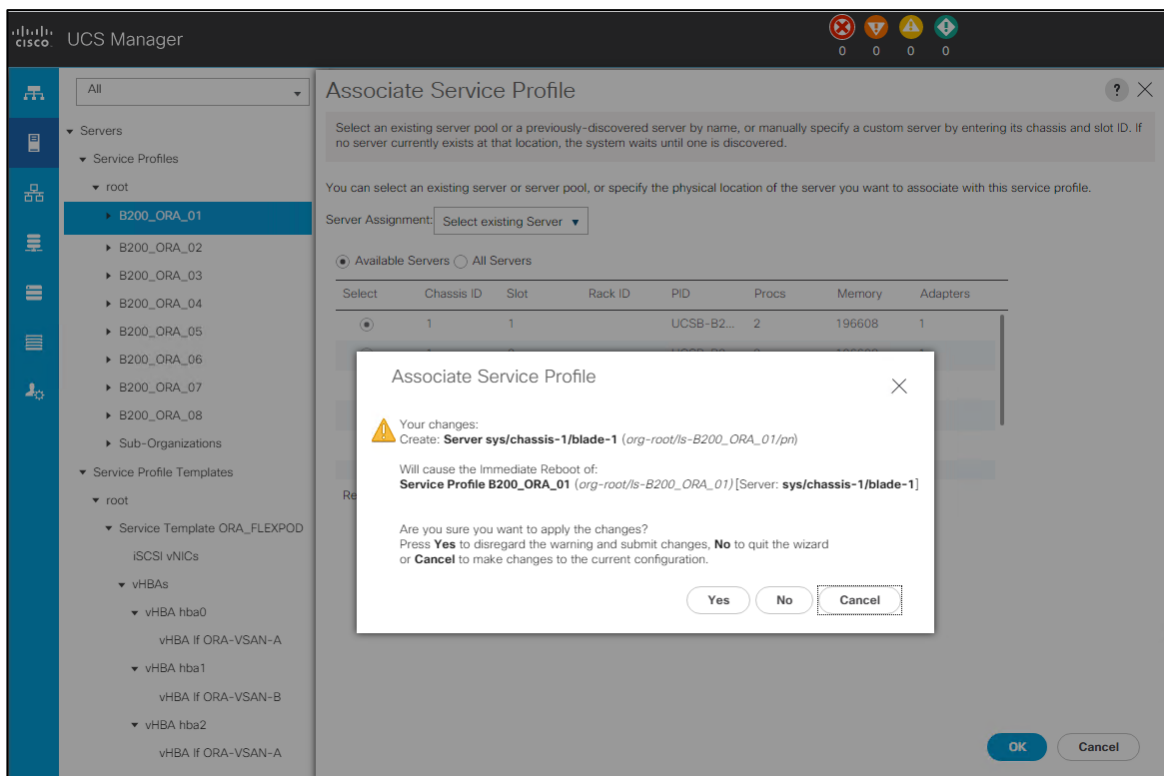
Now the ORA_FLEXPOD service profile template covers four vHBAs and two vNICs.

4. Create service profiles from a template and associate to servers. For all eight Oracle RAC nodes, create eight service profiles: b200_ora_01, b200_ora_02, b200_ora_03, b200_ora_04, b200_ora_05, b200_ora_06, b200_ora_07, and b200_ora_08 from the template ORA_FLEXPOD.
 - a. Go to the tab Servers > Service Profiles > Root > and right-click Create Service Profiles from Template.
 - b. Select the service profile template ORA_FLEXPOD that you created earlier and name the service profile FLEX.
 - c. To create eight service profiles, enter the number of instances as four as shown below. This process creates the following service profiles: B200_ORA_01, B200_ORA_02, B200_ORA_03, B200_ORA_04, B200_ORA_05, B200_ORA_06, B200_ORA_07, and B200_ORA_08.



5. Associate service profiles to the servers.

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



6. Assign service profiles B200_ORA_01 to Chassis 1 Server 1 and service profile B200_ORA_04 to Chassis 1 Server 4. Similarly, assign, service profiles B200_ORA_05 to Chassis 2 Server 1 and service profile B200_ORA_08 to Chassis 2 Server 4.

This concludes the UCS setup configuration for FlexPod. The Oracle blades servers are ready for SAN boot.

Cisco MDS setup

The MDS switch provides connectivity between the fabric interconnects and the NetApp AFF A800 controllers. Table 11 and Table 12 show the port connections between MDS, FI, and AFF A800.

Table 11) Port connections between MDS, FI, and AFF A800.

MDS Switch	MDS Switch port	FI port	Fabric Interconnect
MDS-ORA-01	5	1	FI-ORA-01
MDS-ORA-01	6	2	FI-ORA-01
MDS-ORA-01	7	3	FI-ORA-01
MDS-ORA-01	8	4	FI-ORA-01
MDS-ORA-02	5	1	FI-ORA-02
MDS-ORA-02	6	2	FI-ORA-02
MDS-ORA-02	7	3	FI-ORA-02
MDS-ORA-02	8	4	FI-ORA-02

Table 12) Port connections between MDS, FI, and AFF A800, continued.

MDS Switch	MDS Switch Port	AFF A800 FC Port	AFF A800 controller node
MDS-ORA-01	1	2a	Node 1
MDS-ORA-01	2	2b	Node 1
MDS-ORA-01	3	2c	Node 2
MDS-ORA-01	4	2d	Node 2
MDS-ORA-02	1	2a	Node 2
MDS-ORA-02	2	2b	Node 2
MDS-ORA-02	3	2c	Node 1
MDS-ORA-02	4	2d	Node 1

Initial setup

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter power on auto provisioning. Enter y to get to the system admin account setup. Complete the following on Switch A and B.

```
---- System Admin Account Setup ----

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

```

Mgmt0 IPv4 netmask : <mds-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <mds-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter
Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter
Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter
Enable the http-server? (yes/no) [y]: Enter
Configure clock? (yes/no) [n]: Enter
Configure timezone? (yes/no) [n]: Enter
Configure summertime? (yes/no) [n]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <nexus-A-mgmt0-ip>
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: Enter
Configure default zone mode (basic/enhanced) [basic]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter

```

Enable licenses

Enable the following features on both switches.

```

cie-c9132t-g0139(config)# config terminal
cie-c9132t-g0139(config)# feature npiv
cie-c9132t-g0139(config)# feature telnet
cie-c9132t-g0139(config)# fport-channel-trunk
cie-c9132t-g0139(config)# sshServer
cie-c9132t-g0139(config)# isapi
cie-c9132t-g0139(config)# switchname MDS-ORA-02
MDS-ORA-02(config)# copy running-config startup-config
[#####] 100%
Copy complete.

```

Add second NTP server and local time configuration

```

ntp server <nexus-B-mgmt0-ip>
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-day> <end-month> <end-time> <offset-minutes>
e.g.
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60

```

Configure VSANs for MDS Switch A and MDS Switch B

1. Log in as the admin user into MDS Switch A.
2. Create VSAN 101 for storage traffic:

```

MDS-ORA-01(config)# configure terminal
MDS-ORA-01(config)# VSAN database
DS-ORA-01(config-vsan-db)# vsan 101
MDS-ORA-01(config-vsan-db)# vsan 101 interface fc 1/1-8
Traffic on fc1/1 may be impacted. Do you want to continue? (y/n) [n] y
Traffic on fc1/2 may be impacted. Do you want to continue? (y/n) [n] y
MDS-ORA-01(config-vsan-db)# exit
MDS-ORA-01(config)# interface fc 1/1-8
MDS-ORA-01(config-if)# switchport trunk allowed vsan 101

```



```
Warning: This command will remove all VSANs currently being trunked and trunk only the specified VSANs.
Do you want to continue? (y/n) [n] y
MDS-ORA-01(config-if)# switchport trunk mode off
MDS-ORA-01(config-if)# port-license acquire
MDS-ORA-01(config-if)# no shutdown
MDS-ORA-01(config-if)# exit
MDS-ORA-01(config)# copy running-config startup-config
[#####] 100%
Copy complete.
```

3. Log in as the admin user into MDS Switch B.

4. Create VSAN 102 for storage traffic:

```
MDS-ORA-02# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
MDS-ORA-02(config)# VSAN database
MDS-ORA-02(config-vsan-db)# vsan 102
MDS-ORA-02(config-vsan-db)# vsan 102 interface fc 1/1-8
Traffic on fc1/1 may be impacted. Do you want to continue? (y/n) [n] y
Traffic on fc1/2 may be impacted. Do you want to continue? (y/n) [n]
MDS-ORA-02(config-vsan-db)# exit
MDS-ORA-02(config)# interface fc1/1-8
MDS-ORA-02(config-if)# switchport trunk allowed vsan 102
Warning: This command will remove all VSANs currently being trunked and trunk only the specified VSANs.
Do you want to continue? (y/n) [n] y
MDS-ORA-02(config-if)# switchport trunk mode off
MDS-ORA-02(config-if)# port-license acquire
MDS-ORA-02(config-if)# no shutdown
MDS-ORA-02(config-if)# exit
MDS-ORA-02(config)# copy running-config startup-config
[#####] 100%
```

Configure port-channel 101

Log into MDS switch A as admin user and execute following commands.

```
interface port-channel101
switchport mode F
switchport trunk allowed vsan 101
switchport description ucs-fi-a
switchport rate-mode dedicated
switchport speed 32000
channel mode active
no shutdown

interface fc1/5
switchport speed 32000
switchport mode F
port-license acquire
channel-group 101 force
no shutdown

interface fc1/6
switchport speed 32000
switchport mode F
port-license acquire
channel-group 101 force
no shutdown

interface fc1/7
switchport speed 32000
switchport mode F
port-license acquire
channel-group 101 force
no shutdown

interface fc1/8
```

```
switchport speed 32000
switchport mode F
port-license acquire
channel-group 101 force
no shutdown
```

```
MDS-ORA-01# show port-channel summary
-----
Interface                Total Ports      Oper Ports      First Oper Port
-----
port-channel 101          4                4                fc1/6
MDS-ORA-01#
```

Configure port-channel 102

Login to MDS switch B as the admin user and run the following commands.

```
interface port-channel102
switchport mode F
switchport trunk allowed vsan 102
switchport description ucs-fi-b
switchport rate-mode dedicated
switchport speed 32000
channel mode active
no shutdown

interface fc1/5
switchport speed 32000
switchport mode F
port-license acquire
channel-group 102 force
no shutdown

interface fc1/6
switchport speed 32000
switchport mode F
port-license acquire
channel-group 102 force
no shutdown

interface fc1/7
switchport speed 32000
switchport mode F
port-license acquire
channel-group 102 force
no shutdown

interface fc1/8
switchport speed 32000
switchport mode F
port-license acquire
channel-group 102 force
no shutdown
```

```
MDS-ORA-02# show port-channel summary
-----
Interface                Total Ports      Oper Ports      First Oper Port
-----
port-channel 102          4                4                fc1/8
MDS-ORA-02#
```

Create and configure FC zoning

This procedure sets up the FC connections between the Cisco MDS 9132T switches, the Cisco UCS fabric interconnects, and the NetApp AFF storage systems.

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

To create and configure FC zoning, complete the following steps:

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

Equipment / Chassis / Chassis 1 / Servers / Server 1						
General	Inventory	Virtual Machines	Installed Firmware	CIMC Sessions	SEL Logs	VIF Paths
Motherboard	CIMC	CPUs	GPUs	Memory	Adapters	HBAs
NICs						
iSCSI vNICs						
Security						
Storage						
Persistent Memory						
Advanced Filter Export Print						
Name	Adapter ID	vHBA	Vendor	PID	Operability	WWPN
HBA 1	1	hba0	Cisco Systems Inc	UCSB-MLOM-40G-04	Operable	20:00:00:25:B5:8A:A0:00
HBA 2	1	hba1	Cisco Systems Inc	UCSB-MLOM-40G-04	Operable	20:00:00:25:B5:88:B0:00
HBA 3	1	hba2	Cisco Systems Inc	UCSB-MLOM-40G-04	Operable	20:00:00:25:B5:8A:A0:01
HBA 4	1	hba3	Cisco Systems Inc	UCSB-MLOM-40G-04	Operable	20:00:00:25:B5:88:B0:01

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

Network Interfaces							
Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols
ora19c_svm_data_fcp_lif_1_2a	✓	ora19c_svm		20:0b:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2a	FC
ora19c_svm_data_fcp_lif_1_2b	✓	ora19c_svm		20:0c:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2b	FC
ora19c_svm_data_fcp_lif_2_2a	✓	ora19c_svm		20:0d:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2a	FC
ora19c_svm_data_fcp_lif_2_2b	✓	ora19c_svm		20:12:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2b	FC
ora19c_svm_data_fcp_lif_1_2c	✓	ora19c_svm		20:17:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2c	FC
ora19c_svm_data_fcp_lif_1_2d	✓	ora19c_svm		20:18:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2d	FC
ora19c_svm_data_fcp_lif_2_2d	✓	ora19c_svm		20:19:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2d	FC
ora19c_svm_data_fcp_lif_2_2c	✓	ora19c_svm		20:1a:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2c	FC

The screenshot above shows the network interface, WWPN and ports connectivity configured for NetApp AFF A800 controller. Four FC logical interfaces (LIFs) are created on storage controller cluster node 1 (ora19c_svm_data_fcp_lif_1_2a, ora19c_svm_data_fcp_lif_1_2b, ora19c_svm_data_fcp_lif_1_2c, and ora19c_svm_data_fcp_lif_1_2d), and four FC LIFs are created on storage controller cluster node 2 (ora19c_svm_data_fcp_lif_2_2a, ora19c_svm_data_fcp_lif_2_2b, ora19c_svm_data_fcp_lif_2_2c, and ora19c_svm_data_fcp_lif_2_2d).

Note: You can also obtain this information by logging into the storage cluster and running the network interface show command.

The NetApp Storage AFF A800s have eight active FC connections to the Cisco MDS switches. Four FC ports are connected to Cisco MDS-A, and another four FC ports are connected to the Cisco MDS-B switches. The SAN ports 2a and 2b for NetApp AFF A800s Controller node 1 are connected to Cisco

MDS Switch A, and SAN ports 2c and 2d of node 1 are connected to Cisco MDS Switch B. Similarly, SAN ports 2a and 2b of NetApp AFF A800s controller node 2 are connected to Cisco MDS Switch B, and SAN ports 2c and 2d of node 2 are connected to Cisco MDS Switch A.

3. Create device aliases for FC zoning on Cisco MDS Switch A.

To configure device aliases and zones for the SAN boot paths and data paths for MDS switch A, log in as the admin user and run the following commands:

```
configure terminal
device-alias database
device-alias name B200_ORA_01_hba0 pwwn 20:00:00:25:B5:8A:A0:00
device-alias name B200_ORA_01_hba2 pwwn 20:00:00:25:B5:8A:A0:01
device-alias name B200_ORA_02_hba0 pwwn 20:00:00:25:B5:8A:A0:02
device-alias name B200_ORA_02_hba2 pwwn 20:00:00:25:B5:8A:A0:03
device-alias name B200_ORA_03_hba0 pwwn 20:00:00:25:B5:8A:A0:04
device-alias name B200_ORA_03_hba2 pwwn 20:00:00:25:B5:8A:A0:05
device-alias name B200_ORA_04_hba0 pwwn 20:00:00:25:B5:8A:A0:06
device-alias name B200_ORA_04_hba2 pwwn 20:00:00:25:B5:8A:A0:07
device-alias name B200_ORA_05_hba0 pwwn 20:00:00:25:B5:8A:A0:08
device-alias name B200_ORA_05_hba2 pwwn 20:00:00:25:B5:8A:A0:09
device-alias name B200_ORA_06_hba0 pwwn 20:00:00:25:B5:8A:A0:0A
device-alias name B200_ORA_06_hba2 pwwn 20:00:00:25:B5:8A:A0:0B
device-alias name B200_ORA_07_hba0 pwwn 20:00:00:25:B5:8A:A0:0C
device-alias name B200_ORA_07_hba2 pwwn 20:00:00:25:B5:8A:A0:0D
device-alias name B200_ORA_08_hba0 pwwn 20:00:00:25:B5:8A:A0:0E
device-alias name B200_ORA_08_hba2 pwwn 20:00:00:25:B5:8A:A0:0F
device-alias name NetApp-A800-01-2A pwwn 20:0b:d0:39:ea:20:ba:d3
device-alias name NetApp-A800-01-2B pwwn 20:0c:d0:39:ea:20:ba:d3
device-alias name NetApp-A800-02-2C pwwn 20:1a:d0:39:ea:20:ba:d3
device-alias name NetApp-A800-02-2D pwwn 20:19:d0:39:ea:20:ba:d3
device-alias name A800-01_infra_svm_2A pwwn 20:03:d0:39:ea:20:ba:d3
device-alias name A800-01_infra_svm_2B pwwn 20:05:d0:39:ea:20:ba:d3
device-alias name A800-02_infra_svm_2C pwwn 20:15:d0:39:ea:20:ba:d3
device-alias name A800-02_infra_svm_2D pwwn 20:16:d0:39:ea:20:ba:d3
device-alias commit
copy run start
```

4. Create device aliases for FC zoning on Cisco MDS Switch B. To configure device aliases and zones for the SAN boot paths as well as data paths of MDS switch B, complete the following steps:

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

```
configure terminal
device-alias database
device-alias name B200_ORA_01_hba1 pwwn 20:00:00:25:B5:8B:B0:00
device-alias name B200_ORA_01_hba3 pwwn 20:00:00:25:B5:8B:B0:01
device-alias name B200_ORA_02_hba1 pwwn 20:00:00:25:B5:8B:B0:02
device-alias name B200_ORA_02_hba3 pwwn 20:00:00:25:B5:8B:B0:03
device-alias name B200_ORA_03_hba1 pwwn 20:00:00:25:B5:8B:B0:04
device-alias name B200_ORA_03_hba3 pwwn 20:00:00:25:B5:8B:B0:05
device-alias name B200_ORA_04_hba1 pwwn 20:00:00:25:B5:8B:B0:06
device-alias name B200_ORA_04_hba3 pwwn 20:00:00:25:B5:8B:B0:07
device-alias name B200_ORA_05_hba1 pwwn 20:00:00:25:B5:8B:B0:08
device-alias name B200_ORA_05_hba3 pwwn 20:00:00:25:B5:8B:B0:09
device-alias name B200_ORA_06_hba1 pwwn 20:00:00:25:B5:8B:B0:0A
device-alias name B200_ORA_06_hba3 pwwn 20:00:00:25:B5:8B:B0:0B
device-alias name B200_ORA_07_hba1 pwwn 20:00:00:25:B5:8B:B0:0C
device-alias name B200_ORA_07_hba3 pwwn 20:00:00:25:B5:8B:B0:0D
device-alias name B200_ORA_08_hba1 pwwn 20:00:00:25:B5:8B:B0:0E
device-alias name B200_ORA_08_hba3 pwwn 20:00:00:25:B5:8B:B0:0F
device-alias name NetApp-A800-01-2C pwwn 20:17:d0:39:ea:20:ba:d3
device-alias name NetApp-A800-01-2D pwwn 20:18:d0:39:ea:20:ba:d3
device-alias name NetApp-A800-02-2A pwwn 20:0d:d0:39:ea:20:ba:d3
device-alias name NetApp-A800-02-2B pwwn 20:12:d0:39:ea:20:ba:d3
device-alias name A800-1_infra_svm_2C pwwn 20:13:d0:39:ea:20:ba:d3
device-alias name A800-1_infra_svm_2D pwwn 20:14:d0:39:ea:20:ba:d3
device-alias name A800-2_infra_svm_2A pwwn 20:08:d0:39:ea:20:ba:d3
device-alias name A800-2_infra_svm_2B pwwn 20:09:d0:39:ea:20:ba:d3
```

```
device-alias commit
copy run start
```

5. Create zoning on MDS Switch A. To configure zones for the Cisco MDS Switch A, complete the following steps:

- a. Create a zone for each service profile as shown below. Log in as the admin user into MDS Switch A and run these commands to create the zone for each compute node for data and boot LUN access. Data and boot LUNs are created separately in ora19c_svm and infra_svm. Therefore, two zones must be created for each node.

```
configure terminal
zone name B200_ORA_01 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:00
  member pwwn 20:00:00:25:B5:8A:A0:01
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_02 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:02
  member pwwn 20:00:00:25:B5:8A:A0:03
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_03 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:04
  member pwwn 20:00:00:25:B5:8A:A0:05
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_04 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:06
  member pwwn 20:00:00:25:B5:8A:A0:07
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_05 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:08
  member pwwn 20:00:00:25:B5:8A:A0:09
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_06 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:0A
  member pwwn 20:00:00:25:B5:8A:A0:0B
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_07 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:0C
  member pwwn 20:00:00:25:B5:8A:A0:0D
  member pwwn 20:0b:d0:39:ea:20:ba:d3
  member pwwn 20:0c:d0:39:ea:20:ba:d3
  member pwwn 20:1a:d0:39:ea:20:ba:d3
  member pwwn 20:19:d0:39:ea:20:ba:d3

zone name B200_ORA_08 vsan 101
  member pwwn 20:00:00:25:B5:8A:A0:0E
```

```

member pwnn 20:00:00:25:B5:8A:A0:0F
member pwnn 20:0b:d0:39:ea:20:ba:d3
member pwnn 20:0c:d0:39:ea:20:ba:d3
member pwnn 20:1a:d0:39:ea:20:ba:d3
member pwnn 20:19:d0:39:ea:20:ba:d3

zone name BOOT_ORA_01 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:00
member pwnn 20:00:00:25:B5:8A:A0:01
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_02 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:02
member pwnn 20:00:00:25:B5:8A:A0:03
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_03 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:04
member pwnn 20:00:00:25:B5:8A:A0:05
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_04 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:06
member pwnn 20:00:00:25:B5:8A:A0:07
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_05 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:08
member pwnn 20:00:00:25:B5:8A:A0:09
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_06 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:0A
member pwnn 20:00:00:25:B5:8A:A0:0B
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_07 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:0C
member pwnn 20:00:00:25:B5:8A:A0:0D
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

zone name BOOT_ORA_08 vsan 101
member pwnn 20:00:00:25:B5:8A:A0:0E
member pwnn 20:00:00:25:B5:8A:A0:0F
member pwnn 20:03:d0:39:ea:20:ba:d3
member pwnn 20:05:d0:39:ea:20:ba:d3
member pwnn 20:15:d0:39:ea:20:ba:d3
member pwnn 20:16:d0:39:ea:20:ba:d3

```

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

```
zoneset name B200_ORA vsan 101
member B200_ORA_01
member B200_ORA_02
member B200_ORA_03
member B200_ORA_04
member B200_ORA_05
member B200_ORA_06
member B200_ORA_07
member B200_ORA_08
member BOOT_ORA_01
member BOOT_ORA_02
member BOOT_ORA_03
member BOOT_ORA_04
member BOOT_ORA_05
member BOOT_ORA_06
member BOOT_ORA_07
member BOOT_ORA_08
```

- c. Activate the zone set by running following commands.

```
zoneset activate name B200_ORA vsan 101
copy run start
```

6. Create zoning on Cisco MDS Switch B. To configure zones for the Cisco MDS Switch B, complete the following steps:

- a. Create a zone for each service profile as shown below. Log in as the admin user to MDS Switch A and run these commands to create the zone:

```
configure terminal
zone name B200_ORA_01 vsan 102
member pwwn 20:00:00:25:B5:8B:B0:00
member pwwn 20:00:00:25:B5:8B:B0:01
member pwwn 20:17:d0:39:ea:20:ba:d3
member pwwn 20:18:d0:39:ea:20:ba:d3
member pwwn 20:0d:d0:39:ea:20:ba:d3
member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_02 vsan 102
member pwwn 20:00:00:25:B5:8B:B0:02
member pwwn 20:00:00:25:B5:8B:B0:03
member pwwn 20:17:d0:39:ea:20:ba:d3
member pwwn 20:18:d0:39:ea:20:ba:d3
member pwwn 20:0d:d0:39:ea:20:ba:d3
member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_03 vsan 102
member pwwn 20:00:00:25:B5:8B:B0:04
member pwwn 20:00:00:25:B5:8B:B0:05
member pwwn 20:17:d0:39:ea:20:ba:d3
member pwwn 20:18:d0:39:ea:20:ba:d3
member pwwn 20:0d:d0:39:ea:20:ba:d3
member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_04 vsan 102
member pwwn 20:00:00:25:B5:8B:B0:06
member pwwn 20:00:00:25:B5:8B:B0:07
member pwwn 20:17:d0:39:ea:20:ba:d3
member pwwn 20:18:d0:39:ea:20:ba:d3
member pwwn 20:0d:d0:39:ea:20:ba:d3
member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_05 vsan 102
member pwwn 20:00:00:25:B5:8B:B0:08
member pwwn 20:00:00:25:B5:8B:B0:09
member pwwn 20:17:d0:39:ea:20:ba:d3
member pwwn 20:18:d0:39:ea:20:ba:d3
member pwwn 20:0d:d0:39:ea:20:ba:d3
```

```

    member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_06 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:0A
    member pwwn 20:00:00:25:B5:8B:B0:0B
    member pwwn 20:17:d0:39:ea:20:ba:d3
    member pwwn 20:18:d0:39:ea:20:ba:d3
    member pwwn 20:0d:d0:39:ea:20:ba:d3
    member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_07 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:0C
    member pwwn 20:00:00:25:B5:8B:B0:0D
    member pwwn 20:17:d0:39:ea:20:ba:d3
    member pwwn 20:18:d0:39:ea:20:ba:d3
    member pwwn 20:0d:d0:39:ea:20:ba:d3
    member pwwn 20:12:d0:39:ea:20:ba:d3

zone name B200_ORA_08 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:0E
    member pwwn 20:00:00:25:B5:8B:B0:0F
    member pwwn 20:17:d0:39:ea:20:ba:d3
    member pwwn 20:18:d0:39:ea:20:ba:d3
    member pwwn 20:0d:d0:39:ea:20:ba:d3
    member pwwn 20:12:d0:39:ea:20:ba:d3

zone name BOOT_ORA_01 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:00
    member pwwn 20:00:00:25:B5:8B:B0:01
    member pwwn 20:13:d0:39:ea:20:ba:d3
    member pwwn 20:14:d0:39:ea:20:ba:d3
    member pwwn 20:08:d0:39:ea:20:ba:d3
    member pwwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_02 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:02
    member pwwn 20:00:00:25:B5:8B:B0:03
    member pwwn 20:13:d0:39:ea:20:ba:d3
    member pwwn 20:14:d0:39:ea:20:ba:d3
    member pwwn 20:08:d0:39:ea:20:ba:d3
    member pwwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_03 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:04
    member pwwn 20:00:00:25:B5:8B:B0:05
    member pwwn 20:13:d0:39:ea:20:ba:d3
    member pwwn 20:14:d0:39:ea:20:ba:d3
    member pwwn 20:08:d0:39:ea:20:ba:d3
    member pwwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_04 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:06
    member pwwn 20:00:00:25:B5:8B:B0:07
    member pwwn 20:13:d0:39:ea:20:ba:d3
    member pwwn 20:14:d0:39:ea:20:ba:d3
    member pwwn 20:08:d0:39:ea:20:ba:d3
    member pwwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_05 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:08
    member pwwn 20:00:00:25:B5:8B:B0:09
    member pwwn 20:13:d0:39:ea:20:ba:d3
    member pwwn 20:14:d0:39:ea:20:ba:d3
    member pwwn 20:08:d0:39:ea:20:ba:d3
    member pwwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_06 vsan 102
    member pwwn 20:00:00:25:B5:8B:B0:0a
    member pwwn 20:00:00:25:B5:8B:B0:0b
    member pwwn 20:13:d0:39:ea:20:ba:d3
    member pwwn 20:14:d0:39:ea:20:ba:d3

```



```

member pwn 20:08:d0:39:ea:20:ba:d3
member pwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_07 vsan 102
member pwn 20:00:00:25:B5:8B:B0:0c
member pwn 20:00:00:25:B5:8B:B0:0d
member pwn 20:13:d0:39:ea:20:ba:d3
member pwn 20:14:d0:39:ea:20:ba:d3
member pwn 20:08:d0:39:ea:20:ba:d3
member pwn 20:09:d0:39:ea:20:ba:d3

zone name BOOT_ORA_08 vsan 102
member pwn 20:00:00:25:B5:8B:B0:0e
member pwn 20:00:00:25:B5:8B:B0:0f
member pwn 20:13:d0:39:ea:20:ba:d3
member pwn 20:14:d0:39:ea:20:ba:d3
member pwn 20:08:d0:39:ea:20:ba:d3
member pwn 20:09:d0:39:ea:20:ba:d3

```

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

```

zoneset name B200_ORA vsan 102
member B200_ORA_01
member B200_ORA_02
member B200_ORA_03
member B200_ORA_04
member B200_ORA_05
member B200_ORA_06
member B200_ORA_07
member B200_ORA_08
member BOOT_ORA_01
member BOOT_ORA_02
member BOOT_ORA_03
member BOOT_ORA_04
member BOOT_ORA_05
member BOOT_ORA_06
member BOOT_ORA_07
member BOOT_ORA_08

```

- c. Activate the zone set by running following commands:

```

zoneset activate name B200_ORA vsan 102
copy run start

```

7. Verify host HBA connectivity to MDS switch. To verify the host HBAs connections to the MDS switch, complete the following steps:
- Log in as the admin user to MDS Switch A and run the `show flogi database vsan 101` command to verify that hosts hba0 and hba2 are connected to MDS Switch A.

```

port-channel101 101 0x070080 24:65:00:3a:9c:ad:4b:00 20:65:00:3a:9c:ad:4b:01
port-channel101 101 0x070081 20:00:00:25:b5:8a:a0:00 20:00:00:25:b5:7a:00:00
[B200_ORA_01_hba0]
port-channel101 101 0x070082 20:00:00:25:b5:8a:a0:02 20:00:00:25:b5:7a:00:01
[B200_ORA_02_hba0]
port-channel101 101 0x070083 20:00:00:25:b5:8a:a0:04 20:00:00:25:b5:7a:00:02
[B200_ORA_03_hba0]
port-channel101 101 0x070084 20:00:00:25:b5:8a:a0:08 20:00:00:25:b5:7a:00:04
[B200_ORA_05_hba0]
port-channel101 101 0x070085 20:00:00:25:b5:8a:a0:06 20:00:00:25:b5:7a:00:03
[B200_ORA_04_hba0]
port-channel101 101 0x070086 20:00:00:25:b5:8a:a0:0a 20:00:00:25:b5:7a:00:05
[B200_ORA_06_hba0]
port-channel101 101 0x070087 20:00:00:25:b5:8a:a0:0c 20:00:00:25:b5:7a:00:06
[B200_ORA_07_hba0]
port-channel101 101 0x070088 20:00:00:25:b5:8a:a0:0e 20:00:00:25:b5:7a:00:07
[B200_ORA_08_hba0]
port-channel101 101 0x070089 20:00:00:25:b5:8a:a0:01 20:00:00:25:b5:7a:00:00
[B200_ORA_01_hba2]
port-channel101 101 0x07008a 20:00:00:25:b5:8a:a0:05 20:00:00:25:b5:7a:00:02
[B200_ORA_03_hba2]
port-channel101 101 0x07008b 20:00:00:25:b5:8a:a0:03 20:00:00:25:b5:7a:00:01
[B200_ORA_02_hba2]
port-channel101 101 0x07008c 20:00:00:25:b5:8a:a0:07 20:00:00:25:b5:7a:00:03
[B200_ORA_04_hba2]
port-channel101 101 0x07008d 20:00:00:25:b5:8a:a0:09 20:00:00:25:b5:7a:00:04
[B200_ORA_05_hba2]
port-channel101 101 0x07008e 20:00:00:25:b5:8a:a0:0b 20:00:00:25:b5:7a:00:05
[B200_ORA_06_hba2]
port-channel101 101 0x07008f 20:00:00:25:b5:8a:a0:0f 20:00:00:25:b5:7a:00:07
[B200_ORA_08_hba2]
port-channel101 101 0x070090 20:00:00:25:b5:8a:a0:0d 20:00:00:25:b5:7a:00:06
[B200_ORA_07_hba2]

```

- b. Log in as the admin user to MDS Switch B and run the `show flogi database vsan 101` command to verify that hosts hdb1 and hba3 are connected to MDS Switch B.

```

port-channel102 102 0x090080 24:66:00:3a:9c:ad:49:e0 20:66:00:3a:9c:ad:49:e1
port-channel102 102 0x090081 20:00:00:25:b5:8b:b0:00 20:00:00:25:b5:7a:00:00
[B200_ORA_01_hba1]
port-channel102 102 0x090082 20:00:00:25:b5:8b:b0:02 20:00:00:25:b5:7a:00:01
[B200_ORA_02_hba1]
port-channel102 102 0x090083 20:00:00:25:b5:8b:b0:04 20:00:00:25:b5:7a:00:02
[B200_ORA_03_hba1]
port-channel102 102 0x090084 20:00:00:25:b5:8b:b0:08 20:00:00:25:b5:7a:00:04
[B200_ORA_05_hba1]
port-channel102 102 0x090085 20:00:00:25:b5:8b:b0:06 20:00:00:25:b5:7a:00:03
[B200_ORA_04_hba1]
port-channel102 102 0x090086 20:00:00:25:b5:8b:b0:0a 20:00:00:25:b5:7a:00:05
[B200_ORA_06_hba1]
port-channel102 102 0x090087 20:00:00:25:b5:8b:b0:0c 20:00:00:25:b5:7a:00:06
[B200_ORA_07_hba1]
port-channel102 102 0x090088 20:00:00:25:b5:8b:b0:0e 20:00:00:25:b5:7a:00:07
[B200_ORA_08_hba1]
port-channel102 102 0x090089 20:00:00:25:b5:8b:b0:01 20:00:00:25:b5:7a:00:00
[B200_ORA_01_hba3]
port-channel102 102 0x09008a 20:00:00:25:b5:8b:b0:05 20:00:00:25:b5:7a:00:02
[B200_ORA_03_hba3]
port-channel102 102 0x09008b 20:00:00:25:b5:8b:b0:03 20:00:00:25:b5:7a:00:01
[B200_ORA_02_hba3]
port-channel102 102 0x09008c 20:00:00:25:b5:8b:b0:07 20:00:00:25:b5:7a:00:03
[B200_ORA_04_hba3]
port-channel102 102 0x09008d 20:00:00:25:b5:8b:b0:09 20:00:00:25:b5:7a:00:04
[B200_ORA_05_hba3]
port-channel102 102 0x09008e 20:00:00:25:b5:8b:b0:0b 20:00:00:25:b5:7a:00:05
[B200_ORA_06_hba3]
port-channel102 102 0x09008f 20:00:00:25:b5:8b:b0:0f 20:00:00:25:b5:7a:00:07
[B200_ORA_08_hba3]
port-channel102 102 0x090090 20:00:00:25:b5:8b:b0:0d 20:00:00:25:b5:7a:00:06
[B200_ORA_07_hba3]

```

Storage AFF A800 deployment

In version 9.7, the ONTAP System Manager UI interface is updated and simplified for many configuration tasks. When applicable, we use both the CLI and System Manager UI to configure the AFF A800 storage controller HA pair for Oracle 19c solution deployment.

Cluster setup preparation

1. Complete the configuration worksheet.

Before running the setup script, complete the [cluster setup worksheet](#) in the ONTAP 9 Documentation Center. You must have access to the [NetApp Support site](#) to open the cluster setup worksheet.

2. Collect the ONTAP cluster node details.

Before running the setup script, review the configuration worksheets in the software setup section of the ONTAP 9 Documentation Center to learn about configuring ONTAP. Table 13 lists the information needed to configure two ONTAP nodes. Customize the cluster-detail values with the information applicable to your deployment.

Table 13) ONTAP node configuration details.

Cluster detail	Cluster detail value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
ONTAP 9.7 URL	<url-boot-software>

Configure Node 1

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when the following message displays:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

Note: If ONTAP 9.7 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.7 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M.

```
<node01-mgmt-ip> <node01-mgmt-mask> <node01-mgmt-gateway>
```

9. Enter the URL where the software can be found.

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.
13. Press Ctrl-C when the following message displays:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.

15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.

Note: The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node 02 configuration while the disks for node 01 are zeroing.

Configure Node 2

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when the following message displays:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

Note: If ONTAP 9.7 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.7 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14).

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M.

```
<node02-mgmt-ip> <node02-mgmt-mask> <node02-mgmt-gateway>
```

9. Enter the URL where the software can be found.

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.
13. Press Ctrl-C when the following message displays:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.

Note: The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Set up node and cluster

1. Follow the prompts to set up node 01.

```
Welcome to node setup.
```

```
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
```

"exit" or "quit" - if you want to quit the setup wizard.
Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your system.

For further information on AutoSupport, see:
<http://support.netapp.com/autosupport/>

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created

Use your web browser to complete cluster setup by accessing <https://<node01-mgmt-ip>>

Otherwise press Enter to complete cluster setup using the command line interface:

2. To complete cluster setup, open a web browser and navigate to <https://<node1-mgmt-ip>>.

Before proceeding to cluster setup, you need to have cluster details information listed in Table 14 ready.

Table 14) Cluster details.

Cluster detail	Cluster detail value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-sp-ip>
Node 01 service processor network mask	<node01-sp-mask>
Node 01 service processor gateway	<node01-sp-gateway>
Node 02 service processor IP address	<node02-sp-ip>
Node 02 service processor network mask	<node02-sp-mask>
Node 02 service processor gateway	<node02-sp-gateway>
Node 01 node name	<st-node01>
Node 02 node name	<st-node02>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server A IP address	<switch-a-ntp-ip>

Cluster detail	Cluster detail value
NTP server B IP address	<switch-b-ntp-ip>

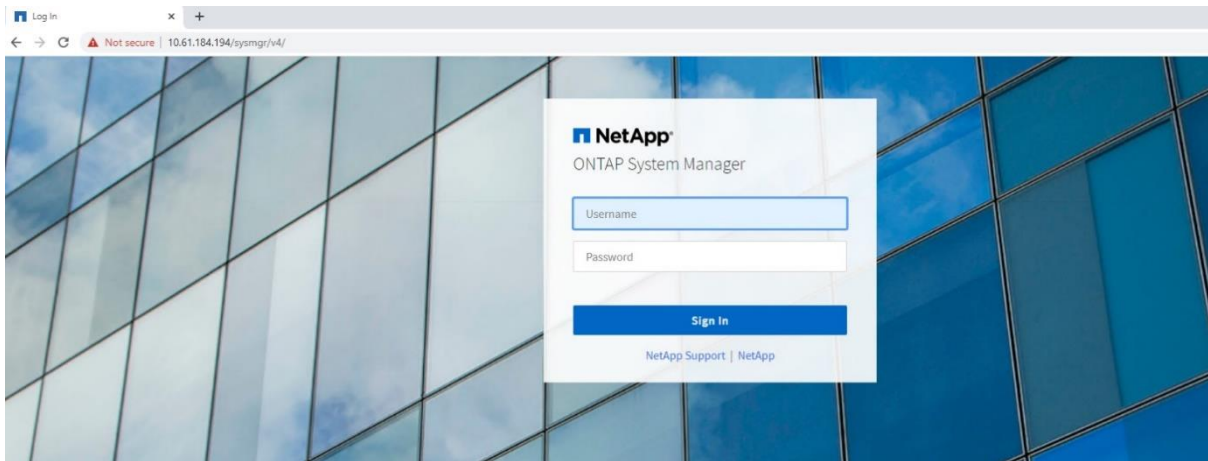
Note: Cluster setup can also be performed using the CLI. This document describes the cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup on the Welcome screen.
4. In the Cluster screen, follow these steps:
 - a. Enter the cluster and node names.
 - b. Select the cluster configuration.
 - c. Enter and confirm the password.
 - d. (Optional) Enter the cluster base and feature licenses.

Note: The nodes should be discovered automatically; if they are not, click the Refresh link. By default, the cluster interfaces are created on all the new factory shipping storage controllers. If all the nodes are not discovered, then configure the cluster using the command line. Cluster license and feature licenses can also be installed after completing cluster creation.
5. Click Submit and Continue.
6. In the network page, complete the following steps:
 - a. Cluster Management. Enter the IP address, netmask, gateway, and port details.
 - b. Node Management. Enter the node management IP addresses and port details for all the nodes.
 - c. Service Processor Management. Enter the IP addresses for all nodes.
 - d. DNS Details. Enter the DNS domain names and server address.
 - e. NTP Details. Enter the primary and alternate NTP server.
7. Click Submit and Continue.
8. In the Support page, configure the AutoSupport and Event Notifications sections.
9. Click Submit and Continue.
10. In the Storage page, review the configuration details. Two data aggregates are to be created as part of setup.
11. Click Submit and Continue.
12. In the SVM Tab, you can create your own SVM, or you can skip this step.
13. Click Skip this Step.
14. In the Summary page, review the configuration details if needed and click Manage Your Cluster.

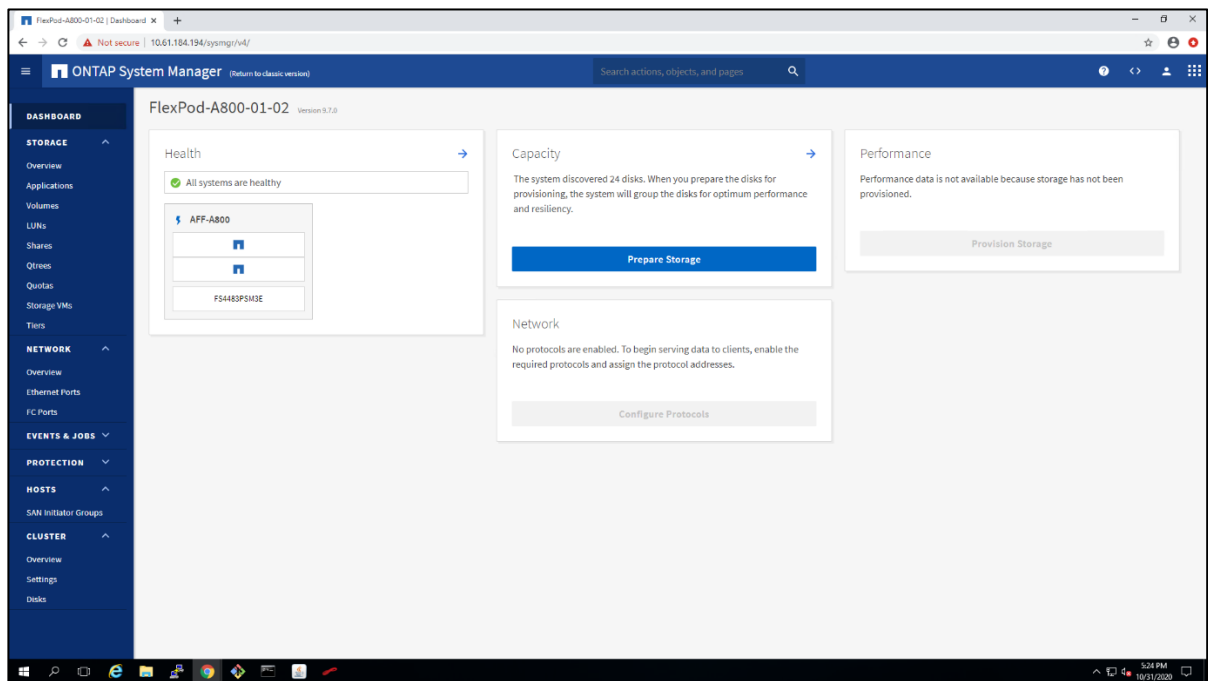
Log in to cluster

Open an SSH connection to either the cluster IP or the host name. Log in to the admin user with the password you provided earlier.

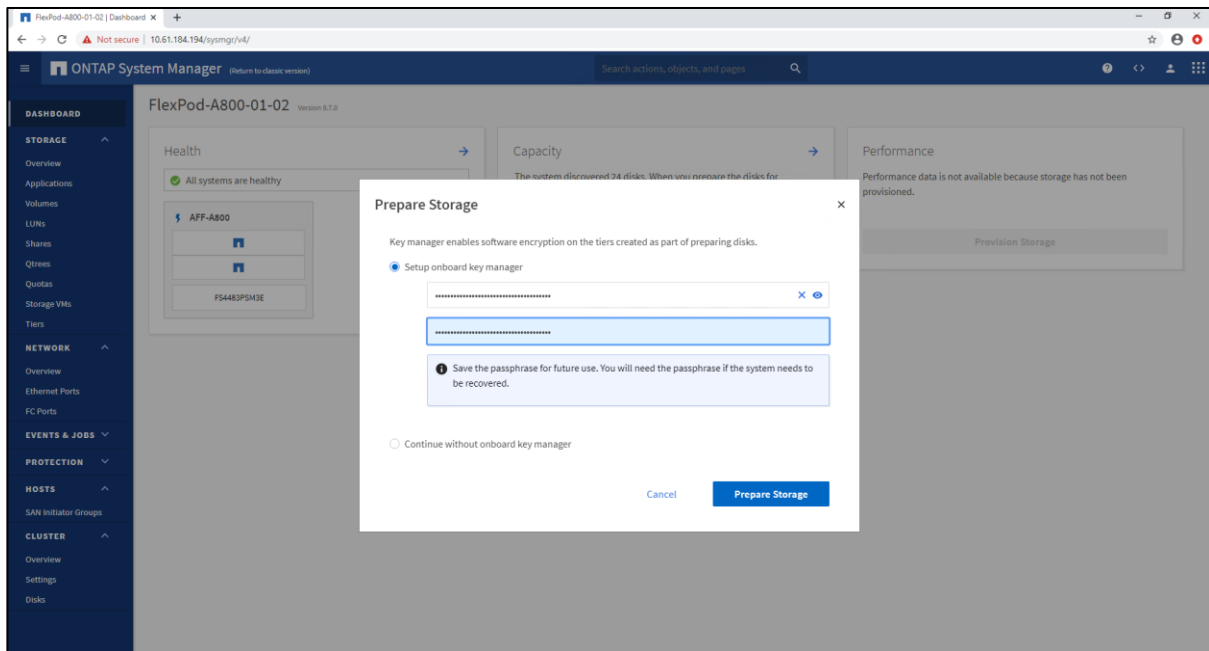


Prepare storage with disk encryption

In the ONTAP System Manager dashboard, click Prepare Storage to turn on NetApp Storage Encryption (NSE) with onboard key manager. NSE provides full disk encryption that protects data at rest with no operational effect.

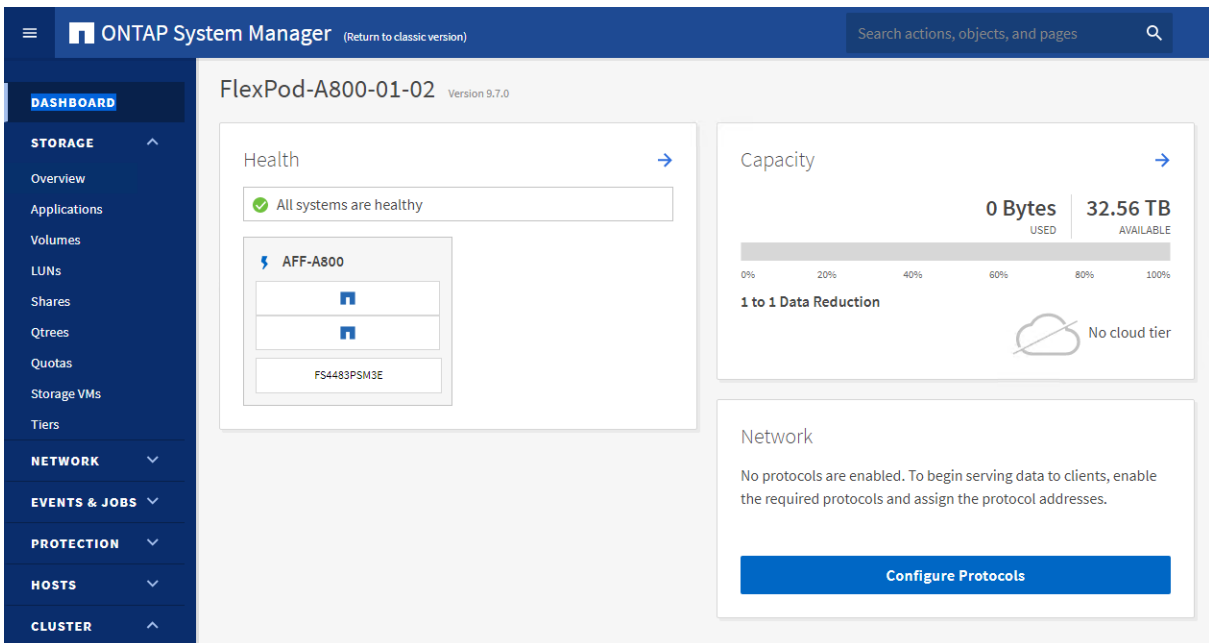


Enter a passphrase to enable onboard key manager.



Enable FC protocol for storage access

Click Prepare Protocol to turn on the FC protocol for cluster data access. ONTAP supports multiple protocols such as NFS, iSCSI, SMB/CIFS, and FC. For this solution, only the FC protocol is needed and enabled on FC ports 2a, 2b, 2c, and 2d.



Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, log into one of the ONTAP cluster nodes via SSH and run the following command:

```
node run -node * options cdpd.enable on
```


Configure Network Time Protocol

1. Set the time zone for the cluster.

```
timezone <timezone>
FlexPod-A800-01-02::> timezone America/New_York
1 entry modified
Warning: Using value "America/New_York" for parameter "-timezone". A future release may require
"-timezone" to be specified before the value.
```

2. Set the date for the cluster.

```
date <ccyyymmddhhmm.ss>
FlexPod-A800-01-02::> date 202011071249.00
FlexPod-A800-01-02::> cluster date show
Node          Date          Time zone
-----
FlexPod-A800-01-02-01
11/7/2020 12:49:05 -05:00 America/New_York
FlexPod-A800-01-02-02
11/7/2020 12:49:05 -05:00 America/New_York
2 entries were displayed.
```

3. Configure the Network Time Protocol (NTP) servers for the cluster.

```
cluster time-service ntp server create -server <nexus-A-mgmt0-ip>
FlexPod-A800-01-02::> cluster time-service ntp server show
This table is currently empty.
FlexPod-A800-01-02::> cluster time-service ntp server create -server 10.61.185.177
```

Verify storage failover

1. Verify the status of storage failover.

```
FlexPod-A800-01-02::> storage failover show
Node          Partner          Possible State Description
-----
FlexPod-A800-01-02-01
FlexPod-A800-01-02-02- true    Connected to FlexPod-A800-01-02-02
FlexPod-A800-01-02-02
FlexPod-A800-01-02-01- true    Connected to FlexPod-A800-01-02-01
2 entries were displayed.
```

2. Enable failover on one of the two nodes if it was not completed during the installation.

```
storage failover modify -node <cluster-node01> -enabled true
```

3. Verify the HA status for a two-node cluster.

```
FlexPod-A800-01-02::> cluster ha show
High Availability Configured: true
```

4. Enable HA status if not done. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

5. Verify that hardware assist is correctly configured.

```
FlexPod-A800-01-02::> storage failover hwassist show
Node
-----
FlexPod-A800-01-02-01
Partner: FlexPod-A800-01-02-02
Hwassist Enabled: true
Hwassist IP: 192.0.2.84
Hwassist Port: 162
```

```

        Monitor Status: active
        Inactive Reason: -
        Corrective Action: -
        Keep-Alive Status: healthy
FlexPod-A800-01-02-02
        Partner: FlexPod-A800-01-02-01
        Hwassist Enabled: true
        Hwassist IP: 192.0.2.85
        Hwassist Port: 162
        Monitor Status: active
        Inactive Reason: -
        Corrective Action: -
        Keep-Alive Status: healthy
2 entries were displayed.

```

Set up service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```

system service-processor network modify -node <cluster-node01> -address-family IPv4 -enable true
-dhcp none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>
system service-processor network modify -node <cluster-node02> -address-family IPv4 -enable true
-dhcp none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
FlexPod-A800-01-02::> system service-processor show

```

Node	Type	Status	IP Configured	Firmware Version	IP Address
FlexPod-A800-01-02-01	BMC	online	true	10.3P1	192.168.103.34
FlexPod-A800-01-02-02	BMC	online	true	10.3P1	192.168.103.36

2 entries were displayed.

Disable auto-negotiate on FC ports

Disable each FC adapter in the controllers with the `fc adapter modify` command.

```

fc adapter modify -node <cluster-node01> -adapter 2a -status-admin down
fc adapter modify -node <cluster-node01> -adapter 2b -status-admin down
fc adapter modify -node <cluster-node01> -adapter 2c -status-admin down
fc adapter modify -node <cluster-node01> -adapter 2d -status-admin down
fc adapter modify -node <cluster-node02> -adapter 2a -status-admin down
fc adapter modify -node <cluster-node02> -adapter 2b -status-admin down
fc adapter modify -node <cluster-node02> -adapter 2c -status-admin down
fc adapter modify -node <cluster-node02> -adapter 2d -status-admin down

```

Set the desired speed on the adapter and return it to the online state.

```

fc adapter modify -node <cluster-node01> -adapter 2a -speed 32 -status-admin up
fc adapter modify -node <cluster-node01> -adapter 2b -speed 32 -status-admin up
fc adapter modify -node <cluster-node01> -adapter 2c -speed 32 -status-admin up
fc adapter modify -node <cluster-node01> -adapter 2d -speed 32 -status-admin up
fc adapter modify -node <cluster-node02> -adapter 2a -speed 32 -status-admin up
fc adapter modify -node <cluster-node02> -adapter 2b -speed 32 -status-admin up
fc adapter modify -node <cluster-node02> -adapter 2c -speed 32 -status-admin up
fc adapter modify -node <cluster-node02> -adapter 2d -speed 32 -status-admin up

```

Create SVM

We created one infrastructure SVM `infra_svm` to host boot LUNs and one data SVM `ora19c_svm` to host Oracle data LUNs.

Log into ONTAP System Manager, navigate to Storage > Storage VMs to add a storage SVM.

ONTAP System Manager (Return to classic version)

Search actions, objects, and pages

Add Storage VM

STORAGE VM NAME
ora19c_svm

Access Protocol
SMB/CIFS and NFS iSCSI **FC**

☒ Enable FC

CONFIGURE FC PORTS ⓘ

Nodes	2a	2b	2c	2d
FlexPod-...1-02-01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
FlexPod-...1-02-02	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Storage VM Administration
☐ Manage administrator account

Save Cancel

You can only add two ports to an SVM initially, and two additional ports can be enabled after an SVM is created by clicking the pencil icon.

ONTAP System Manager (Return to classic version)

Search actions, objects, and pages

Storage VMs

+ Add Filter

Name
ora19c_svm All Storage VMs Edit More

Overview Settings SnapMirror (Local or Remote)

Protocols

NFS Not configured	SMB/CIFS Not configured	iSCSI Not configured	FC STATUS Enabled WWPN: 20:09:d0:39:ea:20:ba:d3 WWPN: 20:0a:d0:39:ea:20:ba:d3 20:09:d0:39:ea:20:ba:d3
-----------------------	----------------------------	-------------------------	--

Create a new SVM enables FC service on the SVM.

Add or modify SVM LIFs

Click the Network Tab, and then click Overview to go to the Network Interfaces page to add, delete, or modify an SVM LIF.

Overview

IPspaces

Cluster	Broadcast Domains
Default	Storage VMs infra_svm, ora19c_svm Broadcast Domains Default

Broadcast Domains

Cluster	9000 MTU	IPspace: Cluster
Default	1500 MTU	IPspace: Default FlexPod-A800-01-02-01 #01 FlexPod-A800-01-02-02 #01

Network Interfaces

Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Protocols	Type
ora19c_svm_data_fc_ip_1_2b	✓	ora19c_svm		20:0c:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2b	FC	Data
ora19c_svm_data_fc_ip_2_2a	✓	ora19c_svm		20:0d:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2a	FC	Data
ora19c_svm_data_fc_ip_2_2b	✓	ora19c_svm		20:12:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2b	FC	Data
ora19c_svm_data_fc_ip_1_2c	✓	ora19c_svm		20:17:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2c	FC	Data
ora19c_svm_data_fc_ip_1_2d	✓	ora19c_svm		20:18:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-01	2d	FC	Data
ora19c_svm_data_fc_ip_2_2d	✓	ora19c_svm		20:19:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2d	FC	Data
ora19c_svm_data_fc_ip_1_2c	✓	ora19c_svm		20:1a:d0:39:ea:20:ba:d3	FlexPod-A800-01-02-02	2c	FC	Data

Create load-sharing mirrors of SVM root

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
FlexPod-A800-01-02::> volume create -vserver infra_svm -volume infra_svm_root_m01 -aggregate
FlexPod_A800_01_02_01_NVME_SSD_1 -size 1GB -type DP
[Job 2281] Job succeeded: Successful
```

```
FlexPod-A800-01-02::> volume create -vserver infra_svm -volume infra_svm_root_m02 -aggregate
FlexPod_A800_01_02_02_NVME_SSD_1 -size 1GB -type DP
[Job 2282] Job succeeded: Successful
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
FlexPod-A800-01-02::> job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
FlexPod-A800-01-02::> snapmirror create -source-path infra_svm:infra_svm_root -destination-path
infra_svm:infra_svm_root_m01 -type LS -schedule 15min
[Job 2283] Job succeeded: SnapMirror: done
```

```
FlexPod-A800-01-02::> snapmirror create -source-path infra_svm:infra_svm_root -destination-path
infra_svm:infra_svm_root_m02 -type LS -schedule 15min
[Job 2285] Job succeeded: SnapMirror: done
```

4. Initialize the mirroring relationship.

```
FlexPod-A800-01-02::> snapmirror initialize-ls-set -source-path infra_svm:infra_svm_root
[Job 2286] Job is queued: snapmirror initialize-ls-set for source "FlexPod-A800-01-
02://infra_svm/infra_svm_root".
FlexPod-A800-01-02::> snapmirror show
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Progress Last Updated
FlexPod-A800-01-02://infra_svm/infra_svm_root	LS	FlexPod-A800-01-02://infra_svm/infra_svm_root_m01	Snapmirrored	Idle	-	true -
		FlexPod-A800-01-02://infra_svm/infra_svm_root_m02	Snapmirrored	Idle	-	true -

2 entries were displayed.

5. Create the same load sharing mirror for ora19c_svm root.

Create SAN initiator groups for Oracle hosts

Extract four HBA WWPNs from UCS Manager for each Oracle host and create an initiator group (igroup) for each host for boot and data LUN access.

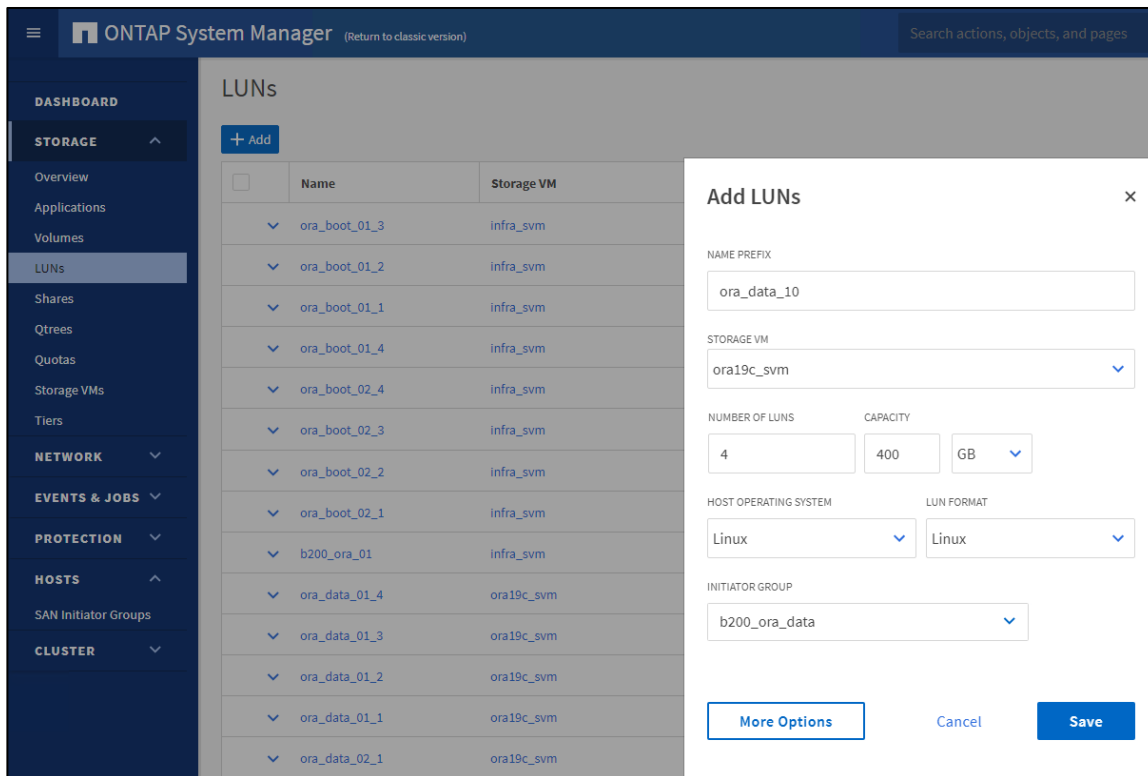
The igroups are SVM specific. Specify the correct SVM, host operating system, and all four HBA WWPNs for the host. If boot LUNs are to be created in `infra_svm`, then an igroup should be created in `infra_svm` to access the boot LUNs.

Although an igroup can be created either from ONTAP System Manager or from the command line, NetApp highly recommends creating an igroup to access the boot LUN hosted in `infra_svm` from the command line instead of from ONTAP System Manager:

```
lun igroup create -vserver infra_svm -igroup b200-ora-01 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:00, 20:00:00:25:B5:8A:A0:01, 20:00:00:25:B5:8B:B0:00,
20:00:00:25:B5:8B:B0:01
lun igroup create -vserver infra_svm -igroup b200-ora-03 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:04, 20:00:00:25:B5:8A:A0:05, 20:00:00:25:B5:8B:B0:04,
20:00:00:25:B5:8B:B0:05
lun igroup create -vserver infra_svm -igroup b200-ora-04 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:06, 20:00:00:25:B5:8A:A0:07, 20:00:00:25:B5:8B:B0:06,
20:00:00:25:B5:8B:B0:07
lun igroup create -vserver infra_svm -igroup b200-ora-05 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:08, 20:00:00:25:B5:8A:A0:09, 20:00:00:25:B5:8B:B0:08,
20:00:00:25:B5:8B:B0:09
lun igroup create -vserver infra_svm -igroup b200-ora-06 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:0A, 20:00:00:25:B5:8A:A0:0B, 20:00:00:25:B5:8B:B0:0A,
20:00:00:25:B5:8B:B0:0B
lun igroup create -vserver infra_svm -igroup b200-ora-07 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:0C, 20:00:00:25:B5:8A:A0:0D, 20:00:00:25:B5:8B:B0:0C,
20:00:00:25:B5:8B:B0:0D
lun igroup create -vserver infra_svm -igroup b200-ora-08 -protocol fcp -ostype linux -initiator
20:00:00:25:B5:8A:A0:0E, 20:00:00:25:B5:8A:A0:0F, 20:00:00:25:B5:8B:B0:0E,
20:00:00:25:B5:8B:B0:0F
```

Create LUNs

From ONTAP System Manager Storage – LUNs, click Add to add all LUNs as shown in storage layout configuration Table 10 (xref).



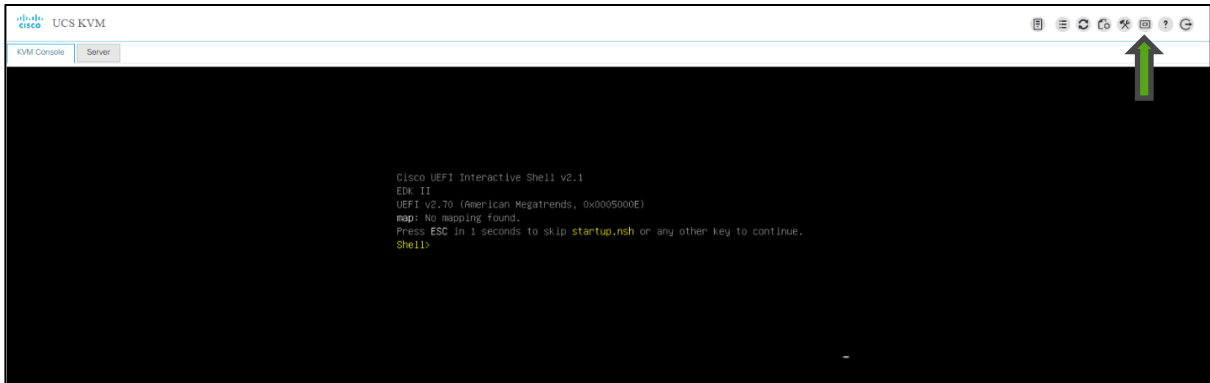
This creates four LUNs: ora_data_10_1, ora_data_10_2, ora_data_10_3, and ora_data_10_4 under volume ora_data_10. We assigned igroup b200_ora_data, which includes all HBA WWPNs from all eight Oracle RAC nodes for shared RAC storage access. The volume size is further expanded to the designed sizing as shown in Table 10 (xref) after the LUNs are created.

Through the same procedure, we created 60 shared Oracle data LUNs and eight hosts boot LUNs and eight Oracle binary LUNs. Each LUN is assigned with respective host access as configured. The AFF A800 storage is now fully configured for Oracle RAC cluster setup.

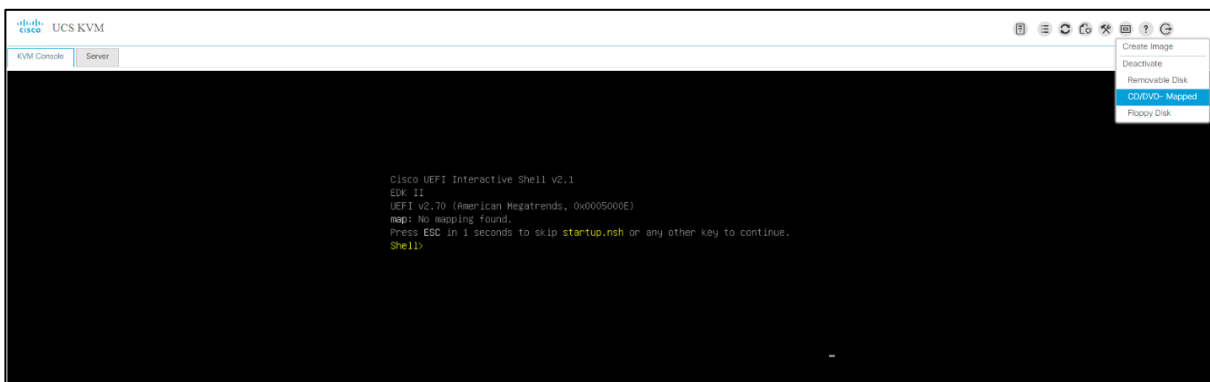
Operating system setup

Operating system installation

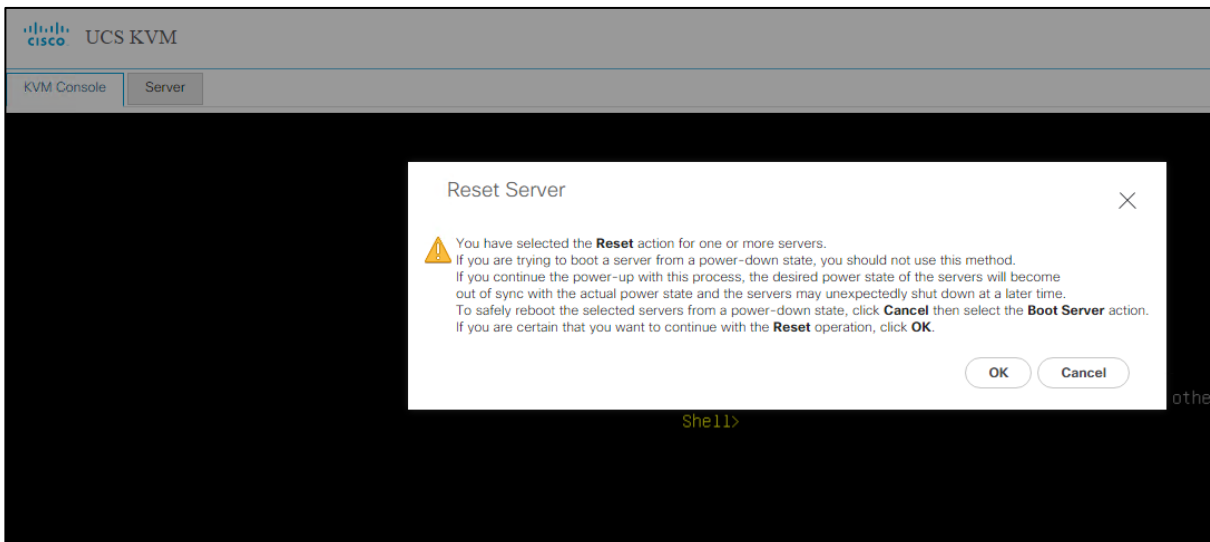
1. Download the Oracle Linux 8.2 OS image from <https://edelivery.oracle.com/linux>.
2. Launch the KVM console on the desired server by going to the tab Equipment > Chassis > Chassis 1 > Servers > Server 1 > from the right-side window General and select KVM Console to open KVM.
3. Click Accept Security and open KVM. Click Enable Virtual Media below to map the ISO drive.



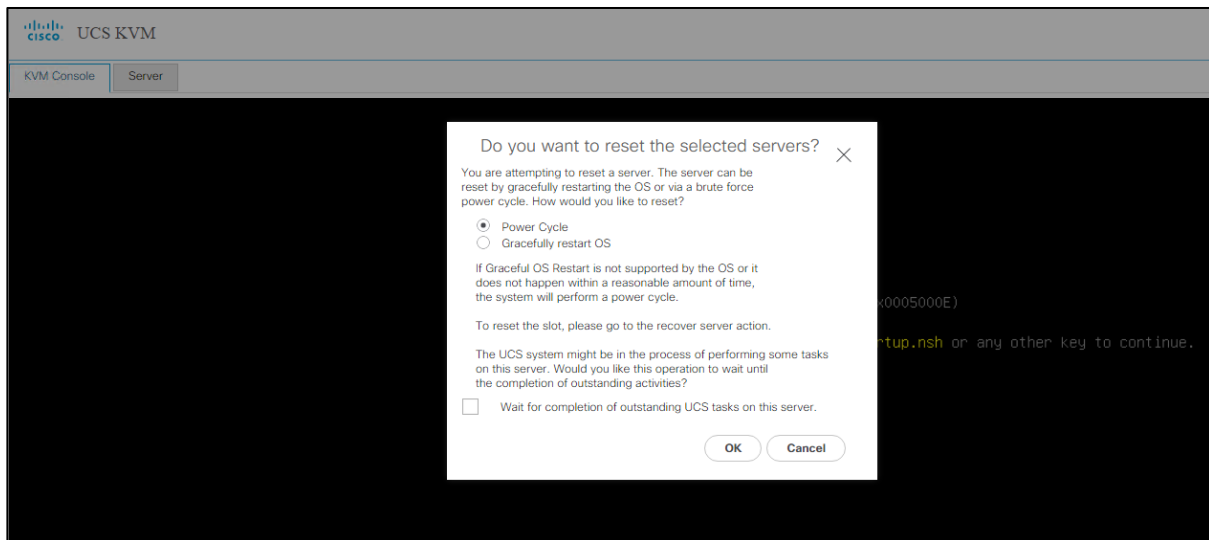
4. Enabling virtual media enables the CD/DVD drive. Click CD/DVD to choose the location of the Linux 8.2 ISO image. CD/DVD now shows that it is mapped.



5. After mapping the Oracle Linux ISO image, reset the server.

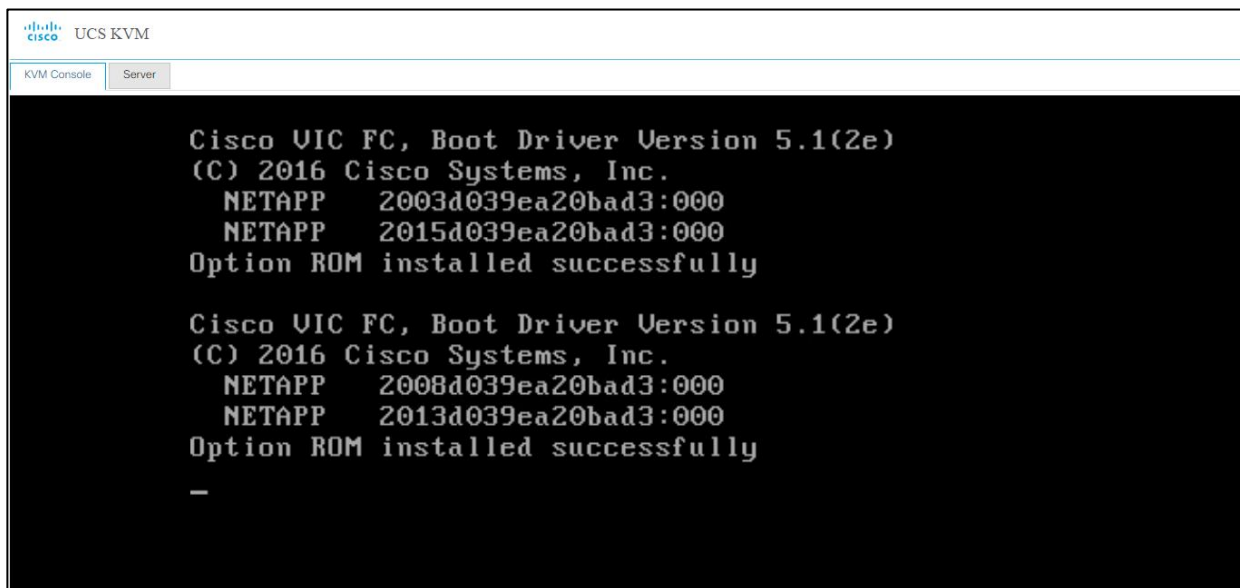


6. Click OK and choose to power cycle:



SAN boot validation

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



Note: If you created a host-access igroup on the AFF A800 controller using the GUI with a mixed FC/iSCSI protocol, the server might not boot after OS installation. In that case, you need to manually create an igroup with FC only from the CLI. That resolves the issue.

```
FlexPod-A800-01-02::> lun igroup create -vserver infra_svm -igroup b200-ora-01 -protocol fcp -
ostype linux -initiator 20:00:00:25:B5:8A:A0:00, 20:00:00:25:B5:8A:A0:01,
20:00:00:25:B5:8B:B0:00, 20:00:00:25:B5:8B:B0:01
FlexPod-A800-01-02::> lun igroup show -protocol fcp
Vserver   Igroup      Protocol OS Type  Initiators
-----
infra_svm b200-ora-01 fcp      linux   20:00:00:25:b5:8a:a0:00
                                     20:00:00:25:b5:8a:a0:01
                                     20:00:00:25:b5:8b:b0:00
```


1. After server installation, log into the server and set the default boot kernel to the Red Hat compatible kernel or the RHCK kernel.

```
[root@b200-ora-08 ~]# ls -l /boot/vmlinuz*
-rwxr-xr-x. 1 root root 9226480 Nov 14 01:43 /boot/vmlinuz-0-rescue-600f517ac51441a0b9adc549788e0b5a
-rwxr-xr-x. 1 root root 9226480 Apr 29 2020 /boot/vmlinuz-4.18.0-193.el8.x86_64
-rwxr-xr-x. 1 root root 8923392 Apr 21 2020 /boot/vmlinuz-5.4.17-2011.1.2.el8uek.x86_64
[root@b200-ora-08 ~]# grubby --default-kernel
/boot/vmlinuz-5.4.17-2011.1.2.el8uek.x86_64
[root@b200-ora-08 ~]# grubby --set-default /boot/vmlinuz-4.18.0-193.el8.x86_64
The default is /boot/loader/entries/600f517ac51441a0b9adc549788e0b5a-4.18.0-193.el8.x86_64.conf with index 1 and
kernel /boot/vmlinuz-4.18.0-193.el8.x86_64
[root@b200-ora-08 ~]# reboot
```

2. Verify the change after reboot:

```
[admin@b200-ora-08 ~]$ uname -a
Linux b200-ora-08.cie.netapp.com 4.18.0-193.el8.x86_64 #1 SMP Wed Apr 29 11:11:17 PDT 2020 x86_64 x86_64 x86_64
GNU/Linux
[admin@b200-ora-08 ~]$ sudo grubby --default-kernel

We trust you have received the usual lecture from the local System
Administrator. It usually boils down to these three things:

#1) Respect the privacy of others.
#2) Think before you type.
#3) With great power comes great responsibility.

[sudo] password for admin:
/boot/vmlinuz-4.18.0-193.el8.x86_64
```

3. Complete these changes on the rest of the nodes.

This solution is designed to run on an Oracle Linux 8.2 Red Hat compatible kernel or on RHCK. Therefore all RAC nodes boot into kernel `4.18.0-193.el8.x86_64` as versus the `el8uek` version of kernel.

Implement operating system prerequisites for Oracle software installation

See the [“Configuring Operating Systems for Oracle Database on Linux”](#) section of the Oracle 19c Database Installation Guide for Linux for prerequisite requirements. This is a summary of the check list for prerequisite configuration.

1. Oracle Preinstallation RPM. This completes many prerequisites for Oracle installation such as the packages required to run Oracle 19c.
2. Configure BIOS for OLTP workloads. This should have been completed as part of UCS server profile setup and configuration. See Cisco UCS setup section 19 for details.
3. Disable transparent huge page. Transparent HugePages memory is enabled by default with Oracle Linux 6 and later but should be disabled for performance.
4. Configure huge page for large database. Configure the HugePages memory allocation to a size large enough to accommodate the sum of the SGA sizes of all the databases you intend to install on the cluster, as well as the Grid Infrastructure Management Repository (GIMR). GIMR is configured to use HugePages and always starts before database instances.
5. Set the disk I/O scheduler to [deadline] on Linux hosts.
6. Disable SELinux.
7. Disable the firewall.
8. Stop and disable `avahi-daemo`.
9. Configure the NTP server.
10. Install the NetApp host utility (`netapp_linux_unified_host_utilities-7-1.x86_64.rpm`).
11. Install the following required packages before installing the NetApp host utility.

- libhbaapi-2.2.9-13.el8.x86_64.rpm
- libhbalinux-1.0.17-7.el8.x86_64.rpm

12. Configure resource limits for Oracle software installation users.

Checking resources limits for Oracle software installation users are met as listed in the installation guide.

In the optional Automated Deployment section, most of these tasks have been automated, and each task can be called on all RAC nodes by their respective tags. We have leveraged those automated Ansible tasks in setting up the Oracle 19c environment.

Configure FC LUNs for Oracle

1. Configure multipath setup.

Oracle Linux RHCK 8.2 is a Red Hat compatible version of the kernel that is compiled with all settings required to recognize and correctly manage ONTAP LUNs. There should be two groups of paths with different priorities. The paths with the higher priorities are active/optimized, meaning I/O is serviced by the controller where the LUN is located. The paths with lower priorities are enabled but not active because the LUN is not local and I/O is served from a peer controller with non-optimized paths. The non-optimized paths are only used when no optimized paths are available.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

```
3600a0980383145374b2451445133574d dm-33 NETAPP,LUN C-Mode
size=400G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
| |- 6:0:0:5 sdfc 129:224 active ready running
| |- 6:0:4:5 sdsd 135:272 active ready running
| |- 7:0:1:5 sdut 67:592 active ready running
| |- 7:0:5:5 sdain 129:944 active ready running
| |- 8:0:1:5 sdbo 68:32 active ready running
| |- 8:0:5:5 sdhv 134:80 active ready running
| |- 9:0:0:5 sdju 65:384 active ready running
| `-- 9:0:4:5 sdyc 128:704 active ready running
`+- policy='service-time 0' prio=10 status=enabled
| |- 6:0:1:5 sdkk 66:384 active ready running
| |- 6:0:5:5 sdze 130:640 active ready running
| |- 7:0:0:5 sdni 71:320 active ready running
| |- 7:0:4:5 sdacd 135:592 active ready running
| |- 8:0:0:5 sdf 8:80 active ready running
| |- 8:0:4:5 sddz 128:16 active ready running
| |- 9:0:1:5 sdra 133:320 active ready running
| `-- 9:0:5:5 sdaet 67:912 active ready running
```

2. Add alias names for NetApp LUNs in the multipath file.

```
[oracle@b200-ora-01 SLOB]$ cat /etc/multipath.conf
defaults {
    find_multipaths yes
    user_friendly_names yes
}

blacklist {
}

multipaths {
    multipath {
        wwid          3600a0980383145374b24514451335767
        alias          ora_bin_01_1
    }
    multipath {
        wwid          3600a0980383145374b2451445133572f
        alias          ora_crs_01_1
    }
    multipath {
        wwid          3600a0980383145374b24514451335761
    }
}
```

```

        alias                ora_crs_01_2
    }
    multipath {
        wwid                  3600a0980383145374b24514451335763
        alias                  ora_crs_02_1
    }
    multipath {
        wwid                  3600a0980383145374b24514451335762
        alias                  ora_crs_02_2
    }
    multipath {
        wwid                  3600a0980383145373524514438424557
        alias                  ora_data_01_1
    }
    multipath {
        wwid                  3600a0980383145373524514438424556
        alias                  ora_data_01_2
    }
    .....

```

3. Create UDEV rules for Oracle devices.
4. Configure UDEV rules to assign permissions in all of the Oracle RAC nodes to access NetApp Storage LUNs. This includes the device details along with required permissions to enable grid and Oracle users to have read/write privileges on these devices. Configure UDEV rules on all the Oracle nodes as shown below:
5. Create a new file named `/etc/udev/rules.d/99-oracle-asmdevices.rules` with the following entries on all nodes:

```

cat /etc/udev/rules.d/99-oracle-asmdevices.rules
ENV{DM_NAME}=="ora_crs*", GROUP=="asmadmin", OWNER=="grid", MODE=="660"
ENV{DM_NAME}=="ora_data*", GROUP=="asmadmin", OWNER=="grid", MODE=="660"
ENV{DM_NAME}=="ora_redo*", GROUP=="asmadmin", OWNER=="grid", MODE=="660"

```

Configure networking for RAC

1. Configure public and private NICs on each RAC node.

If you are running an RAC cluster from two UCS chassis, you might need to rename designated public or private interfaces to have consistent naming across the cluster. We used the Network Manager CLI to rename the interfaces.

```

[root@b200-ora-08 mapper]# nmcli connection show "enp103s0f0" | grep 802-3-ethernet.mac-address:
802-3-ethernet.mac-address:      00:25:B5:89:BB:9F
[root@b200-ora-08 mapper]# ip link show enp103s0f0
3: enp103s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP mode DEFAULT group
default qlen 1000 link/ether 00:25:b5:89:bb:9f brd ff:ff:ff:ff:ff:ff
[root@b200-ora-08 mapper]# nmcli connection modify "enp103s0f0" connection.interface-name
"ens5f1"

```

2. Reboot the host after renaming to ensure that the change are persistent between host reboots.
3. Configure `/etc/hosts` on each RAC node.

```

[oracle@b200-ora-06 ~]$ cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
10.61.180.151 b200-ora-01.cie.netapp.com b200-ora-01
10.61.180.152 b200-ora-02.cie.netapp.com b200-ora-02
10.61.180.153 b200-ora-03.cie.netapp.com b200-ora-03
10.61.180.154 b200-ora-04.cie.netapp.com b200-ora-04
10.61.180.155 b200-ora-05.cie.netapp.com b200-ora-05
10.61.180.156 b200-ora-06.cie.netapp.com b200-ora-06
10.61.180.157 b200-ora-07.cie.netapp.com b200-ora-07
10.61.180.158 b200-ora-08.cie.netapp.com b200-ora-08
172.21.101.151 b200-ora-01-pri.cie.netapp.com b200-ora-01-pri
172.21.101.152 b200-ora-02-pri.cie.netapp.com b200-ora-02-pri
172.21.101.153 b200-ora-03-pri.cie.netapp.com b200-ora-03-pri
172.21.101.154 b200-ora-04-pri.cie.netapp.com b200-ora-04-pri

```

```

172.21.101.155      b200-ora-05-pri.cie.netapp.com b200-ora-05-pri
172.21.101.156      b200-ora-06-pri.cie.netapp.com b200-ora-06-pri
172.21.101.157      b200-ora-07-pri.cie.netapp.com b200-ora-07-pri
172.21.101.158      b200-ora-08-pri.cie.netapp.com b200-ora-08-pri
10.61.180.159       b200-ora-01-vip.cie.netapp.com b200-ora-01-vip
10.61.180.160       b200-ora-02-vip.cie.netapp.com b200-ora-02-vip
10.61.180.161       b200-ora-03-vip.cie.netapp.com b200-ora-03-vip
10.61.180.162       b200-ora-04-vip.cie.netapp.com b200-ora-04-vip
10.61.180.163       b200-ora-05-vip.cie.netapp.com b200-ora-05-vip
10.61.180.164       b200-ora-06-vip.cie.netapp.com b200-ora-06-vip
10.61.180.165       b200-ora-07-vip.cie.netapp.com b200-ora-07-vip
10.61.180.166       b200-ora-08-vip.cie.netapp.com b200-ora-08-vip
10.61.180.167       rtprac-scan.cie.netapp.com rtprac-scan
10.61.180.168       rtprac-scan.cie.netapp.com rtprac-scan
10.61.180.169       rtprac-scan.cie.netapp.com rtprac-scan

```

4. Setup the DNS SCAN address for the RAC cluster.

```

[root@b200-ora-02 oracle]# nslookup rtprac-scan
Server:          10.61.184.251
Address:         10.61.184.251#53
Name:            rtprac-scan.cie.netapp.com
Address: 10.61.180.169
Name:            rtprac-scan.cie.netapp.com
Address: 10.61.180.167
Name:            rtprac-scan.cie.netapp.com
Address: 10.61.180.168

```

Oracle 19c deployment

Oracle 19c grid infrastructure installation

After completing RHCK 8.2 operating-system kernel, storage, and networking configuration. You can launch the Oracle 19c grid installer to set up the Oracle 19c grid infrastructure.

Although Oracle has officially announced Linux 8 support for Oracle 19c, the base installer 19.3 does not recognize Linux 8 as a supported operating system. More importantly, Oracle 19c is the only official production supportable on 19.7 or above on Linux 8. We created a work-around for the installer to work with RHCK 8.2 as well as patching up the installation to 19.8 in the process. See the following details:

1. Set the work around for the grid installer:
 - a. Export `CV_ASSUME_DISTID=OEL7.6`.
 - b. Export `ORACLE_SRVM_REMOTECOPY=/tmp/archive/scp`.

```

[grid@b200-ora-01 OPatch]$ cat /tmp/archive/scp
#!/bin/sh

/usr/bin/scp -T $*

```

2. Download Oracle grid 19.8 patch [31305339](#) and the latest version of [opatch](#).
3. Extract patch 31305339 to a tmp location and update opatch in grid home to the latest version with `p6880880_190000_LINUX-x86-64.zip`.

```

[oracle@b200-ora-01 ntap]$ OPatch/opatch version
OPatch Version: 12.2.0.1.23

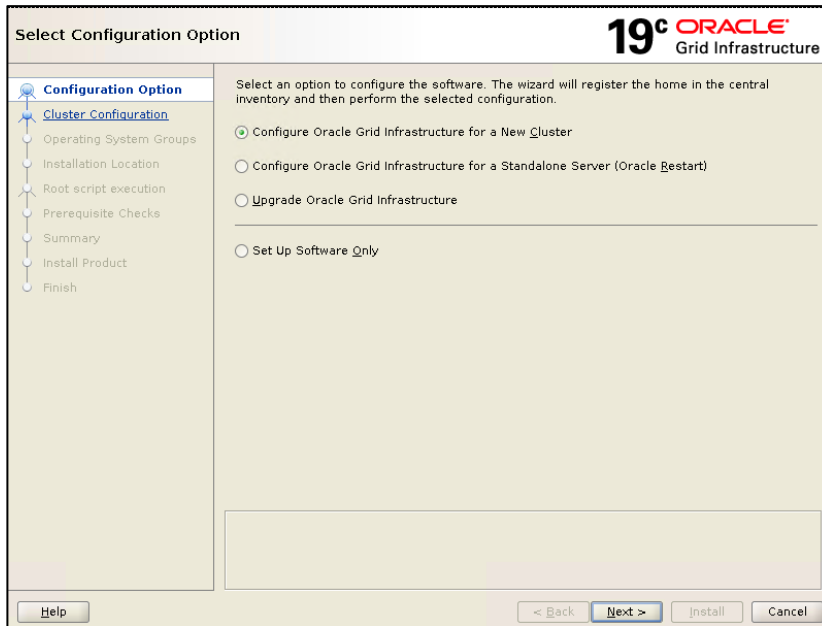
OPatch succeeded

```

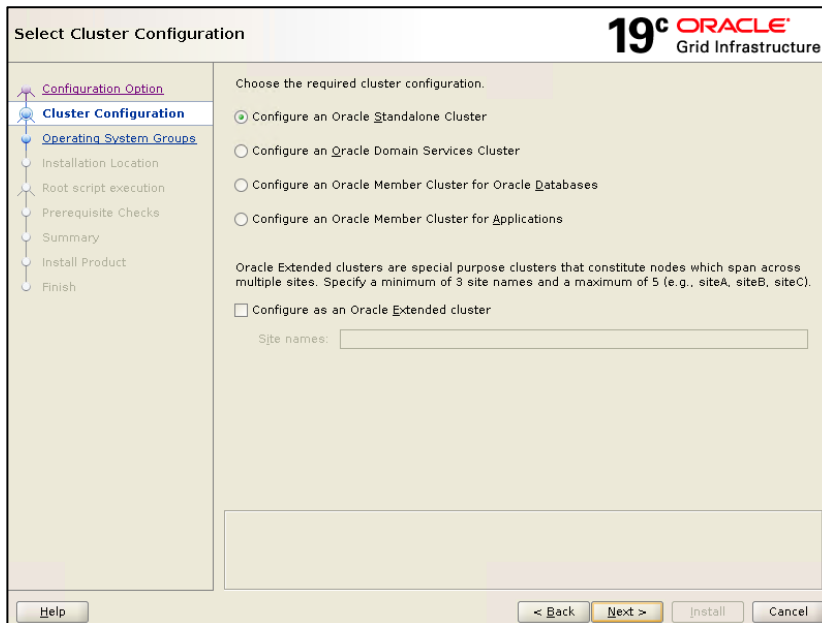
4. Launch grid setup with the `-applyRU` option. The grid installer first applies patches and launches the setup wizard in one command step.

```
[grid@ora01 ~]$ pwd
/tmp/app/19300/grid
[grid@ora01 ~]$ ls /tmp/archive/patch/31305339/
31281355 31304218 31305087 31335188 automation bundle.xml README.html README.txt
[grid@ora01 ~]$ ./gridSetup.sh -applyAll /tmp/archive/patch/31305339/ -applyOneOffs 31281355,31304218,31305087,31335188
[MSG-22832] The path (31281355) provided in the (-applyOneOffs) argument value does not exist. Provide the complete absolute paths of the patches separated by comma.
[grid@ora01 ~]$ ./gridSetup.sh -applyAll /tmp/archive/patch/31305339/ -applyOneOffs /tmp/archive/patch/31305339/31281355,/tmp/archive/patch/31305339/31304218,/tmp/archive/patch/31305339/31305087,/tmp/archive/patch/31305339/31335188
Preparing the home to patch...
Applying the patch /tmp/archive/patch/31305339/...
Successfully applied the patch.
Applying the patch /tmp/archive/patch/31305339/31281355...
Successfully applied the patch.
Applying the patch /tmp/archive/patch/31305339/31304218...
Successfully applied the patch.
Applying the patch /tmp/archive/patch/31305339/31305087...
Successfully applied the patch.
Applying the patch /tmp/archive/patch/31305339/31335188...
Successfully applied the patch.
The log can be found at: /tmp/GridSetupActions2020-12-01_09-09-03PM/installerPatchActions_2020-12-01_09-09-03PM.log
Launching Oracle Grid Infrastructure Setup Wizard...
```

This image shows the setup wizard launching.



5. Choose the option to configure a standalone cluster.



6. This image depicts the cluster node configuration.

Cluster Node Information

Provide the list of nodes to be managed by Oracle Grid Infrastructure with their Public Hostname and Virtual Hostname.

Public Hostname	Virtual Hostname
b200-ora-01.cie.netapp.com	b200-ora-01-vip.cie.netapp.com
b200-ora-02.cie.netapp.com	b200-ora-02-vip.cie.netapp.com
b200-ora-03.cie.netapp.com	b200-ora-03-vip.cie.netapp.com
b200-ora-04.cie.netapp.com	b200-ora-04-vip.cie.netapp.com
b200-ora-05.cie.netapp.com	b200-ora-05-vip.cie.netapp.com
b200-ora-06.cie.netapp.com	b200-ora-06-vip.cie.netapp.com
b200-ora-07.cie.netapp.com	b200-ora-07-vip.cie.netapp.com
b200-ora-08.cie.netapp.com	b200-ora-08-vip.cie.netapp.com

SSH connectivity... Use Cluster Configuration File... Add... Edit... Remove

< Back Next > Install Cancel

7. This image depicts the network interfaces configuration.

Specify Network Interface Usage

Private interfaces are used by Oracle Grid Infrastructure for internode traffic.

Interface Name	Subnet	Use for
ens5f0	10.61.180.0	Public
*	172.21.101.0	ASM & Private

< Back Next > Install Cancel

8. Specify OCR storage.

Storage Option Information **19c ORACLE**
Grid Infrastructure

You can place Oracle Cluster Registry (OCR) files and voting disk files on Oracle ASM storage, or on a file system.

☒ **Use Oracle Flex ASM for storage**
Choose this option to configure OCR and voting disks on ASM storage. ASM instance will be configured on reduced number of cluster nodes.

☐ **Use Shared File System**
Choose this option to configure OCR and voting disk files on an existing shared file system.

Configuration Option
Cluster Configuration
 Grid Plug and Play
 Cluster Node Information
 Network Interface Usage
Storage Option
 Create Grid Infrastructure Management Repository
 Grid Infrastructure Management
 Create ASM Disk Group
 ASM Password
 Operating System Groups
 Installation Location
 Root script execution
 Prerequisite Checks
 Summary
 Install Product
 Finish

Help < Back Next > Install Cancel

9. Create an OCR and VOTE disk group with normal redundancy.

Create ASM Disk Group **19c ORACLE**
Grid Infrastructure

OCR and Voting disk data will be stored in the following ASM Disk group. Select disks and characteristics of this Disk group.

Disk group name:

Redundancy: ☐ Flex ☐ High ☒ Normal ☐ External

Allocation Unit Size: MB

Select Disks Show All Disks

	Disk Path	Size (in MB)	Status	Failure Group
<input checked="" type="checkbox"/>	/dev/mapper/ora_crs_01_1	51200	Candidate	
<input checked="" type="checkbox"/>	/dev/mapper/ora_crs_01_2	51200	Candidate	
<input checked="" type="checkbox"/>	/dev/mapper/ora_crs_02_1	51200	Candidate	
<input checked="" type="checkbox"/>	/dev/mapper/ora_crs_02_2	51200	Candidate	

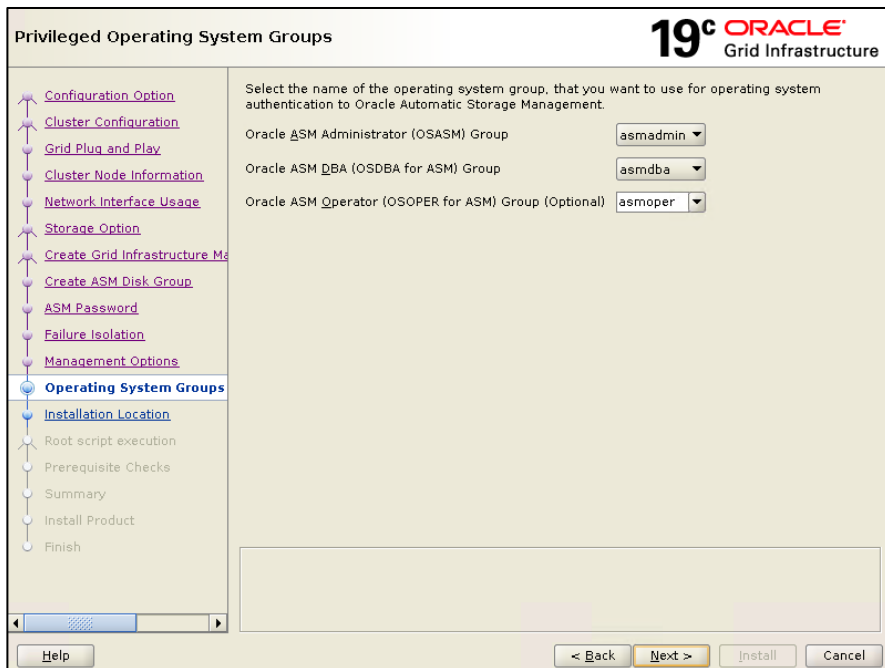
Disk Discovery Path: /dev/mapper/ora_crs*

☐ **Configure Oracle ASM Filter Driver**
Select this option to configure ASM Filter Driver (AFD) to simplify configuration and management of disk devices by Oracle ASM.

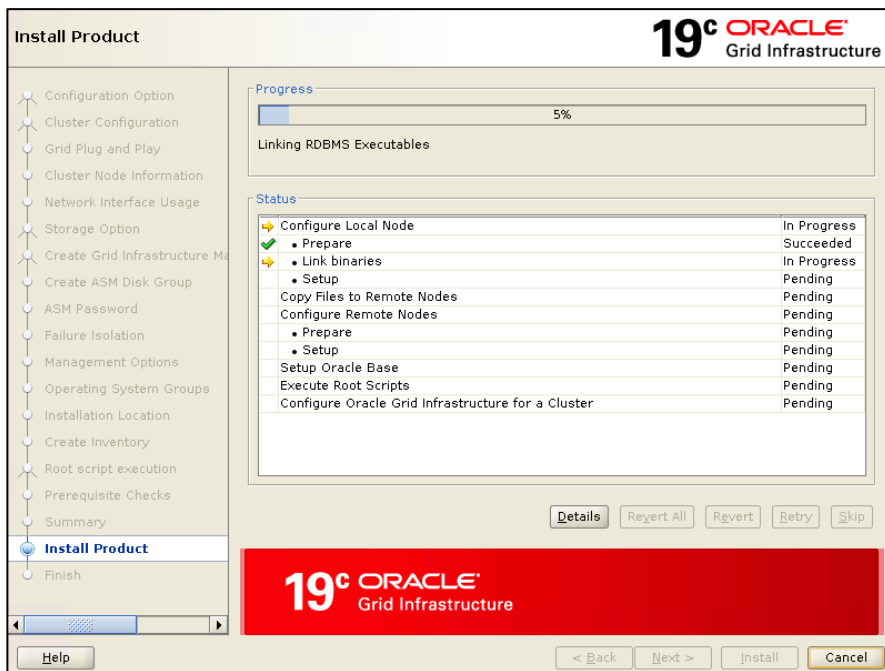
Configuration Option
Cluster Configuration
 Grid Plug and Play
 Cluster Node Information
 Network Interface Usage
 Storage Option
Create ASM Disk Group
 ASM Password
 Operating System Groups
 Installation Location
 Root script execution
 Prerequisite Checks
 Summary
 Install Product
 Finish

Help < Back Next > Install Cancel

10. Define privileged operating system groups:



11. Make sure there is no failed prerequisite check, and then launch installation.



12. Check the grid infrastructure state after installation.

```
[grid@b200-ora-01 bin]$ ./crsctl stat res -t
```

Name	Target	State	Server	State details
Local Resources				
ora.LISTENER.lsnr	ONLINE	ONLINE	b200-ora-01	STABLE

	ONLINE	ONLINE	b200-ora-02	STABLE
	ONLINE	ONLINE	b200-ora-03	STABLE
	ONLINE	ONLINE	b200-ora-04	STABLE
	ONLINE	ONLINE	b200-ora-05	STABLE
	ONLINE	ONLINE	b200-ora-06	STABLE
	ONLINE	ONLINE	b200-ora-07	STABLE
	ONLINE	ONLINE	b200-ora-08	STABLE
ora.chad				
	ONLINE	ONLINE	b200-ora-01	STABLE
	ONLINE	ONLINE	b200-ora-02	STABLE
	ONLINE	ONLINE	b200-ora-03	STABLE
	ONLINE	ONLINE	b200-ora-04	STABLE
	ONLINE	ONLINE	b200-ora-05	STABLE
	ONLINE	ONLINE	b200-ora-06	STABLE
	ONLINE	ONLINE	b200-ora-07	STABLE
	ONLINE	ONLINE	b200-ora-08	STABLE
ora.net1.network				
	ONLINE	ONLINE	b200-ora-01	STABLE
	ONLINE	ONLINE	b200-ora-02	STABLE
	ONLINE	ONLINE	b200-ora-03	STABLE
	ONLINE	ONLINE	b200-ora-04	STABLE
	ONLINE	ONLINE	b200-ora-05	STABLE
	ONLINE	ONLINE	b200-ora-06	STABLE
	ONLINE	ONLINE	b200-ora-07	STABLE
	ONLINE	ONLINE	b200-ora-08	STABLE
ora.ons				
	ONLINE	ONLINE	b200-ora-01	STABLE
	ONLINE	ONLINE	b200-ora-02	STABLE
	ONLINE	ONLINE	b200-ora-03	STABLE
	ONLINE	ONLINE	b200-ora-04	STABLE
	ONLINE	ONLINE	b200-ora-05	STABLE
	ONLINE	ONLINE	b200-ora-06	STABLE
	ONLINE	ONLINE	b200-ora-07	STABLE
	ONLINE	ONLINE	b200-ora-08	STABLE

Cluster Resources				

ora.ASMNET1LSNR_ASM.lsnr(ora.asmgroup)				
1	ONLINE	ONLINE	b200-ora-01	STABLE
2	ONLINE	ONLINE	b200-ora-08	STABLE
3	ONLINE	ONLINE	b200-ora-02	STABLE
ora.LISTENER_SCAN1.lsnr				
1	ONLINE	ONLINE	b200-ora-03	STABLE
ora.LISTENER_SCAN2.lsnr				
1	ONLINE	ONLINE	b200-ora-02	STABLE
ora.LISTENER_SCAN3.lsnr				
1	ONLINE	ONLINE	b200-ora-04	STABLE
ora.VOTE.dg(ora.asmgroup)				
1	ONLINE	ONLINE	b200-ora-01	STABLE
2	ONLINE	ONLINE	b200-ora-08	STABLE
3	ONLINE	ONLINE	b200-ora-02	STABLE
ora.asm(ora.asmgroup)				
1	ONLINE	ONLINE	b200-ora-01	Started,STABLE
2	ONLINE	ONLINE	b200-ora-08	Started,STABLE
3	ONLINE	ONLINE	b200-ora-02	Started,STABLE
ora.asmnet1.asmnetwork(ora.asmgroup)				
1	ONLINE	ONLINE	b200-ora-01	STABLE
2	ONLINE	ONLINE	b200-ora-08	STABLE
3	ONLINE	ONLINE	b200-ora-02	STABLE
ora.b200-ora-01.vip				
1	ONLINE	ONLINE	b200-ora-01	STABLE
ora.b200-ora-02.vip				
1	ONLINE	ONLINE	b200-ora-02	STABLE
ora.b200-ora-03.vip				
1	ONLINE	ONLINE	b200-ora-03	STABLE
ora.b200-ora-04.vip				
1	ONLINE	ONLINE	b200-ora-04	STABLE
ora.b200-ora-05.vip				
1	ONLINE	ONLINE	b200-ora-05	STABLE
ora.b200-ora-06.vip				
1	ONLINE	ONLINE	b200-ora-06	STABLE

```

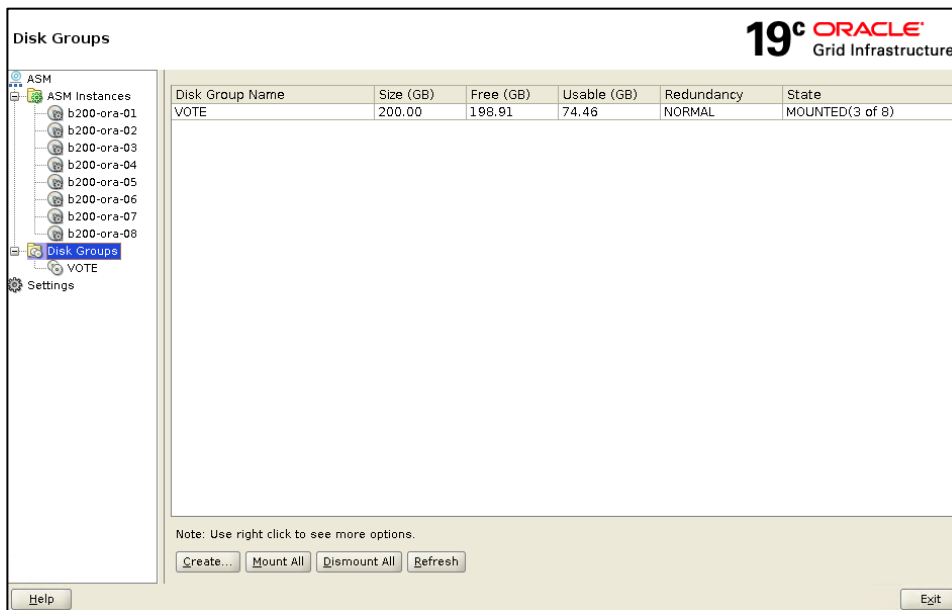
ora.b200-ora-07.vip
1          ONLINE  ONLINE      b200-ora-07          STABLE
ora.b200-ora-08.vip
1          ONLINE  ONLINE      b200-ora-08          STABLE
ora.cvu
1          ONLINE  ONLINE      b200-ora-01          STABLE
ora.qosmsserver
1          ONLINE  ONLINE      b200-ora-01          STABLE
ora.scan1.vip
1          ONLINE  ONLINE      b200-ora-03          STABLE
ora.scan2.vip
1          ONLINE  ONLINE      b200-ora-02          STABLE
ora.scan3.vip
1          ONLINE  ONLINE      b200-ora-04          STABLE
-----
[grid@b200-ora-01 bin]$

[grid@b200-ora-01 admin]$ crsctl query css votedisk
##  STATE      File Universal Id                        File Name                     Disk group
--  -
1.  ONLINE     939a77d6e0a64fd2bf4c9d84f75731ee (/dev/mapper/ora_crs_01_1) [VOTE]
2.  ONLINE     a0980f72fdad4fb3bf5889843fd36b53 (/dev/mapper/ora_crs_01_2) [VOTE]
3.  ONLINE     a9b03178bdfd4fa9bf3024f8b1e2f0a3 (/dev/mapper/ora_crs_02_1) [VOTE]
Located 3 voting disk(s).

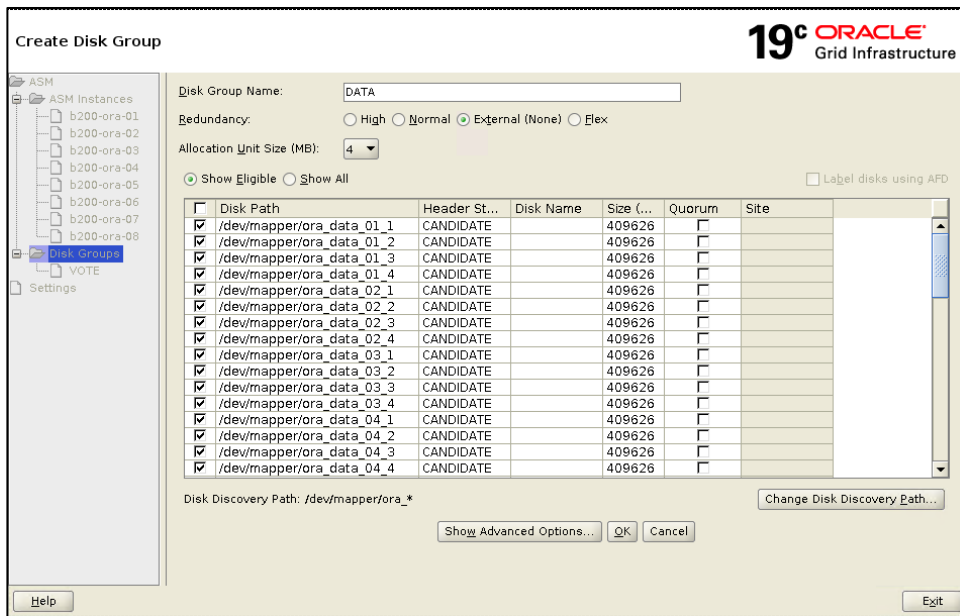
```

13. Create +DATA and +REDO disk groups.

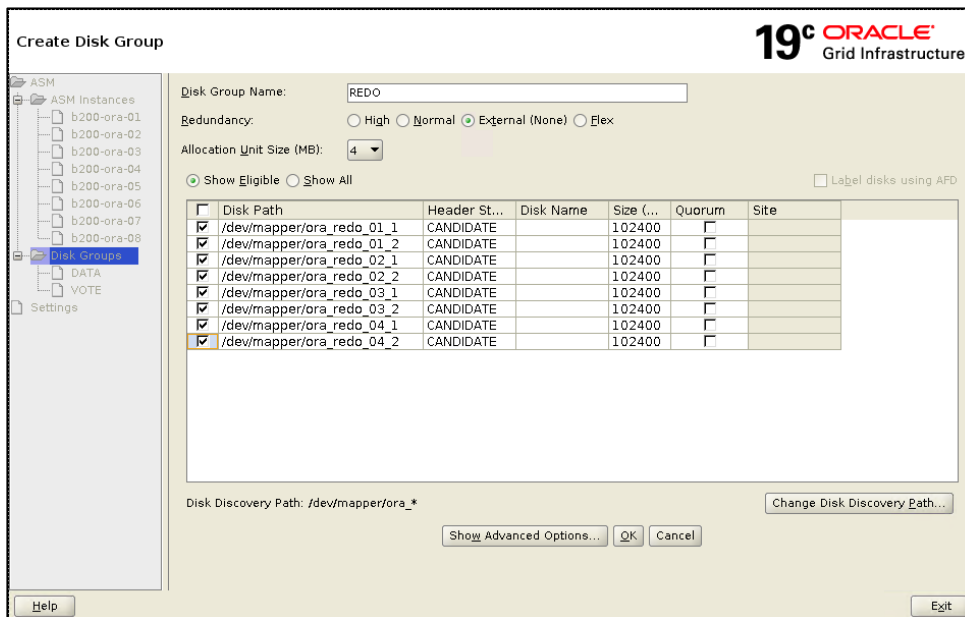
14. Launch ASMCA to configure +DATA and +REDO disk groups for database install.



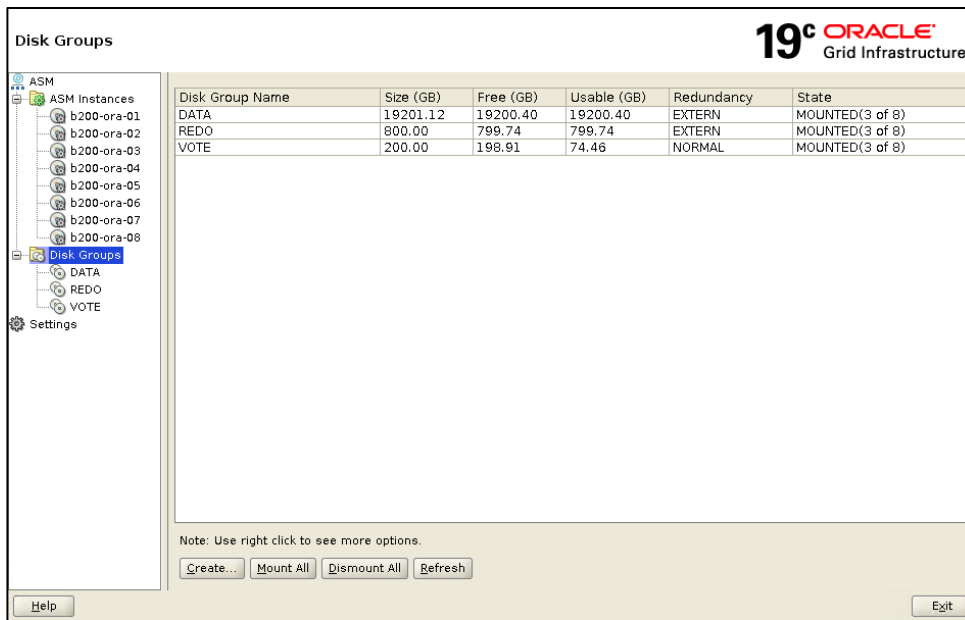
15. Click Create and change the disk discovery path to the right path for the data and redo disks to be discovered by ASMCA.



16. Choose all disks to be tagged for the DATA disk group with external redundancy.
17. Similarly, create a REDO disk group for flash recovery area.



18. The following screenshot depicts the disk groups summary:



Oracle 19c Database software only installation

For Oracle 19c database software installation, we also applied the 19.8 database release update to match the grid infrastructure patch level.

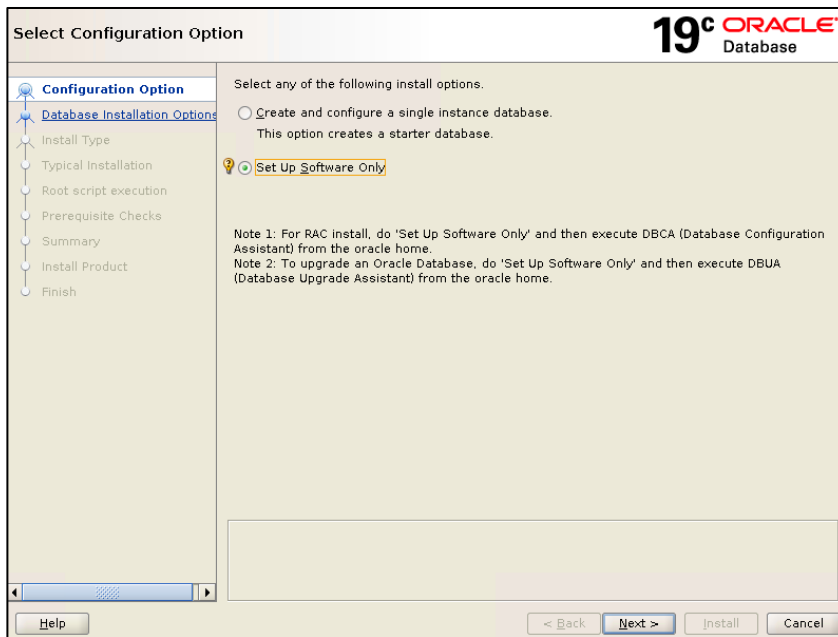
1. Download Oracle database release update [31281355](#) for 19.8 patching.
2. Launch the installer with the RU patch update.

```
[oracle@b200-ora-01 ntap]$ ls /tmp/archive/oracle
31281355  p31281355_190000_Linux-x86-64.zip  PatchSearch.xml
[oracle@b200-ora-01 ntap]$ ./runInstaller -applyRU /tmp/archive/oracle/31281355
Preparing the home to patch...
Applying the patch /tmp/archive/oracle/31281355...
Successfully applied the patch.
The log can be found at: /u01/app/oraInventory/logs/InstallActions2020-12-02_03-17-21PM/installerPatchActions_2020-12-02_03-17-21PM.log
Launching Oracle Database Setup Wizard...

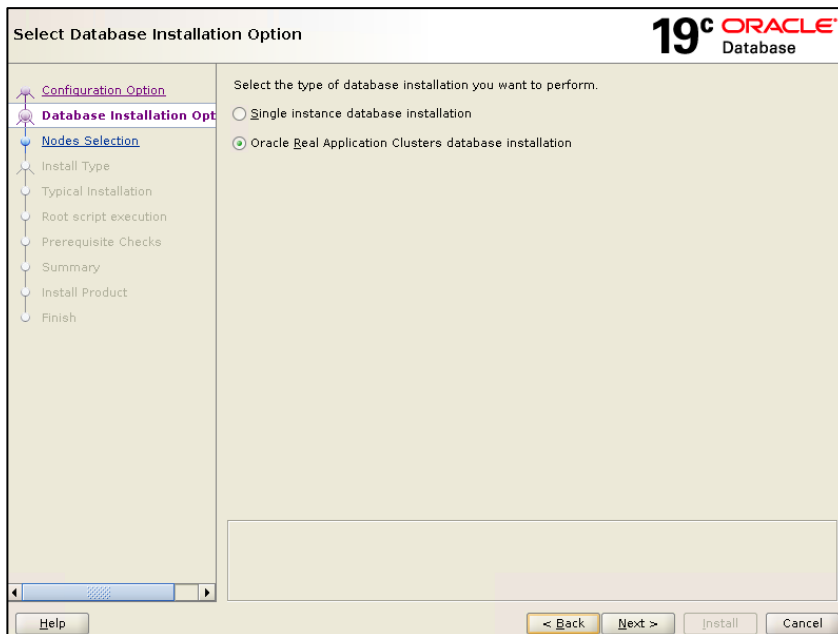
The response file for this session can be found at:
/u01/app/oracle/product/19300/ntap/install/response/db_2020-12-02_03-17-21PM.rsp

You can find the log of this install session at:
/u01/app/oraInventory/logs/InstallActions2020-12-02_03-17-21PM/installActions2020-12-02_03-17-21PM.log
```

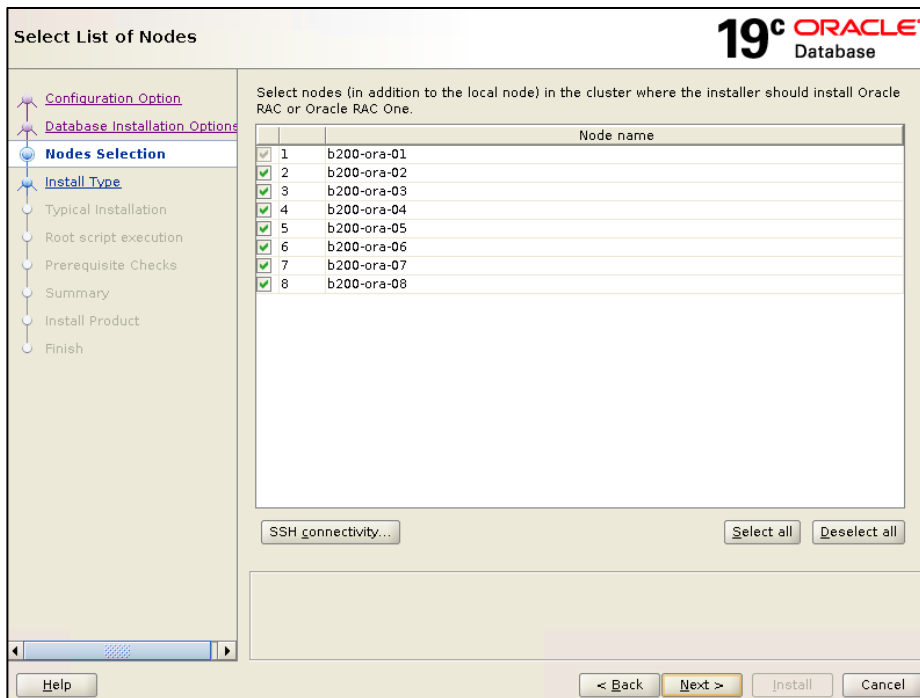
3. The installer launches after the release update is applied in one command step.



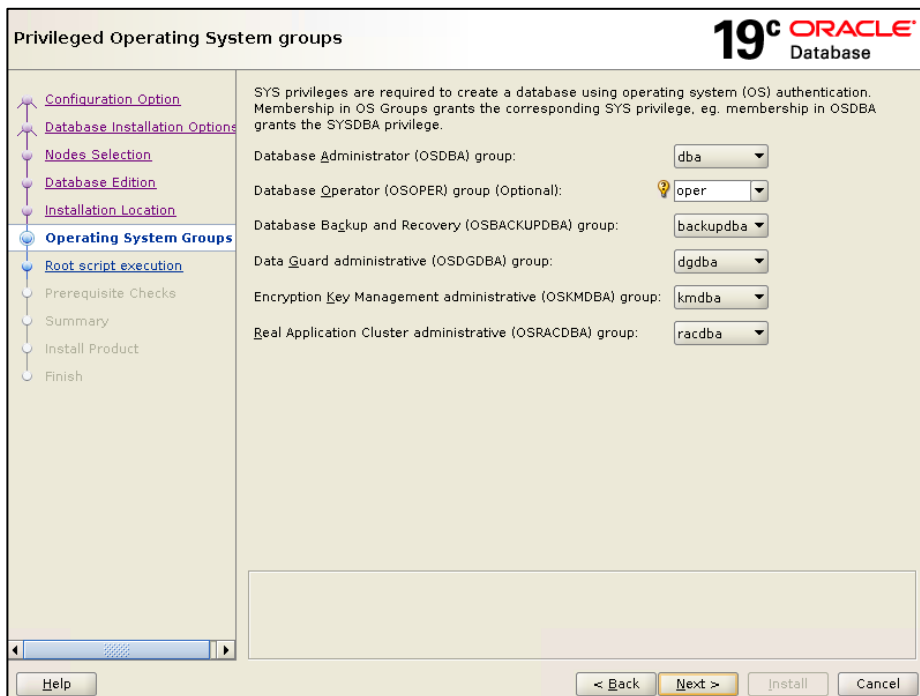
4. Select the Set Up Software Only and Oracle Real Application Clusters Database Installation option.



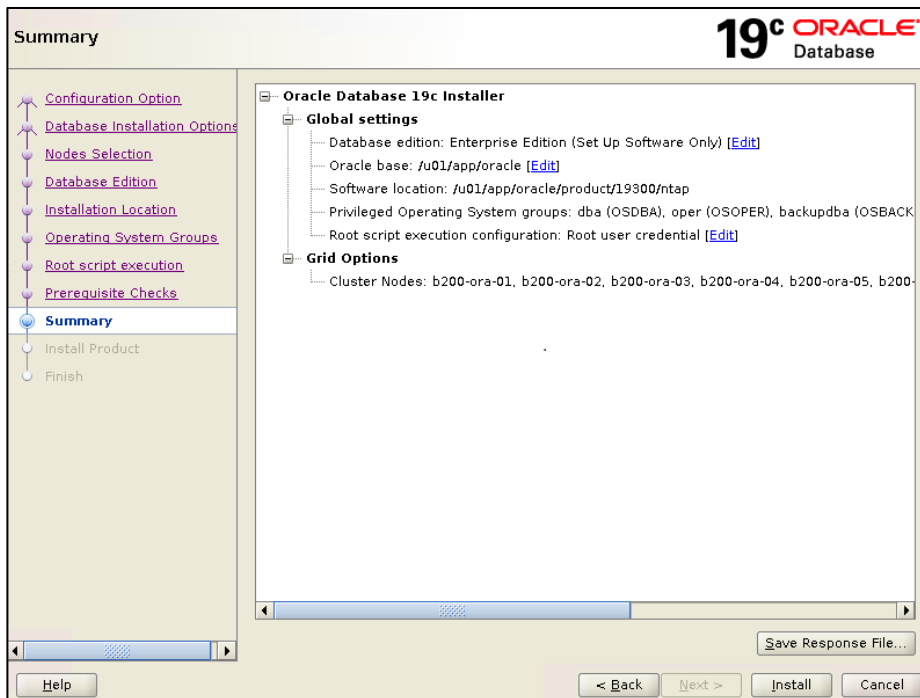
5. Select all the cluster nodes detected.



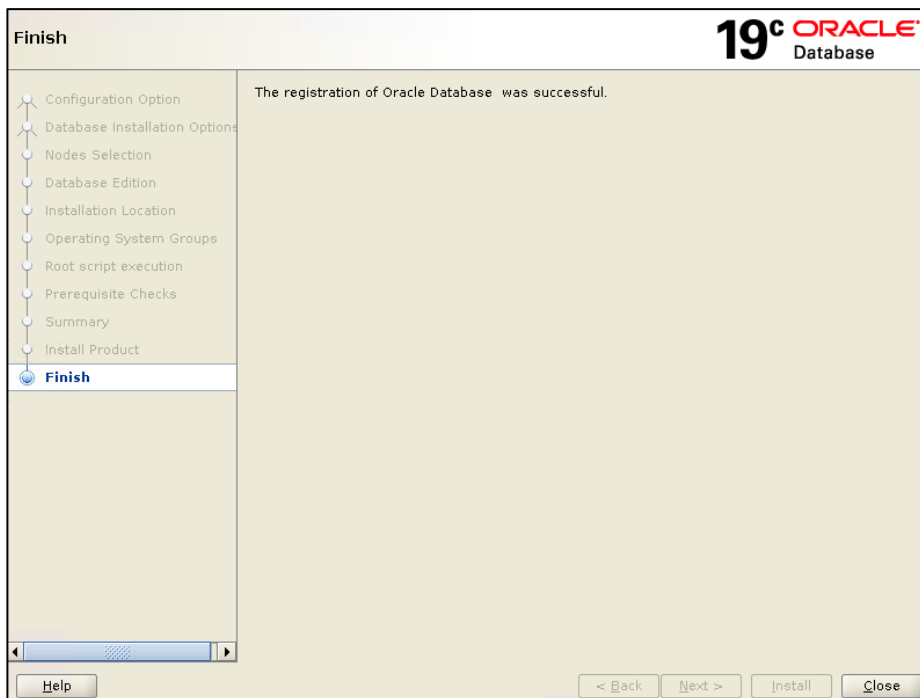
6. Select the privileged operating system group if different the OS groups are defined for separating management tasks.



7. Ignore the NTP time synchronization warning as the NTPD has been replaced by Chronyd in Linux 8.



8. Click Install to finish the software only installation.



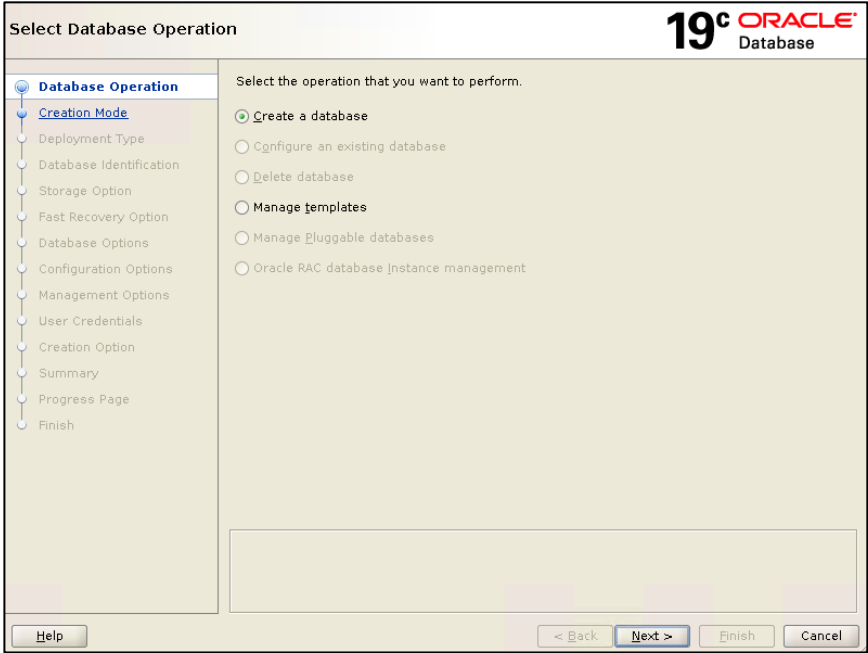
Oracle 19c database deployment with CDB/PDB model

The future of Oracle database deployment is the CDB/PDB model as signified by Oracle in phasing out the standard single instance from release 20. CDB/PDB provides an option for resource sharing and database consolidation without relying on virtualization, which adds overhead as well as additional

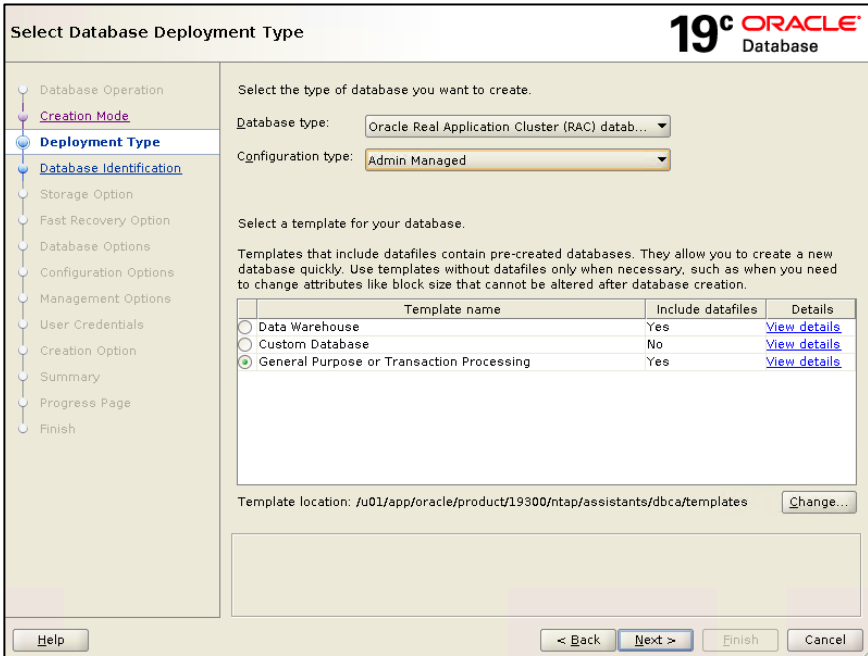
virtualization license costs. You can host as many as 252 PDBs in a container in theory with the enterprise edition.

The following screenshot demonstrates how to create PDBs in a CDB using the Oracle utility DBCA in an RAC configuration.

1. Launch DBCA.



2. Select the type of database to create using templates.



3. Select the RAC node to run the database.

Select List of Nodes **19^c ORACLE[®]**
Database

Select the nodes on which you want to create the cluster database. The local node "b200-ora-01" should always be selected.

	Node name
<input checked="" type="checkbox"/>	1 b200-ora-01
<input checked="" type="checkbox"/>	2 b200-ora-02
<input checked="" type="checkbox"/>	3 b200-ora-03
<input checked="" type="checkbox"/>	4 b200-ora-04
<input checked="" type="checkbox"/>	5 b200-ora-05
<input checked="" type="checkbox"/>	6 b200-ora-06
<input checked="" type="checkbox"/>	7 b200-ora-07
<input checked="" type="checkbox"/>	8 b200-ora-08

4. Create a container database and the number of pluggable databases.

Specify Database Identification Details **19^c ORACLE[®]**
Database

Provide a unique database identifier information. An Oracle database is uniquely identified by a Global database name, typically of the form "name.domain".

Global database name:

SID Prefix:

Service name:

☒ **Create as Container database**

A Container database can be used for consolidating multiple databases into a single database, and it enables database virtualization. A Container database (CDB) can have zero or more pluggable databases (PDB).

☒ Use Local Undo tablespace for PDBs

☐ Create an empty Container database

☒ Create a Container database with one or more PDBs

Number of PDBs:

PDB name prefix:

5. Select the storage option.

Select Database Storage Option **19c** ORACLE Database

- Database Operation
- Creation Mode
- Deployment Type
- Nodes Selection
- Database Identification
- Storage Option**
- Fast Recovery Option
- Database Options
- Configuration Options
- Management Options
- User Credentials
- Creation Option
- Prerequisite Checks
- Summary
- Progress Page
- Finish

☐ Use template file for database storage attributes

Storage type and location for database files will be picked up from the specified template (General Purpose or Transaction Processing).

☒ Use following for the database storage attributes

All the database files will be put at the specified location below. You can customize the name and location of each datafile in the subsequent screen.

Database files storage type: Automatic Storage Management (ASM)

Database files location: +DATA/{DB_UNIQUE_NAME} Browse...

Oracle Managed files option will enable Oracle to automatically generate the names of the datafiles for simplified database management.

☒ Use Oracle-Managed Files (OMF) Multiplex redo logs and control files...

File location variables...

Help
< Back
Next >
Finish
Cancel

Ideally create a sperate disk group for each container database so that QoS can be individually set at the container database volume level.

6. Define the flash recovery area and log archiving.

Select Fast Recovery Option **19c** ORACLE Database

- Database Operation
- Creation Mode
- Deployment Type
- Nodes Selection
- Database Identification
- Storage Option
- Fast Recovery Option**
- Database Options
- Configuration Options
- Management Options
- User Credentials
- Creation Option
- Prerequisite Checks
- Summary
- Progress Page
- Finish

Choose the recovery options for the database.

☒ Specify Fast Recovery Area

Recovery files storage type: Automatic Storage Management (ASM)

Fast Recovery Area: +REDO Browse...

Fast Recovery Area size: 800 GB

☐ Enable archiving Edit archive mode parameters...

Help
< Back
Next >
Finish
Cancel

7. Define the container database configuration in sizing, character sets, and connection mode.

Specify Configuration Options **19^c ORACLE[®]**
Database

- Database Operation
- Creation Mode
- Deployment Type
- Nodes Selection
- Database Identification
- Storage Option
- Fast Recovery Option
- Data Vault Option
- Configuration Options**
- Management Options
- User Credentials
- Creation Option
- Prerequisite Checks
- Summary
- Progress Page
- Finish

Memory
Sizing
Character sets
Connection mode

☒ **Use Automatic Shared Memory Management**

SGA size: MB 490 **65541** 191816

PGA Size: MB

☐ **Use Manual Shared Memory Management**

Shared pool size: MB

Buffer cache size: MB

Java pool size: MB

Large pool size: MB

PGA size: MB

Total memory for database 0 MB

☐ **Use Automatic Memory Management**

Memory target: MB 490 **76726** 191816 39%

Help
< Back
Next >
Finish
Cancel

8. This image shows the database creation summary.

Summary **19^c ORACLE[®]**
Database

- Database Operation
- Creation Mode
- Deployment Type
- Nodes Selection
- Database Identification
- Storage Option
- Fast Recovery Option
- Data Vault Option
- Configuration Options
- Management Options
- User Credentials
- Creation Option
- Prerequisite Checks
- Summary**
- Progress Page
- Finish

Database Configuration Assistant

Global Settings

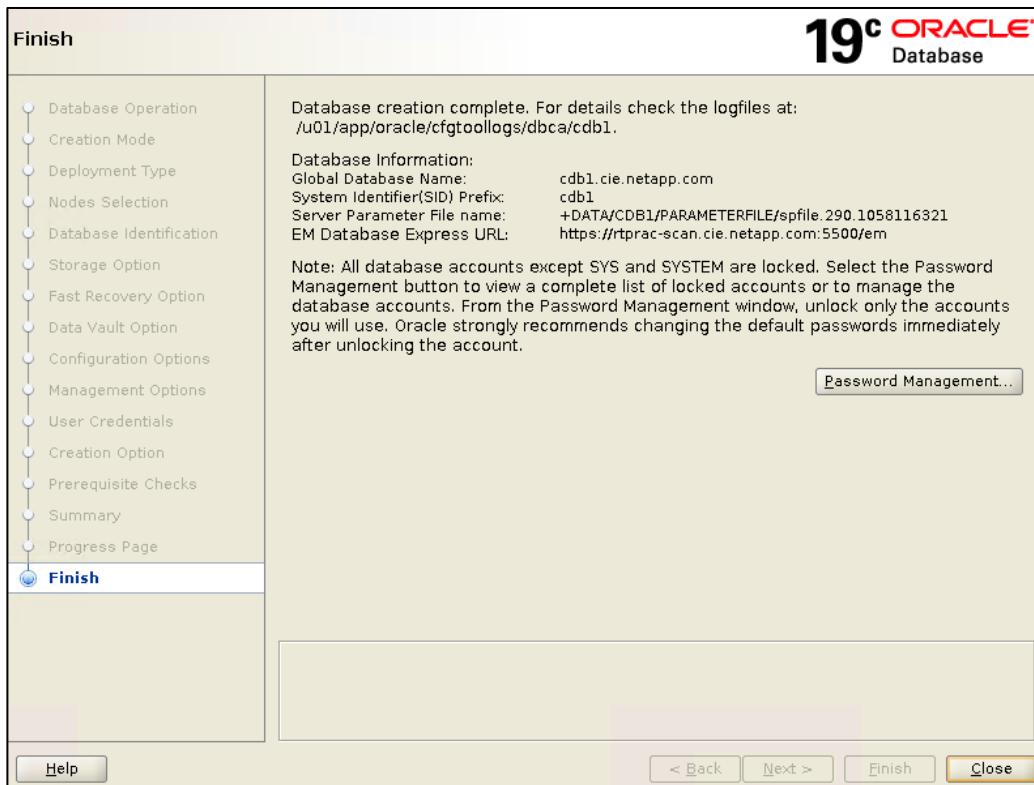
- Global database name: cdb1.cie.netapp.com
- Node List: b200-ora-01,b200-ora-02,b200-ora-03,b200-ora-04,b200-ora-05,b200-ora-06
- Configuration type: Oracle Real Application Cluster (RAC) database - Admin Managed
- SID: cdb1
- Create as Container database: Yes
- Pluggable Database Name: cdb1_pdb
- Number of Pluggable Databases: 2
- Use Local Undo tablespace for PDBs: Yes
- Database Files Storage Type: Automatic Storage Management (ASM)
- Memory Configuration Type: Automatic Shared Memory Management
- Template name: General Purpose

Initialization Parameters

- audit_file_dest: {ORACLE_BASE}/admin/{DB_UNIQUE_NAME}/adump
- audit_trail: db
- cluster_database: true
- compatible: 19.0.0
- db_block_size: 8192 BYTES
- db_create_file_dest: +DATA/{DB_UNIQUE_NAME}/
- db_domain: cie.netapp.com
- db_name: cdb1

Help
< Back
Next >
Finish
Cancel

9. Database creation finish.

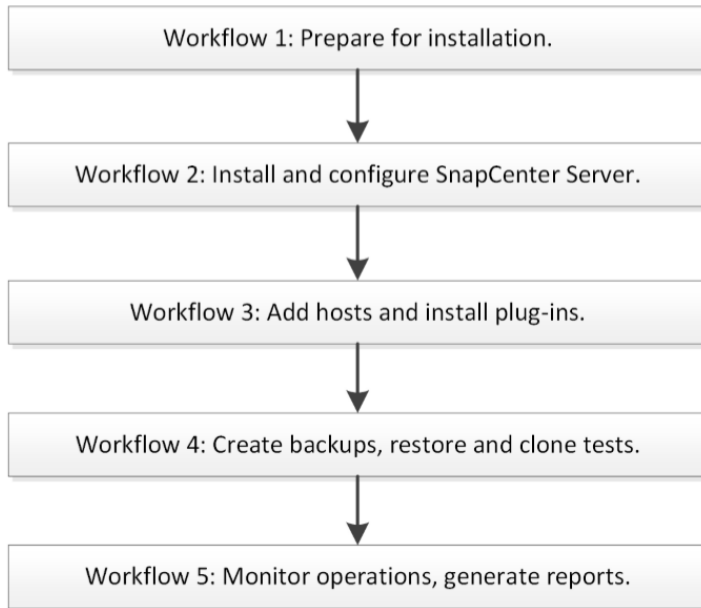


SnapCenter deployment

The instructions that follow highlight the key steps for setting up SnapCenter to manage an Oracle database environment. For comprehensive details for installation and configuration, see the [Installing and setting up SnapCenter Guide](#).

Getting started

To get started with SnapCenter, you must install the SnapCenter Server, add a host that automatically installs the appropriate plug-in, and then run an Oracle database backup. See the following getting started workflow:



Prepare for installation

1. Check host requirements.
 - a. Check domain and workgroup requirements.
 - b. If you are using an Active Directory domain, you must use a domain user with local administrator rights. The domain user must be a member of the local administrator group on the Windows host.
 - c. If you are using workgroups, you must use a local account that has local administrator rights.
 - d. Operating system: Microsoft Windows Server 2012 and above.
 - e. Minimum CPU count: 4 cores.
 - f. Minimum RAM: 8G.
 - g. Minimum storage: software, log, and repository – 10G.
 - h. Required software packages:
 - Microsoft .NET Framework 4.5.2 or later
 - Windows Management Framework (WMF) 4.0 or later
 - PowerShell 4.0 or later
2. Check supported storage systems, applications, browsers, and ports.
 - a. SnapCenter supports ONTAP 8.3.0 and later to protect your data.
 - b. SnapCenter supports protection of different applications and databases.
 - c. For detailed information about the supported Oracle databases version, see the Interoperability Matrix Tool (IMT) - [NetApp Interoperability Matrix Tool](#).
 - d. Supported browsers.
 - Chrome. If you are using v66, you might fail to launch the SnapCenter GUI. For information about the issue and the solution, see the [NetApp Knowledge Base](#).
 - Internet Explorer. Only default-level security is supported. Making changes to Internet Explorer security settings results in significant browser display issues. Internet Explorer compatibility view must be disabled.
 - Microsoft Edge. For latest supported browsers, check the NetApp IMT tool.

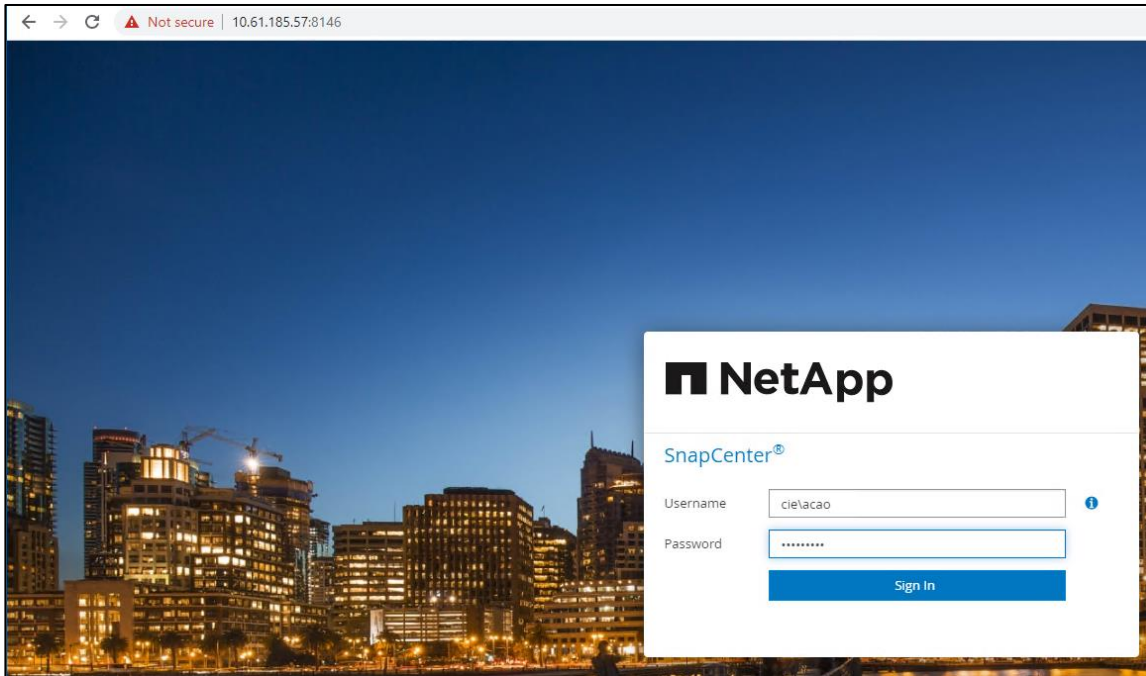
3. Check licensing requirements. Several licenses are required for data protection operations. The [SnapCenter Software Release Notes](#) contain details about the licenses required for your environment.
4. Connections and ports. You should make sure that connection and port requirements are met before installing the SnapCenter Server and application or database plug-ins. If there is a firewall between the SnapCenter host and target hosts, set up rules to allow traffic flow through the firewall. See the following list of ports that must be open for Oracle deployment:

Type of port	Default port
SnapCenter port	8146 (HTTPS), bidirectional, customizable, as in the URL <code>https://server:8146</code> . Used for communication between the SnapCenter client (the SnapCenter user) and the SnapCenter server. Also used for communication from the plug-in hosts to the SnapCenter server.
SnapCenter SMCore communication port	8145 (HTTPS), bidirectional, customizable. The port is used for communication between the SnapCenter server and the hosts where the SnapCenter plug-ins are installed.
Linux or AIX plug-in hosts	Twenty-two (SSH) ports are used for communication between the SnapCenter server and the host where the plug-in is being installed. The ports are used by SnapCenter to copy plug-in package binaries to Linux or AIX plug-in hosts and should be open or excluded from the firewall or iptables.
SnapCenter Plug-in for Oracle Database	27216, customizable. The default JDBC port is used by the plug-in for Oracle for connecting to the Oracle database.
ONTAP cluster or SVM communication port	443 (HTTPS), bidirectional 80 (HTTP), bidirectional. The port is used by the storage abstraction layer (SAL) for communication between the host running the SnapCenter server and the SVM. The port is currently also used by the SAL on SnapCenter for Windows plug-in hosts for communication between the SnapCenter plug-in host and the SVM.
Domain controller communication port	See the Microsoft documentation to identify the ports that should be opened in the firewall on a domain controller for authentication to work properly. It is necessary to open the Microsoft required ports on the domain controller so that the SnapCenter server, plug-in hosts, or other Windows clients can authenticate users.

Install and configure SnapCenter server

1. Install the SnapCenter server using the Install wizard.
2. Log in to SnapCenter.

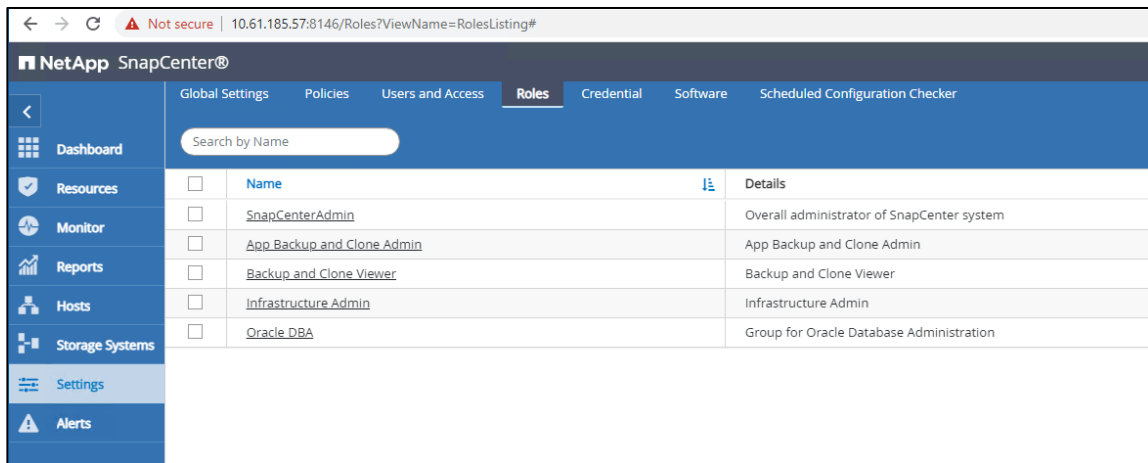
The default GUI URL is a secure connection to default port 8146 on the server where the SnapCenter Server is installed (`https://server:8146`). If you provided a different server port during SnapCenter installation, that port is used instead.



3. Configure role-based access control (RBAC).

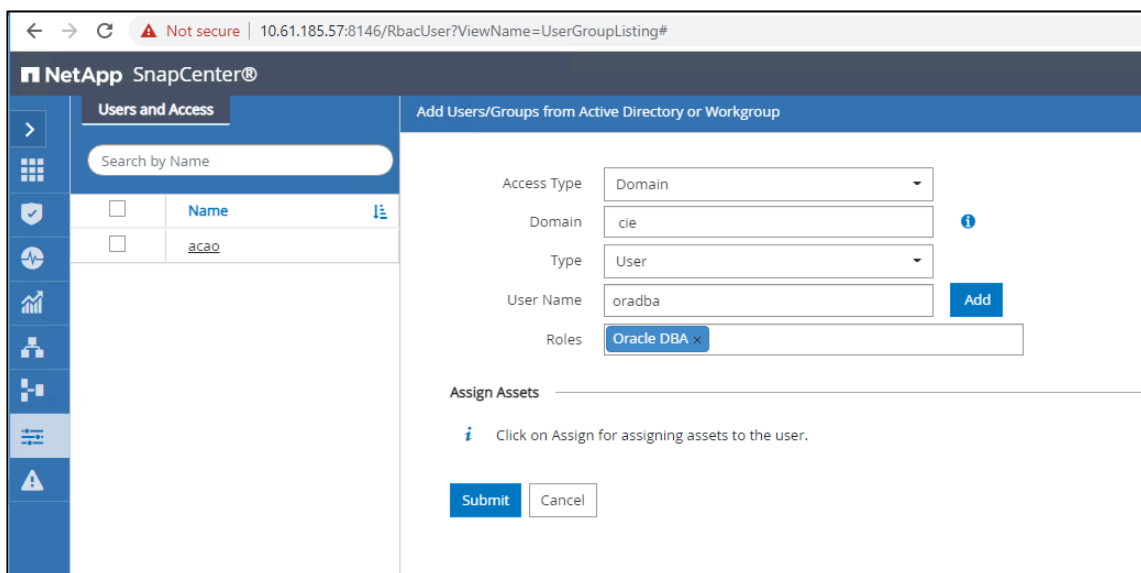
Roles can be created to assign relevant permissions for the role to operate on with different type of assets managed by SnapCenter. The roles can then be assigned to users intended for the roles or functions.

Asset Name	Allow	Create	Read	Update	Delete
Resource Group	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Policy	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Backup	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Host	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Storage Connection	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clone	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>



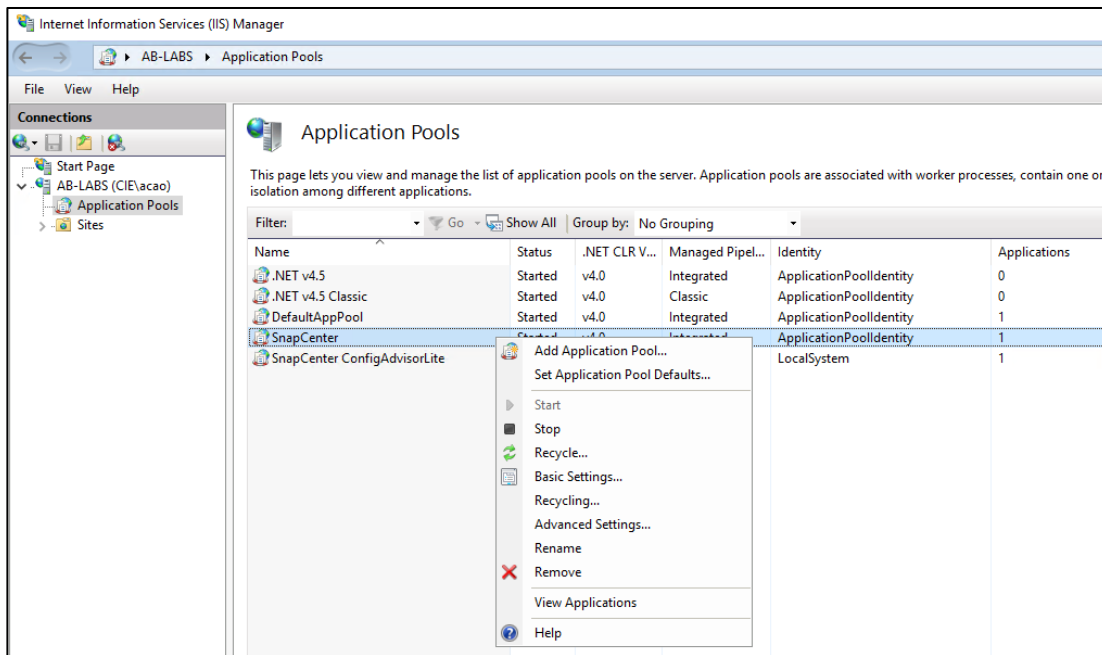
4. Add a user or group and assigning role and assets.

Assets must be created beforehand to assign proper assets to new users.

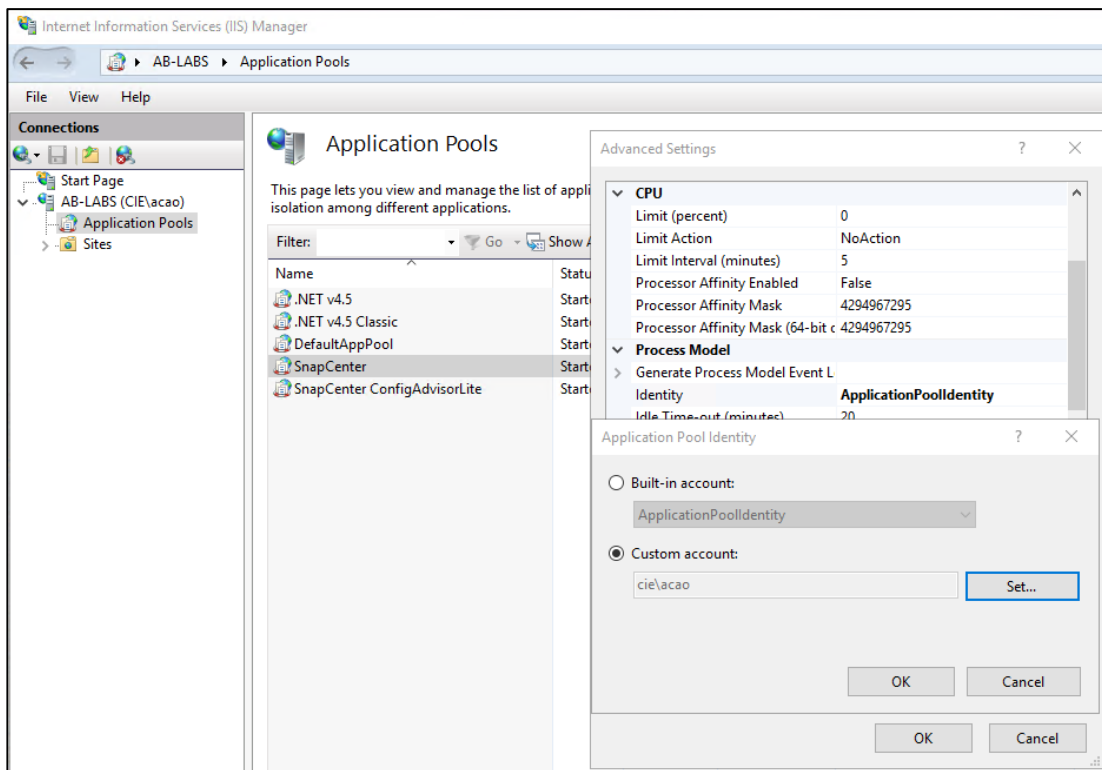


5. Configure IIS application pools to enable Active Directory read permissions.

Open IIS Manager on the Windows Server where SnapCenter is installed. In the left navigation pane, click Application Pools. Select SnapCenter in the Application Pools list, and then click Advanced Settings in the Actions pane.



6. Select Identity, and then click ... to edit the SnapCenter application pool identity. In the Custom Account field, enter a domain user or domain admin account name with Active Directory read permission. Click OK. The custom account replaces the built-in ApplicationPoolIdentity account for the SnapCenter application pool.

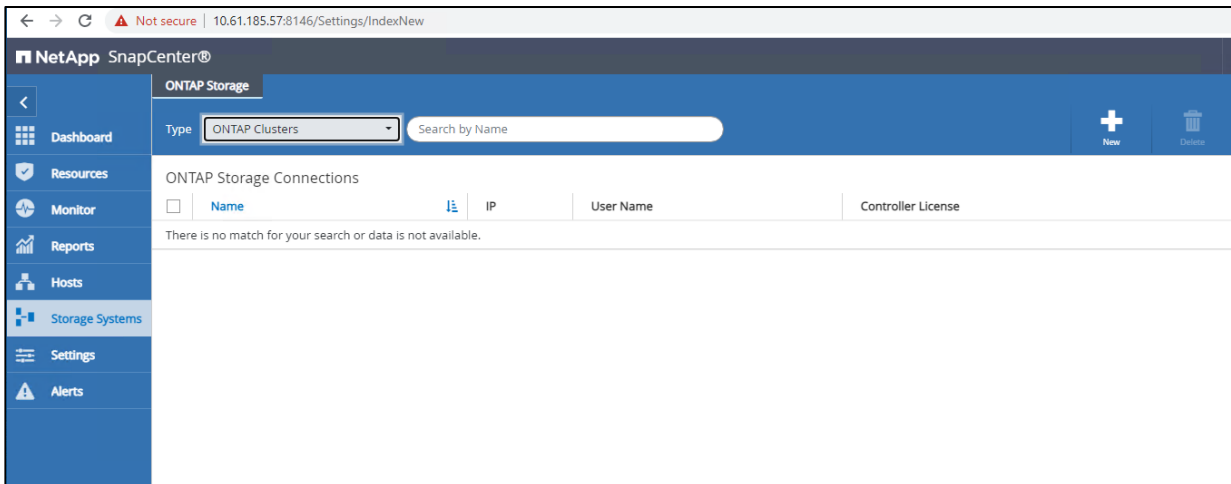


7. Click on OK to set the custom account.

8. Adding storage systems. You should set up the storage system that gives SnapCenter access to ONTAP storage to perform data protection and provisioning operations. You can either add a standalone SVM or a cluster comprising of multiple SVMs.

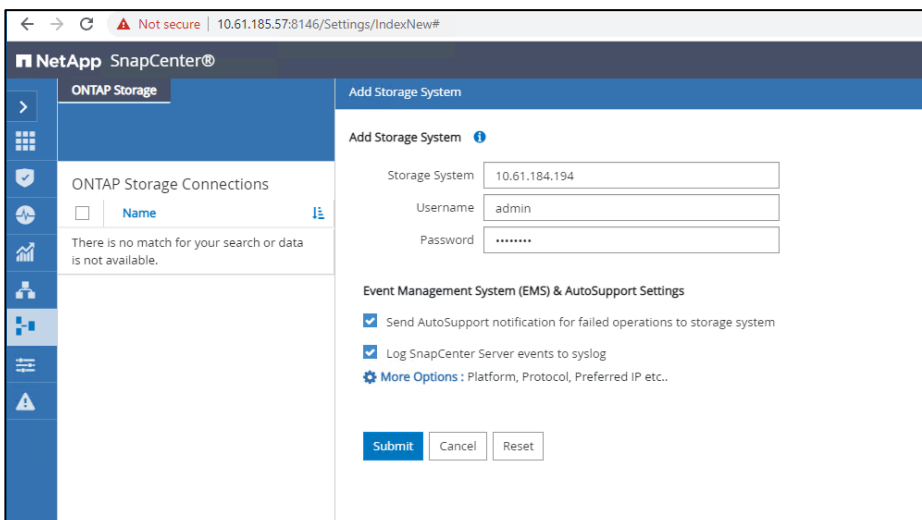
Note:

- You should have the required permissions in the Infrastructure Admin role to create storage connections.
 - You should ensure that the plug-in installations are not in progress. Host plug-in installations must not be in progress while adding a storage system connections because the host cache might not be updated, and the database status might be displayed in the SnapCenter GUI as Not Available for Backup or Not On NetApp Storage.
 - Storage system names should be unique. SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.
- a. In the left navigation pane, click Storage Systems.
 - b. On the Storage Systems page, click New.

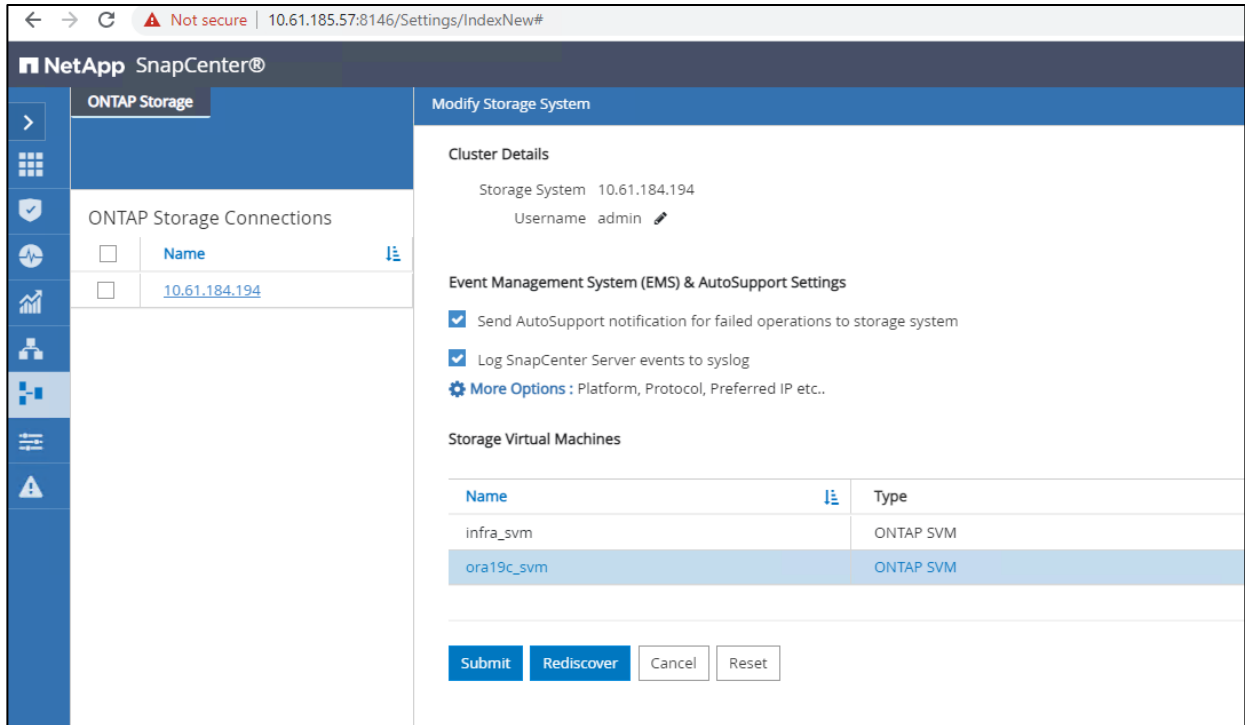


You can choose to either add an SVM or a cluster. In our testing, we chose to add a cluster by changing the type to ONTAP Clusters from the drop-down window.

- c. On the Add Storage System page, provide the following information:



- d. Enter the cluster virtual IP as storage system address and admin/password for authentication. Click Submit to add the storage system.
9. After the storage system is added, click Storage System. You can see that the two SVMs (infra_svm and ora19c_svm) that we created for the Oracle 19c solution have been added to SnapCenter.



- a. If a new SVM is added to the ONTAP cluster using the ONTAP GUI, click Rediscover to view the newly added SVM.
- b. Enable AutoSupport on each storage system. A cluster administrator must enable AutoSupport on each storage system node to send email notifications from all storage systems to which SnapCenter has access by running the following command from the storage system command line:

```
autosupport trigger modify -node nodename -autosupport-message client.app.info enable -noteto enable
```

10. Secure the SnapCenter web server by disabling SSL 3.0.
 - a. To launch Registry Editor on the SnapCenter web server host, click Start > Run, and then enter regedit.
 - b. In the Registry Editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 3.0\.
 - If the server key already exists, select the enabled DWORD, and then click Edit > Modify. Then change the value to 0, and click OK.
 - If the Server key does not exist, click Edit > New > Key, and then name the key Server. With the new Server key selected, click Edit > New > DWORD. Name the new DWORD enabled, and then enter 0 as the value.
 - c. Close Registry Editor.
11. Add SnapCenter licenses. If you already have a SnapManagerSuite license on your controller, SnapCenter Standard controller-based license entitlement is provided automatically. The names

SnapManagerSuite license and SnapCenter Standard controller-based license are used interchangeably, but they refer to the same license.

You can use the SnapCenter GUI or command line to view whether a SnapManager Suite license is installed on FAS or AFF primary storage systems and to identify which storage systems might require SnapManager Suite licenses. SnapManager Suite licenses apply only to FAS and AFF SVMs or clusters on primary storage systems.

```
FlexPod-A800-01-02::> license show
(system license show)
```

Serial Number: 1-80-000011			
Owner: FlexPod-A800-01-02			
Package	Type	Description	Expiration
Base	site	Cluster Base License	-

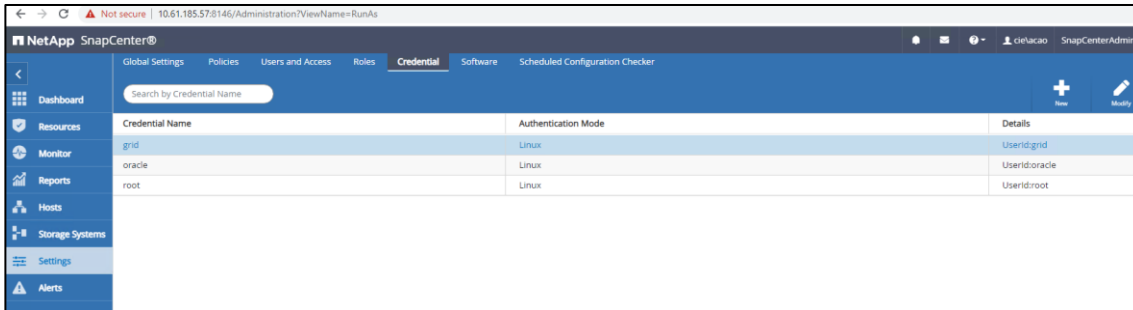
```
Serial Number: 1-81-0000000000000952029000220
Owner: FlexPod-A800-01-02-01
```

Package	Type	Description	Expiration
NFS	license	NFS License	-
CIFS	license	CIFS License	-
iSCSI	license	iSCSI License	-
FCP	license	FCP License	-
SnapRestore	license	SnapRestore License	-
SnapMirror	license	SnapMirror License	-
FlexClone	license	FlexClone License	-
SnapVault	license	SnapVault License	-
SnapManagerSuite	license	SnapManagerSuite License	-
TPM	license	Trusted Platform Module License	-
VE	license	Volume Encryption License	-
SnapMirror_Sync	license	SnapMirror Synchronous License	-

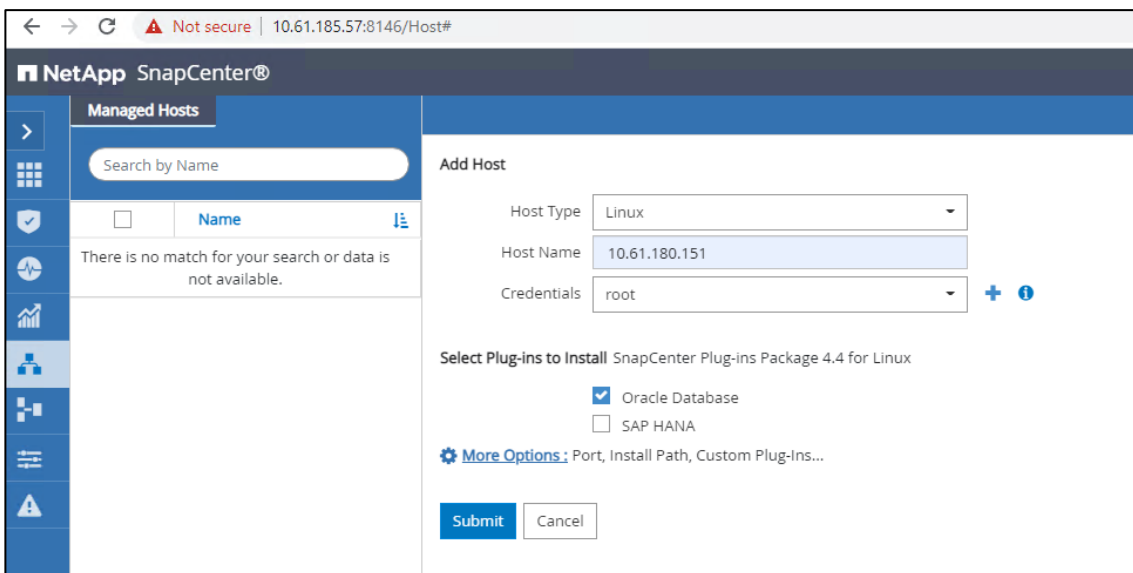
```
Serial Number: 1-81-0000000000000952029000447
Owner: FlexPod-A800-01-02-02
```

Package	Type	Description	Expiration
NFS	license	NFS License	-
CIFS	license	CIFS License	-
iSCSI	license	iSCSI License	-
FCP	license	FCP License	-
SnapRestore	license	SnapRestore License	-
SnapMirror	license	SnapMirror License	-
FlexClone	license	FlexClone License	-
SnapVault	license	SnapVault License	-
SnapManagerSuite	license	SnapManagerSuite License	-
TPM	license	Trusted Platform Module License	-
VE	license	Volume Encryption License	-
SnapMirror_Sync	license	SnapMirror Synchronous License	-

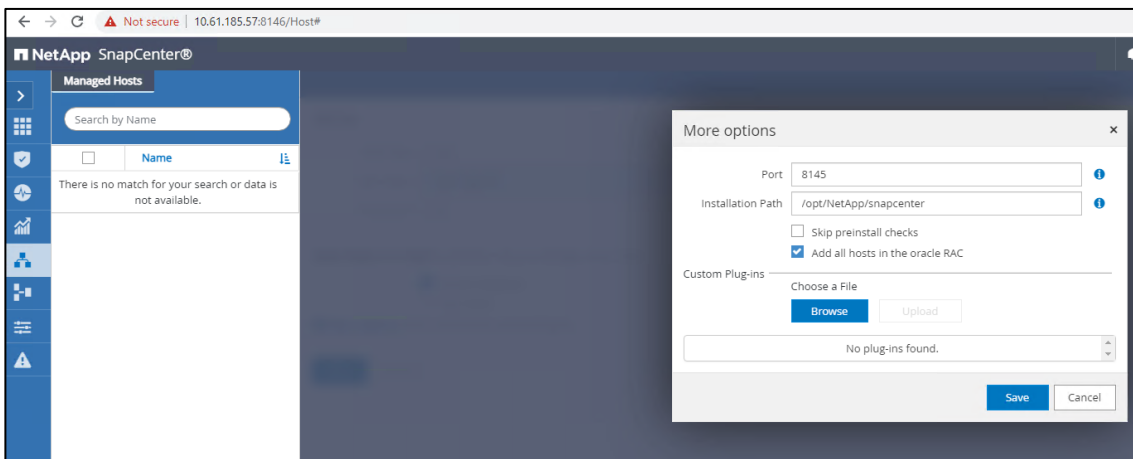
25 entries were displayed.




- c. Add hosts and install the plug-in.
- d. From the left pane, click the hosts tab, and then click Add to open the Add Hosts screen shown below:



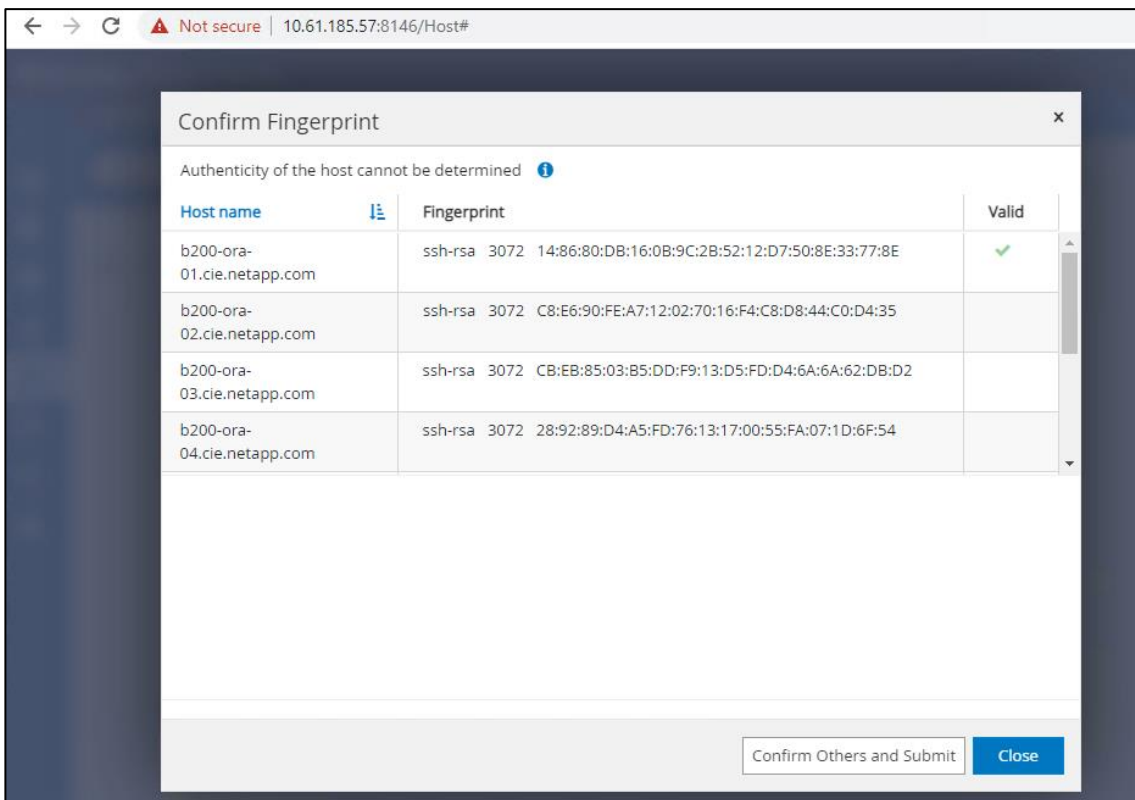
- e. Enter the first node IP address of the RAC cluster and choose Root Credential for the installation. Click More Options to check Add All Hosts in the Oracle RAC.



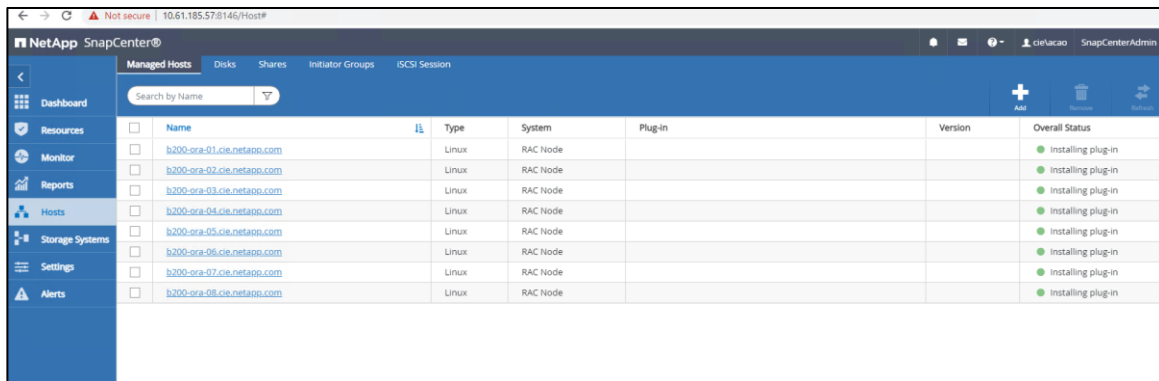
- f. Click Submit to start installation. You are prompted to confirm the fingerprint right after the start.



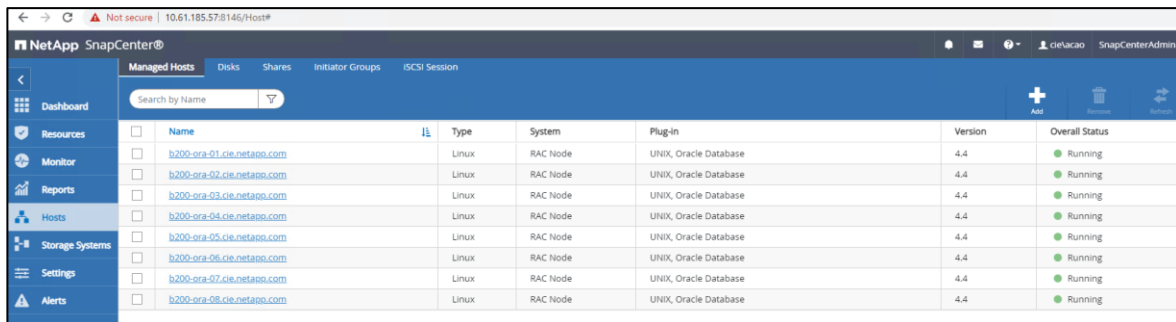
g. Click Confirm and Submit.



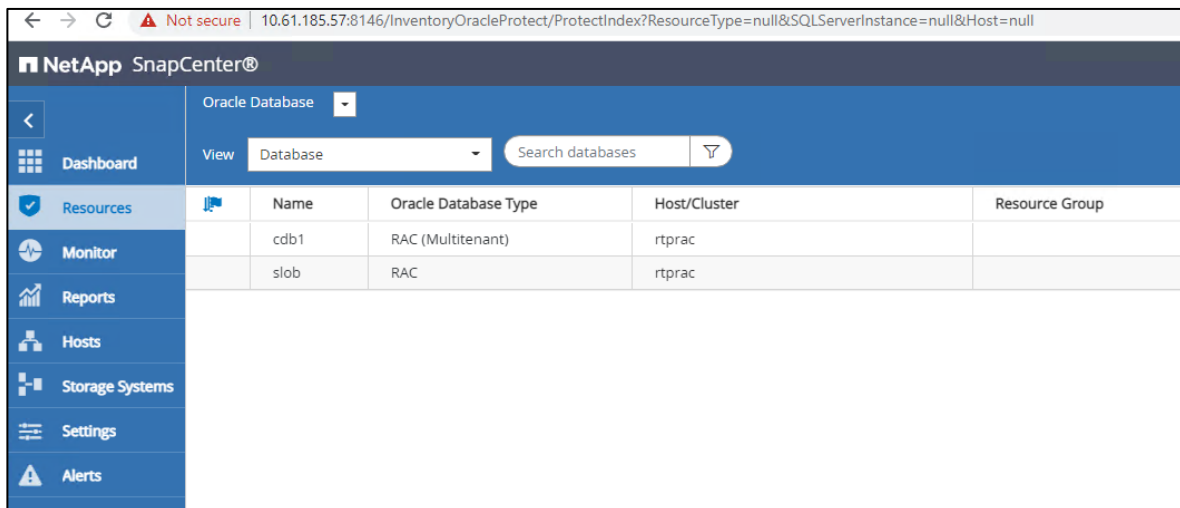
h. Click Confirm Others and Submit to trigger host registration and installation.



i. After successful installation, the Overall Status shows as Running.



j. As part of Oracle database plug-in installation process, the existing DBs are discovered and listed under the Resources tab:



3. Disable alias and user-friendly names in `multipath.conf`. During Oracle RAC setup host configuration, we use alias and user-friendly names in the SAN multipath configuration. Because SnapCenter does not support aliases with user-friendly names for the LUNs, complete the following steps to disable alias names and user-friendly naming in multipath configuration.

a. Log into the ASM instance as `sysasm` to check or set `asm_diskstring` to `/dev/mapper/*`.

```
SQL> show parameter asm_diskstring
```

NAME	TYPE	VALUE
------	------	-------


```
-----  
asm_diskstring          string          /dev/mapper/*  
-----
```

b. Shutdown the RAC cluster as the root user.

```
./crsctl stop cluster -all
```

c. Remove alias names in the multipath configuration file and set `user_friendly_names` to `no`.

```
[root@b200-ora-02 oracle]# cat /etc/multipath.conf  
defaults {  
    find_multipaths yes  
    user_friendly_names no  
}  
  
blacklist {  
}
```

d. Modify Oracle device udev rules as follows:

```
[oracle@b200-ora-02 ~]$ cat /etc/udev/rules.d/99-oracle-asmdevices.rules  
ENV{DM_NAME}=="3600a0980383145374b2451445133572f", SYMLINK+="ora_crs_01_1", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b24514451335761", SYMLINK+="ora_crs_01_2", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b24514451335763", SYMLINK+="ora_crs_02_1", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b24514451335762", SYMLINK+="ora_crs_02_2", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145373524514438424557", SYMLINK+="ora_data_01_1", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145373524514438424556", SYMLINK+="ora_data_01_2", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145373524514438424555", SYMLINK+="ora_data_01_3", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145373524514438424558", SYMLINK+="ora_data_01_4", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b2451445133574d", SYMLINK+="ora_data_02_1", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b2451445133574e", SYMLINK+="ora_data_02_2", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b2451445133574c", SYMLINK+="ora_data_02_3", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b2451445133574b", SYMLINK+="ora_data_02_4", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
ENV{DM_NAME}=="3600a0980383145374b24514451335750", SYMLINK+="ora_data_03_1", GROUP:="asmadmin",  
OWNER:="grid", MODE:="0660"  
...
```

e. Reload and trigger the udev rules and restart multipathd.

```
udevadm control --reload-rules  
udevadm trigger  
systemctl restart multipathd
```

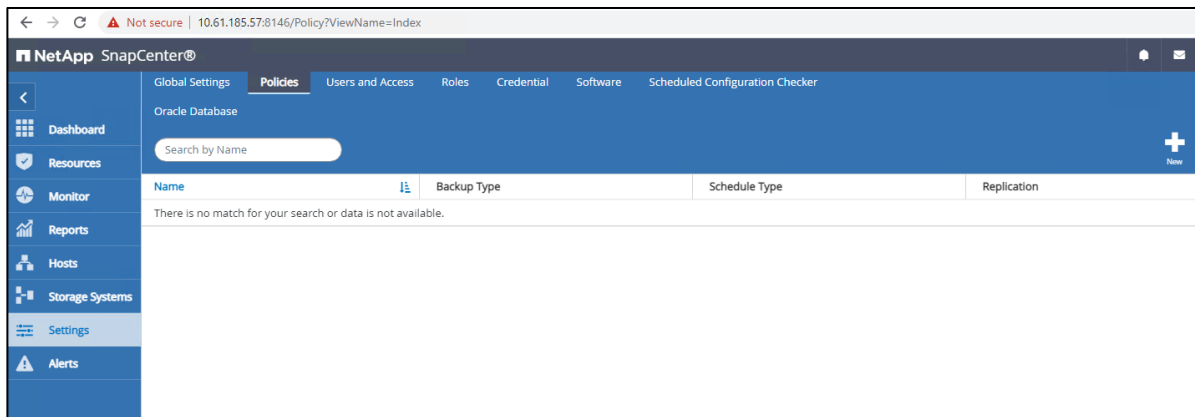
f. Restart the Oracle RAC cluster.

```
./crsctl start cluster -all
```

Back up Oracle Database with SnapCenter

1. Create a backup policy:

- a. From the Settings > Policies screen, click New to add a database backup policy.



b. Create the desired backup policy with a detailed description.

The screenshot shows the 'New Oracle Database Backup Policy' form. The sidebar on the left lists steps: 1 Name, 2 Backup Type, 3 Retention, 4 Replication, 5 Script, 6 Verification, and 7 Summary. The 'Name' step is active. The main form area has a section titled 'Provide a policy name' with two input fields: 'Policy name' (containing 'Daily Online Database Backup') and 'Details' (containing 'Daily backup database files and archive log files.'). At the bottom right, there are 'Previous' and 'Next' buttons.

c. Select the backup options and schedule the frequency. You can also choose whether the backup should be registered with RMAN catalog and archive log pruning options.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select Oracle database backup options

Choose backup type

☒ Online backup

☒ Datafiles, control files, and archive logs

☐ Datafiles and control files

☐ Archive logs

☐ Offline backup ⓘ

☒ Mount

☐ Shutdown

☐ Save state of PDBs ⓘ

Choose schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

☐ On demand

☐ Hourly

☒ Daily

☐ Weekly

☐ Monthly

Previous Next

- d. In the Retention page, specify the retention settings for the backup type and the schedule type selected in the Backup Type page. For ONTAP 9.7, the maximum number of NetApp Snapshot copies that can be retained is 1018.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Retention settings ⓘ

Daily retention settings

Data backup retention settings ⓘ

☒ Total Snapshot copies to keep 7

☐ Keep Snapshot copies for 14 days

Archive Log backup retention settings

☒ Total Snapshot copies to keep 7

☐ Keep Snapshot copies for 14 days

Previous Next

- e. Secondary replication for backup protection is highly recommended for database protection in the case of site failure. In this solution, we did not setup secondary site for the solution validation.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Select secondary replication options ⓘ

☐ Update SnapMirror after creating a local Snapshot copy.

☐ Update SnapVault after creating a local Snapshot copy.

Secondary policy label Choose ⓘ

Error retry count 3 ⓘ

Previous Next

- f. You have the option to run scripts before and after database backup. Otherwise, you can confirm the summary page to finish the database backup policy configuration.

New Oracle Database Backup Policy

1 Name

2 Backup Type

3 Retention

4 Replication

5 Script

6 Verification

7 Summary

Summary

Policy name	Daily Online Database Backup
Details	Daily backup database files and archive log files.
Backup type	Online backup
Schedule type	Daily
RMAN catalog backup	Enabled
Archive log pruning	None
On demand data backup retention	None
On demand archive log backup retention	None
Hourly data backup retention	None
Hourly archive log backup retention	None
Daily data backup retention	Total snapshot copies to retain : 7
Daily archive log backup retention	Total snapshot copies to retain : 7
Weekly data backup retention	None
Weekly archive log backup retention	None
Monthly data backup retention	None
Monthly archive log backup retention	None
Replication	None
Backup prescript settings	undefined
Prescript arguments:	
Backup postscript settings	undefined

Previous Finish

- g. Similarly, create an archive log backup policy if needed on an hourly basis to regularly purge archived logs for highly active databases to manage REDO space usage.

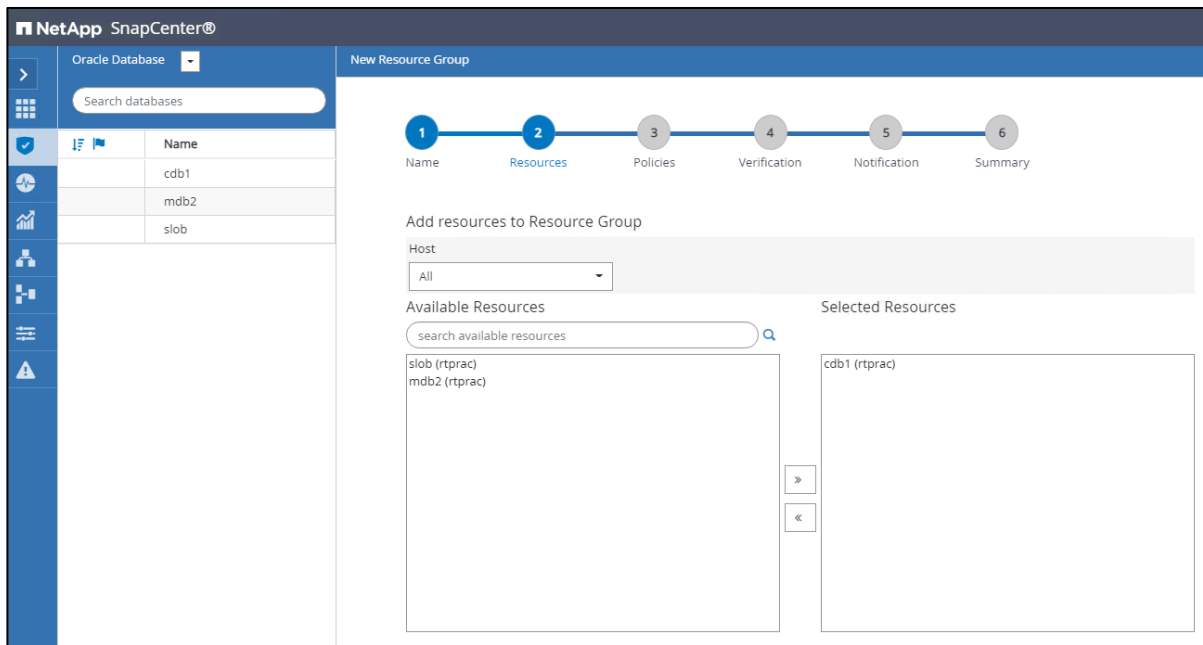
The screenshot shows the NetApp SnapCenter interface. The left sidebar contains navigation links: Dashboard, Resources, Monitor, Reports, Hosts, Storage Systems, Settings, and Alerts. The main content area is titled 'Oracle Database' and has a 'Policies' tab selected. Below the tab is a search bar labeled 'Search by Name'. A table lists the policies:

Name	Backup Type	Schedule Type
Archive log Backup	LOG, ONLINE	Hourly
Daily Online Database Backup	FULL, ONLINE	Daily

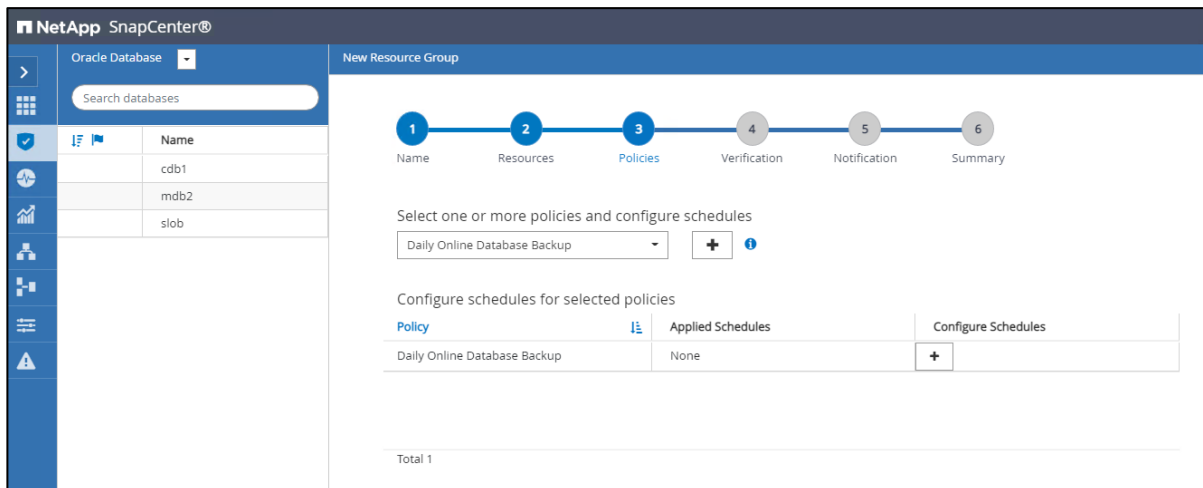
2. Create resource groups and attach policies for Oracle databases. A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously.

The screenshot shows the 'New Resource Group' wizard in NetApp SnapCenter. The left sidebar is the same as the previous screenshot. The main content area is titled 'New Resource Group' and shows a progress bar with six steps: 1. Name, 2. Resources, 3. Policies, 4. Verification, 5. Notification, and 6. Summary. Step 2, 'Resources', is currently active. Below the progress bar, there are input fields for 'Name' (CD81) and 'Tags' (CD81). A checkbox labeled 'Use custom name format for Snapshot copy' is checked, and a text field below it contains '\$ResourceGroup'. Under the 'Backup settings' section, there is a field for 'Exclude archive log destinations from backup' with a dropdown menu and a plus sign button.

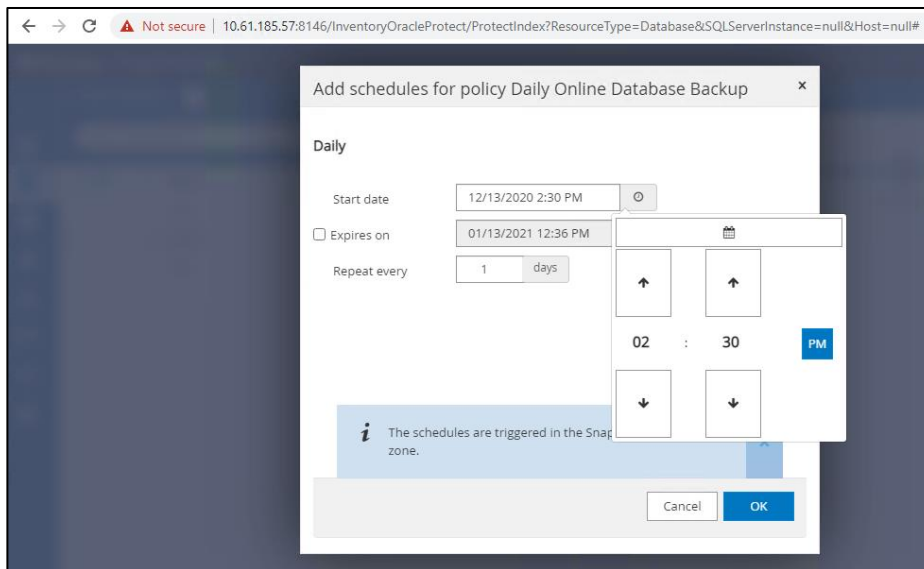
3. Add DBs to a resource group:



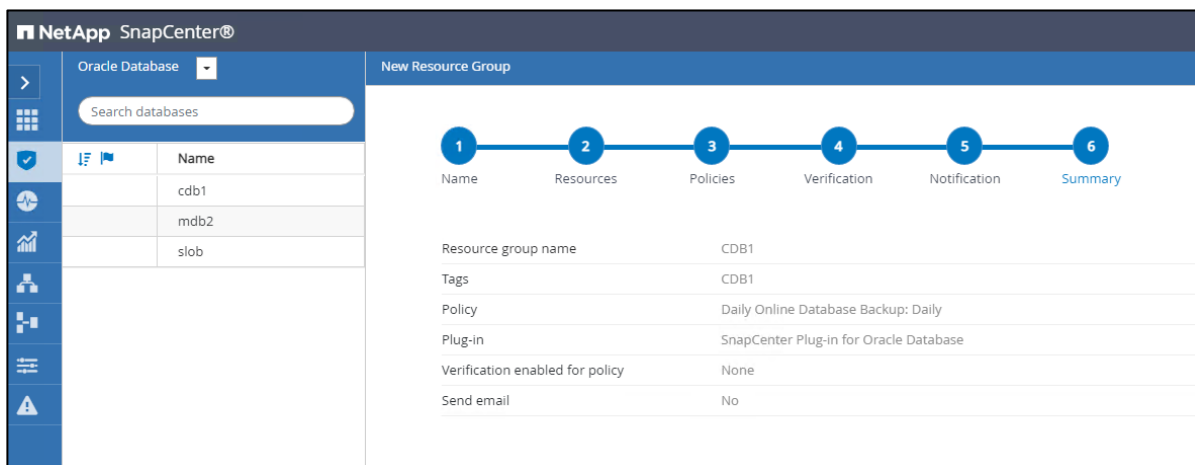
4. Select a backup policy for the resource group.



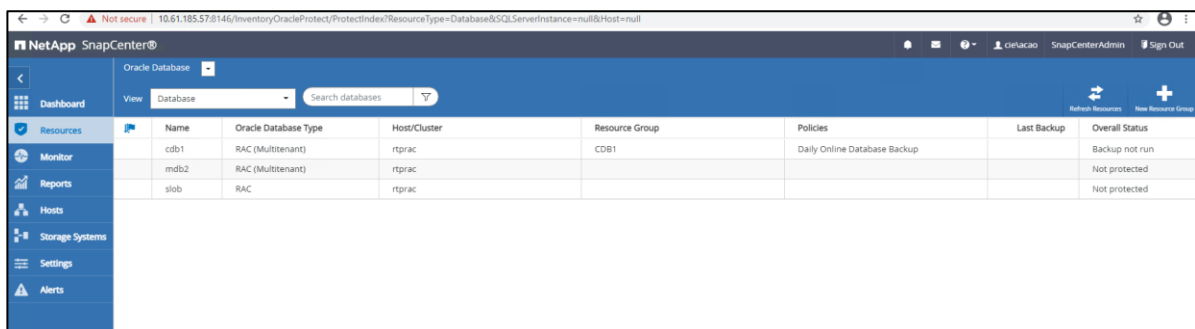
5. Set backup schedule timing.



- If you set verification for the policy and you configured email notification, it is populated in the Verification and Notification tab. Otherwise, the summary page shows the details of the new resource group.



- Click Finish to complete the addition of a new resource group. Now CDB1 is showing as a protected resource group, but backups have yet to run.



8. After a successful database back, you can check that snapshots have been created on RAC +DATA and +REDO disk groups volumes:

```
FlexPod-A800-01-02::> volume snapshot show ora_*
```

				---Blocks---
--				
Vserver	Volume	Snapshot	Size	Total% Used%
oral9c_svm	ora_data_01	CDB1_12-13-2020_14.30.32.0837_0	1.97MB	0% 0%
	ora_data_02	CDB1_12-13-2020_14.30.32.0837_0	1.85MB	0% 0%
	ora_data_03	CDB1_12-13-2020_14.30.32.0837_0	0B	0% 0%
	ora_data_04	CDB1_12-13-2020_14.30.32.0837_0	0B	0% 0%
	ora_data_05	CDB1_12-13-2020_14.30.32.0837_0	1.83MB	0% 0%
	ora_data_06	CDB1_12-13-2020_14.30.32.0837_0	0B	0% 0%
	ora_data_07	CDB1_12-13-2020_14.30.32.0837_0	1.88MB	0% 0%
	ora_data_08	CDB1_12-13-2020_14.30.32.0837_0	0B	0% 0%
	ora_data_09	CDB1_12-13-2020_14.30.32.0837_0	2.02MB	0% 0%
	ora_data_10	CDB1_12-13-2020_14.30.32.0837_0	1.90MB	0% 0%
	ora_data_11	CDB1_12-13-2020_14.30.32.0837_0	2.07MB	0% 0%
	ora_data_12	CDB1_12-13-2020_14.30.32.0837_0	1.98MB	0% 0%
	ora_redo_01	CDB1_12-13-2020_14.30.32.0837_1	0B	0% 0%
	ora_redo_02	CDB1_12-13-2020_14.30.32.0837_1	580KB	0% 0%
	ora_redo_03	CDB1_12-13-2020_14.30.32.0837_1	0B	0% 0%
	ora_redo_04	CDB1_12-13-2020_14.30.32.0837_1	580KB	0% 0%

16 entries were displayed.

9. Create an RMAN catalog and register databases to be protected in the RMAN catalog. Procedures for creating an RMAN catalog are beyond scope of this solution. See the Oracle recovery manager documentation for setting up an RMAN catalog for your Oracle environment. The following steps show how to register a database with an RMAN catalog.

```
[oracle@b200-ora-01 admin]$ rman target / catalog rman@slob

Recovery Manager: Release 19.0.0.0.0 - Production on Mon Dec 14 14:36:25 2020
Version 19.8.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.

connected to target database: CDB1 (DBID=1031832871)
recovery catalog database Password:
connected to recovery catalog database

RMAN> register database;

database registered in recovery catalog
starting full resync of recovery catalog
full resync complete
```

Note: The databases must first be registered with the RMAN catalog before configuring catalog settings in SnapCenter.

10. Configure databases in SnapCenter to register backups with an RMAN catalog.

We have created a credential for an RMAN catalog owner called `rman` to be used for connection to RMAN catalog via a TNS name `slob` on RAC node 1: `b200-ora-01`. The RMAN catalog is hosted in the database `slob`. This must be configured for all databases for your Oracle environment to be protected by SnapCenter with RMAN cataloging capability.

11. Validate SnapCenter backup copies registered with RMAN. After you configure database connections to the RMAN catalog in the database configuration setting, SnapCenter can now register data file backup copies with RMAN as an option. This can be easily turned on or off in your backup policy.

SnapCenter backups are registered as data file copies in the RMAN catalog as shown below:

```
[oracle@b200-ora-01 admin]$ rman target / catalog rman@slob
```

Recovery Manager: Release 19.0.0.0.0 - Production on Thu Dec 17 15:18:45 2020
Version 19.8.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.

connected to target database: CDB1 (DBID=1031832871)
recovery catalog database Password:
connected to recovery catalog database

RMAN> list copy;

List of Datafile Copies
=====

Key	File S	Completion Time	Ckp SCN	Ckp Time	Sparse
1210	1	A 17-DEC-20	27398594	17-DEC-20	NO
	Name: +DATA_CDB1_29/CDB1/DATAFILE/system.257.1058115327				
	Tag: SCO_CDB1_29				
1224	3	A 17-DEC-20	27398594	17-DEC-20	NO
	Name: +DATA_CDB1_29/CDB1/DATAFILE/sysaux.258.1058115363				
	Tag: SCO_CDB1_29				
1206	4	A 17-DEC-20	27398594	17-DEC-20	NO
	Name: +DATA_CDB1_29/CDB1/DATAFILE/undotbs1.259.1058115377				
	Tag: SCO_CDB1_29				
1211	5	A 17-DEC-20	2547452	02-DEC-20	NO
	Name: +DATA_CDB1_29/CDB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/system.265.1058115727				
	Tag: SCO_CDB1_29				
	Container ID: 2, PDB Name: PDB\$SEED				
1216	6	A 17-DEC-20	2547452	02-DEC-20	NO
	Name: +DATA_CDB1_29/CDB1/86B637B62FE07A65E053F706E80A27CA/DATAFILE/sysaux.266.1058115727				
	Tag: SCO_CDB1_29				
	Container ID: 2, PDB Name: PDB\$SEED				
1197	7	A 17-DEC-20	27398594	17-DEC-20	NO
	Name: +DATA_CDB1_29/CDB1/DATAFILE/users.260.1058115379				
	Tag: SCO_CDB1_29				
...					

Oracle Database restore and recovery - PDB PITR

With SnapCenter 4.4, you can restore and recover individual PDBs in a CDB (container database) to any desired point in time. We demonstrate a PDB PITR or PDB point-in-time recovery in the following exercise.

1. Create a test table in a PDB.

To test the SnapCenter PDB point-in-time recovery functionality, we created a test table called `restore_check` in container database CDB1 and PDB `cdb1_pdb1` and inserted a row into the table.

```
SQL> conn pdbadmin@cdb1_pdb1
Enter password:
Connected.
SQL> show con_name

CON_NAME
-----
CDB1_PDB1
SQL> create table restore_check(
event varchar(100),
dt timestamp default sysdate);

Table created.
insert into restore_check values('test SnapCenter PDB PITR', sysdate)
SQL> /
```

```

1 row created.
commit;
SQL> select * from restore_check;
EVENT
-----

```

```

DT
-----

```

```

test SnapCenter PDB PITR
16-DEC-20 03.32.19.000000 PM

```

2. Drop PDB cdb1_pdb1 after a backup. We performed an online backup for the container database, and then dropped the PDB cdb1_pdb1.

```

SQL> alter pluggable database cdb1_pdb1 close immediate instances=all;
Pluggable database altered.
SQL> drop pluggable database cdb1_pdb1 including datafiles;
Pluggable database dropped.

```

3. Recover the dropped PDB and test table by recovering PDB to a point in time. Now we try to recover the table by restoring PDB cdb1_pdb1 to its point in time before the table was dropped.
 - a. From SnapCenter in the Resources page, select either Database or Resource Group from the View list. Select the database from either the database details view, or the resource group details view. The database topology page is displayed.

The screenshot shows the NetApp SnapCenter interface. On the left is a navigation sidebar with icons for search, databases, topology, performance, alerts, and a warning icon. The main content area is titled 'cdb1 Topology'. It includes a 'Search databases' input field and a table with columns 'Name' and 'Status'. The table lists 'cdb1', 'mdb2', and 'slob'. Below this is a 'Manage Copies' section showing '18 Backups' and '0 Clones'. A 'Primary Backup(s)' section contains a search input and a table with columns: Backup Name, Count, Type, and End Date. The table lists several backup entries with their respective counts, types (Log or Data), and end dates.

Backup Name	Count	Type	End Date
CDB1_12-17-2020_14.30.31.8735_1	1	Log	12/17/2020 2:31:33 PM
CDB1_12-17-2020_14.30.31.8735_0	1	Data	12/17/2020 2:31:11 PM
RTPRAC_12-17-2020_13.52.48.9196_1	1	Log	12/17/2020 1:53:55 PM
RTPRAC_12-17-2020_13.52.48.9196_0	1	Data	12/17/2020 1:53:30 PM
RTPRAC_12-16-2020_17.47.27.2139_1	1	Log	12/16/2020 5:48:28 PM
RTPRAC_12-16-2020_17.47.27.2139_0	1	Data	12/16/2020 5:48:06 PM

- b. From the Manage Copies view, select Backups from either the primary or the secondary (mirrored or replicated) storage systems.
We used the backup copy RTPRAC_12-17-2020_13.52.48.9196_0 to recover the PDB cdb1_pdb1, which was backup copy before table drop.
 - c. Mount the archive log backup that was backed up with the backup job.

NetApp SnapCenter®

Oracle Database | cdb1 Topology

Search databases

Manage Copies

16 Backups
0 Clones
Local copies

Summary Card

18 Backups
9 Data Backups
9 Log Backups
0 Clones

Primary Backup(s)

search

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
CDB1_12-17-2020_14.30.31.8735_1	1	Log		12/17/2020 2:31:33 PM	Not Applicable	False	Not Cataloged	27472978
CDB1_12-17-2020_14.30.31.8735_0	1	Data		12/17/2020 2:31:11 PM	Unverified	False	Not Cataloged	27472878
RTPRAC_12-17-2020_13.52.48.9196_1	1	Log		12/17/2020 1:53:55 PM	Not Applicable	False	Cataloged	27399225
RTPRAC_12-17-2020_13.52.48.9196_0	1	Data		12/17/2020 1:53:30 PM	Unverified	False	Cataloged	27398631

Mount backups

Choose the host to mount the backup: b200-ora-01.cie.netapp.com

Mount path: +REDO_cdb1_31

ASM instance: None

Credential name: +

ASM Port: 1521

Mount Cancel

After the archive log backup is mounted, it shows the Mounted status of True as shown below:

NetApp SnapCenter®

Oracle Database | cdb1 Topology

Search databases

Manage Copies

18 Backups
0 Clones
Local copies

Summary Card

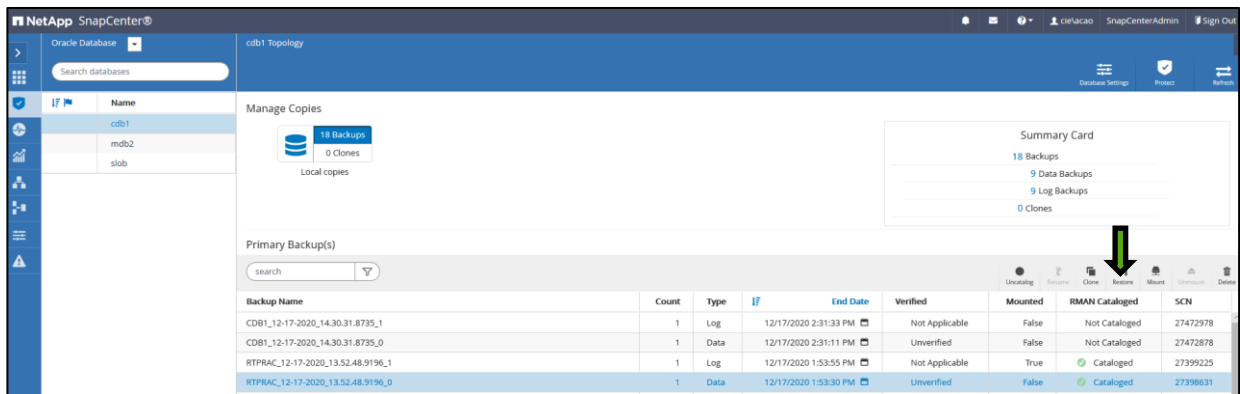
18 Backups
9 Data Backups
9 Log Backups
0 Clones

Primary Backup(s)

search

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
CDB1_12-17-2020_14.30.31.8735_1	1	Log		12/17/2020 2:31:33 PM	Not Applicable	False	Not Cataloged	27472978
CDB1_12-17-2020_14.30.31.8735_0	1	Data		12/17/2020 2:31:11 PM	Unverified	False	Not Cataloged	27472878
RTPRAC_12-17-2020_13.52.48.9196_1	1	Log		12/17/2020 1:53:55 PM	Not Applicable	True	Cataloged	27399225
RTPRAC_12-17-2020_13.52.48.9196_0	1	Data		12/17/2020 1:53:30 PM	Unverified	False	Cataloged	27398631

- d. Select the backup from the table, and then click Restore.



e. Choose the RAC node for the restore and specify the restore scope.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select RAC Node

b200-ora-01.cie.netapp.com

Restore Scope

All Datafiles

Pluggable databases (PDBs)

CDB1_PDB1

Pluggable database(PDB) tablespaces

Control files

Database State

Change database state if needed for restore and recovery

Restore Mode

Force in place restore

If this check box is not selected and if any of the in place restore criteria is not met, restore will be performed using the connect and copy method. The connect and copy restore method might take time based on the files being restored.

Previous

Next

- Check the change database state if needed for restore and recovery. Check Force In Place Restore if you want to perform in-place restore in the scenarios where new datafiles are added after backup or when LUNs are added, deleted, or recreated to a disk group.
- Set the recovery scope.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Choose Recovery Scope

☐ All Logs
 ☒ Until SCN (System Change Number)

SCN

27398631

Auxiliary destination

+DATA

☐ Date and Time
 ☐ No recovery

Specify external archive log files locations

Previous

Next

- h. Specify the SCN when backup is completed for the desired SCN for point-in-time recovery. Log backup needs to be mounted and specified in external archive log files locations.
4. PreOps, PostOps, and Notification:
 - a. You can bypass the PreOps, PostOps, and notification if you do not intend to run any scripts before or after you receive notification about the restore. You can also choose to open the database after restore and recovery.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run after performing a restore job

Postscript full path

/var/opt/snapcenter/spl/scripts/Enter Postscript path

Arguments

☒ Open the pluggable databases in READ-WRITE mode after recovery.
 If required, the container database and the associated pluggable database will be opened in READ-WRITE mode to bring the restored and recovered tablespaces ONLINE.

Previous

Next

b. Restore summary.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Summary

Restore node	b200-ora-01.cie.netapp.com
Backup name	RTPRAC_12-17-2020_13.52.48.9196_0
Backup date	12/17/2020 1:53:30 PM
Restore scope	1PDB Selected
Recovery scope	Until SCN -27398631
Auxiliary destination	+DATA
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

Previous

Finish

- c. Click Finish to start the recovery.
- d. Validate the restore and recovery of PDB cdb1_pdb1.
- e. Check the status of the restore job for completion.

The screenshot shows the NetApp SnapCenter web interface. The 'Jobs' tab is selected, displaying a table of restore jobs. The job 'Restore 'tpraccdb1'' is shown with a status of 'Success' and a green checkmark. The start date is 12/17/2020 7:28:44 PM and the end date is 12/17/2020 7:54:45 PM.

ID	Status	Name	Start date	End date
83	✓	Restore 'tpraccdb1'	12/17/2020 7:28:44 PM	12/17/2020 7:54:45 PM

- f. Log into PDB cdb1_pdb1

```
[oracle@b200-ora-01 ~]$ sqlplus pdbadmin@//b200-ora-01/cdb1_pdb1.cie.netapp.com
SQL> show con_name
CON_NAME
-----
CDB1_PDB1
SQL> show user
USER is "PDBADMIN"
SQL> select sysdate from dual;
SYSDATE
-----
21-DEC-20
SQL> select table_name from user_tables;
TABLE_NAME
-----
RESTORE_CHECK
SQL> select table_name from user_tables;
TABLE_NAME
-----
RESTORE_CHECK
SQL> select * from RESTORE_CHECK;
EVENT
-----
DT
-----
test SnapCenter PDB PITR
16-DEC-20 03.32.19.000000 PM
```

This validates that the dropped PDB and test table has been recovered to its point-in-time image.

Why this matters

In a consolidated Oracle applications environment, the ability to restore an individual PDB to a point-in-time image allows you to rollback an application to a previous state without affecting other applications or tenants within the same container database. When a few applications or tenants are consolidated into a container database, the flexibility to restore and recover individual pluggable databases and their associated applications might be the mandate and service-level requirements for workload consolidation.

PDB tablespace PITR

An Oracle tablespace is a subset of data within a pluggable database in a container database. With SnapCenter 4.4, you can now restore and recover a tablespace in a pluggable database to any desired point in time using options like date/time or the latest system change number (SCN). The following demonstration shows how Oracle tablespace point-in-time recovery works in a pluggable database with SnapCenter.

1. Create a new tablespace and a test table in PDB cdb1_pdb2.

```
SQL> sho con_name
CON_NAME
-----
CDB1_PDB2
```

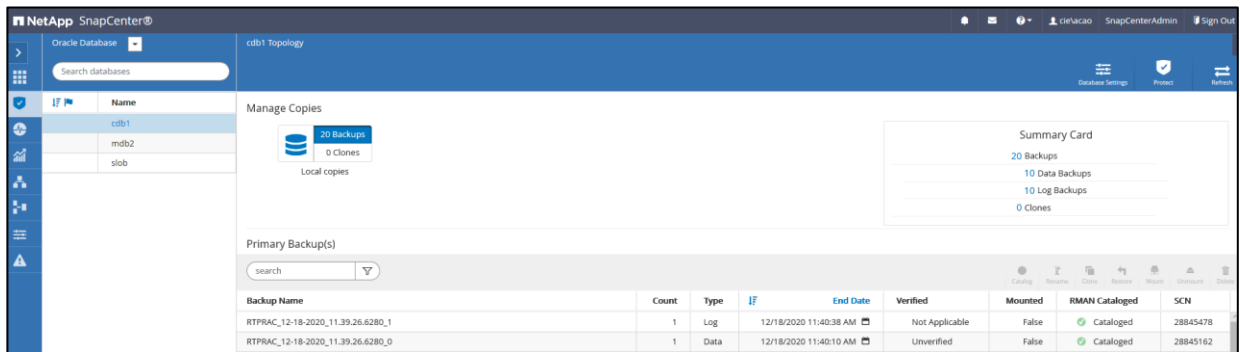


```

SQL> show user
USER is "PDBADMIN"
SQL> create tablespace hr;
Tablespace created.
SQL> create table hr_test(event varchar(100), dt timestamp default sysdate) tablespace hr;
Table created.
SQL> insert into hr_test values('pdb tablespace Pitr test', sysdate);
1 row created.
SQL> select * from hr_test;
EVENT
-----
DT
-----

pdb tablespace Pitr test
18-DEC-20 11.28.02.000000 AM
SQL> select table_name, tablespace_name from user_tables;
TABLE_NAME                                TABLESPACE_NAME
-----
RESTORE_CHECK_CDB1_PDB2                  USERS
HR_TEST                                  HR
Run database backup to backup cdb1

```



2. Drop table hr_test after backup.

```

SQL> drop table HR_TEST;
Table dropped.
SQL> select table_name, tablespace_name from user_tables;
TABLE_NAME                                TABLESPACE_NAME
-----
RESTORE_CHECK_CDB1_PDB2                  USERS
SQL> commit;
Commit complete.

```

3. Mount the backup.

Mount backups

Choose the host to mount the backup
b200-ora-01.cie.netapp.com

Mount path :
+DATA_cdb1_37

ASM instance
None

Credential name

ASM Port
1521

Mount Cancel

4. Restore the hr tablespace to the point in time before the table drop to the recovery table `hr_test`.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select RAC Node **b200-ora-01.cle.netapp.com**

Restore Scope

☐ All Datafiles

☐ Pluggable databases (PDBs)

☒ Pluggable database(PDB) tablespaces

CDB1_PDB2

HR

☐ Control files

Database State

☒ Change database state if needed for restore and recovery

Restore Mode

☐ Force in place restore

If this check box is not selected and if any of the in place restore criteria is not met, restore will be performed using the connect and copy method. The connect and copy restore method might take time based on the files being restored.

Previous

Next

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Choose Recovery Scope

☐ All Logs

☒ Until SCN (System Change Number)

SCN

28845162

Auxiliary destination

+DATA

☐ Date and Time

☐ No recovery

Specify external archive log files locations

Previous

Next

Restore cdb1 x

1 Restore Scope
2 Recovery Scope
3 PreOps
4 PostOps
5 Notification
6 Summary

Summary

Restore node	b200-ora-01.cie.netapp.com
Backup name	RTPRAC_12-18-2020_11.39.26.6280_0
Backup date	12/18/2020 11:40:10 AM
Restore scope	CDB1_PDB2PDB selected with1 Tablespaces
Recovery scope	Until SCN -28845162
Auxiliary destination	+DATA
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

Previous
Finish

5. Validate that the hr_test table was recovered.

During restore, SnapCenter uses Oracle RMAN to restore the HR tablespace to the point in time. After a successful restore, we validated that the dropped hr_test table was recovered.

```

Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.8.0.0.0
SQL> show con_name
CON_NAME
-----
CDB1_PDB2
SQL> show user
USER is "PDBADMIN"
SQL> col name form a50
SQL> col table_name form a50
SQL> set lin 200
SQL> select table_name, tablespace_name from user_tables;
TABLE_NAME                                TABLESPACE_NAME
-----
RESTORE_CHECK_CDB1_PDB2                   USERS
HR_TEST                                   HR

```

Why this matters

As a workload consolidation option, multiple applications can be consolidated into a pluggable database and segregated into different Oracle tablespaces. The ability to restore a tablespace to a point in time enables rolling back an application to a previous point in time without affecting other applications hosted in the same pluggable database.

The ability to restore a tablespace in a PDB to a point in time also allows user errors or logical application data corruption recovery in place without affecting any other unassociated users or applications.

CDB restore/recovery - CDB recovery to last log

1. Insert a new row to the restore_check table in cdb1_pdb1.

```
SQL> show con_name
CON_NAME
-----
CDB1_PDB1
SQL> select table_name from user_tables;
TABLE_NAME
-----
RESTORE_CHECK
SQL> select * from restore_check;
EVENT
-----
DT
-----
test SnapCenter PDB PITR
16-DEC-20 03.32.19.000000 PM
SQL> insert into restore_check values('test CDB recovery to last log', sysdate);
1 row created.
SQL> commit;
Commit complete.

SQL> select * from restore_check;
EVENT
-----
DT
-----
test SnapCenter PDB PITR
16-DEC-20 03.32.19.000000 PM
test CDB recovery to last log
04-JAN-21 11.30.04.000000 AM
```

2. Restore CDB1 to a previous day backup copy on 1/3/2021.

The screenshot shows the NetApp SnapCenter interface for an Oracle Database. The left sidebar contains navigation icons. The main area displays the 'cdb1 Topology' with a table of databases:

IF	Name
<input checked="" type="checkbox"/>	cdb1
<input type="checkbox"/>	mdb2
<input type="checkbox"/>	slsb

Below the table, it shows 'Manage Copies' with '56 Backups' and '1 Clone' for 'Local copies'. A 'Summary Card' on the right provides a quick overview: 56 Backups, 28 Data Backups, 28 Log Backups, and 1 Clone.

The 'Primary Backup(s)' section contains a table of backups:

Backup Name	Count	Type	IF	End Date	Verified	Mounted	RMAN Cataloged	SCN
CDB1_01-03-2021_18.30.20.1406_1	1	Log	<input checked="" type="checkbox"/>	01/03/2021 6:31:24 PM	Not Applicable	False	<input checked="" type="checkbox"/> Cataloged	51475788
CDB1_01-03-2021_18.30.20.1406_0	1	Data	<input checked="" type="checkbox"/>	01/03/2021 6:30:57 PM	Unverified	False	<input checked="" type="checkbox"/> Cataloged	51475529
CDB1_01-02-2021_18.30.19.6269_1	1	Log	<input checked="" type="checkbox"/>	01/02/2021 6:31:27 PM	Not Applicable	False	<input checked="" type="checkbox"/> Cataloged	50095237
CDB1_01-02-2021_18.30.19.6269_0	1	Data	<input checked="" type="checkbox"/>	01/02/2021 6:30:58 PM	Unverified	False	<input checked="" type="checkbox"/> Cataloged	50093135
CDB1_01-01-2021_18.30.29.2633_1	1	Log	<input checked="" type="checkbox"/>	01/01/2021 6:31:33 PM	Not Applicable	False	<input checked="" type="checkbox"/> Cataloged	48759106

3. CDB1 restore scope.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Select RAC Node

b200-ora-01.cie.netapp.com

Restore Scope

All Datafiles

Pluggable databases (PDBs)

Pluggable database(PDB) tablespaces

Control files

Database State

Change database state if needed for restore and recovery

Restore Mode

Force in place restore

If this check box is not selected and if any of the in place restore criteria is not met, restore will be performed using the connect and copy method. The connect and copy restore method might take time based on the files being restored.

Previous

Next

4. CDB1 recovery scope.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Choose Recovery Scope

All Logs

Until SCN (System Change Number)

Date and Time

No recovery

Specify external archive log files locations

Previous

Next

5. PostOps to open restored database.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify optional scripts to run after performing a restore job ⓘ

Postscript full path Enter Postscript path

Arguments

☒ Open the database or container database in READ-WRITE mode after recovery

Previous

Next

6. Summary.

Restore cdb1

1 Restore Scope

2 Recovery Scope

3 PreOps

4 PostOps

5 Notification

6 Summary

Summary

Restore node	b200-ora-01.cie.netapp.com
Backup name	CDB1_01-03-2021_18.30.20.1406_0
Backup date	01/03/2021 6:30:57 PM
Restore scope	All DataFiles
Recovery scope	All Logs
Options	Change database state if necessary , Open the database or container database in READ-WRITE mode after recovery
Prescript full path	None
Prescript arguments	
Postscript full path	None
Postscript arguments	
Send email	No

Previous

Finish

7. Restore job ID 203 execution.



ID	Status	Name	Start date	End date	Owner
203	Success	Restore vtpack/cdb1	01/04/2021 12:01:53 PM		CIElacao
202	Success	Storage savings computation	01/04/2021 12:00:01 AM	01/04/2021 12:00:08 AM	CIElacao
201	Success	Validate license	01/03/2021 11:59:25 PM	01/03/2021 11:59:26 PM	CIElacao
200	Success	Configuration Check for the SnapCenterServer 'AB-Labs.cie.netapp.com'	01/03/2021 11:59:13 PM	01/03/2021 11:59:42 PM	CIElacao
197	Success	Backup of Resource Group 'CDB1' with policy 'Daily Online Database Backup'	01/03/2021 6:30:03 PM	01/03/2021 7:34:40 PM	CIElacao

8. CDB1 restore and recovery check.

CDB1 was restored using an online backup copy on 1/3/2021, and then CDB1 was rolled forward to the last available log, which would recover the new row we just inserted on 1/4/2021.

After completion of restore and recovery, CDB1 was left in the Mount state. We manually opened CDB1 and validated that the new row in the test table was indeed recovered after logs from 1/4/2021 were applied after recovery.

```
SQL> show pdbs
  CON_ID CON_NAME                                OPEN MODE RESTRICTED
-----
      2 PDB$SEED                                READ ONLY NO
      3 CDB1_PDB1                                MOUNTED
      4 CDB1_PDB2                                READ WRITE NO
SQL> alter pluggable database cdb1_pdb1 open instances=all;
Pluggable database altered.
SQL> show con_name

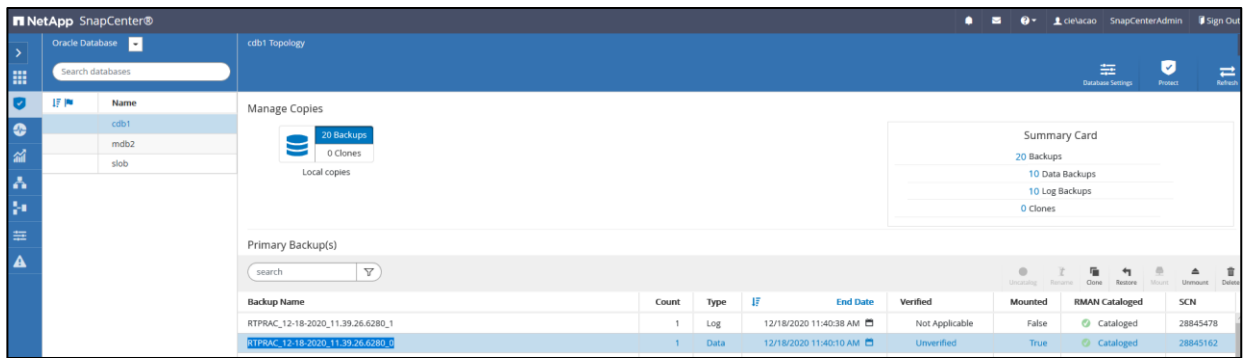
CON_NAME
-----
CDB1_PDB1
SQL> select table_name from user_tables;
TABLE_NAME
-----
RESTORE_CHECK
SQL> select * from restore_check;
EVENT
-----
DT
-----
test SnapCenter PDB PITR
16-DEC-20 03.32.19.000000 PM
test CDB recovery to last log
04-JAN-21 11.30.04.000000 AM
```

PDB clone

SnapCenter supports cloning of different types of backups of Oracle databases. This includes online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration.

In the following exercise, we clone a PDB from one CDB to another from an online full backup.

1. Identify the backup for the clone. We clone off a full online backup.



2. Click Clone to proceed.
3. Chose the clone option and name the cloned PDB.

Clone from cdb1

1 Name

2 Locations

3 PreOps

4 PostOps

5 Notification

6 Summary

☐ Complete Database Clone

☒ PDB Clone

PDB

CDB1_PDB2 x

PDB Clone Name

clone_cdb1_p2

Previous

Next

4. Chose the clone host as well as the target CDB. We are cloning cdb1_pdb2 to CDB mdb2.

Clone from cdb1

1 Name

2 Locations

3 PreOps

4 PostOps

5 Notification

6 Summary

Select the host to create a clone

Clone host

b200-ora-01.cie.netapp.com

Target CDB

mdb2

Database State

☒ Open the PDB to be cloned in READ-WRITE mode.

Datafile locations

+SC_2090922_SCJOBID

Reset

Oracle Home Settings

Oracle Home

/u01/app/oracle/product/19300/ntap

Oracle OS User

oracle

Oracle OS Group

oinstall

Previous

Next

5. PreOps.

Clone from cdb1

1 Name

2 Locations

3 PreOps

4 PostOps

5 Notification

6 Summary

Specify scripts to run before clone operation

Prescript full path

/var/opt/snapcenter/spl/scripts/

Enter Prescript path

Arguments

Script timeout

60

secs

Auxiliary CDB clone database parameter settings

Previous

Next

6. After Ops.

Clone from cdb1

1 Name

2 Locations

3 PreOps

4 PostOps

5 Notification

6 Summary

PostOps for Auxiliary CDB Clone

Recover Database

☐ Until Cancel

☐ Date and Time

1

Date-time format: MM/DD/YYYY hh:mm:ss

☒ Until SCN (System Change Number)

1

Specify external archive log locations

☒ Create new DBID

1

☐ Enter SQL queries to apply when clone is created

1

☐ Enter scripts to run after clone operation

1

Previous

Next

7. Clone summary.

Clone from cdb1

1 Name

2 Locations

3 PreOps

4 PostOps

5 Notification

6 Summary

Summary

Clone from backup	RTPRAC_12-18-2020_11.39.26.6280_0
PDB	CDB1_PDB2
PDB Clone Name	clonecdb1p2
Clone server	b200-ora-01.cie.netapp.com
Target CDB	mdb2
Options	Open the PDB to be cloned in READ-WRITE mode.
Oracle home	/u01/app/oracle/product/19300/ntap
Oracle OS user	oracle
Oracle OS group	oinstall
Datafile mountpaths	+SC_2090922_SCJOBID
Recovery scope	Until SCN -28845162
Prescript full path	none

Previous

Finish

8. Validate the clone.

In this clone validation process, we have created a cloned PDB called clonecdb1p2 in the container database mdb2 from a backup copy of PDB cdb1_pdb2 in the container database cdb1. We can log into the cloned PDB clonecdb1p2 in CDB mdb2 and see the two tables we have created for testing tablespace Pitr in the cloned PDB under user pdbadmin.

```
SQL> select name, cdb, con_id from v$databases;
NAME          CDB      CON_ID
-----
MDB2          YES        0
SQL> show pdbs
CON_ID CON_NAME                                OPEN MODE  RESTRICTED
-----
      2 PDB$SEED                                READ ONLY  NO
      3 MDB2_PDB1                                MOUNTED
      4 CLONECDB1P2                             READ WRITE NO
SQL> alter session set container=CLONECDB1P2;
Session altered.
SQL> select table_name from dba_tables where owner='PDBADMIN';
TABLE_NAME
-----
RESTORE_CHECK_CDB1_PDB2
HR_TEST
SQL> select * from pdbadmin.HR_TEST;
EVENT
-----
DT
-----

pdb tablespace Pitr test
18-DEC-20 11.28.02.000000 AM
```

9. Clone refresh.

A cloned database can be refreshed using two methods. First, drop the clone and repeat the same procedure from 1) through 7) using a newer backup copy. Second, from command line, the existing clone refresh can be triggered using the following command at an Oracle host:

```
sccli Refresh-SmClone -OracleCloneSpecificationFile
'/var/opt/snapcenter/sco/clone_specs/oracle_clonespec_dgp_CLSID1_2018-06-04_02.35.05.605.xml'
-PolicyName polSecondary -CloneDatabaseSID sid2
```

SID2 is the cloned database SID. With the command, SnapCenter drops the existing clone, makes a backup of the source database, and creates a new clone at the target from the backup.

Why this matters

The ability to quickly create a clone of a production database serves many purposes. First, the cloned database can be used for a dev/test data refresh for an application coding test against a full production data copy. It could also be used for data recovery in the case of data corruption or a logical data error, or it could be used for Oracle patching or to upgrade before production deployment. The cloned copy uses minimal storage space and can be created very quickly.

Monitor operations or generate reports

1. Monitor activities within SnapCenter.

ID	Status	Name	Start date	End date	Owner
75		Backup of Resource Group 'RTPRAC' with policy 'One Time Hot Backup'	12/17/2020 1:52:32 PM		CIElacao
74		Storage savings computation	12/17/2020 12:00:01 AM	12/17/2020 12:00:03 AM	CIElacao
73		Validate license	12/16/2020 11:59:50 PM	12/16/2020 11:59:51 PM	CIElacao
69		Backup of Resource Group 'RTPRAC' with policy 'One Time Hot Backup'	12/16/2020 5:47:12 PM	12/16/2020 6:19:33 PM	CIElacao
63		Backup of Resource Group 'RTPRAC' with policy 'One Time Hot Backup'	12/16/2020 3:39:07 PM	12/16/2020 4:20:54 PM	CIElacao
62		Update Policy 'One Time Hot Backup'	12/16/2020 3:27:02 PM	12/16/2020 3:27:03 PM	CIElacao
61		Backup of Resource Group 'CDB1' with policy 'Daily Online Database Backup'	12/16/2020 2:30:03 PM	12/16/2020 2:31:33 PM	CIElacao
60		Storage savings computation	12/16/2020 12:00:01 AM	12/16/2020 12:00:03 AM	CIElacao
59		Validate license	12/15/2020 11:59:52 PM	12/15/2020 11:59:53 PM	CIElacao
58		Backup of Resource Group 'CDB1' with policy 'Daily Online Database Backup'	12/15/2020 2:30:02 PM	12/15/2020 2:31:27 PM	CIElacao
56		Cataloging Backups(RTPRAC_12-13-2020_17.51.37.3685_0	12/15/2020 9:25:44 AM	12/15/2020 9:25:03 AM	CIElacao
55		Storage savings computation	12/15/2020 12:00:01 AM	12/15/2020 12:00:03 AM	CIElacao
54		Validate license	12/14/2020 11:59:53 PM	12/14/2020 11:59:54 PM	CIElacao

For each job, there is option to view the details, generate job specific reports, download logs for the job, cancel a running job if needed.

- View a job log while the job is running. Click the job ID, details, and view logs to review the detailed job log. The log can be downloaded if needed for trouble shooting.

Source	Log Level	Message
spl_98_20201218_0359.log	INFO	2020-12-18T16:08:53.0000596-05:00 INFO qtp1761307678-24664 c.n.s.l.r.SMCoreResourceManager - plugin Name () sco operation name GetLogs
spl_98_20201218_0359.log	INFO	2020-12-18T16:08:53.0000596-05:00 INFO qtp1761307678-24664 c.n.s.l.r.SMCoreResourceManager - Incoming input type com.netapp.smcore.plugin.contracts.SmGetLogRequest
spl_98_20201218_0359.log	INFO	2020-12-18T16:08:53.0000596-05:00 INFO qtp1761307678-24664 c.n.s.l.r.SMCoreResourceManager - Work flow context () Context collection parameters ()
spl_98_20201218_0359.log	INFO	2020-12-18T16:08:53.0000596-05:00 INFO qtp1761307678-24664 c.n.s.l.r.SMCoreResourceManager - Plugin Execution metadata
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:51.0435639-05:00 INFO SCUPlugin: SCU:Operation:RescanDevices:run(): Running Rescan of devices on the host
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:51.0000141-05:00 INFO qtp1761307678-24910 c.n.p.s.m.RescanDevicesPluginCoreMapper - SCU-LOG-00502: Completed convertRequest for RescanDevices operation.
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:51.0000141-05:00 INFO qtp1761307678-24910 c.n.p.s.m.RescanDevicesPluginCoreMapper - SCU-LOG-00501: Started convertRequest for RescanDevices operation.
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:51.0000142-05:00 INFO qtp1761307678-24910 c.n.p.s.p.PerfExecutor - SCU-LOG-00519: Executing command : /opt/NetApp/snapcenter/spl/lib/.plugins/scu/scuore/bin/scu-perf-client --infile=/var/opt/snapcenter/scu/json/RescanDevicesGlobalInput_98_0.json --outfile=/var/opt/snapcenter/scu/json/RescanDevicesGlobalOutput_98_0.json
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:50.0000929-05:00 INFO qtp1761307678-24910 c.n.s.l.r.SMCoreResourceManager - Plugin Execution metadata createClonecom.netapp.plugin.scu.operation.CloneOperation
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:50.0000929-05:00 INFO qtp1761307678-24910 c.n.s.l.r.SMCoreResourceManager - Incoming input type com.netapp.smcore.plugin.contracts.SmCreateCloneRequest
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:50.0000929-05:00 INFO qtp1761307678-24910 c.n.s.l.r.SMCoreResourceManager - Work flow context (DataCollectionEmmsEnabled=True) Context collection parameters ()
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:50.0000929-05:00 INFO qtp1761307678-24910 c.n.s.l.r.SMCoreResourceManager - plugin Name () sco operation name CreateClone
spl_98_20201218_0359.log	INFO	2020-12-18T16:00:50.0000929-05:00 INFO qtp1761307678-24910 c.n.p.s.o.CreateClone - SCU-LOG-00504: SCU CreateClone operation started.
SnapManagerWeb_98_20201218_0359.log	INFO	2020-12-18T16:00:19.3781124-05:00 INFO SnapManagerWeb_98 PID=[9152] TID=[71] Exit JobManagerProvider: UpdateJobStatus
SnapManagerWeb_98_20201218_0359.log	INFO	2020-12-18T16:00:19.3771109-05:00 INFO SnapManagerWeb_98 PID=[9152] TID=[71] Exit UpdateJobStatus
SnapManagerWeb_98_20201218_0359.log	INFO	2020-12-18T16:00:19.3691136-05:00 INFO SnapManagerWeb_98 PID=[9152] TID=[71] Enter UpdateJobStatus
SnapManagerWeb_98_20201218_0359.log	INFO	2020-12-18T16:00:19.3631129-05:00 INFO SnapManagerWeb_98 PID=[9152] TID=[71] Enter JobManagerProvider: UpdateJobStatus
SMCore_98_20201218_0359.log	INFO	2020-12-18T16:00:13.1661929-05:00 INFO SMCore_98 PID=[4660] TID=[30] Activity - FileSystem Clone, Starting

It is particularly useful for debugging that you can filter the log level by category such as ERROR to view all the error message that may have caused the job failure.

around any potential pitfalls. Traditionally, DBAs are overwhelmed with dev and test refresh data using RMAN, which produces significant overhead for daily DBA activities.

SnapCenter can help with application code testing and validation by easily and quickly cloning production database copies. Cloned production databases can then be used for testing and validating of new application code. Since NetApp SnapCenter clones databases so quickly and easily, you can refresh test data as quickly and as frequently as needed to boost your application code testing and migration preparation.

For many customers who are still sitting on an aging, or even de-supported Oracle version, the migration journey to Oracle 19c is particularly challenging. You might not have a one-step path to Oracle 19c. Rather it might be a multistep process. With a multistep migration approach, it is advisable to set up checkpoints and validate applications at each checkpoint so that any failure at a step would not nullify the entire migration effort. To fall back to a checkpoint, a database backup or a data dump is required. However, for many business applications, the window of migration is most likely limited to 48 hours on the weekend. You might not have the luxury to run database backups or data dumps, especially for large databases. It simply takes too long to be viable option with a limited migration window. Therefore, your choice of checkpoint and fallback planning is very limited.

For NetApp SnapCenter protected databases, you are free of above limitation. You can set up as many checkpoints as you wish in your migration go-live planning. Because it only takes SnapCenter a few minutes to do an offline full-database snapshot backup, and it takes a similar amount of time to restore and recover the data copy. A database snapshot at each migration step then becomes a very reliable fallback point when needed along a long migration journey. This can help mitigate many database and application migration risks and ensure a successful migration.

Automated deployment (optional)

Manual Oracle RAC deployment is a time consuming and tedious process. The benefits of automation using Ansible are not limited to the initial RAC setup in that you can execute the same commands on multiple hosts and use playbook to automate many RAC deployment tasks. This reduces time to value from days to hours.

Ansible automation is also a meaningful tool for day-to-day DBA activities. It is particularly efficient for configuring and supporting large Oracle RAC cluster with a large number of nodes.

Ansible controller setup

For this solution, we use open-source Centos 8.2 to run the automation controller with Ansible. Complete the following steps to set up the ansible controller:

1. Install python 3.

CentOS 8 comes with Python 3 installed by default. However, if, for whatever reason, Python3 is not installed, install it using the following DNF commands:

```
[admin@ansiblectl ~]$ sudo dnf install python3
[sudo] password for admin:
Last metadata expiration check: 1:18:09 ago on Wed 25 Nov 2020 08:59:41 AM EST.
Package python36-3.6.8-2.module_el8.1.0+245+c39af44f.x86_64 is already installed.
Dependencies resolved.
Nothing to do.
Complete!
[admin@ansiblectl ~]$ python3 -V
Python 3.6.8
```

2. Install PIP, the Python package installer.

```
sudo dnf install python3-pip
[sudo] password for admin:
Last metadata expiration check: 1:34:11 ago on Wed 18 Nov 2020 10:03:51 AM EST.
Package python3-pip-9.0.3-16.el8.noarch is already installed.
```

```
Dependencies resolved.
Nothing to do.
Complete!
```

3. Install Ansible.

```
[admin@ansiblectl ~]$ pip3 install ansible --user
```

4. Check the Ansible version after installation.

```
[admin@ansiblectl ~]$ ansible --version
ansible 2.10.3
  config file = None
  configured module search path = ['/home/admin/.ansible/plugins/modules',
'/usr/share/ansible/plugins/modules']
  ansible python module location = /home/admin/.local/lib/python3.6/site-packages/ansible
  executable location = /home/admin/.local/bin/ansible
  python version = 3.6.8 (default, Apr 16 2020, 01:36:27) [GCC 8.3.1 20191121 (Red Hat 8.3.1-5)]
```

5. Generate the SSH key.

```
[admin@ansiblectl ~]$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/home/admin/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/admin/.ssh/id_rsa.
Your public key has been saved in /home/admin/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:2zuduu8ljfQkp5A9ScaJGUhl/7WebzT8Zh0yds1RuOc admin@ansiblectl
The key's randomart image is:
+---[RSA 3072]-----+
|      ..o+      . |
|      .. * .. . |
|      o *  .o |
|      = o.o |
|      S o * **o |
|      o &.BE |
|      . ..=.B+= |
|      ..oo B |
|      +*o  +. |
+----[SHA256]-----+
```

6. Copy the SSH ID to the Oracle hosts.

```
[admin@ansiblectl ~]$ ssh-copy-id admin@b200-ora-01
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/admin/.ssh/id_rsa.pub"
The authenticity of host 'b200-ora-01 (10.61.180.151)' can't be established.
ECDSA key fingerprint is SHA256:DbTp9GZlxOzSDcId1cnpFAxpfe72Thsx5TXaxEnnG3Q.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are
already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to
install the new keys
admin@b200-ora-01's password:
Number of key(s) added: 1
Now try logging into the machine, with: "ssh 'admin@b200-ora-01'"
and check to make sure that only the key(s) you wanted were added.
```

7. Validate password-less authentication.

```
[admin@ansiblectl ~]$ ssh admin@b200-ora-01
Activate the web console with: systemctl enable --now cockpit.socket
Last failed login: Mon Nov 23 13:26:58 EST 2020 from 10.61.186.241 on ssh:notty
There were 4 failed login attempts since the last successful login.
Last login: Sat Nov 21 11:48:09 2020 from 10.61.180.175
[admin@b200-ora-01 ~]$
```

8. Complete the copy ID to all Oracle hosts to enable SSH access without a password.

9. Ping remote Oracle hosts to validate connectivity.

```
[admin@ansiblectl ansible]$ ansible -i hosts b200-ora-01 -m ping
[WARNING]: Platform linux on host b200-ora-01 is using the discovered Python interpreter at
/usr/bin/python3.6, but future installation of another Python interpreter could change the
meaning of that path. See
https://docs.ansible.com/ansible/2.10/reference_appendices/interpreter_discovery.html for more
information.
b200-ora-01 | SUCCESS => {
  "ansible_facts": {
    "discovered_interpreter_python": "/usr/bin/python3.6"
  },
  "changed": false,
  "ping": "pong"
}
```

The warning can be safely ignored.

10. Uncomment %wheel group to allow the admin user to run sudo without a password on Oracle hosts.

11. Run visudo and remove the # sign from the following line in front of the %wheel. Add the admin user at the Oracle hosts to the wheel group if it is not there already.

```
## Allows people in group wheel to run all commands
%wheel  ALL=(ALL)        ALL
## Same thing without a password
%wheel  ALL=(ALL)        NOPASSWD: ALL
```

12. Set ANSIBLE_PYTHON_INTERPRETER=auto_silent in the shell to suppress the python interpreter warning.

```
export ANSIBLE_PYTHON_INTERPRETER=auto_silent
```

Role-based solution automation playbook

NetApp Solution Engineering maintains a GitHub repository for role-based Ansible playbooks for FlexPod infrastructure automation as well as applications deployment on FlexPod, such as Oracle Database. For this solution, selective deployment tasks are automated with Ansible roles. If interested, you can clone a copy of the automation code at: [NetApp-Automation/na_NVA-1155: na_NVA-1155 \(github.com\)](https://github.com/NetApp-Automation/na_NVA-1155).

Deployment summary

In this deployment demonstration, we first deployed FlexPod infrastructure by setting up and configuring Nexus switches, a UCS fabric interconnect, MDS SAN switches, and a NetApp AFF A800 controller. We deployed UCS Manager to set up an Oracle server service profile template. We applied a service profile template to eight Oracle cluster nodes. We created FC SAN-based boot LUNs as well as data, redo, CRS, and binary LUNs on the AFF A800 controller in accordance with a predefined Oracle RAC storage layout plan and granted LUNs access as appropriate to Oracle hosts.

Through the UCS Manager KVM interface, Oracle server blade nodes were booted from the Oracle Linux 8.2 ISO, and operating systems were installed to the boot LUNs. The Oracle Linux 8.2 RHCK kernel was chosen as the operating kernel, and it was subsequently configured for Oracle grid infrastructure and Oracle database installation. An Oracle database was deployed using the CDB/PDB model. We created two CDBs with each hosting up to three PDB as an example of possible Oracle 19c database deployment.

The FlexPod infrastructure setup was validated against an FIO benchmark workload, and the Oracle installation was validated against a SLOB benchmark workload for performance verification.

We deployed SnapCenter 4.4 to protect the FlexPod Oracle database environment. We demonstrated the capability and efficiency of SnapCenter to backup, restore, and clone Oracle databases. We particularly showcased how to recover a subset of an Oracle database hosted in a CDB using PDB and PDB

tablespace point-in-time recovery. We further explained how SnapCenter can be employed as a viable tool for Oracle 19c migration assistance.

The solution deployment demonstration primarily used manual steps. As an option, we demonstrated how to set up an Ansible automation controller to run automated deployment tasks. For customers who are interested in solution deployment using Ansible automation, NetApp has built and published role-based automation playbooks for authorized access, which can be applied to FlexPod infrastructure deployment as well as follow-on applications deployment such as Oracle database for the benefits of time to value, implementation of best practices enforced through automation, as well as minimizing the possibility of user errors during solution deployment.

Validation results

We tested the performance of Oracle RAC deployment on eight cluster nodes with bare metal configuration. We have tested performance scalability of the infrastructure platform as well as the Oracle database using popular the FIO and SLOB benchmark tools as explained below.

FIO benchmark validation


FIO is an I/O workload generator that can spawn concurrent threads or processes doing a particular type of I/O operation as specified by the user. For our solution testing, we ran FIO against 56 shared NetApp AFF A800 storage LUNs (46 data and 8 redo) from eight RAC nodes concurrently and ramped up the concurrent workload further by increase the number of jobs or threads to push the back-end AFF A800 storage to the limit.

We ran an FIO test to simulate an OLTP workload with a random 4k and 8k block I/O size and measured achievable IOPS and latency. We also tested simulated 4k, 8k workloads in a mixed read/write ratio with 100/0 percent read/write, 90/10 percent read/write, 70/30 percent read/write, 50/50 percent read/write, and 0/100 read/write.


For a DSS workload, we ran a FIO test with a large block size I/O in 512k and 1024k to measure sequential I/O performance in throughput and latency. Similarly, we tested simulated 512k and 1024k sequential workloads by varying the read/write ratio with 100/0 percent read/write, 90/10 percent read/write, 70/30 percent read/write, 50/50 percent read/write, and 0/100 read/write in a mixed read and write ratio.

Following are the test results.


4k Random I/O performance




8k Random I/O performance



512k Sequential I/O performance



1024k Sequential I/O performance



In all test cases, we achieved high IOPS or throughput with consistent low latency, which is a validation of the infrastructure platform for the support of mission critical, latency-sensitive Oracle workloads.

SLOB benchmark validation

The Silly Little Oracle Benchmark (SLOB) is a toolkit for generating and testing I/O through an Oracle database. SLOB is very effective to test the I/O subsystem with genuine, Oracle SGA-buffered physical I/O.

To test a simulated Oracle workload using the SLOB tool, we created a SLOB database on the same 56 LUNs that we tested with FIO in two disk groups: one +DATA and one +REDO with an 8k block size. The SLOB database was populated with eight schemas. Therefore, each RAC node was running against one schema. The SLOB schema data was populated with the option tag, `OBFUSCATE_COLUMNS=TRUE`, so that each schema data was not highly compressible as in standard SLOB test use cases. Each SLOB schema was configured with a scale of 512G for a total of 4TB of test data.

We launched SLOB tests from RAC node 1 with `SQLNET_SERVICE_MAX` set at 8, so that SLOB spawn connections to each RAC instance using SQL net connections defined in `tns_names.ora` on node 1. The workload was ramped up by the number of threads as configured in the `slob.conf` file. Each thread represented a session to a SLOB database that executed I/O against the database. See the following session count from a SLOB test captured from the SLOB database.

USERNAME	INSTANCE	COUNT (*)
USER1	1	100
USER2	2	100
USER3	3	100
USER4	4	100
USER5	5	100
USER6	6	100
USER7	7	100
USER8	8	100
8 rows selected.		

The connections were evenly distributed among the eight RAC nodes with each connected to a database user ID when threads count was set at 100 in the `slob.conf` file. We ramped up to 200 threads for each RAC node for a combined total of 1,600 threads in our test cases.

From the AWR report generated from the test, we can see that the I/O was evenly balanced with SLOB test cases.


System Statistics - Per Second			DB/Inst: SLOB/slob6 Snaps: 247-248							
I#	Logical Reads/s	Physical Reads/s	Physical Writes/s	Redo Size (k)/s	Block Changes/s	User Calls/s	Execs/s	Parses/s	Logons/s	Txns/s
1	166,228.49	161,720.1	2.7	6.3	26.2	0.8	1,713.2	7.1	0.10	0.1
2	128,322.81	122,957.2	2.2	5.4	26.0	0.8	1,936.1	2.6	0.10	0.1
3	128,484.18	123,365.1	1.9	5.2	16.2	0.8	1,943.5	1.3	0.09	0.0
4	128,142.06	123,074.2	1.8	5.1	15.5	0.8	1,939.0	1.4	0.09	0.0
5	129,552.15	124,411.4	1.6	5.2	12.6	0.8	1,960.1	1.2	0.09	0.0
6	131,169.28	126,052.5	1.8	5.1	15.2	0.8	1,984.8	1.4	0.10	0.0
7	129,338.02	124,266.6	1.6	5.0	15.5	0.8	1,956.9	1.4	0.09	0.0
8	128,598.23	123,510.2	1.3	4.5	13.2	1.0	1,945.8	1.4	0.10	0.0

Sum	1,069,835.22	1,029,357.2	14.9	41.8	140.5	6.6	15,379.4	17.7	0.76	0.2
Avg	133,729.40	128,669.7	1.9	5.2	17.6	0.8	1,922.4	2.2	0.10	0.0
Std	13,167.98	13,391.9	0.4	0.5	5.4	0.1	86.0	2.0	0.00	0.0


In the SLOB tests, we measured Oracle I/O performance in terms of IOPS against latency with a read/write I/O ratio of 100 read, 75% read/25% write, and 100% write.

The following diagram demonstrated the SLOB test results.


100% select



75% select, 25% update



100% update



Again, the SLOB test results were consistent with the FIO test results in that the platform was capable of delivering substantial IOPS with minimal latency in less than one millisecond.

Infrastructure high availability and resiliency validation

FlexPod infrastructure is built with high availability and resiliency in mind. All components within the platform have double redundancy to eliminate any single point of failure.

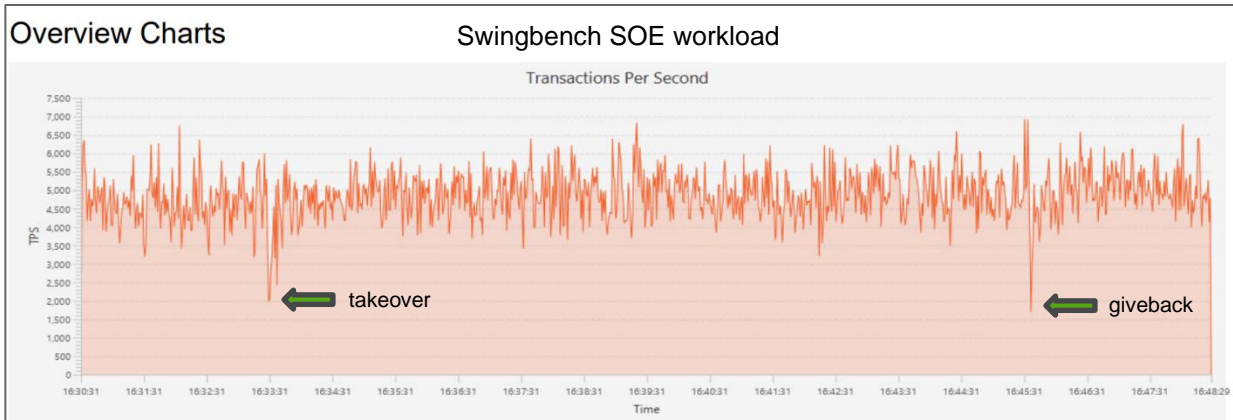
To validate the platform's high availability and resiliency, all major components were tested by simulating a failure with one of each component shut down.

AFF A800 controllers

To validate AFF A800 controller pair HA capability, we deployed active OLTP Swingbench Order Entry (SOE) workload in PDB mdb2_pdb1. With SOE workload stable at around 300K transaction per minute with 64 concurrent users, we performed manual storage takeover and giveback from AFF A800 node2 against node1:

```
FlexPod-A800-01-02:> storage failover takeover -ofnode FlexPod-A800-01-02-01
Warning: A takeover will be initiated. Once the partner node reboots, a giveback will be automatically initiated. Do you want to continue? {y|n}: y
FlexPod-A800-01-02:> storage failover show
Node      Partner      Takeover      Possible State Description
-----
FlexPod-A800-01-02-01  FlexPod-A800-01-02-02  -      Unknown
FlexPod-A800-01-02-02  FlexPod-A800-01-02-01  false  In takeover, Auto giveback will be initiated in 531 seconds
2 entries were displayed.
```

With the takeover, AFF A800 node1 was rebooted and was not available for period of about 12 minutes. We monitored and observed that there was a very brief dip in transaction volume in the SOE workload that quickly bounced back after AFF A800 node2 took over or gave back the storage LUNs on AFF A800 node1.



We also tested storage take over and give back from node1 against node2 and observed similar results. The tests indicated that there was a negligible application performance effect when one storage controller went offline, and the surviving controller picked up the I/O workload almost immediately.

Nexus switches

Nexus switches provide redundancy for end user connections through peer-linked dual switches. Each switch is configured with multiple virtual port channels that are linked to a separate fabric interconnect or LAN switches and each port channel includes at least two ports for redundancy.

To test the Nexus switches high availability, we simulated a failure by rebooting one Nexus switch after an active Swingbench SOE workload stabilized at above 300k transaction per minute. We validated application user connections to database were evenly distributed among the eight RAC nodes.


INST_ID	USERNAME	MACHINE	COUNT (*)
7	SOE	AB-Labs	8
2	SOE	AB-Labs	8
4	SOE	AB-Labs	8
5	SOE	AB-Labs	8
1	SOE	AB-Labs	8
3	SOE	AB-Labs	8
6	SOE	AB-Labs	8
8	SOE	AB-Labs	8

8 rows selected.

SQL>

We observed that the shutdown and reboot of a Nexus switch did not affect the active workload transaction volume. When the rebooted switch came back online, we observed that it initially accepted traffic and routed through the peer link while the virtual port channels were being enabled. There was a small hiccup when the traffic was routed through a port channel versus a peer link when port channels were enabled as shown below.

Overview Charts



The tests of simulated failure at either Nexus switch validated that failure at one switch only takes down one port in a port channel and traffic was routed through the remaining live port to sustain availability.

Fabric interconnects

UCS fabric interconnects serve as a central hub to provide connectivity between blade servers to storage as well as to provide connectivity between end users and Oracle databases or applications. On fabric interconnects, Oracle traffic is segregated as public and private/ASM traffic that routes through fabric A and B. In normal operation, public traffic is routed through fabric A with vLAN 180 and private and ASM traffic is routed through fabric B with vLAN 3357.

```
FI-ORA-01-A(nx-os)# show mac address-table
Legend:
    * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
    age - seconds since last seen, + - primary entry using vPC Peer-Link,
    (T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan
    VLAN    MAC Address    Type    age    Secure NTFY Ports
    -----
* 180      0025.b589.aa8f    static  -      F      F      Veth963
* 180      0025.b589.aa9f    static  -      F      F      Veth971
* 180      0025.b589.aaaf    static  -      F      F      Veth955
* 180      0025.b589.aabf    static  -      F      F      Veth939
* 180      0025.b589.aacf    static  -      F      F      Veth947
* 180      0025.b589.aadf    static  -      F      F      Veth923
* 180      0025.b589.aadf    static  -      F      F      Veth931
* 180      0025.b589.aaff    static  -      F      F      Veth915
```

```
FI-ORA-01-B(nx-os)# show mac address-table
Legend:
    * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
    age - seconds since last seen, + - primary entry using vPC Peer-Link,
    (T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan
    VLAN    MAC Address    Type    age    Secure NTFY Ports
    -----
* 3357     0025.b589.bb8f    static  -      F      F      Veth965
* 3357     0025.b589.bb9f    static  -      F      F      Veth973
* 3357     0025.b589.bbaf    static  -      F      F      Veth957
* 3357     0025.b589.bbbf    static  -      F      F      Veth941
* 3357     0025.b589.bbcf    static  -      F      F      Veth949
* 3357     0025.b589.bbdf    static  -      F      F      Veth925
* 3357     0025.b589.bbdf    static  -      F      F      Veth933
* 3357     0025.b589.bbff    static  -      F      F      Veth917
```


We simulated failure scenarios by rebooting either FI A or FI B. The tests showed that the failure scenarios were disruptive initially as either application connections from the public link or RAC cluster traffic were severed and SOE transaction volume suddenly dropped off.

The vNIC for both public and private traffic on fabric interconnects were configured for failover. The failure on fabric A triggered the vNIC failover to fabric B and vice versa as shown below:

```
FI-ORA-01-B(nx-os)# show mac address-table
```

Legend:

* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen, + - primary entry using vPC Peer-Link,
(T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan

VLAN	MAC Address	Type	age	Secure	NTFY	Ports
* 180	0025.b589.aa8f	static	-	F	F	Veth964
* 180	0025.b589.aa9f	static	-	F	F	Veth972
* 180	0025.b589.aaaf	static	-	F	F	Veth956
* 180	0025.b589.aabf	static	-	F	F	Veth940
* 180	0025.b589.aacf	static	-	F	F	Veth948
* 180	0025.b589.aadf	static	-	F	F	Veth924
* 180	0025.b589.aaef	static	-	F	F	Veth932
* 180	0025.b589.aaff	static	-	F	F	Veth916
* 3357	0025.b589.bb8f	static	-	F	F	Veth965
* 3357	0025.b589.bb9f	static	-	F	F	Veth973
* 3357	0025.b589.bbaf	static	-	F	F	Veth957
* 3357	0025.b589.bbbf	static	-	F	F	Veth941
* 3357	0025.b589.bbcf	static	-	F	F	Veth949
* 3357	0025.b589.bddf	static	-	F	F	Veth925
* 3357	0025.b589.bbef	static	-	F	F	Veth933
* 3357	0025.b589.bbff	static	-	F	F	Veth917

As the application reconnected after vNIC failover, SOE transaction volume started to resume and bounce back. However, since a failed fabric interconnect also brought down half of the SAN paths to storage, the transaction volume did not recover to previous level until all SAN paths were reinstated.

MDS switches

MDS switches provide application connectivity to NetApp AFF A800 storage through the fabric interconnects. FC zoning was enforced on MDS switches. A failure on either MDS switches causes application to lose half of active FC paths to a LUN and can be interruptive briefly.

Our failure tests showed that, at the removal of an MDS switch, I/O was rejected at failed active paths and queued until the failed paths were reinstated. The test results were similar to the failure tests for the fabric interconnects. Upon giveback, the application transaction volume quickly recovered to its previous level.

Conclusion

For many mission-critical Oracle enterprise applications that require high IOPS, high throughput, and low latency, the bare-metal FlexPod Datacenter with the FC protocol is an ideal platform for these workloads.

The FlexPod Datacenter with a NetApp AFF A800 system and a Cisco UCS Gen5 computing system is a converged infrastructure platform that combines technologies from Cisco and NetApp into a powerful converged platform for hosting enterprise Oracle database applications.

The prevalidated FlexPod Datacenter with Oracle 19c RAC databases architecture delivers proven high performance, value, and agility, and this system enables faster deployments, greater flexibility of choice, efficiency, high availability, lower cost, and lower risk.

Where to find additional information

To learn more about the information described in this document, review the following documents and/or websites:

- Learn how to Automate deployment of Oracle 19c RAC on FlexPod with Ansible roles and playbooks
<https://youtu.be/VcQMjIRzhoY>
- Oracle Database Installation Guide for Linux
<https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/index.html>
- Oracle Grid Infrastructure Installation and Upgrade Guide for Linux
<https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/index.html>
- Upgrade and Migrate to Oracle Database 19c
<https://www.oracle.com/sn/a/tech/docs/twp-upgrade-oracle-database-19c.pdf>
- UCS Hardware and Software Compatibility
<https://ucshcltool.cloudapps.cisco.com/public/>
- NetApp Automation / na_NVA-1155
https://github.com/NetApp-Automation/na_NVA-1155
- NetApp Interoperability Matrix Tool
<https://mysupport.netapp.com/matrix/#welcome>
- Best Practices for Modern SAN
<https://www.netapp.com/pdf.html?item=/media/10680-tr4080.pdf>
- FlexPod Datacenter Oracle Database Backup with SnapCenter
<https://www.netapp.com/pdf.html?item=/media/16973-sb-3999.pdf>
- SnapCenter Plug-In for Oracle Database Best Practices
<https://www.netapp.com/pdf.html?item=/media/12403-tr4700pdf.pdf>
- Using Oracle Linux 8.2 with NetApp ONTAP
https://docs.netapp.com/us-en/ontap-sanhost/hu_ol_82.html#installing-the-linux-unified-host-utilities
- TR-3633 Oracle Databases on ONTAP
<https://www.netapp.com/media/8744-tr3633.pdf>

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.

Copyright Information

Copyright © 2021 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-1155