West Virginia
Purchasing Division

2019 Washington Street, East
Charleston, WV 25305
Telephone: 304-558-2306
General Fax: 304-558-6026
Bid Fax: 304-558-3970

The following documentation is an electronically-submitted vendor response to an advertised solicitation from the *West Virginia Purchasing Bulletin* within the Vendor Self-Service portal at *wvOASIS.gov*. As part of the State of West Virginia's procurement process, and to maintain the transparency of the bid-opening process, this documentation submitted online is publicly posted by the West Virginia Purchasing Division at *WVPurchasing.gov* with any other vendor responses to this solicitation submitted to the Purchasing Division in hard copy format.

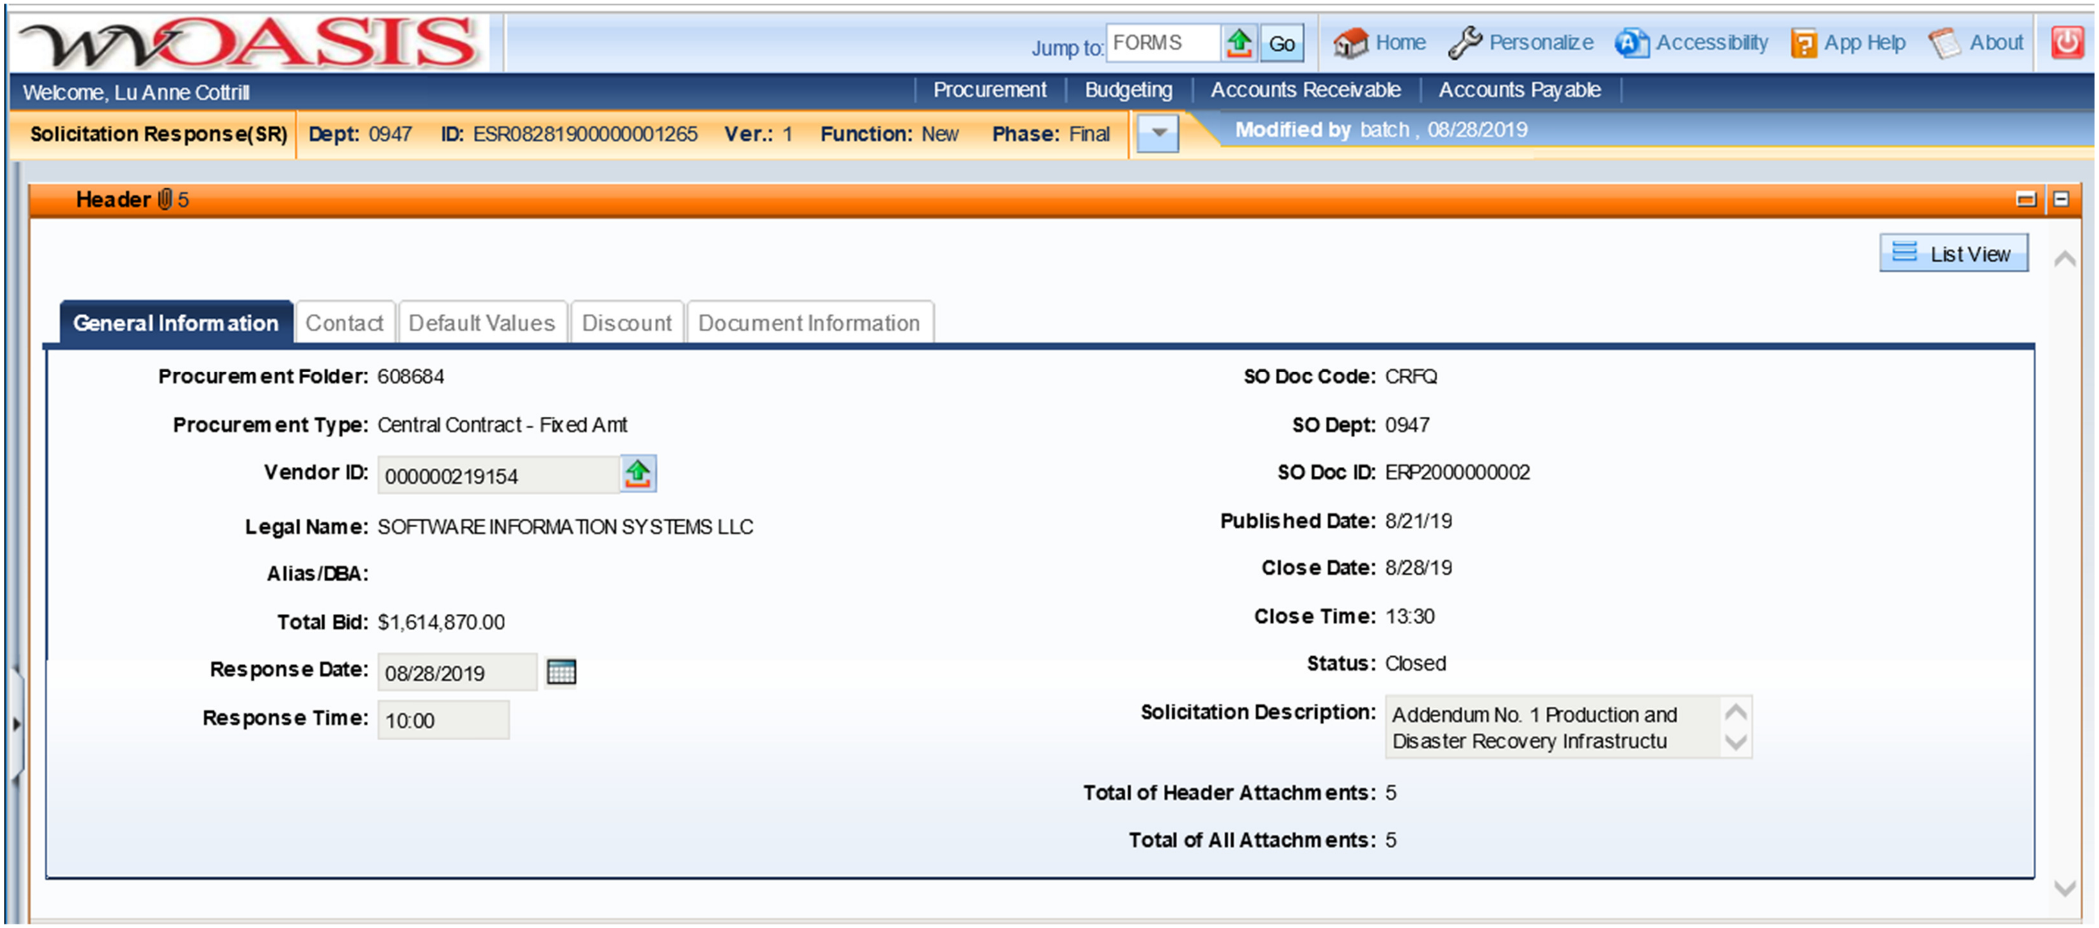

Welcome, Lu Anne Cottrill

Solicitation Response(SR) Dept: 0947 ID: ESR08281900000001265 Ver.: 1 Function: New Phase: Final

Modified by batch , 08/28/2019

Header 5

Procurement Folder: 608684

Procurement Type: Central Contract - Fixed Amt

Vendor ID: 000000219154

Legal Name: SOFTWARE INFORMATION SYSTEMS LLC

Alias/DBA:

Total Bid: $1,614,870.00

Response Date: 08/28/2019

Response Time: 10:00

SO Doc Code: CRFQ

SO Dept: 0947

SO Doc ID: ERP2000000002

Published Date: 8/21/19

Close Date: 8/28/19

Close Time: 13:30

Status: Closed

Solicitation Description: Addendum No. 1 Production and Disaster Recovery Infrastructu

Total of Header Attachments: 5

Total of All Attachments: 5

Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

State of West Virginia
Solicitation Response

Proc Folder : 608684

Solicitation Description : Addendum No. 1 Production and Disaster Recovery Infrastructu

Proc Type : Central Contract - Fixed Amt

Date issued	Solicitation Closes	Solicitation Response	Version
	2019-08-28 13:30:00	SR 0947 ESR08281900000001265	1

VENDOR

000000219154

SOFTWARE INFORMATION SYSTEMS LLC

Solicitation Number: CRFQ 0947 ERP2000000002

Total Bid : $1,614,870.00 **Response Date:** 2019-08-28 **Response Time:** 10:00:56

Comments:

FOR INFORMATION CONTACT THE BUYER

Melissa Pettrey
(304) 558-0094
melissa.k.pettrey@wv.gov

Signature on File **FEIN #** **DATE**

All offers subject to all terms and conditions contained in this solicitation

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
1	Storage Array Solution - Production Site	1.00000	EA	$267,000.000000	$267,000.00

Comm Code	Manufacturer	Specification	Model #
43201802			

Extended Description : Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
2	Storage Array Solution - Disaster Recovery Site	1.00000	EA	$267,000.000000	$267,000.00

Comm Code	Manufacturer	Specification	Model #
43201802			

Extended Description : Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
3	Compute/Networking Solution - Production Site (with VMware)	1.00000	EA	$450,000.000000	$450,000.00

Comm Code	Manufacturer	Specification	Model #
43232907			

Extended Description : Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
4	Compute/Networking Solution - Disaster Rec Site (w. VMware)	1.00000	EA	$288,000.000000	$288,000.00

Comm Code	Manufacturer	Specification	Model #
43232907			

Extended Description : Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
5	Implementation and Migration Services				$100,370.00

Comm Code	Manufacturer	Specification	Model #
81112309			

Extended Description :	
	Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
6	Cabinet - Production Site	1.00000	EA	$7,000.000000	$7,000.00

Comm Code	Manufacturer	Specification	Model #
43223306			

Extended Description :	
	Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
7	Cabinet - Disaster Recovery Site	1.00000	EA	$7,000.000000	$7,000.00

Comm Code	Manufacturer	Specification	Model #
43223306			

Extended Description :	
	Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
8	Firewalls (with Subscription Services)	1.00000	EA	$185,000.000000	$185,000.00

Comm Code	Manufacturer	Specification	Model #
43222501			

Extended Description :	
	Exact specifications can be found on the attached specifications sheet.

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Ln Total Or Contract Amount
9	Knowledge Transfer and Training				$43,500.00

Comm Code	Manufacturer	Specification	Model #
86000000			

Extended Description : Exact specifications can be found on the attached specifications sheet.

SIS

Strategy | Technology | Results

200 Association Drive
Charleston, WV 25311
304 768-1645
304 768-1671 (fax)

cisco NetApp

Ms. Melissa Perry
Department of Administration
Purchasing Division
2019 Washington Street, East
Charleston, WV 25305

August 28, 2019

Dear Ms. Pettrey,

Please find attached the SIS response to CRFQ 0947 ERP2000000003 due August 28, 2019. SIS agrees with the requirements of the RFQ and is submitting a response offering the Cisco UCS with NetApp Storage and Palo Alto Firewalls and VMware software. This response meets or exceeds the mandatory requirements of the RFQ for both the Production and Disaster Recovery sites and all components are included. The prices included are for this one-time purchase - with five years of support for Cisco, NetApp and Palo Alto products and three years of support for VMware software. Future procurements for these products are not subject to the prices quoted.

Thank you for the opportunity to submit this response and we look forward to further discussions.

Sincerely,

Charles D. Arnett
Senior Client Executive
carnett@thinksis.com

Enclosures

Purchasing Divison
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

State of West Virginia
Request for Quotation
21 — Info Technology

Proc Folder: 608684

Doc Description: Production and Disaster Recovery Infrastructure Upgrade

Proc Type: Central Contract - Fixed Amt

Date Issued	Solicitation Closes	Solicitation No	Version
2019-08-14	2019-08-28 13:30:00	CRFQ 0947 ERP2000000002	1

BID RECEIVING LOCATION

BID CLERK
DEPARTMENT OF ADMINISTRATION
PURCHASING DIVISION
2019 WASHINGTON ST E
CHARLESTON WV 25305
US

VENDOR

Vendor Name, Address and Telephone Number:

Software Information Systems, LLC
200 Association Drive, Suite 210
Charleston, WV 25311
304 768-1645

FOR INFORMATION CONTACT THE BUYER

Melissa Pettrey
(304) 558-0094
melissa.k.pettrey@wv.gov

Signature X FEIN # 61-1371685 DATE 8-28-2019

All offers subject to all terms and conditions contained in this solicitation

ADDITIONAL INFORMATION:

REQUEST FOR QUOTATION

The West Virginia Purchasing Division is soliciting bids on behalf of the Agency, the West Virginia Enterprise Resource Planning Board (wvOASIS) to establish a contract for the replacement of its current Production and Disaster Recovery Infrastructure and obtain Services to assist in the installation and configuration of the new hardware per the bid requirements, specifications, terms and conditions attached to this solicitation.

The primary Production Site is in Charleston, WV with a Disaster Recovery Site located in Morgantown, WV.

INVOICE TO	SHIP TO
CONTROLLER ENTERPRISE RESOURCE PLANNING BOARD 1007 BULLITT STREET SUITE 400 CHARLESTON WV25301 US	CONTROLLER ENTERPRISE RESOURCE PLANNING BOARD 1007 BULLITT STREET SUITE 400 CHARLESTON WV 25301 US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
1	Storage Array Solution - Production Site	1.00000	EA	$267,000	$267,000

Comm Code	Manufacturer	Specification	Model #
43201802			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO	SHIP TO
CONTROLLER ENTERPRISE RESOURCE PLANNING BOARD 1007 BULLITT STREET SUITE 400 CHARLESTON WV25301 US	CONTROLLER ENTERPRISE RESOURCE PLANNING BOARD 1007 BULLITT STREET SUITE 400 CHARLESTON WV 25301 US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
2	Storage Array Solution - Disaster Recovery Site	1.00000	EA	$267,000	$267,000

Comm Code	Manufacturer	Specification	Model #
43201802			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO	SHIP TO
CONTROLLER	CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD	ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET	1007 BULLITT STREET
SUITE 400	SUITE 400
CHARLESTON WV25301	CHARLESTON WV 25301
US	US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
3	Compute/Networking Solution - Production Site (with VMware)	1.00000	EA	$450,000	$450,000

Comm Code	Manufacturer	Specification	Model #
43232907			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO	SHIP TO
CONTROLLER	CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD	ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET	1007 BULLITT STREET
SUITE 400	SUITE 400
CHARLESTON WV25301	CHARLESTON WV 25301
US	US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
4	Compute/Networking Solution - Disaster Rec Site (w. VMware)	1.00000	EA	$288,000	$288,000

Comm Code	Manufacturer	Specification	Model #
43232907			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO	SHIP TO
CONTROLLER	CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD	ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET	1007 BULLITT STREET
SUITE 400	SUITE 400
CHARLESTON WV25301	CHARLESTON WV 25301
US	US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
5	Implementation and Migration Services			$100,370	$100,370

Comm Code	Manufacturer	Specification	Model #
81112309			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO	SHIP TO
CONTROLLER	CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD	ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET	1007 BULLITT STREET
SUITE 400	SUITE 400
CHARLESTON WV25301	CHARLESTON WV 25301
US	US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
6	Cabinet - Production Site	1.00000	EA	$7,000	$7,000

Comm Code	Manufacturer	Specification	Model #
43223306			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO	SHIP TO
CONTROLLER	CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD	ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET	1007 BULLITT STREET
SUITE 400	SUITE 400
CHARLESTON WV25301	CHARLESTON WV 25301
US	US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
7	Cabinet - Disaster Recovery Site	1.00000	EA	$7,000	$7000.

Comm Code	Manufacturer	Specification	Model #
43223306			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO

CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET
SUITE 400
CHARLESTON WV25301
US

SHIP TO

CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET
SUITE 400
CHARLESTON WV 25301
US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
8	Firewalls (with Subscription Services)	1.00000	EA	$185,000	$185,000

Comm Code	Manufacturer	Specification	Model #
43222501			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

INVOICE TO

CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET
SUITE 400
CHARLESTON WV25301
US

SHIP TO

CONTROLLER
ENTERPRISE RESOURCE PLANNING BOARD
1007 BULLITT STREET
SUITE 400
CHARLESTON WV 25301
US

Line	Comm Ln Desc	Qty	Unit Issue	Unit Price	Total Price
9	Knowledge Transfer and Training			$43,500	$43,500

Comm Code	Manufacturer	Specification	Model #
86000000			

Extended Description :

Exact specifications can be found on the attached specifications sheet.

REQUEST FOR QUOTATION
Network Infrastructure Upgrade/Refresh

EXHIBIT A -PRICING PAGE

Vendor is to complete this Pricing Page and submit it with its bid response if not responding electronically via wvOASIS. Vendor shall enter the total bid amount (lump sum) per category/line item to provide the Goods and/or Services contemplated and as described in Section 4 of the Specifications for the project in wvOASIS, and/or on the corresponding lines below.

Additionally, Vendor should provide an itemized Build List with the brand, model numbers, etc. for all equipment proposed for each category and location as a separate attachment. Equipment prices for each category/location must include the services as specified in Section 5.2.1. (Hardware Delivery & Base Implementation Services). Vendor must supply this information with their bid response.

Price shall include a five (5) year warranty on all hardware, equipment, software, (except VMware) and support/licensing upgrade of applicable software from the date of Agency acceptance. Prices shall include all costs, fees, including but not limited to shipping, travel, lodging, meals and other related costs.

Item 1: Storage Array Solution -Production Site	$ 267,000
Item 2: Storage Array Solution -Disaster Recovery Site	$ 267,000
Item 3: Compute/Networking Solution -Production Site (with VMware)	$ 450,000
Item 4: Compute/Networking Solution-Disaster Recovery Site (with VMware)	$ 288,000
Item 5: Implementation/Migration Services	$ 100,370
Item 6: Cabinet -Production Site	$ 7,000
Item 7: Cabinet – Disaster Recovery Site	$ 7,000
Item8: Firewalls (with Subscriptions Services)	$ 185,000
Item 9: Knowledge Transfer & Training	$ 43,500

TOTAL BID AMOUNT $ 1,614,870.00
(Items 1-8)

DESIGNATED CONTACT: Vendor appoints the individual identified in this Section as the Contract Administrator and the initial point of contact for matters relating to this Contract.

Charles D. Arnett, Sr. Client Executive
(Name, Title)

(Printed Name and Title)
200 Association Drive Suite 210, Charleston, WV 25311
(Address)
304 768-1645 Fax 304 768-1671
(Phone Number) / (Fax Number)
carnett@ thinksis. com
(email address)

CERTIFICATION AND SIGNATURE: By signing below, or submitting documentation through wvOASIS, I certify that I have reviewed this Solicitation in its entirety; that I understand the requirements, terms and conditions, and other information contained herein; that this bid, offer or proposal constitutes an offer to the State that cannot be unilaterally withdrawn; that the product or service proposed meets the mandatory requirements contained in the Solicitation for that product or service, unless otherwise stated herein; that the Vendor accepts the terms and conditions contained in the Solicitation, unless otherwise stated herein; that I am submitting this bid, offer or proposal for review and consideration; that I am authorized by the vendor to execute and submit this bid, offer, or proposal, or any documents related thereto on vendor's behalf; that I am authorized to bind the vendor in a contractual relationship; and that to the best of my knowledge, the vendor has properly registered with any State agency that may require registration.

Software Information Systems LLC
(Company)
Charles D Arnett, Sr. Client Executive
(Authorized Signature) (Representative Name, Title)
Charles D Arnett Sr. Client Executive
(Printed Name and Title of Authorized Representative)

8-28-2019
(Date)

304 768-1645 Fax 304 768-1671
(Phone Number) (Fax Number)

10. VENDOR DEFAULT:

10.1 The following shall be considered a vendor default under this Contract.

10.1.1 Failure to provide Contract Items in accordance with the requirements contained herein.

10.1.2 Failure to comply with other specifications and requirements contained herein.

10.1.3 Failure to comply with any laws, rules, and ordinances applicable to the Contract Services provided under this Contract.

10.1.4 Failure to remedy deficient performance upon request.

10.2 The following remedies shall be available to Agency upon default.

10.2.1 Immediate cancellation of the Contract.

10.2.2 Immediate cancellation of one or more release orders issued under this Contract.

10.2.3 Any other remedies available in law or equity.

11. MISCELLANEOUS:

11.1 Contract Manager: During its performance of this Contract, Vendor must designate and maintain a primary contract manager responsible for overseeing Vendor's responsibilities under this Contract. The Contract manager must be available during normal business hours to address any customer service or other issues related to this Contract. Vendor should list its Contract manager and his or her contact information below.

Contract Manager: Charles Arnett
(Printed/Typed Name)

Telephone Number: 304 768-1645

Fax Number: 304 768-1671

Email Address: carnett@thinksis.com

West Virginia Ethics Commission

Disclosure of Interested Parties to Contracts

(Required by *W. Va. Code* § 6D-1-2)

Name of Contracting Business Entity: Software Information Systems, LLC **Address:** 165 Barr Street
Lexington, KY 40507

Name of Authorized Agent: Stephen F. Sigg **Address:** 165 Barr Street, Lexington, KY 40507

Contract Number: CRFQ ERP 2000000002 **Contract Description:** Production and Disaster Recovery Infrastructure Upgrade

Governmental agency awarding contract: Enterprise Resource Planning Board

□ **Check here if this is a Supplemental Disclosure**

List the Names of Interested Parties to the contract which are known or reasonably anticipated by the contracting business entity for each category below *(attach additional pages if necessary)*:

1. **Subcontractors or other entities performing work or service under the Contract**

   □ Check here if none, otherwise list entity/individual names below.

   NetApp
   Contact: David Franco

2. **Any person or entity who owns 25% or more of contracting entity (not applicable to publicly traded entities)**

   ☑ Check here if none, otherwise list entity/individual names below.

3. **Any person or entity that facilitated, or negotiated the terms of, the applicable contract (excluding legal services related to the negotiation or drafting of the applicable contract)**

   ☑ Check here if none, otherwise list entity/individual names below.

Signature: Stephen F. Sigg Date Signed: 8/27/2019

Notary Verification

State of Kentucky, County of Fayette:

I, Stephen F. Sigg, the authorized agent of the contracting business entity listed above, being duly sworn, acknowledge that the Disclosure herein is being made under oath and under the penalty of perjury.

Taken, sworn to and subscribed before me this 27th day of August, 2019.

Karen A. Smallwood
Notary Public's Signature

KAREN A. SMALLWOOD
NOTARY PUBLIC
ID NO. 574503
MY COMMISSION EXPIRES
3/27/2021
STATE AT LARGE

To be completed by State Agency:
Date Received by State Agency: ____________
Date submitted to Ethics Commission: ____________
Governmental agency submitting Disclosure: ____________

STATE OF WEST VIRGINIA
Purchasing Division

PURCHASING AFFIDAVIT

CONSTRUCTION CONTRACTS: Under W. Va. Code § 5-22-1(i), the contracting public entity shall not award a construction contract to any bidder that is known to be in default on any monetary obligation owed to the state or a political subdivision of the state, including, but not limited to, obligations related to payroll taxes, property taxes, sales and use taxes, fire service fees, or other fines or fees.

ALL CONTRACTS: Under W. Va. Code §5A-3-10a, no contract or renewal of any contract may be awarded by the state or any of its political subdivisions to any vendor or prospective vendor when the vendor or prospective vendor or a related party to the vendor or prospective vendor is a debtor and: (1) the debt owed is an amount greater than one thousand dollars in the aggregate; or (2) the debtor is in employer default.

EXCEPTION: The prohibition listed above does not apply where a vendor has contested any tax administered pursuant to chapter eleven of the W. Va. Code, workers' compensation premium, permit fee or environmental fee or assessment and the matter has not become final or where the vendor has entered into a payment plan or agreement and the vendor is not in default of any of the provisions of such plan or agreement.

DEFINITIONS:

"Debt" means any assessment, premium, penalty, fine, tax or other amount of money owed to the state or any of its political subdivisions because of a judgment, fine, permit violation, license assessment, defaulted workers' compensation premium, penalty or other assessment presently delinquent or due and required to be paid to the state or any of its political subdivisions, including any interest or additional penalties accrued thereon.

"Employer default" means having an outstanding balance or liability to the old fund or to the uninsured employers' fund or being in policy default, as defined in W. Va. Code § 23-2c-2, failure to maintain mandatory workers' compensation coverage, or failure to fully meet its obligations as a workers' compensation self-insured employer. An employer is not in employer default if it has entered into a repayment agreement with the Insurance Commissioner and remains in compliance with the obligations under the repayment agreement.

"Related party" means a party, whether an individual, corporation, partnership, association, limited liability company or any other form or business association or other entity whatsoever, related to any vendor by blood, marriage, ownership or contract through which the party has a relationship of ownership or other interest with the vendor so that the party will actually or by effect receive or control a portion of the benefit, profit or other consideration from performance of a vendor contract with the party receiving an amount that meets or exceed five percent of the total contract amount.

AFFIRMATION: By signing this form, the vendor's authorized signer affirms and acknowledges under penalty of law for false swearing (*W. Va. Code* §61-5-3) that: (1) for construction contracts, the vendor is not in default on any monetary obligation owed to the state or a political subdivision of the state, and (2) for all other contracts, that neither vendor nor any related party owe a debt as defined above and that neither vendor nor any related party are in employer default as defined above, unless the debt or employer default is permitted under the exception above.

WITNESS THE FOLLOWING SIGNATURE:

Vendor's Name: Software Information Systems, LLC

Authorized Signature: ______________________ Date: 8/27/2019

State of Kentucky

County of Fayette, to-wit:

Taken, subscribed, and sworn to before me this 27th day of August, 209.

My Commission expires March 27, 2021.

AFFIX SEAL HERE

KAREN A. SMALLWOOD
NOTARY PUBLIC
ID NO. 574603
MY COMMISSION EXPIRES 3/27/2021
STATE AT LARGE, KY.

NOTARY PUBLIC Karen A. Smallwood

ADDENDUM ACKNOWLEDGEMENT FORM
SOLICITATION NO.: ERP2000000002

Instructions: Please acknowledge receipt of all addenda issued with this solicitation by completing this addendum acknowledgment form. Check the box next to each addendum received and sign below. Failure to acknowledge addenda may result in bid disqualification.

Acknowledgment: I hereby acknowledge receipt of the following addenda and have made the necessary revisions to my proposal, plans and/or specification, etc.

Addendum Numbers Received:
(Check the box next to each addendum received)

[✓] Addendum No. 1

[ ] Addendum No. 2

[ ] Addendum No. 3

[ ] Addendum No. 4

[ ] Addendum No. 5

[ ] Addendum No. 6

[ ] Addendum No. 7

[ ] Addendum No. 8

[ ] Addendum No. 9

[ ] Addendum No. 10

I understand that failure to confirm the receipt of addenda may be cause for rejection of this bid. I further understand that that any verbal representation made or assumed to be made during any oral discussion held between Vendor's representatives and any state personnel is not binding. Only the information issued in writing and added to the specifications by an official addendum is binding.

Software Information Systems, LLC
Company

Karen Smallwood
Authorized Signature

8/27/2019
Date

NOTE: This addendum acknowledgement should be submitted with the bid to expedite document processing.

WVOasis

mvannette | August 28, 2019

Production

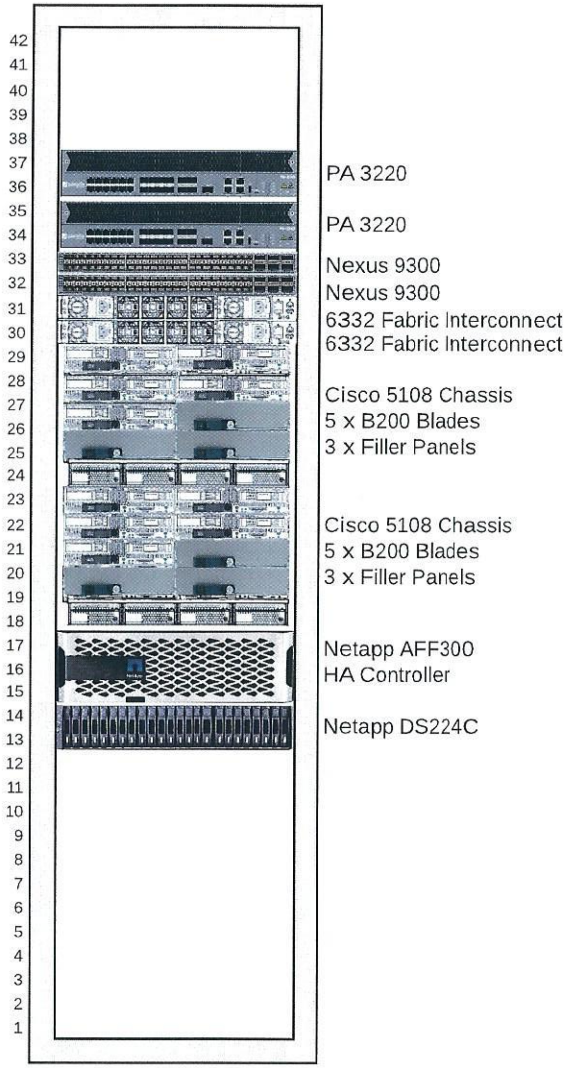

DR

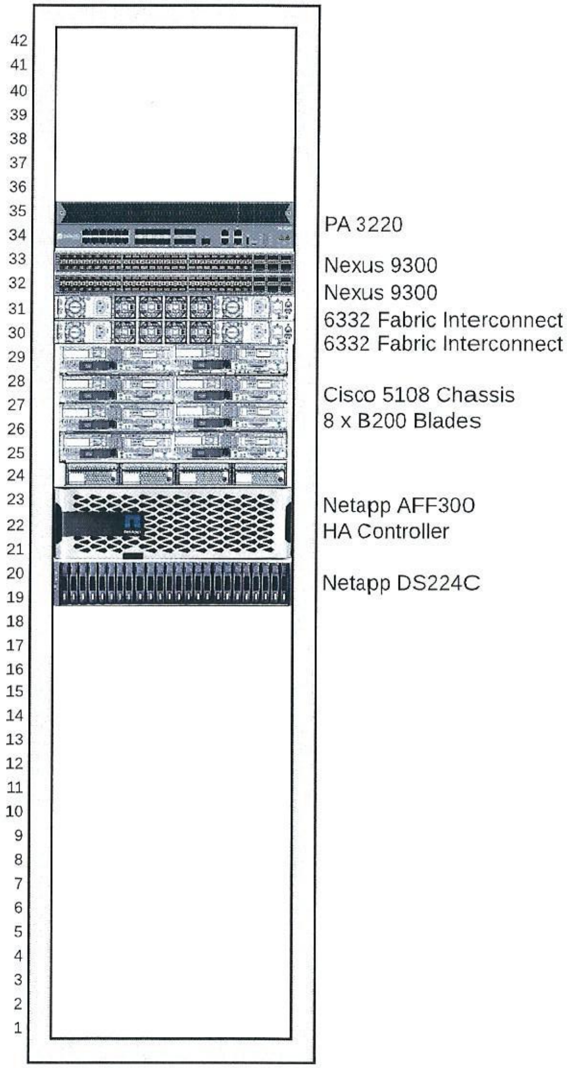

Production Cisco UCS Bill of Materials

System 1

Line No.	Qty	Part Number	Description
1	5	UCSB-B200-M5-U	UCS B200 M5 BLADE W/O CPU, MEM,
2	5	CON-SSC4P-	SOLN SUPP 24X7X4OS UCS B200 M5
3	60	UCS-MR-X32G2RT-H	32GB DDR4-2933-MHZ RDIMM/2RX4/1.2V
4	5	N20-FW016	UCS 5108 BLADE CHASSIS FW PACKAGE
5	5	UCS-SID-INFR-CFP	CONVERGED-FLEXPOD
6	5	UCS-SID-WKL-OW	OTHER WORKLOAD
7	10	UCSB-LSTOR-BK	FLEXSTORAGE BLANKING PANELS W/O
8	5	UCSB-HS-M5-F	CPU HEAT SINK FOR UCS B-SERIES M5
9	60	UCS-DIMM-BLK	UCS DIMM BLANKS
10	5	UCSB-HS-M5-R	CPU HEAT SINK FOR UCS B-SERIES M5
11	10	UCS-CPU-I6240	2.6GHZ/150W 6240 18C/24.75MB 3DX
12	10	UCS-SD-64G-S	64GB SD CARD FOR UCS SERVERS
13	5	UCSX-TPM2-002	TRUSTED PLATFORM MODULE 2.0 FOR
14	5	UCS-MSTOR-SD	MINI STORAGE CARRIER FOR SD
15	5	UCSB-MLOM-40G-04	CISCO UCS VIC 1440 MODULAR LOM
16	5	UCSB-VIC-M84-4P	CISCO UCS VIC 1480 MODULAR LOM
17	5	P177/R1	INTEGRATION SERVICES

System 2

Line No.	Qty	Part Number	Description
1	5	UCSB-B200-M5-U	UCS B200 M5 BLADE W/O CPU, MEM,
2	5	CON-SSC4P-	SOLN SUPP 24X7X4OS UCS B200 M5
3	5	N20-FW016	UCS 5108 BLADE CHASSIS FW PACKAGE
4	5	UCS-SID-INFR-CFP	CONVERGED-FLEXPOD
5	5	UCS-SID-WKL-ORCL	ORACLE
6	10	UCSB-LSTOR-BK	FLEXSTORAGE BLANKING PANELS W/O
7	5	UCSB-HS-M5-F	CPU HEAT SINK FOR UCS B-SERIES M5
8	60	UCS-DIMM-BLK	UCS DIMM BLANKS
9	5	UCSB-HS-M5-R	CPU HEAT SINK FOR UCS B-SERIES M5
10	10	UCS-CPU-I5222	3.8GHZ/125W 5222 4C/16.50MB 3DX DDR4
11	10	UCS-SD-64G-S	64GB SD CARD FOR UCS SERVERS
12	5	UCSX-TPM2-002	TRUSTED PLATFORM MODULE 2.0 FOR
13	5	UCS-MSTOR-SD	MINI STORAGE CARRIER FOR SD
14	60	UCS-MR-X32G2RT-H	32GB DDR4-2933-MHZ RDIMM/2RX4/1.2V
15	5	UCSB-MLOM-40G-04	CISCO UCS VIC 1440 MODULAR LOM
16	5	UCSB-VIC-M84-4P	CISCO UCS VIC 1480 MODULAR LOM
17	5	P177/R1	INTEGRATION SERVICES

System 3

Line No.	Qty	Part Number	Description
1	1	UCS-SP-FI633216-2X	UCS SP SELECT 6332-16UP FI/NO PSU/24
2	2	UCS-SP-FI6332-16UP	(NOT SOLD STANDALONE) UCS 6332-
3	2	CON-SSC4P-SP16UP	SOLN SUPP 24X7X4OS (NOT SOLD
4	4	UCS-PSU-6332-AC	UCS 6332 POWER SUPPLY/100-240VAC
5	4	CAB-C13-C14-2M	POWER CORD JUMPER, C13-C14
6	8	QSFP-H40G-CU3M	40GBASE-CR4 PASSIVE COPPER CABLE,
7	8	QSFP-40G-SR-BD	QSFP40G BIDI SHORT-REACH
8	8	DS-SFP-FC16G-SW	16 GBPS FIBRE CHANNEL SW SFP+, LC
9	2	N10-MGT015	UCS MANAGER V3.2(1)
10	2	UCS-ACC-6332	UCS 6332 CHASSIS ACCESSORY KIT

11	8	UCS-FAN-6332	UCS 6332 FAN MODULE
12	2	P177/O1	INTEGRATION SERVICES
Line No.	**Qty**	**Part Number**	**Description**
1	4	SFP-H10GB-CU1M=	10GBASE-CU SFP+ CABLE 1 METER
System 5			
Line No.	**Qty**	**Part Number**	**Description**
1	2	UCS-SP-FI6332-L-4X	UCS SP SELECT 6300 SERIES FABRIC INT
2	8	UCS-SP-LIC-40GE	SP 3RDGEN FI PER PORT LIC TO
3	8	UCS-LIC-6300-40G-B	3RD GEN FI PER PORT LICENSE TO
System 6			
Line No.	**Qty**	**Part Number**	**Description**
1	2	UCS-SP-5108-AC3	UCS SP SELECT 5108 AC2 CHASSIS
2	2	CON-SSC4P-5108AC3	SOLN SUPP 24X7X4OS, UCS SP SELECT
3	8	UCSB-PSU-2500ACDV	2500W PLATINUM AC HOT PLUG POWER
4	8	CAB-C19-CBN	CABINET JUMPER POWER CORD, 250
5	8	QSFP-H40G-CU3M	40GBASE-CR4 PASSIVE COPPER CABLE,
6	2	N20-FW015	UCS 5108 BLADE CHASSIS FW PACKAGE
7	2	UCSB-5108-PKG-HW	UCS 5108 PACKAGING FOR CHASSIS
8	2	N01-UAC1	SINGLE PHASE AC POWER MODULE FOR
9	2	N20-CAK	ACCESS. KIT FOR 5108 BLADE CHASSIS
10	16	N20-CBLKB1	BLADE SLOT BLANKING PANEL FOR
11	16	N20-FAN5	FAN MODULE FOR UCS 5108
12	4	UCS-IOM-2304	UCS 2304XP I/O MODULE (4 EXTERNAL,
13	2	P177/M1	INTEGRATION SERVICES

Price: $300,000.00

Disaster Recovery Cisco UCS Bill of Materials

System 1

Line No.	Qty	Part Number	Description
1	3	UCSB-B200-M5-U	UCS B200 M5 BLADE W/O CPU, MEM,
2	3	CON-SSC4P-	SOLN SUPP 24X7X4OS UCS B200 M5
3	36	UCS-MR-X32G2RT-H	32GB DDR4-2933-MHZ RDIMM/2RX4/1.2V
4	3	N20-FW016	UCS 5108 BLADE CHASSIS FW PACKAGE
5	3	UCS-SID-INFR-CFP	CONVERGED-FLEXPOD
6	3	UCS-SID-WKL-OW	OTHER WORKLOAD
7	6	UCSB-LSTOR-BK	FLEXSTORAGE BLANKING PANELS W/O
8	3	UCSB-HS-M5-F	CPU HEAT SINK FOR UCS B-SERIES M5
9	36	UCS-DIMM-BLK	UCS DIMM BLANKS
10	3	UCSB-HS-M5-R	CPU HEAT SINK FOR UCS B-SERIES M5
11	6	UCS-CPU-I6240	2.6GHZ/150W 6240 18C/24.75MB 3DX
12	6	UCS-SD-64G-S	64GB SD CARD FOR UCS SERVERS
13	3	UCSX-TPM2-002	TRUSTED PLATFORM MODULE 2.0 FOR
14	3	UCS-MSTOR-SD	MINI STORAGE CARRIER FOR SD
15	3	UCSB-MLOM-40G-04	CISCO UCS VIC 1440 MODULAR LOM
16	3	UCSB-VIC-M84-4P	CISCO UCS VIC 1480 MODULAR LOM
17	3	P177/R1	INTEGRATION SERVICES

System 2

Line No.	Qty	Part Number	Description
1	2	UCSB-B200-M5-U	UCS B200 M5 BLADE W/O CPU, MEM,
2	2	CON-SSC4P-	SOLN SUPP 24X7X4OS UCS B200 M5
3	2	N20-FW016	UCS 5108 BLADE CHASSIS FW PACKAGE
4	2	UCS-SID-INFR-CFP	CONVERGED-FLEXPOD
5	2	UCS-SID-WKL-ORCL	ORACLE
6	4	UCSB-LSTOR-BK	FLEXSTORAGE BLANKING PANELS W/O
7	2	UCSB-HS-M5-F	CPU HEAT SINK FOR UCS B-SERIES M5
8	24	UCS-DIMM-BLK	UCS DIMM BLANKS
9	2	UCSB-HS-M5-R	CPU HEAT SINK FOR UCS B-SERIES M5
10	4	UCS-CPU-I5222	3.8GHZ/125W 5222 4C/16.50MB 3DX DDR4
11	4	UCS-SD-64G-S	64GB SD CARD FOR UCS SERVERS
12	2	UCSX-TPM2-002	TRUSTED PLATFORM MODULE 2.0 FOR
13	2	UCS-MSTOR-SD	MINI STORAGE CARRIER FOR SD
14	24	UCS-MR-X32G2RT-H	32GB DDR4-2933-MHZ RDIMM/2RX4/1.2V
15	2	UCSB-MLOM-40G-04	CISCO UCS VIC 1440 MODULAR LOM
16	2	UCSB-VIC-M84-4P	CISCO UCS VIC 1480 MODULAR LOM
17	2	P177/R1	INTEGRATION SERVICES

System 3

Line No.	Qty	Part Number	Description
1	1	UCS-SP-FI633216-2X	UCS SP SELECT 6332-16UP FI/NO PSU/24
2	2	UCS-SP-FI6332-16UP	(NOT SOLD STANDALONE) UCS 6332-
3	2	CON-SSC4P-SP16UP	SOLN SUPP 24X7X4OS (NOT SOLD
4	4	UCS-PSU-6332-AC	UCS 6332 POWER SUPPLY/100-240VAC
5	4	CAB-C13-C14-2M	POWER CORD JUMPER, C13-C14
6	8	QSFP-H40G-CU3M	40GBASE-CR4 PASSIVE COPPER CABLE,
7	8	QSFP-40G-SR-BD	QSFP40G BIDI SHORT-REACH
8	8	DS-SFP-FC16G-SW	16 GBPS FIBRE CHANNEL SW SFP+, LC
9	2	N10-MGT015	UCS MANAGER V3.2(1)
10	2	UCS-ACC-6332	UCS 6332 CHASSIS ACCESSORY KIT

11	8	UCS-FAN-6332	UCS 6332 FAN MODULE
12	2	P177/O1	INTEGRATION SERVICES
System 4			
Line No.	**Qty**	**Part Number**	**Description**
1	4	SFP-H10GB-CU1M=	10GBASE-CU SFP+ CABLE 1 METER
System 5			
Line No.	**Qty**	**Part Number**	**Description**
1	1	UCS-SP-5108-AC3	UCS SP SELECT 5108 AC2 CHASSIS
2	1	CON-SSC4P-5108AC3	SOLN SUPP 24X7X4OS, UCS SP SELECT
3	4	UCSB-PSU-2500ACDV	2500W PLATINUM AC HOT PLUG POWER
4	4	CAB-C19-CBN	CABINET JUMPER POWER CORD, 250
5	4	QSFP-H40G-CU3M	40GBASE-CR4 PASSIVE COPPER CABLE,
6	1	N20-FW015	UCS 5108 BLADE CHASSIS FW PACKAGE
7	1	UCSB-5108-PKG-HW	UCS 5108 PACKAGING FOR CHASSIS
8	1	N01-UAC1	SINGLE PHASE AC POWER MODULE FOR
9	1	N20-CAK	ACCESS. KIT FOR 5108 BLADE CHASSIS
10	8	N20-CBLKB1	BLADE SLOT BLANKING PANEL FOR
11	8	N20-FAN5	FAN MODULE FOR UCS 5108
12	2	UCS-IOM-2304	UCS 2304XP I/O MODULE (4 EXTERNAL,
13	1	P177/M1	INTEGRATION SERVICES

Price: $235,000.00

NetApp Production Bill of Materials

Line Item	Qty	Part Number	Product
1	1	SW-2-CL-BASE	SW-2,Base,CL,Node
2	1	AFF-A300	
3	2	AFF-A300A-001	AFF A300 HA System,FlashBundle
4	2	SW-2-A300A-NVE-C	SW,Data at Rest Encryption Enabled,A300A,-C
5	2	SW-2-A300A-TPM-C	SW,Trusted Platform Module Enabled,A300A,-C
6	2	X6566B-05-R6-C	Cable,Direct Attach CU SFP+ 10G,0.5M,-C
7	8	X66250-5-C	Cable,LC-LC,OM4,5m,-C
8	4	X66034A-C	Cable,12Gb,Mini SAS HD,5m,-C
9	1	X6235-C	Chassis,FAS8200,AFF-A300,AC PS,-C
10	1	DS224C-S-7.6-24S-2P-C	SSD Shelf,12G,24x7.6TB,2P,-C
11	1	DOC-AFF-A300-C	Documents,AFF-A300,-C
12	8	X6599A-R6-C	SFP+ Optical 10Gb Shortwave,-C
13	2	X-02659-00-C	Rail Kit,4-Post,Rnd/Sq-Hole,Adj,24-32,-C
14	4	X800-42U-R6-C	Power Cable,In-Cabinet,C13-C14,-C
15	2	DATA-AT-REST-ENCRYPTION	Data at Rest Encryption Capable Operating Sys
16	1824	SW-FLASH-BUNDLE-2P-C	ONTAP,Per-0.1TB,FlashBundle,Ult-Perf,2P,-C
17	1	PS-DEPLOY-STAND-AFF-M	PS Deployment,Standard,AFF,Med
18	1	CS-WARRANTY-EXTENSION	Warranty Extension
19	1	CS-O2-4HR	SupportEdge Premium 4hr Onsite

Price: $267,000.00

NetApp Disaster Recovery Bill of Materials

Line Item	Qty	Part Number	Product
1	1	SW-2-CL-BASE	SW-2,Base,CL,Node
2	1	AFF-A300	
3	2	AFF-A300A-001	AFF A300 HA System,FlashBundle
4	2	SW-2-A300A-NVE-C	SW,Data at Rest Encryption Enabled,A300A,-C
5	2	SW-2-A300A-TPM-C	SW,Trusted Platform Module Enabled,A300A,-C
6	2	X6566B-05-R6-C	Cable,Direct Attach CU SFP+ 10G,0.5M,-C
7	8	X66250-5-C	Cable,LC-LC,OM4,5m,-C
8	4	X66034A-C	Cable,12Gb,Mini SAS HD,5m,-C
9	1	X6235-C	Chassis,FAS8200,AFF-A300,AC PS,-C
10	1	DS224C-S-7.6-24S-2P-C	SSD Shelf,12G,24x7.6TB,2P,-C
11	1	DOC-AFF-A300-C	Documents,AFF-A300,-C
12	8	X6599A-R6-C	SFP+ Optical 10Gb Shortwave,-C
13	2	X-02659-00-C	Rail Kit,4-Post,Rnd/Sq-Hole,Adj,24-32,-C
14	4	X800-42U-R6-C	Power Cable,In-Cabinet,C13-C14,-C
15	2	DATA-AT-REST-ENCRYPTION	Data at Rest Encryption Capable Operating Sys
16	1824	SW-FLASH-BUNDLE-2P-C	ONTAP,Per-0.1TB,FlashBundle,Ult-Perf,2P,-C
17	1	PS-DEPLOY-STAND-AFF-M	PS Deployment,Standard,AFF,Med
18	1	CS-WARRANTY-EXTENSION	Warranty Extension
19	1	CS-O2-4HR	SupportEdge Premium 4hr Onsite

Price: $267,000.00

Palo Alto Production and Disaster Recovery Bill of Materials

Line No.	Qty	Part Number	Description
1	3	PAN-PA-3320	Palo Alto Networks PA-3220 with redundant AC power supplies
2	3	PAN-PA-3220-TP-5YR	Threat prevention subscription 5-year prepaid, PA-3220
3	3	PAN-PA-3220-URL4-5YR	PANDB URL filtering subscription 5-year prepaid, PA-3220
4	3	PAN-PA-3220-WF-5YR	WildFire subscription 5-year prepaid, PA-3220
5	3	PAN-SVC-PREM-3220-5YR	Premium support 5 year prepaid, PA-3220 3 $

Price: $185,000.00

VMware Production and Disaster Recovery Bill of Materials

Line No.	Qty	Part Number	Description
1	20	VS6-EPL-C-T3	CUSTOMER PURCHASING PROGRAM T3 VMWARE VSPHERE 6 ENTERPRISE PLUS FOR 1 PROCESSOR
2	20	VS6-EPL-3P-SSS-C	PRODUCTION SUPPORT/SUBSCRIPTION VMWARE VSPHERE 6 ENTERPRISE PLUS FOR 1 PROCESSOR FOR 3 YEAR
3	1	VCS6-STD-C-T3	CUSTOMER PURCHASING PROGRAM T3 VMWARE VCENTER SERVER 6 STANDARD FOR VSPHERE 6 (PER INSTANCE)
4	1	VCS6-STD-3P-SSS-C	PRODUCTION SUPPORT/SUBSCRIPTION VMWARE VCENTER SERVER 6 STANDARD FOR VSPHERE 6 (PER INSTANCE) FOR 3 YEAR
5	2	VC-SRM8-25E-C-T3	CUSTOMER PURCHASING PROGRAM T3 VMWARE SITE RECOVERY MANAGER 8 ENTERPRISE (25 VM PACK)
6	2	VC-SRM8-25E-3PSSS-C	PRODUCTION SUPPORT/SUBSCRIPTION FOR VMWARE SITE RECOVERY MANAGER 8 ENTERPRISE (25 VM PACK) FOR 3 YEARS
7	10	VS6-EPL-C-T3	CUSTOMER PURCHASING PROGRAM T3 VMWARE VSPHERE 6 ENTERPRISE PLUS FOR 1 PROCESSOR
8	10	VS6-EPL-3P-SSS-C	PRODUCTION SUPPORT/SUBSCRIPTION VMWARE VSPHERE 6 ENTERPRISE PLUS FOR 1 PROCESSOR FOR 3 YEAR

Price: $203,000.00

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

Updated: January 10, 2019

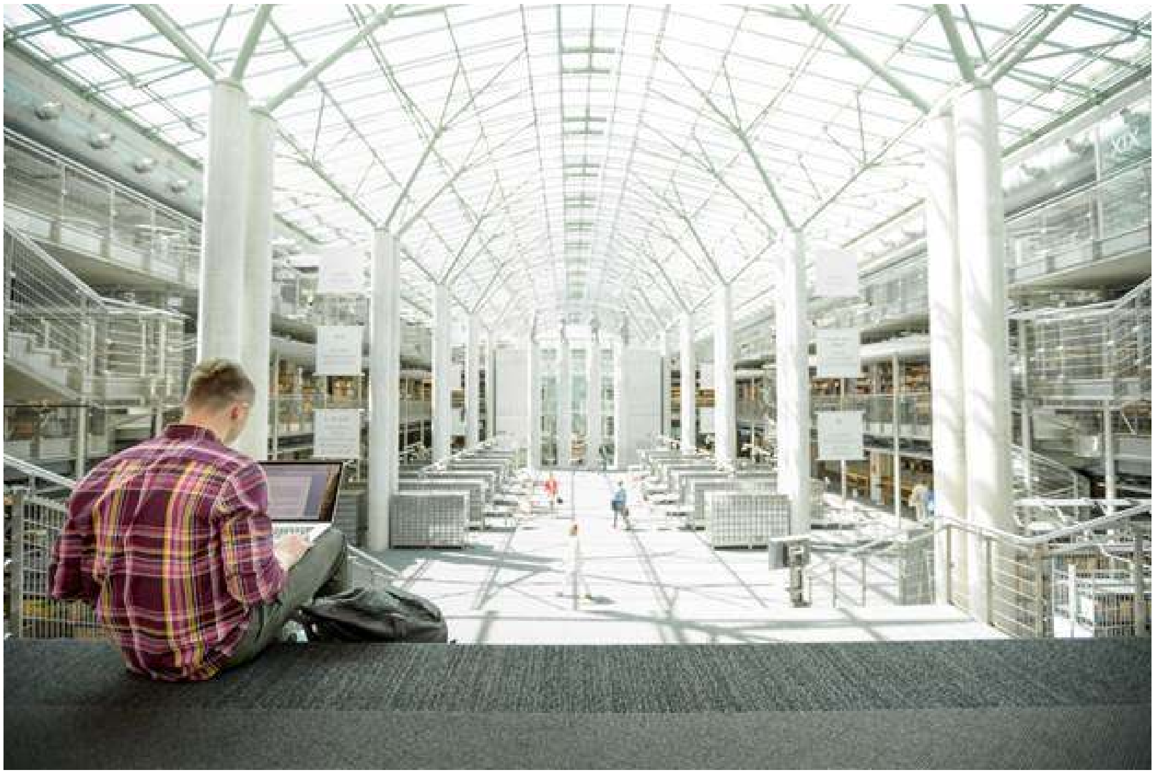

FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series

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

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

Last Updated: January 10, 2019

About the Cisco Validated Design Program

> The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, visit:
>
> http://www.cisco.com/go/designzone.

Table of Contents

Executive Summary

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

Cisco Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

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

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

Solution Overview

Introduction

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

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

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

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, database administrators, IT managers, oracle database architects, and customers who want to deploy Oracle RAC 12cR2 database solution on FlexPod Converged Infrastructure with NetApp clustered Data ONTAP® and the Cisco UCS platform. A working knowledge of

Oracle RAC Database, Linux, Storage technology, and Network is assumed but is not a prerequisite to read this document.

Purpose of this Document

Oracle RAC databases deployments are extremely complicated in nature and customers face enormous challenges in maintaining these landscapes in terms of time, efforts and cost. Oracle RAC databases often manage the mission critical components of a customer's IT department. Ensuring availability while also lowering the IT TCO is always their top priority. This FlexPod solution for Oracle databases delivers industry-leading storage, unprecedented scalability, continuous data access, and automated data management for immediate responses to business opportunities.

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

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

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

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

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

FlexPod System Overview

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

FlexPod provides the following differentiators:

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

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

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

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

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

Figure 1 FlexPod System Overview

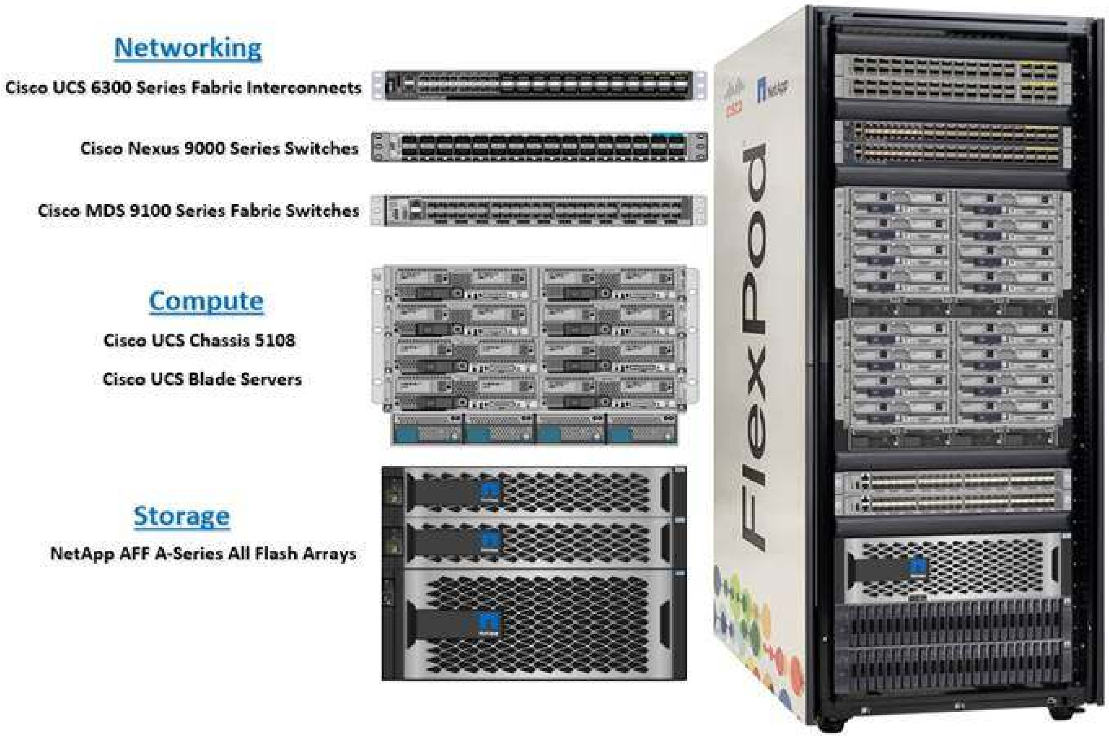
Networking
Cisco UCS 6300 Series Fabric Interconnects
Cisco Nexus 9000 Series Switches
Cisco MDS 9100 Series Fabric Switches
Compute
Cisco UCS Chassis 5108
Cisco UCS Blade Servers
Storage
NetApp AFF A-Series All Flash Arrays
FlexPod

FlexPod datacenter architecture includes three components:

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

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

What’s New in this Release?

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

It incorporates the following features:

- Support for the Cisco UCS 3.2 unified software release and Cisco UCS B200 M5 Blade Servers
- Support for the latest release of NetApp ONTAP® 9.3
- Fibre channel storage design
- Validation of Oracle RAC Database 12c Release 2

Solution Components

Cisco UCS 5108 Blade Server Chassis

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

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

Cisco UCS 2304 Fabric Extender

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

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

Cisco UCS B200 M5 Blade Servers

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

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

Cisco UCS Virtual Interface Card (VIC) 1340

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

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

Cisco UCS 6332-16UP Fabric Interconnect

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

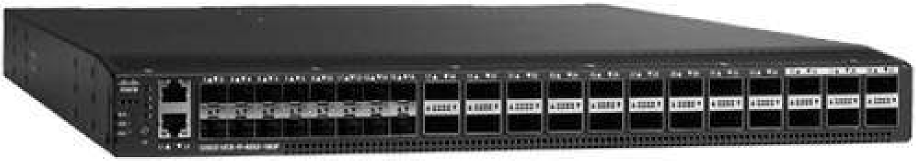

The Fabric Interconnect provides high-speed upstream connectivity to the network, or converged traffic to servers through its 40 Gbps ports, but also allows for Fibre Channel connectivity to SAN switches like the MDS, or alternately directly attached Fibre Channel to storage arrays.

Cisco UCS MDS 9148S Fabric Switch

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

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

Cisco Nexus 93180LC-EX Switch

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

NetApp AFF All Flash Array A700s

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

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

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

- HA and fault tolerance. Non-disruptive operations are achieved through clustering, HA pairing of controllers, hot-swappable components, NetApp RAID-TEC® and disk protection (allowing two independent disk failures without data loss), and network interface redundancy.

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

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

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

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

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

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

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

Solution Design

Physical Infrastructure

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

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

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

Physical Topology

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

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

Figure 2 FlexPod Architecture Design

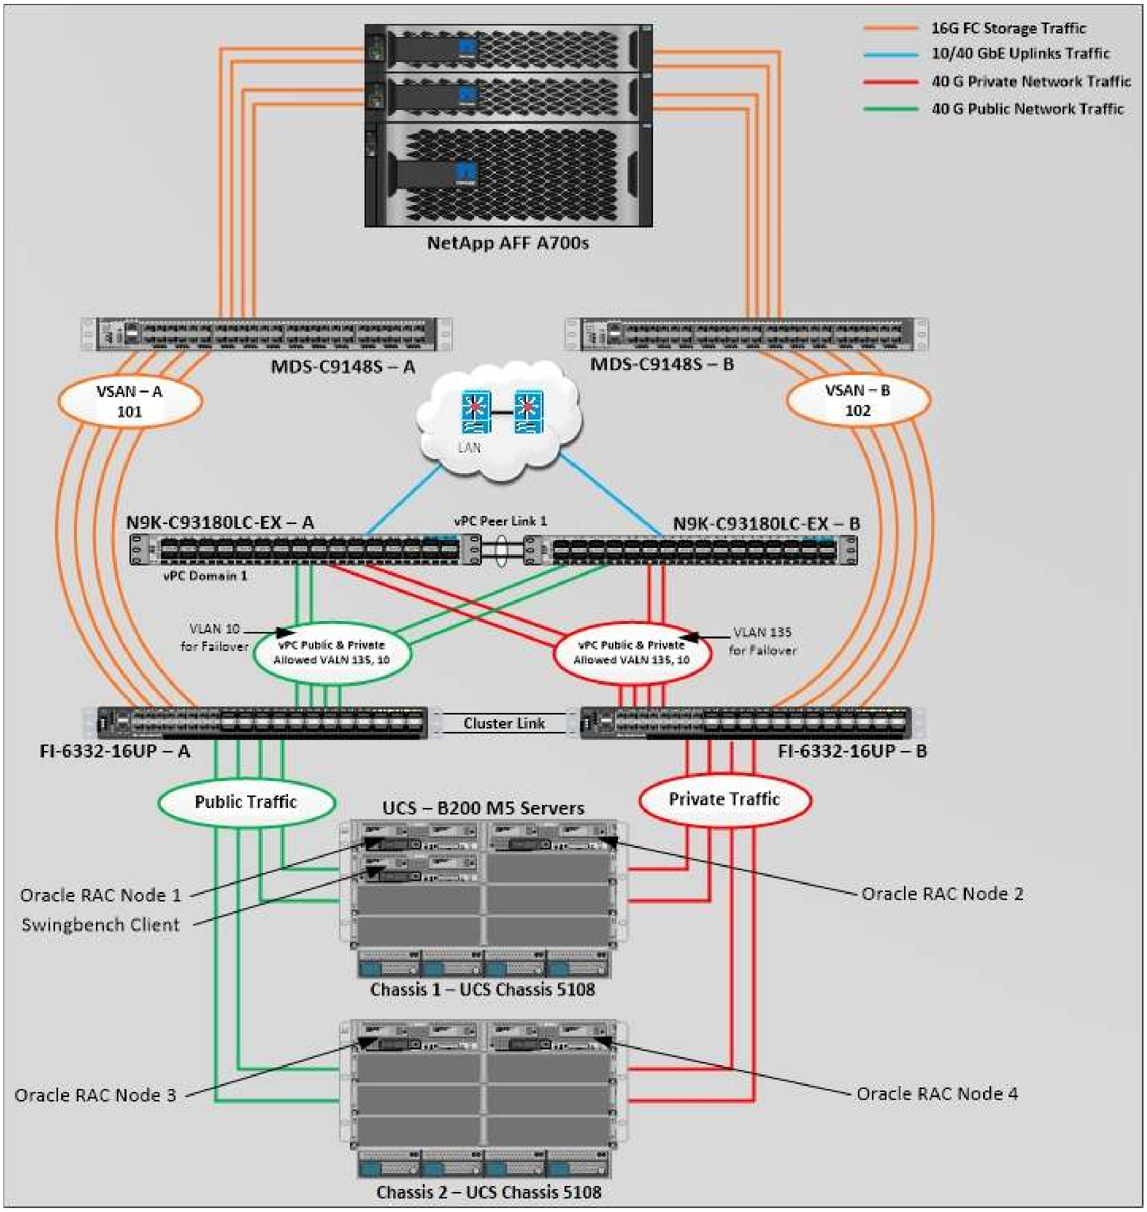

Figure 2 details the cable connections used in the validation lab for the 40Gb end-to-end with Fibre Channel topology based on the Cisco UCS 6332-16UP fabric interconnect. As shown in above architecture, a pair of Cisco UCS 6332-16UP Fabric Interconnects (FI) carries both storage and network traffic from the Cisco UCS B200 M5 server blades with the help of Cisco Nexus 93180LC-EX and Cisco MDS 9148S switches. Both the Fabric Interconnects and the Cisco Nexus switches are clustered with the peer link between them to provide high availability. Two virtual Port-Channels (vPCs) are configured to provide public network and private network paths for the server blades to northbound switches. Each vPC has VLANs created for application network data and management data paths.

As illustrated in the above architecture, four (2 x 40G link per chassis) links go to Fabric Interconnect - A. Similarly, four (2 x 40G link per chassis) links go to Fabric Interconnect - B. Fabric Interconnect - A links are used for Oracle Public network traffic shown as green lines. Fabric Interconnect - B links are used for Oracle private interconnect traffic shown as red lines. FC Storage access from Fabric Interconnect - A and Fabric Interconnect - B shown as orange lines.

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

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

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

The reference architecture includes the following hardware:

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

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

Design Topology

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

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

Table 1 Hardware Components

Name	Model	Description	Quantit
Cisco Nexus 93180LC-EX Switch	N9K-C93180LC-EX	Cisco Nexus 9300 Series Switches	2
Cisco MDS 9148S 16G Fabric Switch	DS-C9148S-12PK9	Cisco MDS 9100 Series Multilayer Fabric Switches	2
Cisco UCS 6332-16UP Fabric Interconnect	UCS-FI-6332-16UP	Cisco 6300 Series Fabric Interconnects	2
Cisco UCS 5108 Blade Server Chassis	UCSB-5108-AC2	Cisco UCS 5100 Series Blade Server AC2 Chassis	2
Cisco UCS Fabric Extender	UCS-IOM-2304	Cisco UCS 2304XP I/O Module (4 External, 8 Internal 40Gb Ports)	4
Cisco UCS B200 M5 Blade Server	UCSB-B200-M5	Cisco UCS B-Series Blade Servers	4
Cisco UCS VIC 1340	UCSB-MLOM-40G-03	Cisco UCS Virtual Interface Card 1340	4
Cisco UCS Port Expander	UCSB-MLOM-	Port Expander Card for Cisco UCS	4

Name	Model	Description	Quantit
NetApp Storage	AFF A700s	AFF A-Series All Flash Arrays	1
NetApp Disk	DS-224C	Disk Shelves and Storage Media for NetApp AFF and FAS systems	2

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

Table 2 Cisco UCS B200 M5 Blade Server Configuration

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

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

Table 3 VLAN and VSAN Configurations

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

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

Table 4 NetApp AFF A700s Storage Configuration

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

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

Table 5 Software and Firmware Revisions

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

Solution Configuration

Configure Cisco Nexus 93180LC-EX Switch

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

Set Up Initial Configuration

Cisco Nexus A

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

1. Configure the switch.

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

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

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

Enter the password for "admin": <password>

Confirm the password for "admin": <password>

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

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

Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [noshut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus B

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

1. Configure the switch.

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

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

```
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [noshut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

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

Configure Global Settings for Cisco Nexus A and Cisco Nexus B

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

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

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

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

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

Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

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

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

```
configure terminal

vlan 135

name Oracle_RAC_Public_Traffic

no shutdown

vlan 10

name Oracle_RAC_Private_Traffic

no shutdown

copy run start
```

3. Log in as admin user into the Nexus Switch B and create VLAN 135 for Public Network Traffic and VLAN 10 for Private Network Traffic.

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

In the Cisco Nexus 93180LC-EX switch topology, a single vPC feature is enabled to provide HA, faster convergence in the event of a failure, and greater throughput. Cisco Nexus vPC configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers is shown in Table 6

Table 6 vPC Summary

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

As listed in Table 6 , a single vPC domain with Domain ID 1 is created across two Cisco Nexus 93180LC-EX member switches to define vPC members to carry specific VLAN network traffic. In this topology, we defined a total number of 3 vPCs.

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

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

Create vPC Peer-Link Between the Two Nexus Switches

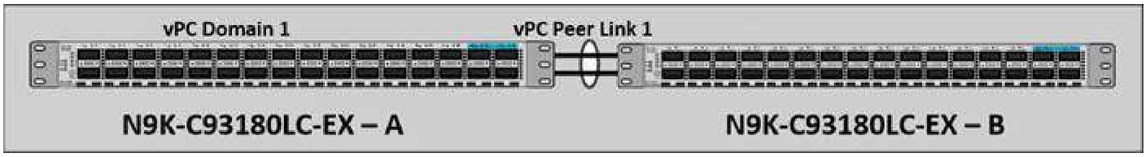

1. Login as admin user into the Nexus Switch A.

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

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

```
configure terminal
```

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

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

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

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

3. Log in as admin user into the Nexus Switch B and repeat the above steps to configure second nexus switch.

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

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

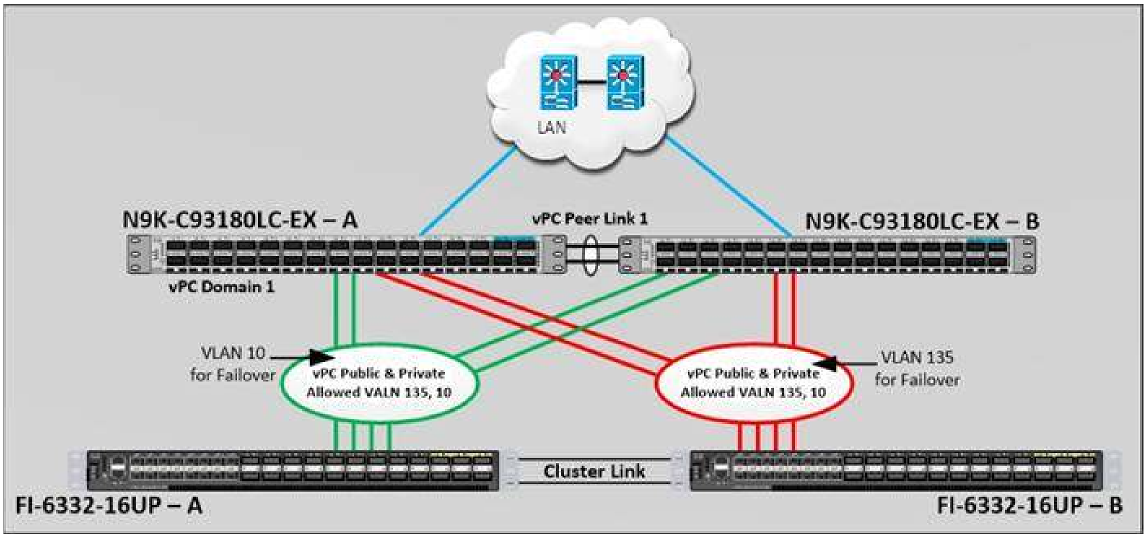

Next, you will create and configure vPC 51 and 52 for Data network between Nexus switches and Fabric Interconnects.

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

Table 7 vPC IDs and VLAN IDs

vPC Description	vPC ID	Fabric Interconnects Ports	Nexus Ports	Allowed VLANs
Port Channel FI-A	51	FI-A Port 1/31	N9K-A Port 1/21	135, 10 Note: VLAN 10 Needed for Failover
		FI-A Port 1/32	N9K-A Port 1/22	
		FI-A Port 1/33	N9K-B Port 1/21	
		FI-A Port 1/34	N9K-B Port 1/22	
Port Channel FI-B	52	FI-B Port 1/31	N9K-A Port 1/23	10, 135 Note: VLAN 135 Needed for Failover
		FI-B Port 1/32	N9K-A Port 1/24	
		FI-B Port 1/33	N9K-B Port 1/23	
		FI-B Port 1/34	N9K-B Port 1/24	

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

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

```
configure terminal

interface port-channel51

description Port-Channel FI-A

switchport mode trunk

switchport trunk allowed vlan 1,10,135

spanning-tree port type edge trunk
```

```
mtu 9216
vpc 51
no shutdown

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

interface Ethernet1/21
  description Fabric-Interconnect-A-31
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active
  no shutdown

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

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

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

copy run start
```

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

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

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

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

interface Ethernet1/22
  description Fabric-Interconnect-A-34
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 51 mode active
  no shutdown

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

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

copy run start
```

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

Verify All vPC Status Is Up on Both the Nexus Switches

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

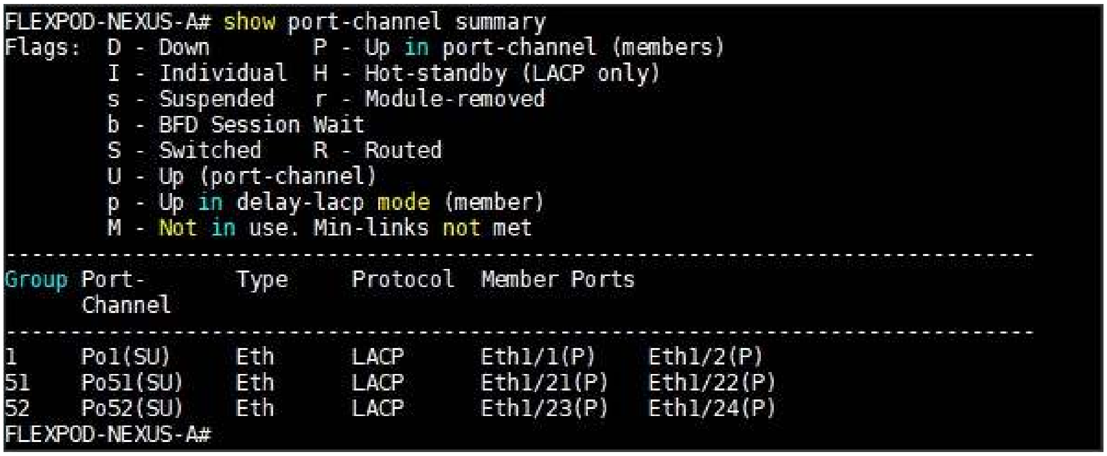

```
FLEXPOD-NEXUS-A# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        b - BFD Session Wait
        S - Switched    R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
51    Po51(SU)    Eth      LACP      Eth1/21(P)   Eth1/22(P)
52    Po52(SU)    Eth      LACP      Eth1/23(P)   Eth1/24(P)
FLEXPOD-NEXUS-A#
```

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

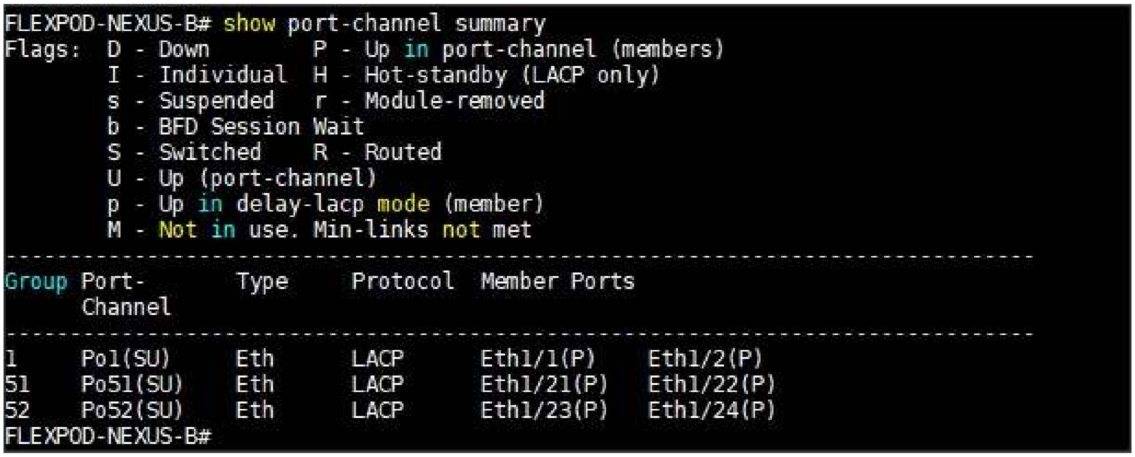

```
FLEXPOD-NEXUS-B# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        b - BFD Session Wait
        S - Switched    R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
51    Po51(SU)    Eth      LACP      Eth1/21(P)   Eth1/22(P)
52    Po52(SU)    Eth      LACP      Eth1/23(P)   Eth1/24(P)
FLEXPOD-NEXUS-B#
```

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

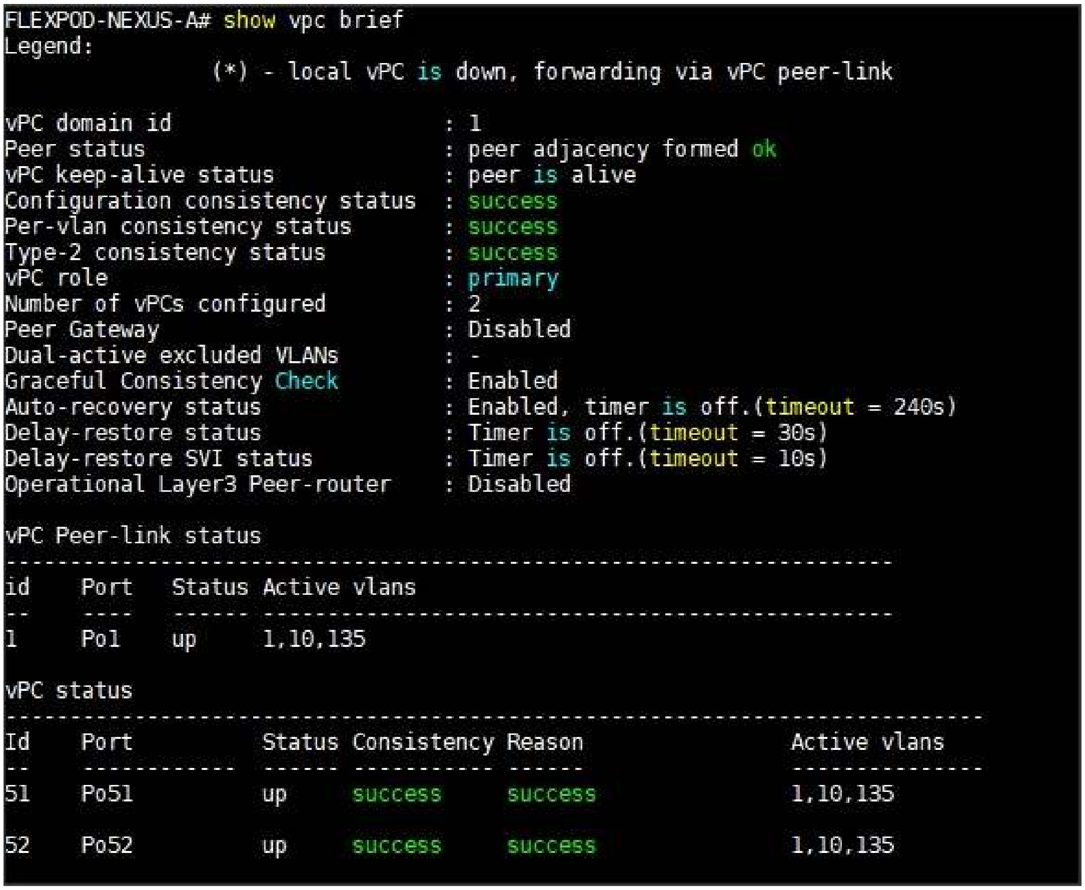

```
FLEXPOD-NEXUS-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled

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

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

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

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

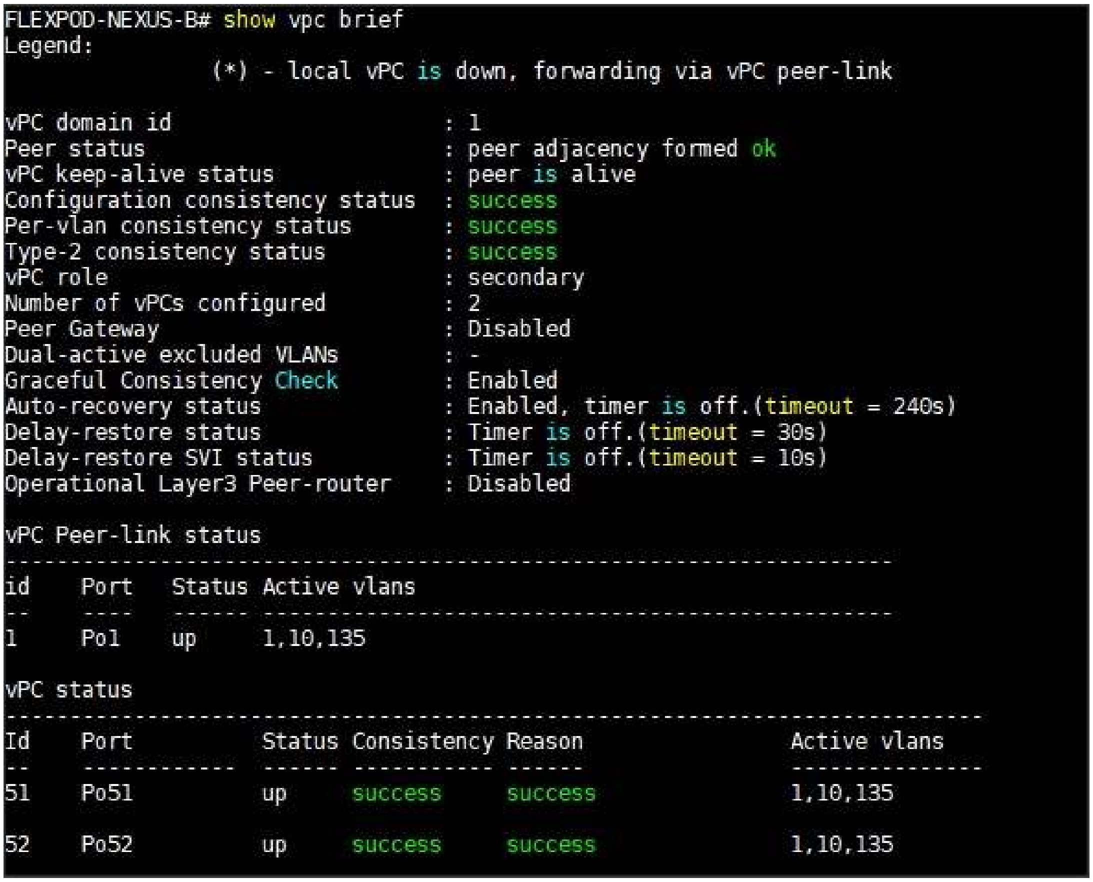

```
FLEXPOD-NEXUS-B# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled

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

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

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

Cisco UCS Configuration Overview

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

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

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

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

Cisco UCS Fabric Interconnect A and Fabric Interconnect B

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

1. Verify the following physical connections on the fabric interconnect:

a. The management Ethernet port (mgmt0) is connected to an external hub, switch, or router

b. The L1 ports on both fabric interconnects are directly connected to each other

c. The L2 ports on both fabric interconnects are directly connected to each other

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

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

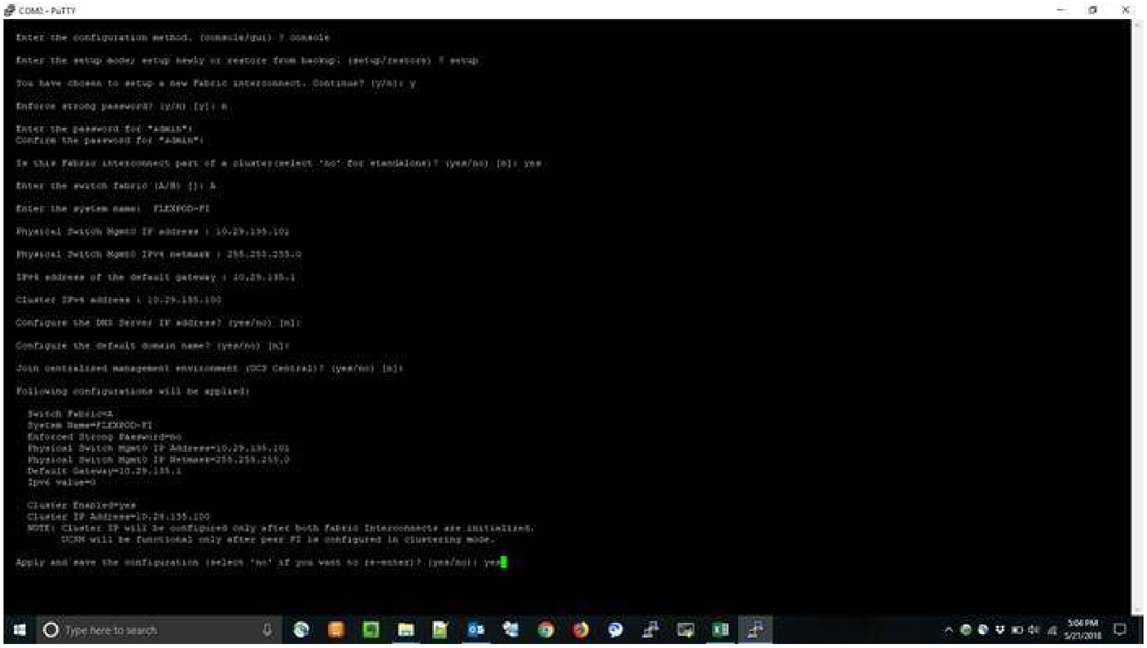

```
Enter the configuration method. (console/gui) ? console

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

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

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

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

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

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

Enter the system name:  FLEXPOD-FI

Physical Switch Mgmt0 IP address : 10.29.135.101

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.29.135.1

Cluster IPv4 address : 10.29.135.100

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

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

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

Following configurations will be applied:

  Switch Fabric=A
  System Name=FLEXPOD-FI
  Enforced Strong Password=no
  Physical Switch Mgmt0 IP Address=10.29.135.101
  Physical Switch Mgmt0 IP Netmask=255.255.255.0
  Default Gateway=10.29.135.1
  Ipv6 value=0

  Cluster Enabled=yes
  Cluster IP Address=10.29.135.100
  NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized.
        UCSM will be functional only after peer FI is configured in clustering mode.

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

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

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

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

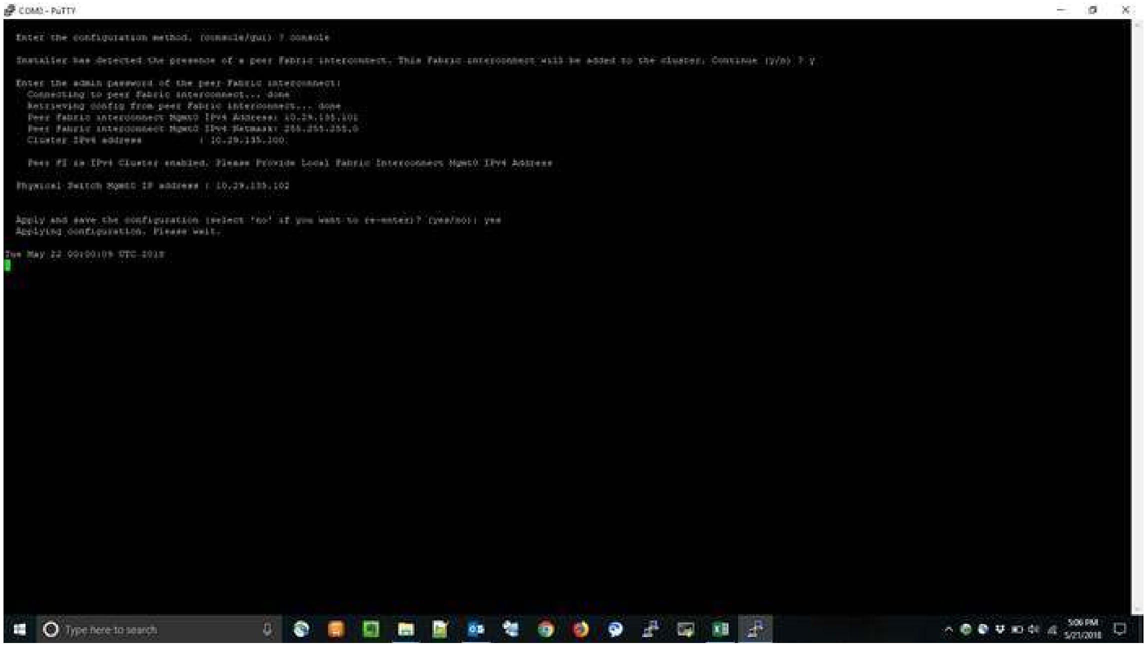

```
Enter the configuration method. (console/gui) ? console

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

Enter the admin password of the peer Fabric interconnect:
  Connecting to peer Fabric interconnect... done
  Retrieving config from peer Fabric interconnect... done
  Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.135.101
  Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
  Cluster IPv4 address          : 10.29.135.100

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

Physical Switch Mgmt0 IP address : 10.29.135.102


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

Tue May 22 00:00:09 UTC 2018
```

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

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

Log in to Cisco UCS Manager

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

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

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

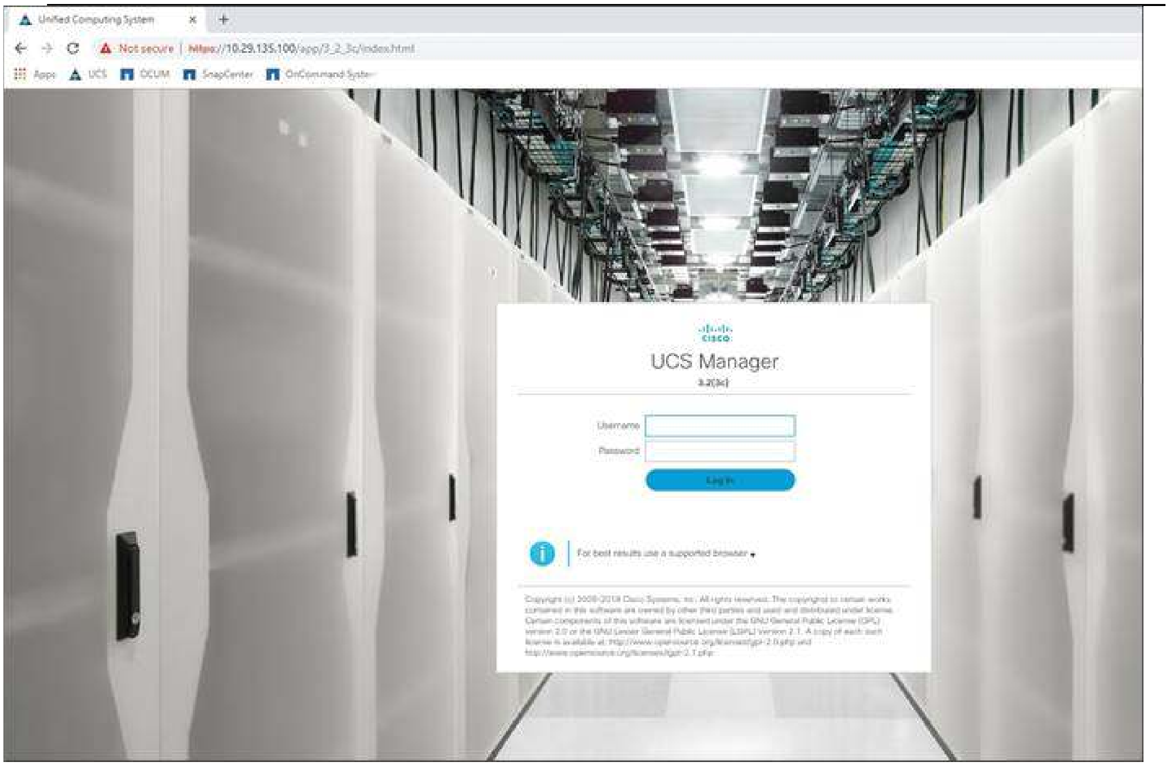

Copyright (c) 2009-2018 Cisco Systems, Inc. All rights reserved. The copyrights to certain works contained in this software are owned by other third parties and used and distributed under license. Certain components of this software are licensed under the GNU General Public License (GPL) version 2.0 or the GNU Lesser General Public License (LGPL) Version 2.1. A copy of each such license is available at: http://www.opensource.org/licenses/gpl-2.0.php and http://www.opensource.org/licenses/lgpl-2.1.php

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

Upgrade Cisco UCS Manager Software to Version 3.2

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

Anonymous Reporting

To create anonymous reporting, complete the following step:

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

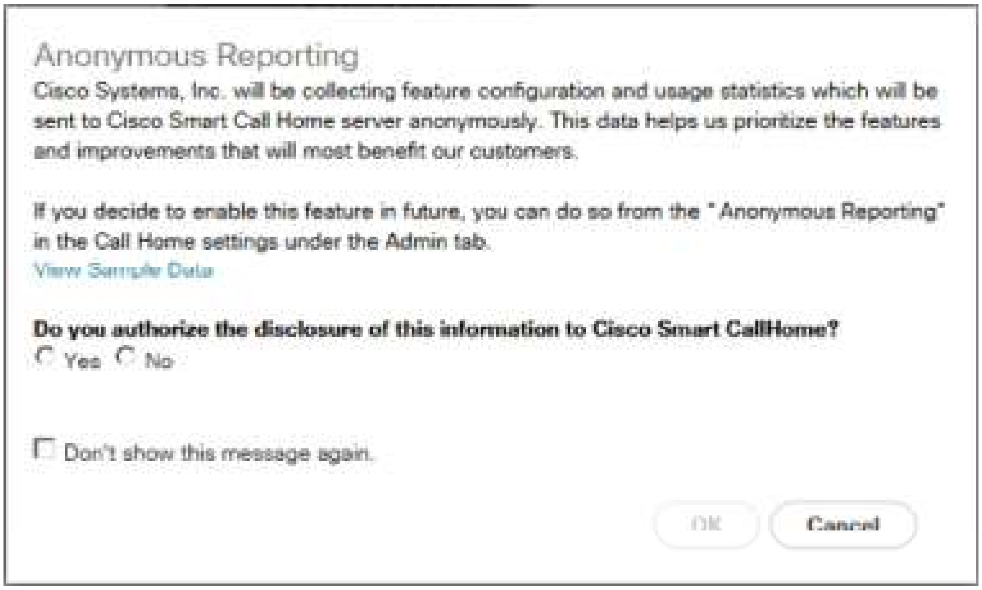

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.

If you decide to enable this feature in future, you can do so from the "Anonymous Reporting" in the Call Home settings under the Admin tab.

View Sample Data

Do you authorize the disclosure of this information to Cisco Smart CallHome?

○ Yes ○ No

□ Don't show this message again.

Configure Cisco UCS Call Home

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

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

High Level Steps to Configure Base Cisco UCS

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

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

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

a. Configure Global Policies

b. Configure Server Ports

4. Configure LAN and SAN on Cisco UCS Manager

a. Configure Ethernet LAN Uplink Ports

b. Create Uplink Port Channels to Cisco Nexus Switches

c. Configure FC SAN Uplink Ports

d. Configure VLAN

e. Configure VSAN

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

a. IP Pool Creation

b. UUID Suffix Pool Creation

c. Server Pool Creation

d. MAC Pool Creation

e. WWNN and WWPN Pool Creation

6. Set Jumbo Frames in both the Cisco Fabric Interconnect
7. Configure Server BIOS Policy
8. Create Adapter Policy
9. Configure Update Default Maintenance Policy
10. Configure vNIC and vHBA Template

a. Create Public and Private vNIC Template

b. Create Storage vHBA Template

11. Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

Set Fabric Interconnects to Fibre Channel End Host Mode

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

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

3. Follow the dialogs to complete the change.

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

Synchronize Cisco UCS to NTP

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

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

Configure Fabric Interconnect for Chassis and Blade Discovery

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

Configure Global Policies

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

To configure Global Policies, follow these steps:

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

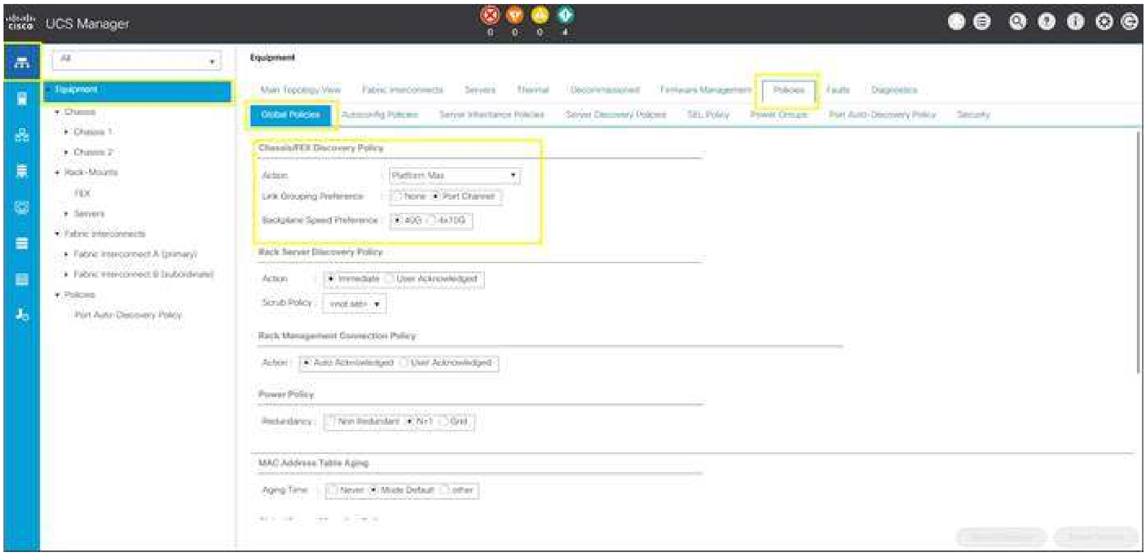

Chassis/FEX Discovery Policy

Action: Platform Max

Link Grouping Preference: ○ None ● Port Channel

Backplane Speed Preference: ● 40G ○ 4x10G

Rack Server Discovery Policy

Action: ● Immediate ○ User Acknowledged

Scrub Policy: <not set>

Rack Management Connection Policy

Action: ● Auto Acknowledged ○ User Acknowledged

Power Policy

Redundancy: ○ Non Redundant ● N+1 ○ Grid

MAC Address Table Aging

Aging Time: ○ Never ● Mode Default ○ other

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

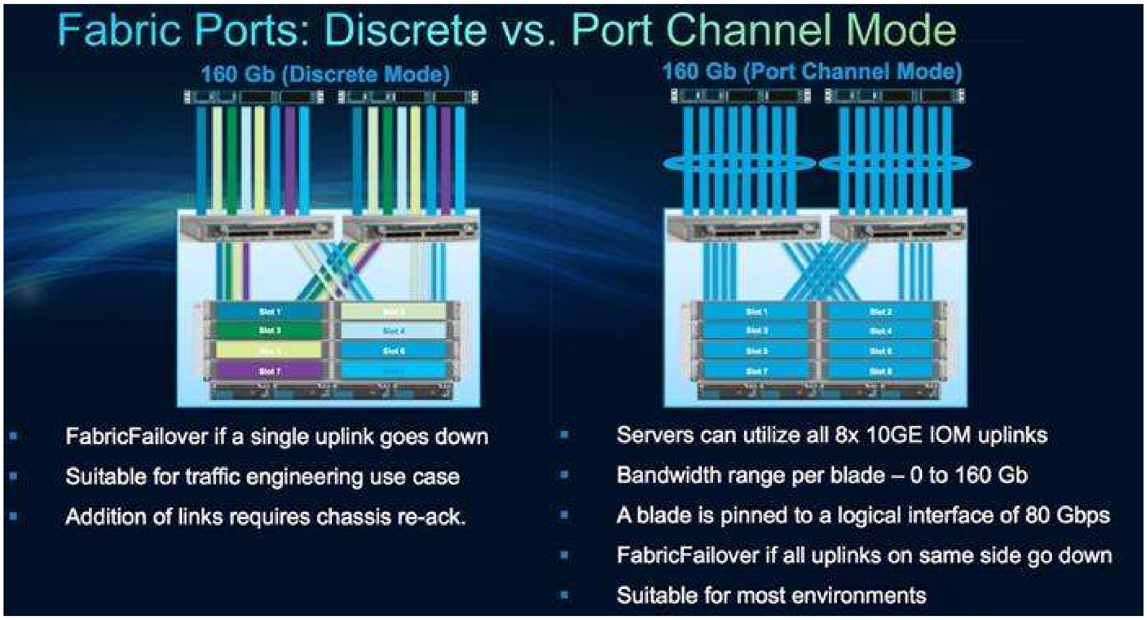

Fabric Ports: Discrete vs. Port Channel Mode

- FabricFailover if a single uplink goes down
- Suitable for traffic engineering use case
- Addition of links requires chassis re-ack.

- Servers can utilize all 8x 10GE IOM uplinks
- Bandwidth range per blade – 0 to 160 Gb
- A blade is pinned to a logical interface of 80 Gbps
- FabricFailover if all uplinks on same side go down
- Suitable for most environments

Configure Server Ports

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

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

2. Select the ports (for this solution ports are 17-20) which are connected to the Cisco IO Modules of the two B-Series 5108 Chassis.

3. Right-click and select “Configure as Server Port.”

4. Click Yes to confirm and click OK.

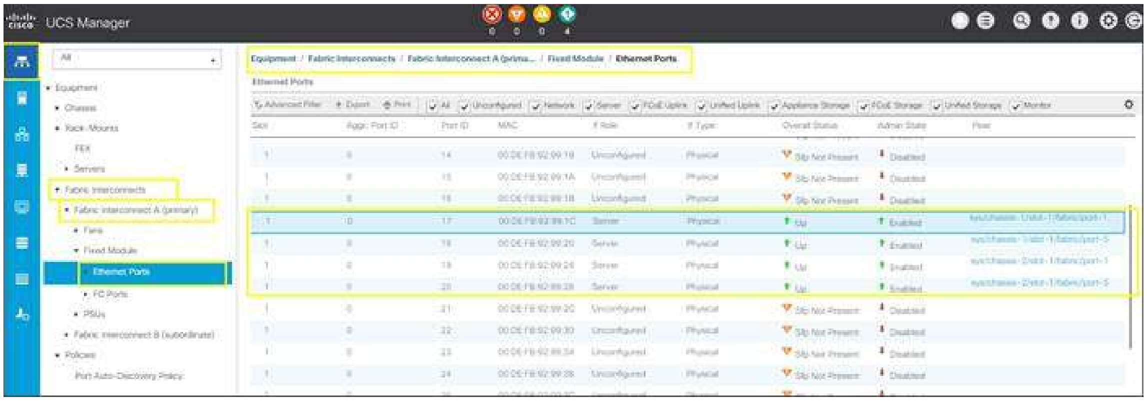

Equipment / Fabric Interconnects / Fabric Interconnect A (prima... / Fixed Module / Ethernet Ports

Ethernet Ports

☑ All ☑ Unconfigured ☑ Network ☑ Server ☑ FCoE Uplink ☑ Unified Uplink ☑ Appliance Storage ☑ FCoE Storage ☑ Unified Storage ☑ Monitor

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	14	00:DE:FB:92:99:19	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	15	00:DE:FB:92:99:1A	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	16	00:DE:FB:92:99:1B	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	17	00:DE:FB:92:99:1C	Server	Physical	Up	Enabled	sys/chassis-1/slot-1/fabric/port-1
1	0	18	00:DE:FB:92:99:20	Server	Physical	Up	Enabled	sys/chassis-1/slot-1/fabric/port-5
1	0	19	00:DE:FB:92:99:24	Server	Physical	Up	Enabled	sys/chassis-2/slot-1/fabric/port-1
1	0	20	00:DE:FB:92:99:28	Server	Physical	Up	Enabled	sys/chassis-2/slot-1/fabric/port-5
1	0	21	00:DE:FB:92:99:2C	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	22	00:DE:FB:92:99:30	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	23	00:DE:FB:92:99:34	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	24	00:DE:FB:92:99:38	Unconfigured	Physical	Sfp Not Present	Disabled	

5. Repeat the same task for Fabric Interconnect B.

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

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

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

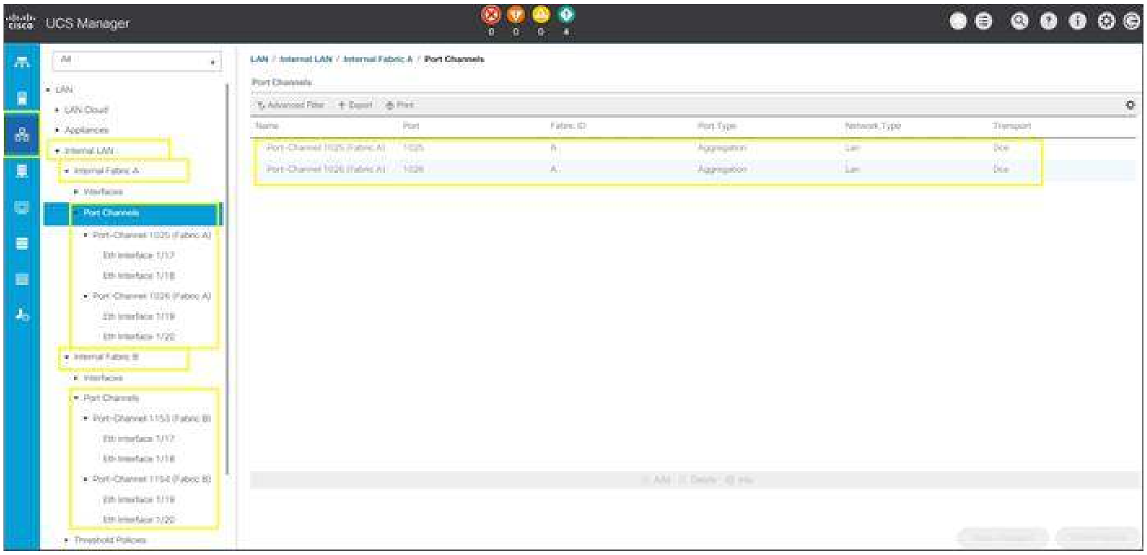

LAN / Internal LAN / Internal Fabric A / Port Channels

Port Channels

Name	Port	Fabric ID	Port Type	Network Type	Transport
Port-Channel 1025 (Fabric A)	1025	A	Aggregation	Lan	Dce
Port-Channel 1026 (Fabric A)	1026	A	Aggregation	Lan	Dce

9. Verify the same for Internal Fabric B.

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

Configure LAN and SAN on Cisco UCS Manager

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

Configure Ethernet LAN Uplink Ports

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

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports (for this solution ports are 31-34) that are connected to the Nexus switches, right-click them, and select Configure as Network Port.
5. Click Yes to confirm ports and click OK.
6. Verify the Ports connected to Cisco Nexus upstream switches are now configured as network ports.
7. Repeat the above steps for Fabric Interconnect B. The screenshot shows the network uplink ports for Fabric A.

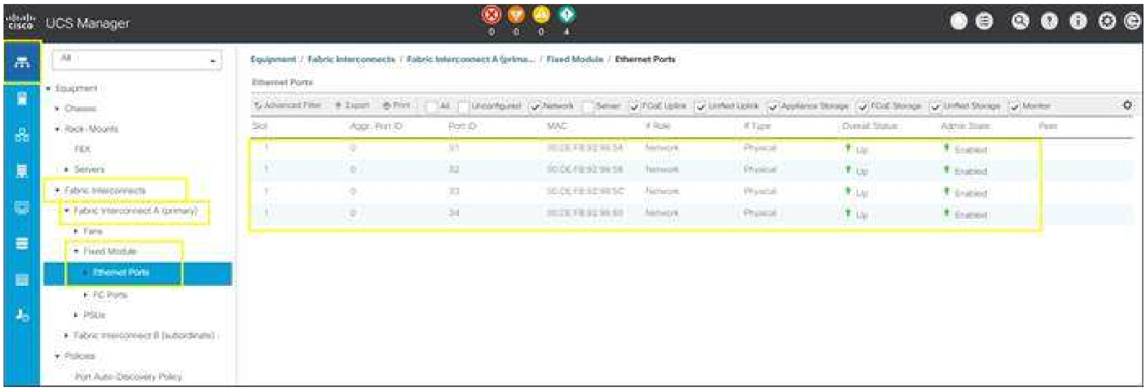

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	31	00:DE:FB:92:98:54	Network	Physical	↑ Up	↑ Enabled	
1	0	32	00:DE:FB:92:98:58	Network	Physical	↑ Up	↑ Enabled	
1	0	33	00:DE:FB:92:98:5C	Network	Physical	↑ Up	↑ Enabled	
1	0	34	00:DE:FB:92:98:60	Network	Physical	↑ Up	↑ Enabled	

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

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

Create Uplink Port Channels to Cisco Nexus Switches

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, expand node Fabric A tree:
3. Right-click Port Channels.
4. Select Create Port Channel.

5. Enter 51 as the unique ID of the port channel.

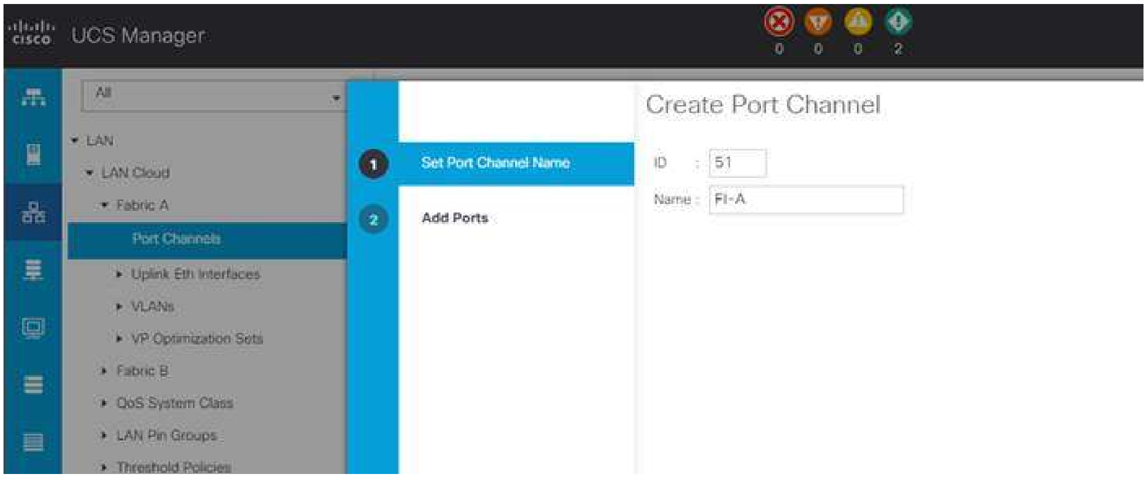

Create Port Channel

ID : 51

Name : FI-A

6. Enter FI-A as the name of the port channel.
7. Click Next.
8. Select Ethernet ports 31-34 for the port channel.
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel and then click OK.

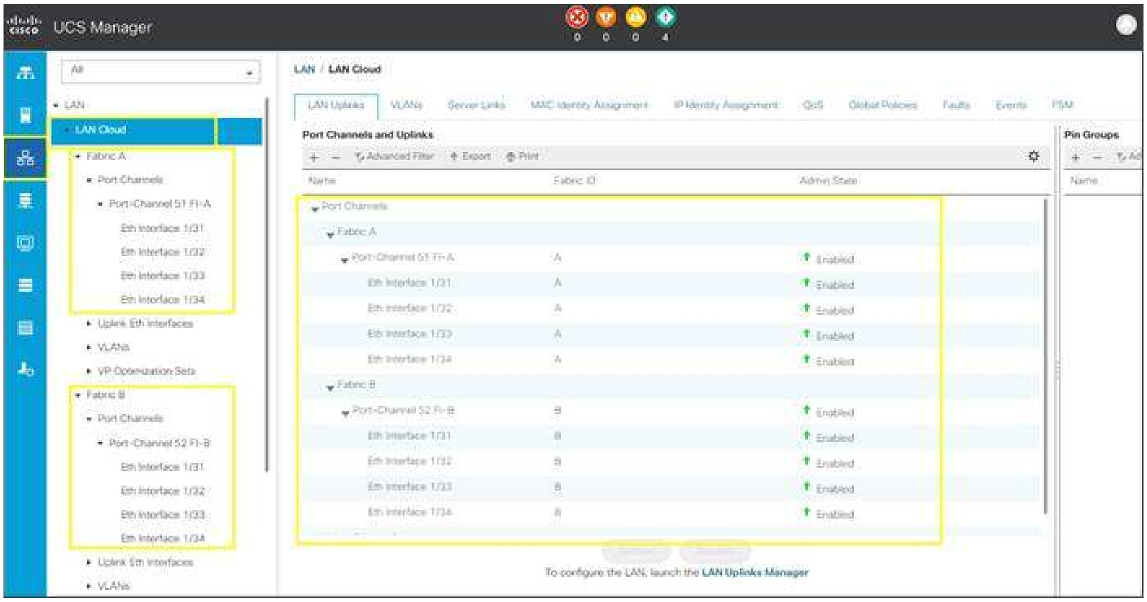

Port Channels and Uplinks

Name	Fabric ID	Admin State
Port Channels		
Fabric A		
Port-Channel 51 FI-A	A	Enabled
Eth Interface 1/31	A	Enabled
Eth Interface 1/32	A	Enabled
Eth Interface 1/33	A	Enabled
Eth Interface 1/34	A	Enabled
Fabric B		
Port-Channel 52 FI-B	B	Enabled
Eth Interface 1/31	B	Enabled
Eth Interface 1/32	B	Enabled
Eth Interface 1/33	B	Enabled
Eth Interface 1/34	B	Enabled

To configure the LAN, launch the LAN Uplinks Manager

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

Configure FC SAN Uplink Ports

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

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

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. Select Configure Unified Ports.

4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.

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

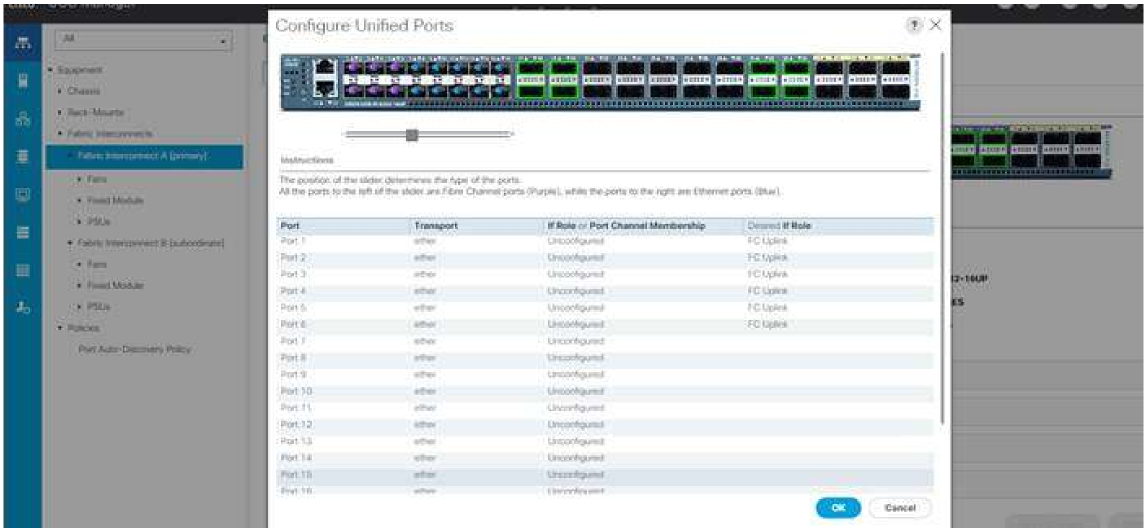

Configure Unified Ports

Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	FC Uplink
Port 6	ether	Unconfigured	FC Uplink
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Unconfigured	

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

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

7. Select Equipment > Fabric Interconnects > Fabric Interconnect B (primary).

8. Select Configure Unified Ports.

9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.

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

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

12. Wait for both Fabric Interconnects to reboot.

13. Log back into Cisco UCS Manager.

Configure VLAN

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

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

> In this solution, we created two VLANs: one for private network (VLAN 10) traffic and one for public network (VLAN 135) traffic. These two VLANs will be used in the vNIC templates that are discussed later.

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

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

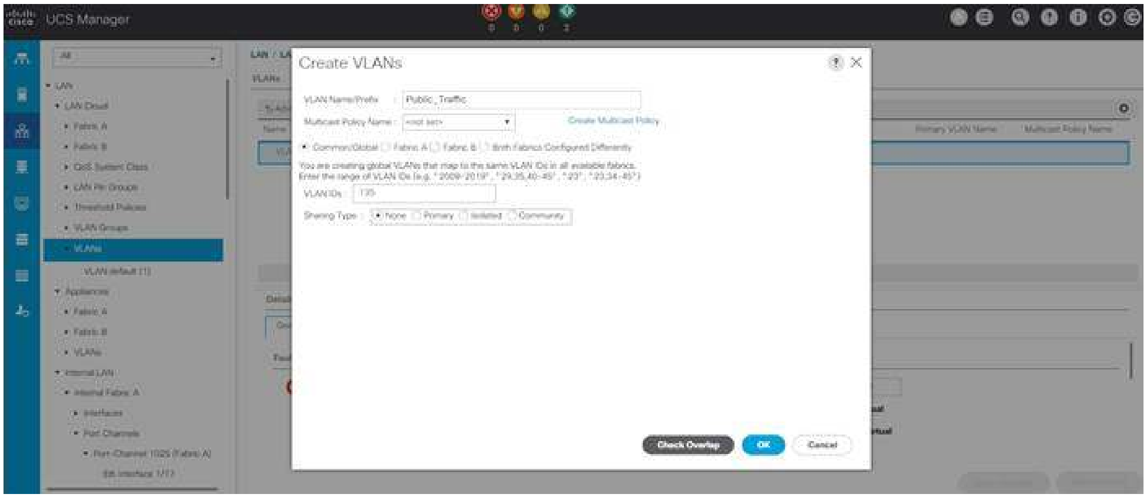

Create VLANs

VLAN Name/Prefix : Public_Traffic

Multicast Policy Name : <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. " 2009-2019" , " 29,35,40-45" , " 23" , " 23,34-45")

VLAN IDs : 135

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

5. Enter Public_Traffic as the name of the VLAN to be used for Public Network Traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 135 as the ID of the VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.

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

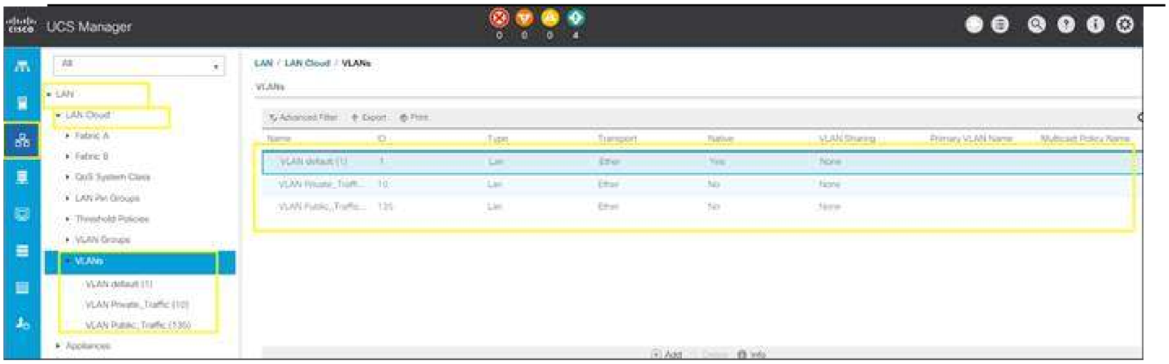

LAN / LAN Cloud / VLANs

VLANs

Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Name	Multicast Policy Name
VLAN default (1)	1	Lan	Ether	Yes	None		
VLAN Private_Traff...	10	Lan	Ether	No	None		
VLAN Public_Traffic...	135	Lan	Ether	No	None		

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

Configure VSAN

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

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

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

2. Select SAN > SAN Cloud.
3. Under VSANs, right-click VSANs.
4. Select Create VSANs.
5. Enter ORA-VSAN-A as the name of the VSAN.
6. Leave FC Zoning set at Disabled.

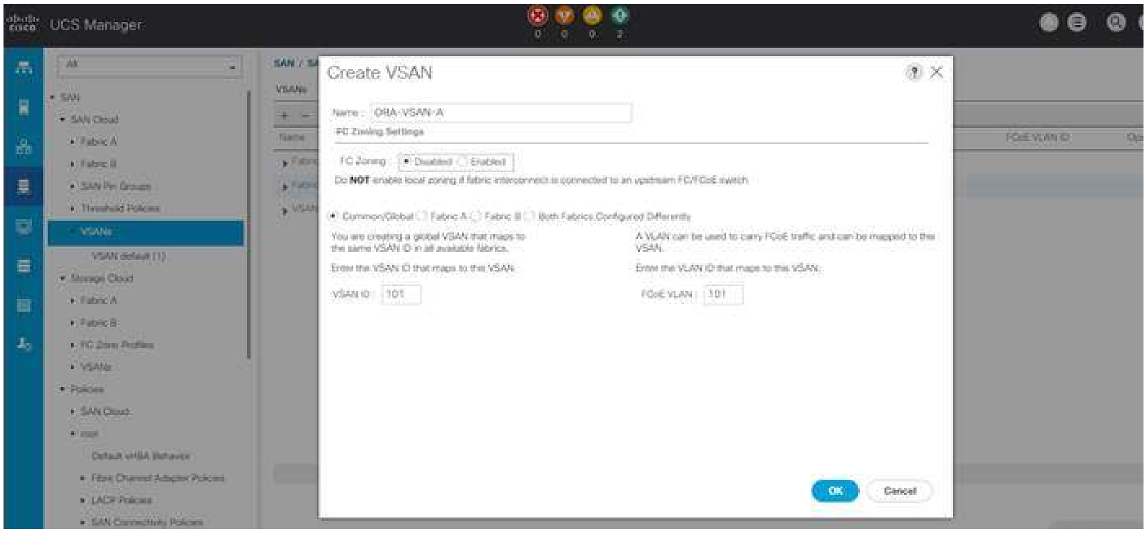

Create VSAN

Name : ORA-VSAN-A

FC Zoning Settings

FC Zoning ● Disabled ○ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating a global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 101

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN : 101

7. Select Fabric A for the scope of the VSAN.
8. Enter 101 as the ID of the VSAN.

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

9. Click OK and then click OK again.

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

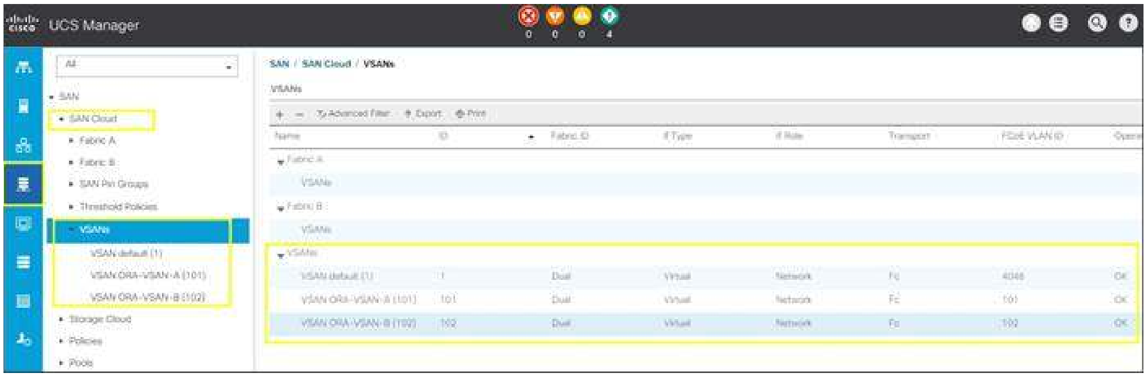

SAN / SAN Cloud / VSANs

VSANs

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

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

IP Pool Creation

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

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

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

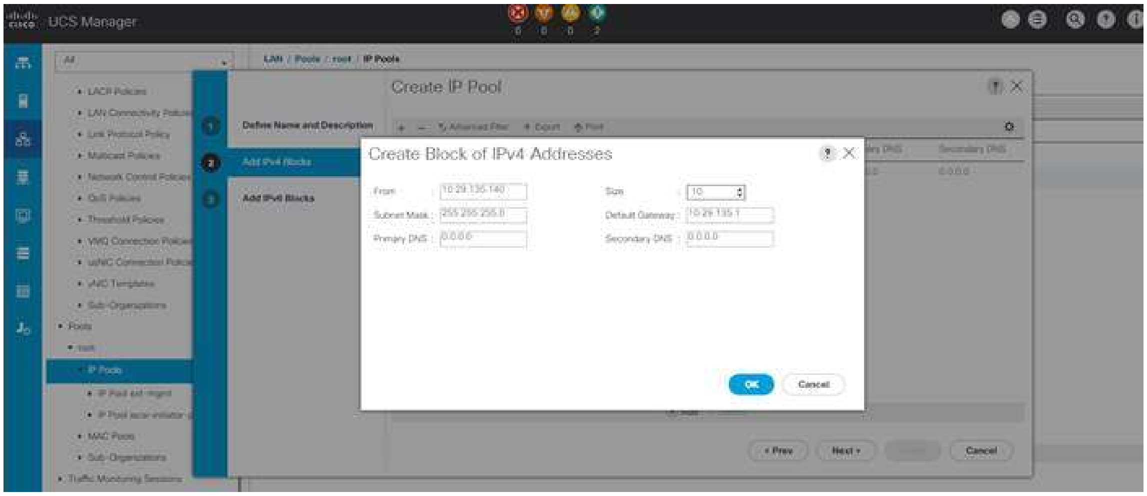

Create Block of IPv4 Addresses

From : 10.29.135.140
Size : 10
Subnet Mask : 255.255.255.0
Default Gateway : 10.29.135.1
Primary DNS : 0.0.0.0
Secondary DNS : 0.0.0.0

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

UUID Suffix Pool Creation

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools and then select Create UUID Suffix Pool.
4. Enter ORA-UUID-Pool as the name of the UUID name.
5. Optional: Enter a description for the UUID pool.
6. Keep the prefix at the derived option and select Sequential in as Assignment Order then click Next.
7. Click Add to add a block of UUIDs
8. Create a starting point UUID as per your environment.

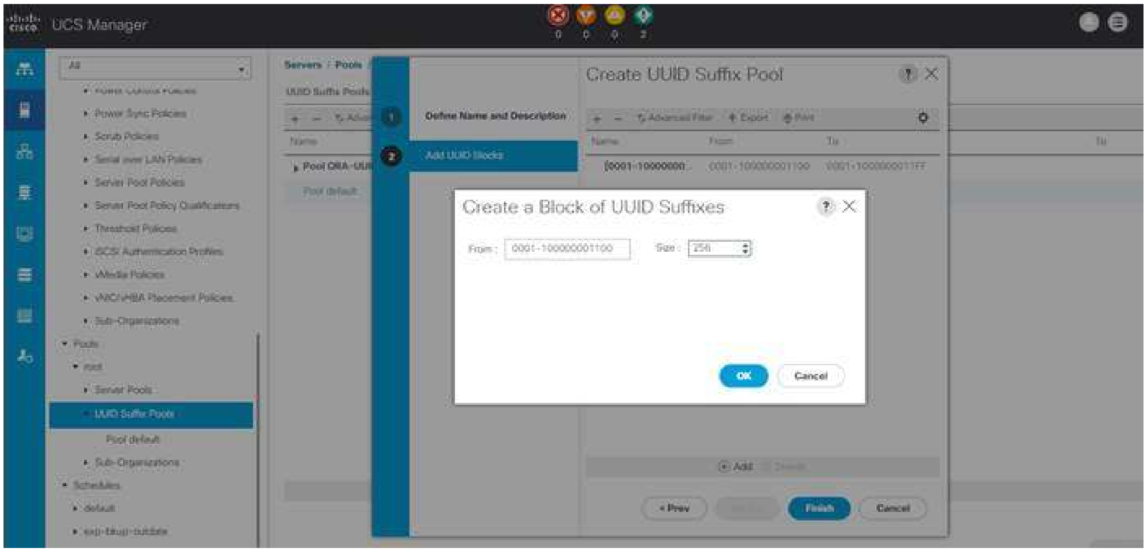

[0001-10000000... 0001-100000001100 0001-1000000011FF

Create a Block of UUID Suffixes

From : 0001-100000001100 Size : 256

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

Server Pool Creation

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

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

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

MAC Pool Creation

To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

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

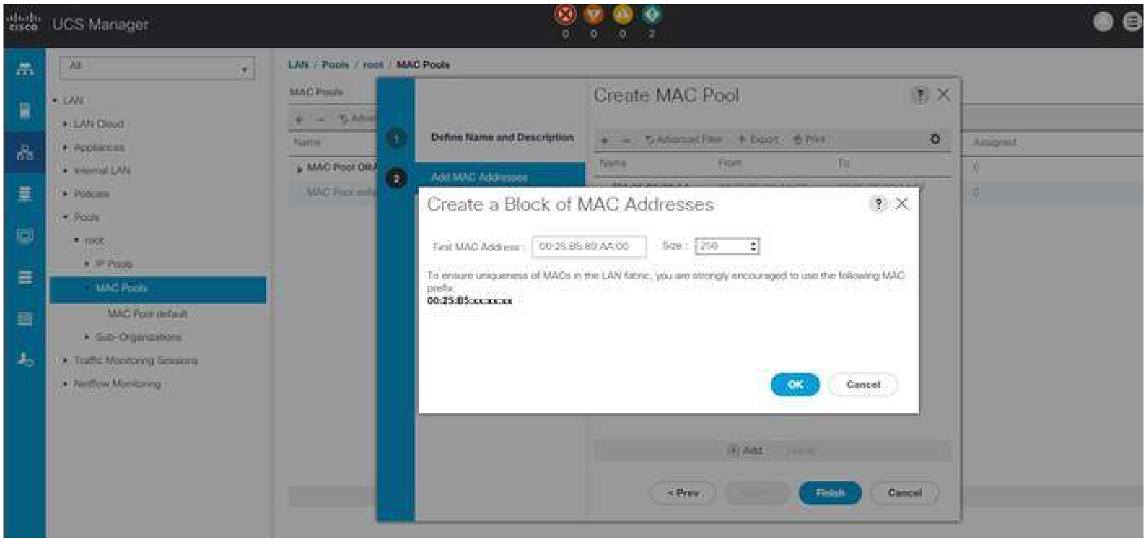

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:89:AA:00 Size : 256

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

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

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

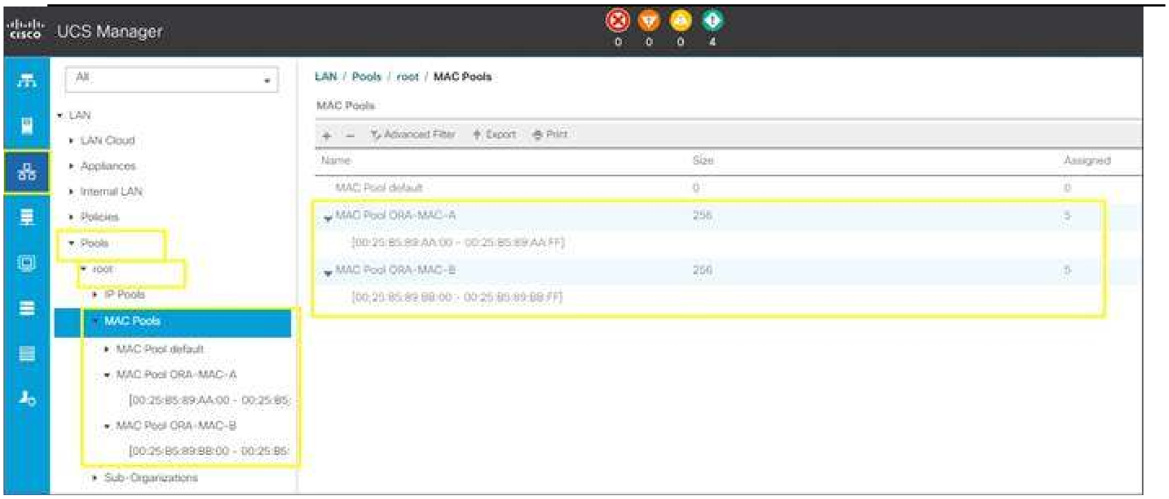

Name	Size	Assigned
MAC Pool default	0	0
MAC Pool ORA-MAC-A	256	5
[00:25:B5:89:AA:00 - 00:25:B5:89:AA:FF]		
MAC Pool ORA-MAC-B	256	0
[00:25:B5:89:BB:00 - 00:25:B5:89:BB:FF]		

WWNN and WWPN Pool Creation

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

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

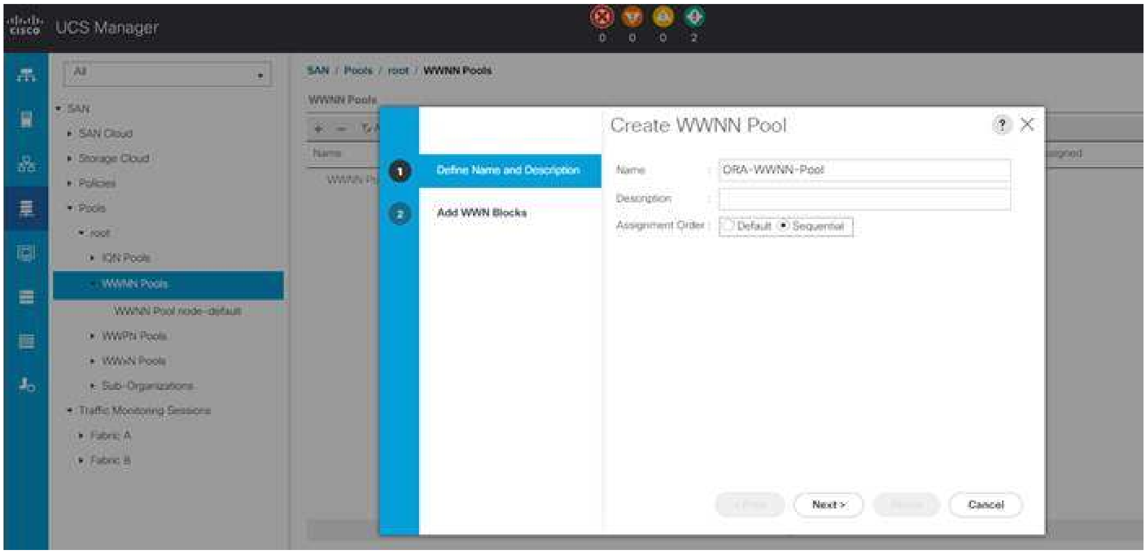

Create WWNN Pool

Name : ORA-WWNN-Pool

Description :

Assignment Order : ○ Default ● Sequential

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

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

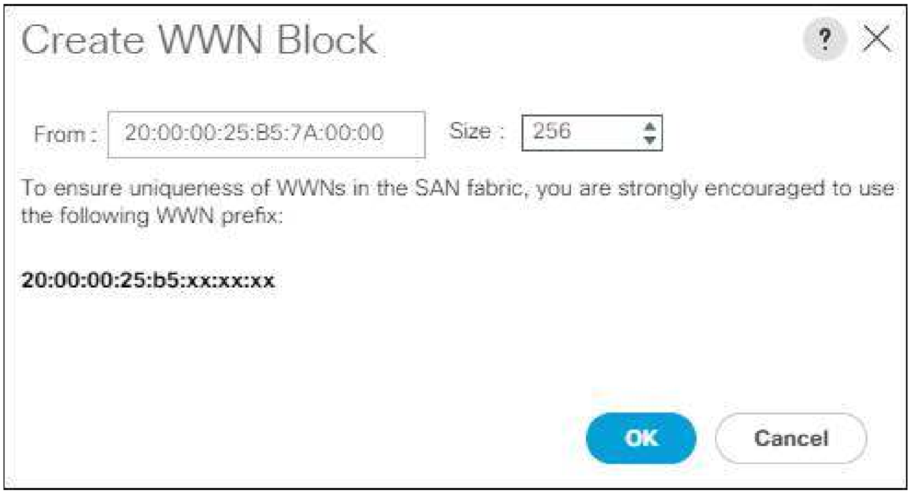

Create WWN Block

From : 20:00:00:25:B5:7A:00:00 Size : 256

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

6. Click OK and then click Finish.

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

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

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

2. Select Pools > Root > WWPN Pools > right-click WWPN Pools > Select Create WWPN Pool.

3. Assign name as ORA-WWPN-A and Assignment Order as sequential.

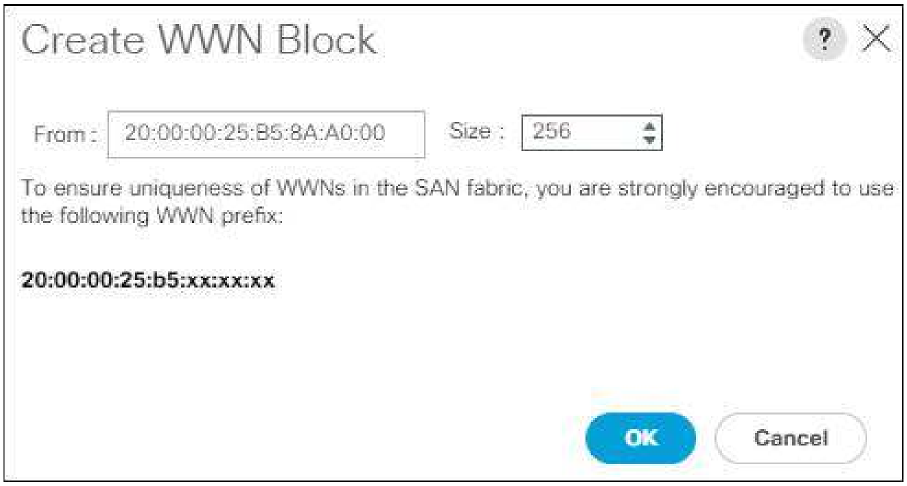

Create WWN Block

From : 20:00:00:25:B5:8A:A0:00 Size : 256

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

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

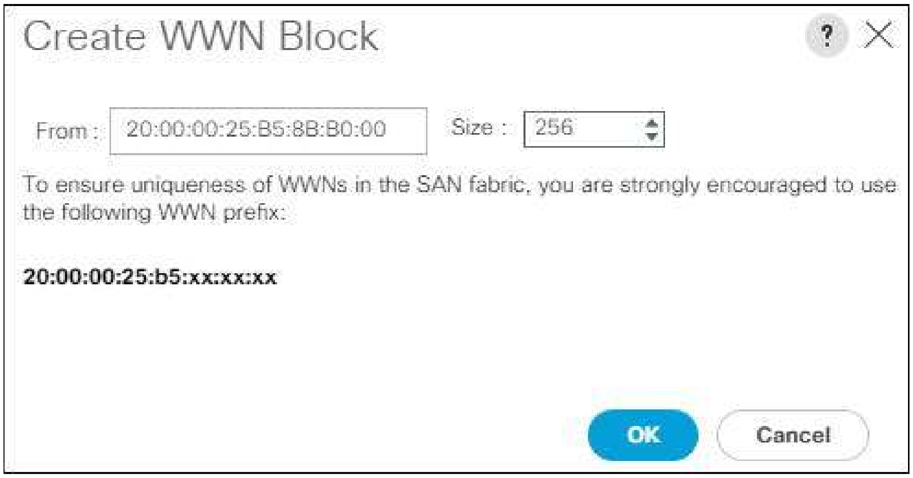

Create WWN Block

From : 20:00:00:25:B5:8B:B0:00 Size : 256

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

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

Set Jumbo Frames in both the Cisco Fabric Interconnect

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

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

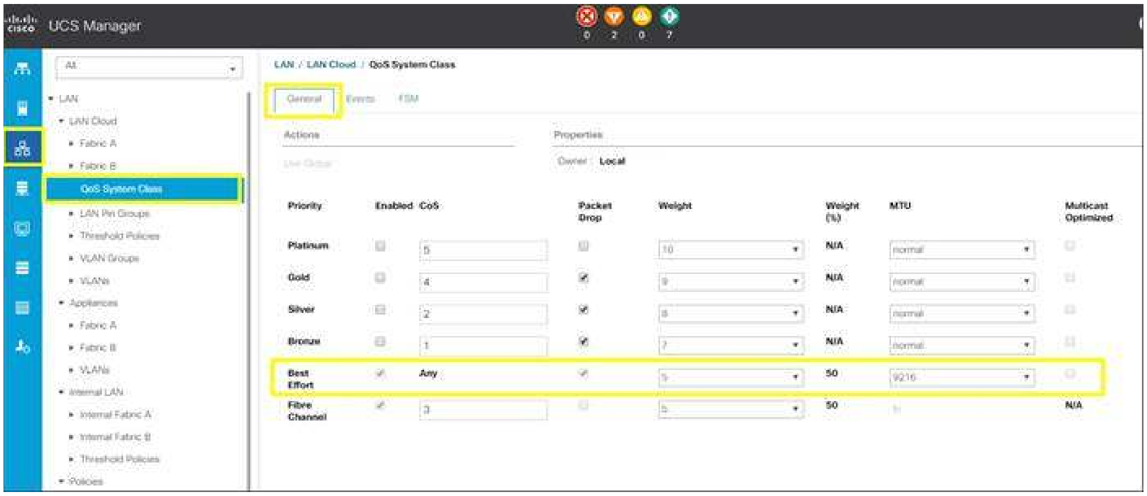

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	□	5	□	10	N/A	normal	□
Gold	□	4	☑	9	N/A	normal	□
Silver	□	2	☑	8	N/A	normal	□
Bronze	□	1	☑	7	N/A	normal	□
Best Effort	☑	Any	☑	5	50	9216	□
Fibre Channel	☑	3	□	5	50	fc	N/A

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

6. Click OK.

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

Configure Server BIOS Policy

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

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

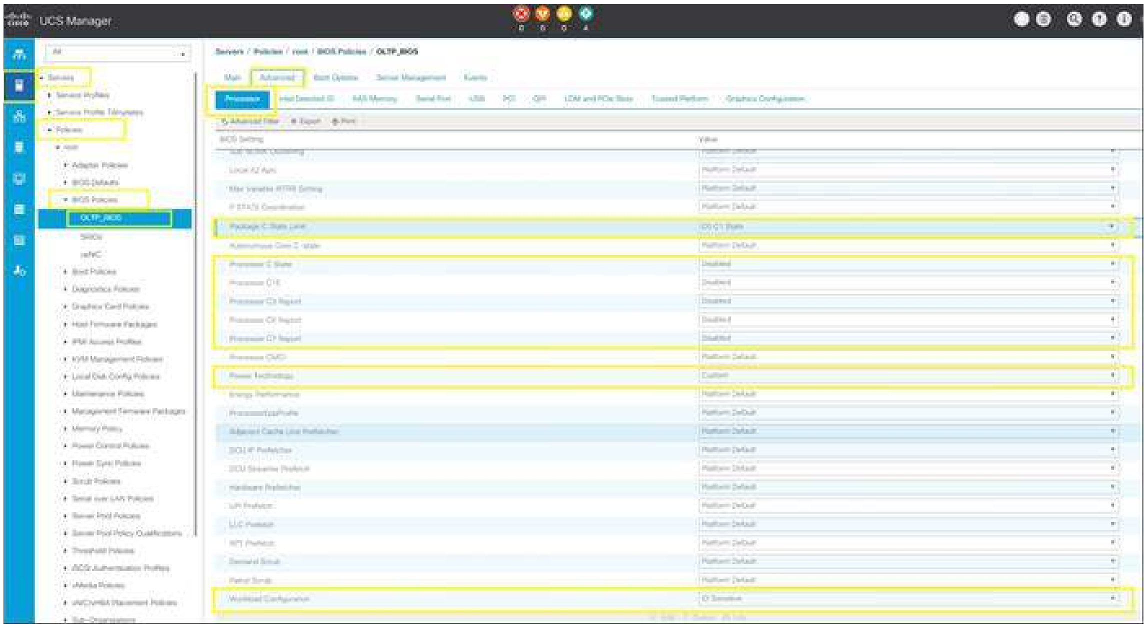

BIOS Setting	Value
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	C0 C1 State
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Custom
Energy Performance	Platform Default
ProcessorEppProfile	Platform Default
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	IO Sensitive

9. Click Save Changes and then click OK

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

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

---

It is recommended to disable C states in the BIOS and in addition, Oracle recommends disabling it from OS level as well by modifying grub entries.

Create Adapter Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > right-click Adapter Policies.
3. Select Create Ethernet Adapter Policy.
4. Provide a name for the Ethernet adapter policy. Change the following fields and click Save Changes when you are finished:

a. Resources

i. Transmit Queues: 8

ii. Ring Size: 4096

iii. Receive Queues: 8

iv. Ring Size: 4096

v. Completion Queues: 16

vi. Interrupts: 32

b. Options

i. Receive Side Scaling (RSS): Enabled

5. Configure adapter policy as shown below.

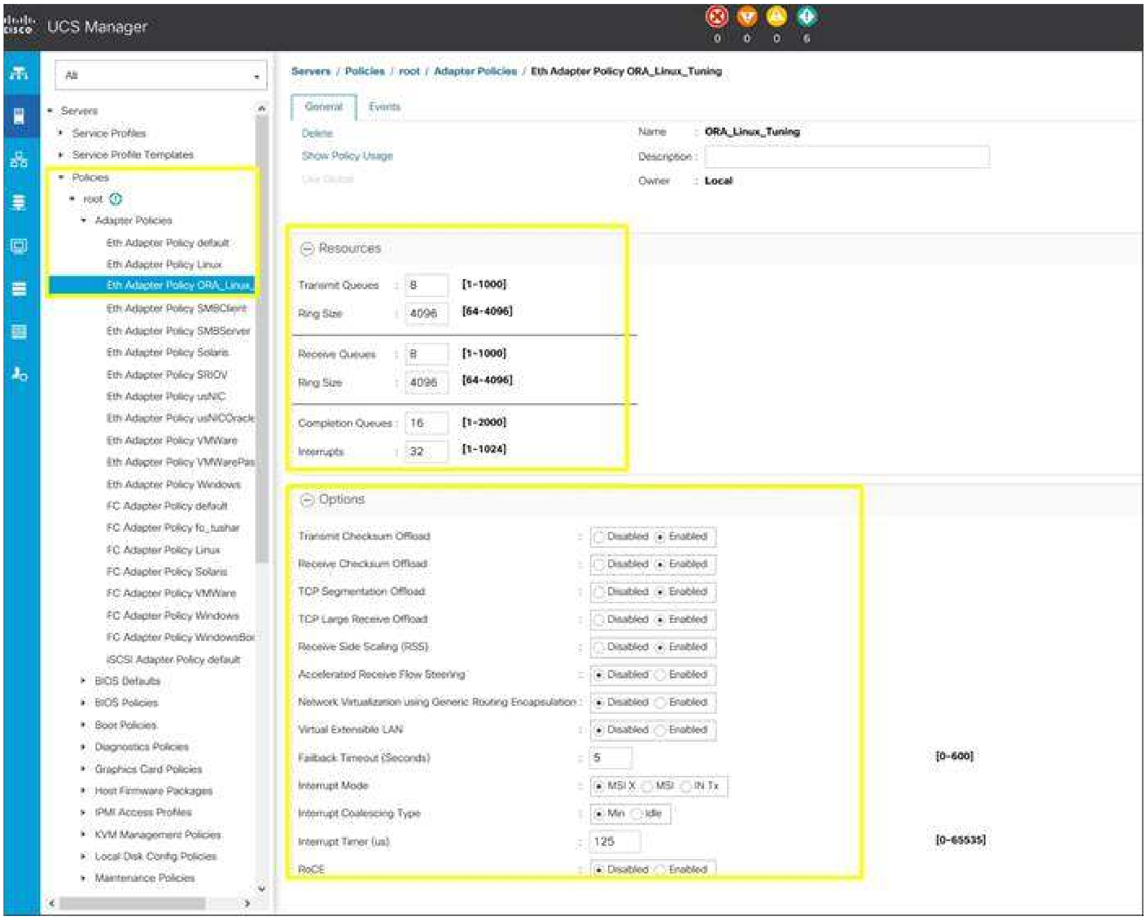
UCS Manager
0 0 0 6
All
Servers
Service Profiles
Service Profile Templates
Policies
root
Adapter Policies
Eth Adapter Policy default
Eth Adapter Policy Linux
Eth Adapter Policy ORA_Linux
Eth Adapter Policy SMBClient
Eth Adapter Policy SMBServer
Eth Adapter Policy Solaris
Eth Adapter Policy SRIOV
Eth Adapter Policy usNIC
Eth Adapter Policy usNICOracle
Eth Adapter Policy VMWare
Eth Adapter Policy VMWarePas
Eth Adapter Policy Windows
FC Adapter Policy default
FC Adapter Policy fc_tushar
FC Adapter Policy Linux
FC Adapter Policy Solaris
FC Adapter Policy VMWare
FC Adapter Policy Windows
FC Adapter Policy WindowsBoo
iSCSI Adapter Policy default
BIOS Defaults
BIOS Policies
Boot Policies
Diagnostics Policies
Graphics Card Policies
Host Firmware Packages
IPMI Access Profiles
KVM Management Policies
Local Disk Config Policies
Maintenance Policies
General
Events
Delete
Show Policy Usage
Use Global

Servers / Policies / root / Adapter Policies / Eth Adapter Policy ORA_Linux_Tuning

Name : ORA_Linux_Tuning
Description :
Owner : Local

Resources

Transmit Queues : 8 [1-1000]
Ring Size : 4096 [64-4096]

Receive Queues : 8 [1-1000]
Ring Size : 4096 [64-4096]

Completion Queues : 16 [1-2000]
Interrupts : 32 [1-1024]

Options

Transmit Checksum Offload : ○ Disabled ● Enabled
Receive Checksum Offload : ○ Disabled ● Enabled
TCP Segmentation Offload : ○ Disabled ● Enabled
TCP Large Receive Offload : ○ Disabled ● Enabled
Receive Side Scaling (RSS) : ○ Disabled ● Enabled
Accelerated Receive Flow Steering : ● Disabled ○ Enabled
Network Virtualization using Generic Routing Encapsulation : ● Disabled ○ Enabled
Virtual Extensible LAN : ● Disabled ○ Enabled
Failback Timeout (Seconds) : 5 [0-600]
Interrupt Mode : ● MSI X ○ MSI ○ IN Tx
Interrupt Coalescing Type : ● Min ○ Idle
Interrupt Timer (us) : 125 [0-65535]
RoCE : ● Disabled ○ Enabled

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

Enabled—Network receive processing is shared across processors whenever possible.

6. Click OK to finish.

Configure Update Default Maintenance Policy

To update the default Maintenance Policy, follow these steps:

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

Configure vNIC and vHBA Template

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

Create Public and Private vNIC Template

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root > vNIC Templates > right-click to vNIC Template and Select " Create vNIC Template"
3. Enter ORA-vNIC-A as the vNIC template name and keep Fabric A selected.
4. Select the Enable Failover checkbox for high availability of the vNIC.

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

5. Select Template Type as Updating Template.

6. Under VLANs, select the checkboxes default and Public_Traffic and set Native-VLAN as the Public_Traffic.

7. Keep MTU value 1500 for Public Network Traffic.

8. In the MAC Pool list, select ORA-MAC-A.

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

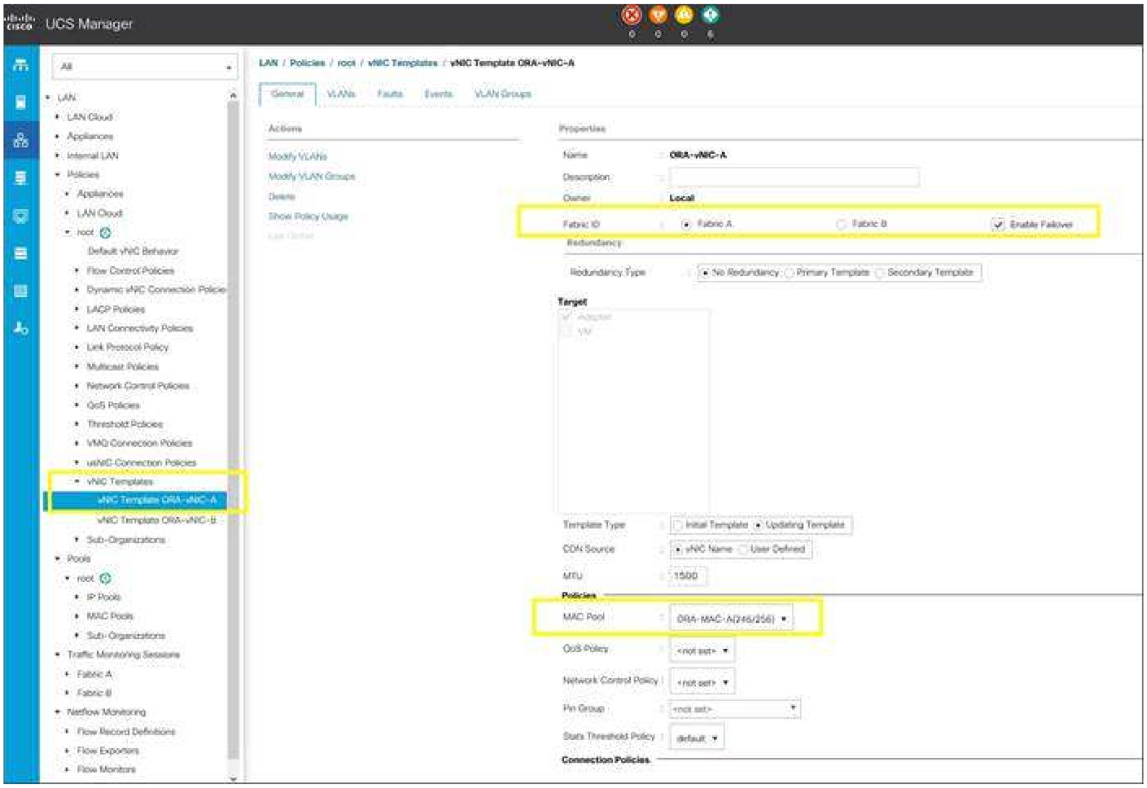

Name : ORA-vNIC-A

Description :

Owner : Local

Fabric ID : ● Fabric A ○ Fabric B ☑ Enable Failover

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Template Type : ○ Initial Template ● Updating Template

CDN Source : ● vNIC Name ○ User Defined

MTU : 1500

MAC Pool : ORA-MAC-A(246/256)

QoS Policy : <not set>

Network Control Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

10. Click OK to finish.

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

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

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

c. Select Template Type as Updating Template.

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

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

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

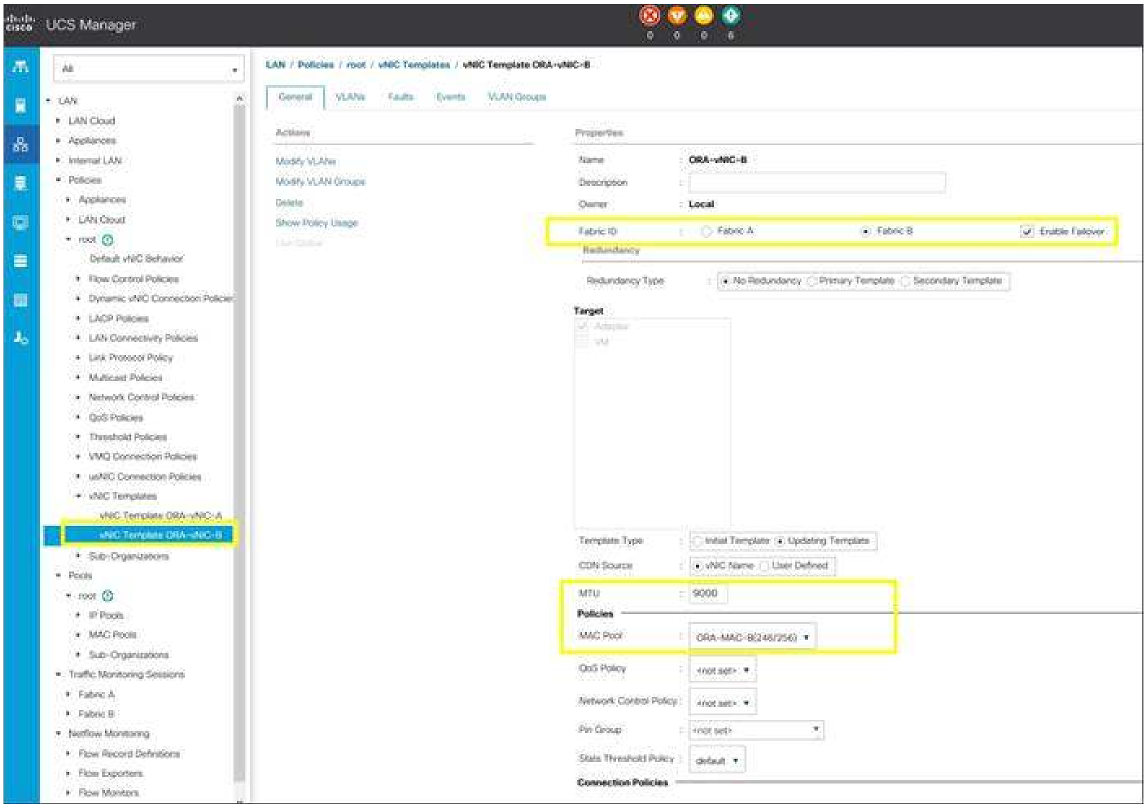

Properties

Name : ORA-vNIC-B

Description :

Owner : Local

Fabric ID : ○ Fabric A ● Fabric B ☑ Enable Failover

Redundancy

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Target

☑ Adapter

□ VM

Template Type : ○ Initial Template ● Updating Template

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

Policies

MAC Pool : ORA-MAC-B(248/256)

QoS Policy : <not set>

Network Control Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Create Storage vHBA Template

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

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

2. Select Policies > root > Right Click vHBA Templates > Select “Create vHBA Template” to create vHBAs.

3. Enter name as ORA-vHBA-A and keep Fabric A selected.

4. Select VSAN as ORA-VSAN-A and template type to Updating Template.

5. Select WWPN Pool as ORA-WWPN-A from the drop-down list as shown below.

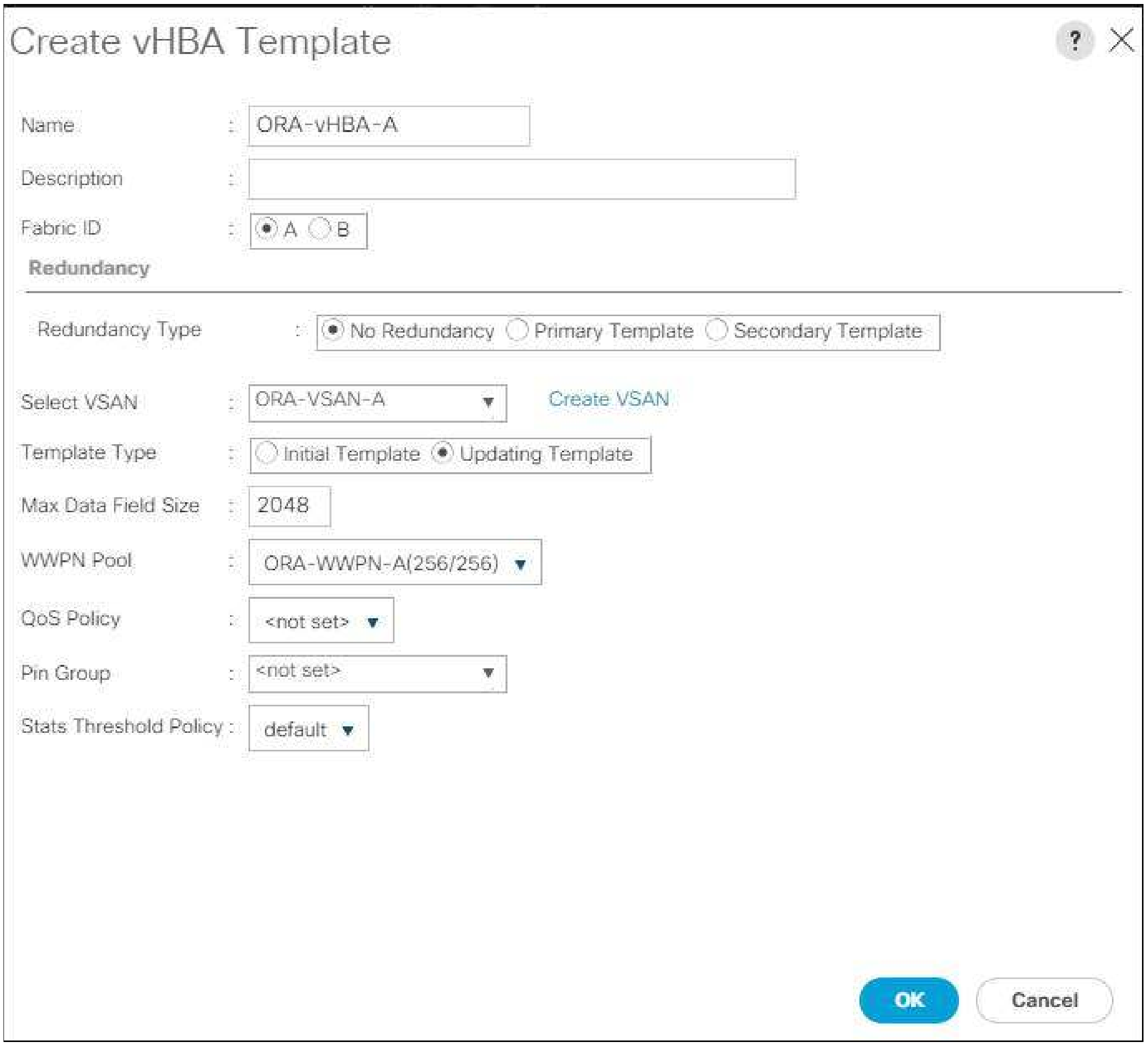

Create vHBA Template

Name : ORA-vHBA-A

Description :

Fabric ID : ● A ○ B

Redundancy

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Select VSAN : ORA-VSAN-A ▾ Create VSAN

Template Type : ○ Initial Template ● Updating Template

Max Data Field Size : 2048

WWPN Pool : ORA-WWPN-A(256/256) ▾

QoS Policy : <not set> ▾

Pin Group : <not set> ▾

Stats Threshold Policy : default ▾

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

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

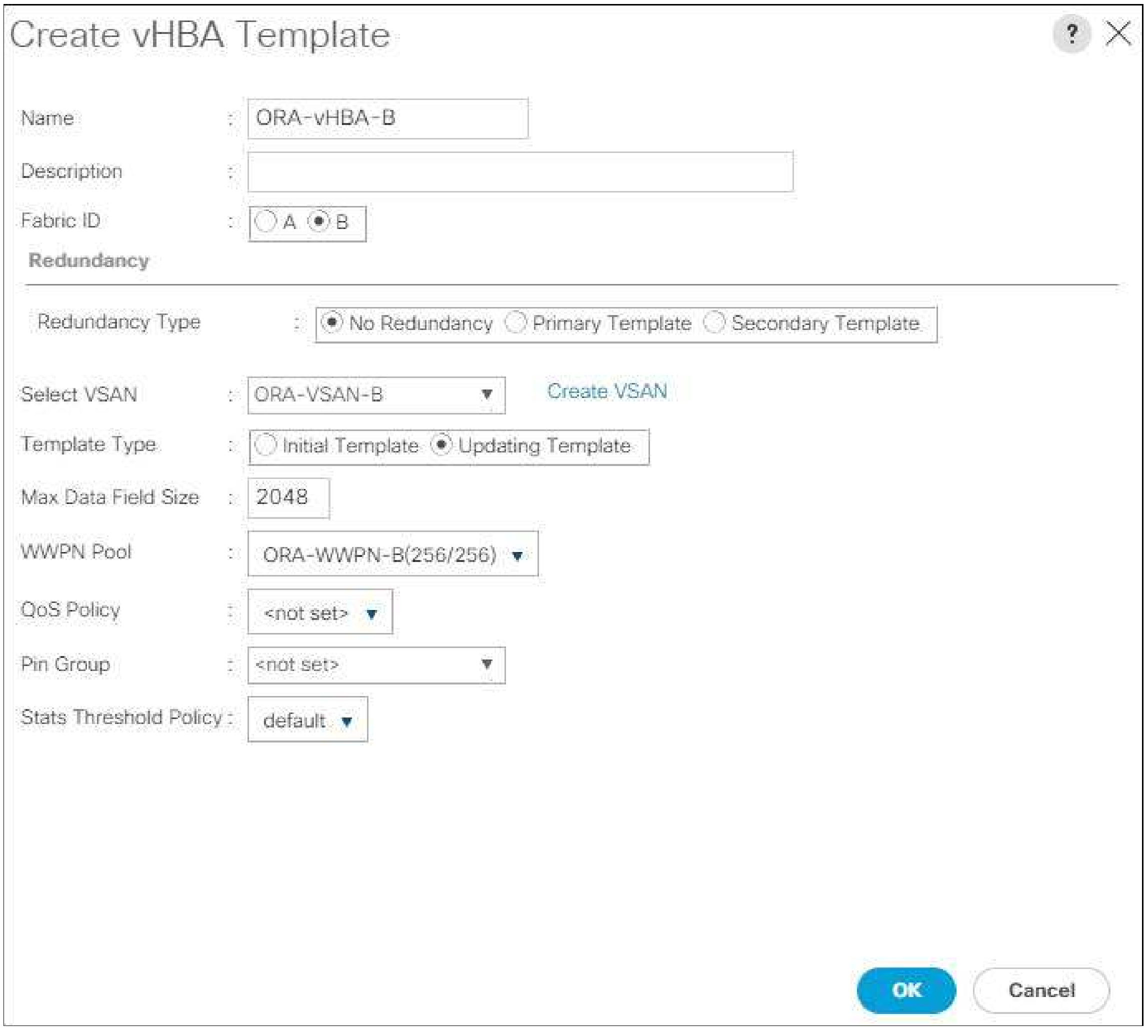

Create vHBA Template

Name : ORA-vHBA-B

Description :

Fabric ID : ○ A ● B

Redundancy

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Select VSAN : ORA-VSAN-B ▼ Create VSAN

Template Type : ○ Initial Template ● Updating Template

Max Data Field Size : 2048

WWPN Pool : ORA-WWPN-B(256/256) ▼

QoS Policy : <not set> ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

Create Server Boot Policy for SAN Boot

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

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

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

Create Local Disk Configuration Policy

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

To configure Local disk policy, follow these steps:

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

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

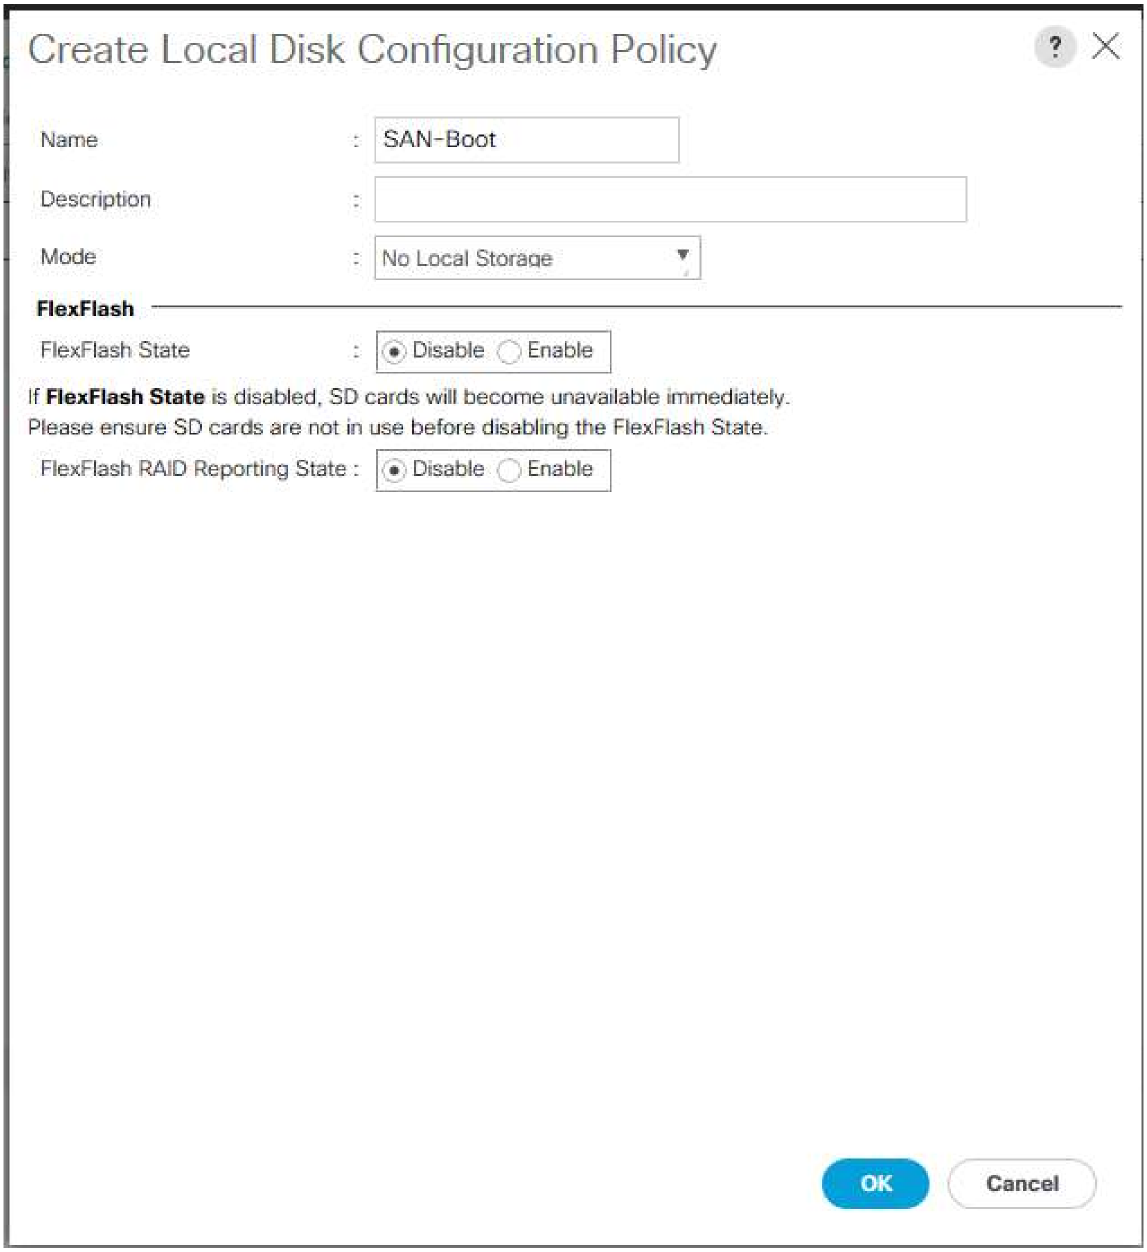

Create Local Disk Configuration Policy

Name	:	SAN-Boot
Description	:	
Mode	:	No Local Storage

FlexFlash

FlexFlash State : ● Disable ○ Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately.
Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : ● Disable ○ Enable

Create SAN Boot Policy

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

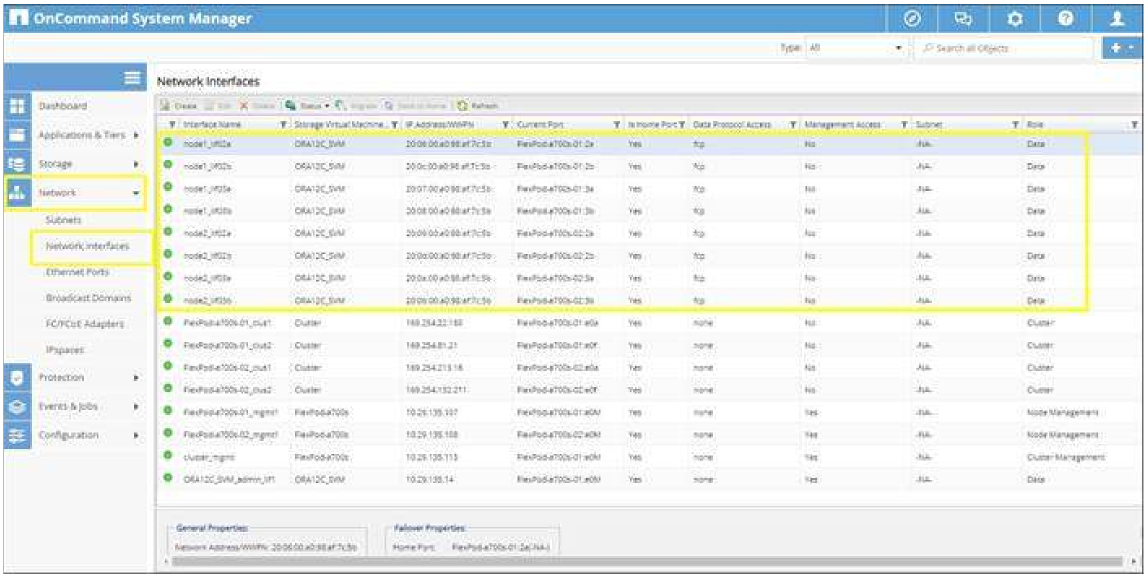

Network Interfaces

Interface Name	Storage Virtual Machine...	IP Address/WWPN	Current Port	Is Home Port	Data Protocol Access	Management Access	Subnet	Role
node1_lif02a	ORA12C_SVM	20:06:00:a0:98:af:7c:5b	FlexPod-a700s-01:2a	Yes	fcp	No	-NA-	Data
node1_lif02b	ORA12C_SVM	20:0c:00:a0:98:af:7c:5b	FlexPod-a700s-01:2b	Yes	fcp	No	-NA-	Data
node1_lif03a	ORA12C_SVM	20:07:00:a0:98:af:7c:5b	FlexPod-a700s-01:3a	Yes	fcp	No	-NA-	Data
node1_lif03b	ORA12C_SVM	20:08:00:a0:98:af:7c:5b	FlexPod-a700s-01:3b	Yes	fcp	No	-NA-	Data
node2_lif02a	ORA12C_SVM	20:09:00:a0:98:af:7c:5b	FlexPod-a700s-02:2a	Yes	fcp	No	-NA-	Data
node2_lif02b	ORA12C_SVM	20:0b:00:a0:98:af:7c:5b	FlexPod-a700s-02:2b	Yes	fcp	No	-NA-	Data
node2_lif03a	ORA12C_SVM	20:0a:00:a0:98:af:7c:5b	FlexPod-a700s-02:3a	Yes	fcp	No	-NA-	Data
node2_lif03b	ORA12C_SVM	20:09:00:a0:98:af:7c:5b	FlexPod-a700s-02:3b	Yes	fcp	No	-NA-	Data
FlexPod-a700s-01_clus1	Cluster	169.254.22.160	FlexPod-a700s-01:e0a	Yes	none	No	-NA-	Cluster
FlexPod-a700s-01_clus2	Cluster	169.254.81.21	FlexPod-a700s-01:e0f	Yes	none	No	-NA-	Cluster
FlexPod-a700s-02_clus1	Cluster	169.254.213.18	FlexPod-a700s-02:e0a	Yes	none	No	-NA-	Cluster
FlexPod-a700s-02_clus2	Cluster	169.254.132.211	FlexPod-a700s-02:e0f	Yes	none	No	-NA-	Cluster
FlexPod-a700s-01_mgmt1	FlexPod-a700s	10.29.135.107	FlexPod-a700s-01:e0M	Yes	none	Yes	-NA-	Node Management
FlexPod-a700s-02_mgmt1	FlexPod-a700s	10.29.135.108	FlexPod-a700s-02:e0M	Yes	none	Yes	-NA-	Node Management
cluster_mgmt	FlexPod-a700s	10.29.135.115	FlexPod-a700s-01:e0M	Yes	none	Yes	-NA-	Cluster Management
ORA12C_SVM_admin_lif1	ORA12C_SVM	10.29.135.14	FlexPod-a700s-01:e0M	Yes	none	Yes	-NA-	Data

General Properties:
Network Address/WWPN: 20:06:00:a0:98:af:7c:5b

Failover Properties:
Home Port: FlexPod-a700s-01:2a(-NA-)

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

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

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

To create Boot Policies for the Cisco UCS environments follow these steps:

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

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

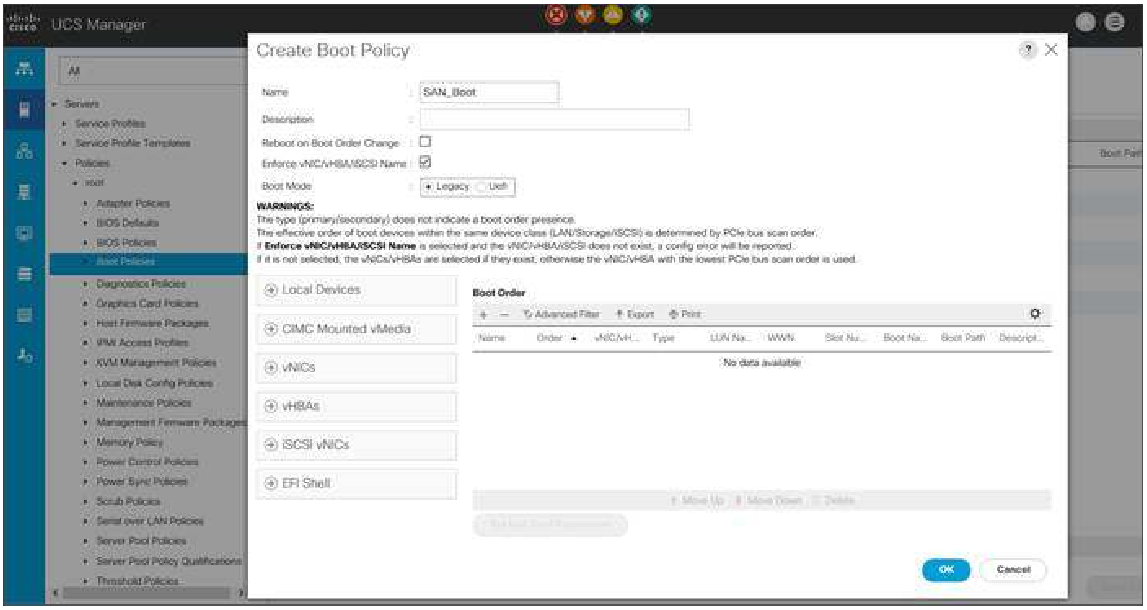

Create Boot Policy

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.

Boot Order

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descript...

No data available

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

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

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

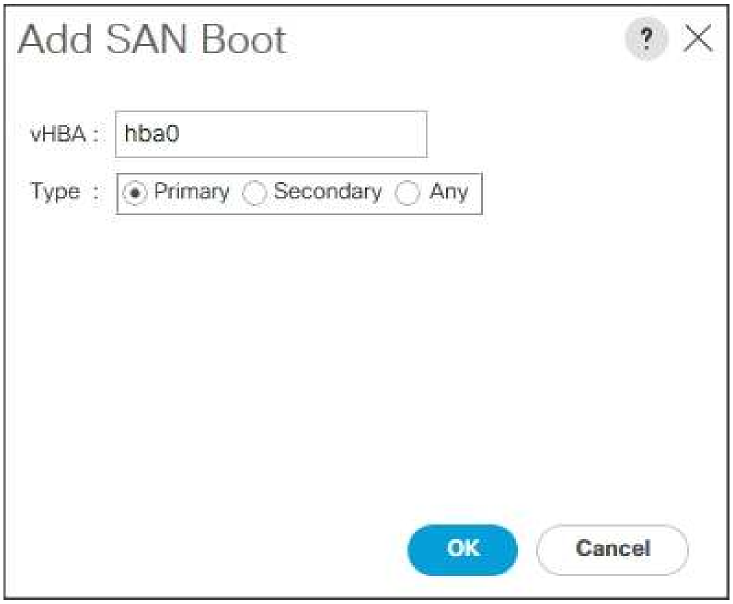

Add SAN Boot

vHBA : hba0

Type : ● Primary ○ Secondary ○ Any

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

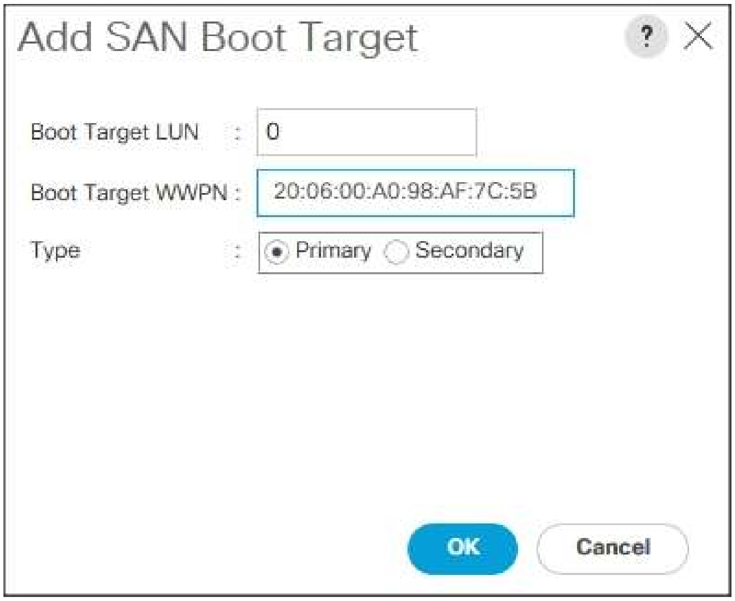

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:06:00:A0:98:AF:7C:5B

Type : ● Primary ○ Secondary

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

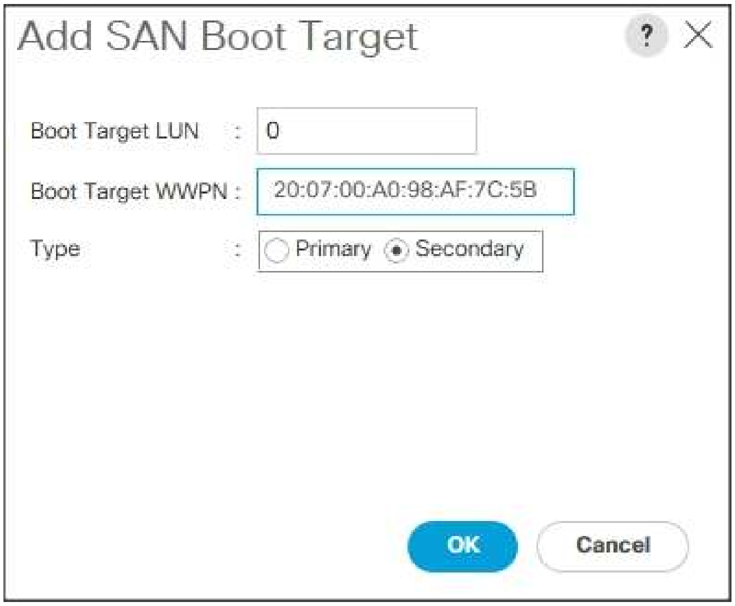

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:07:00:A0:98:AF:7C:5B

Type : ○ Primary ● Secondary

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

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

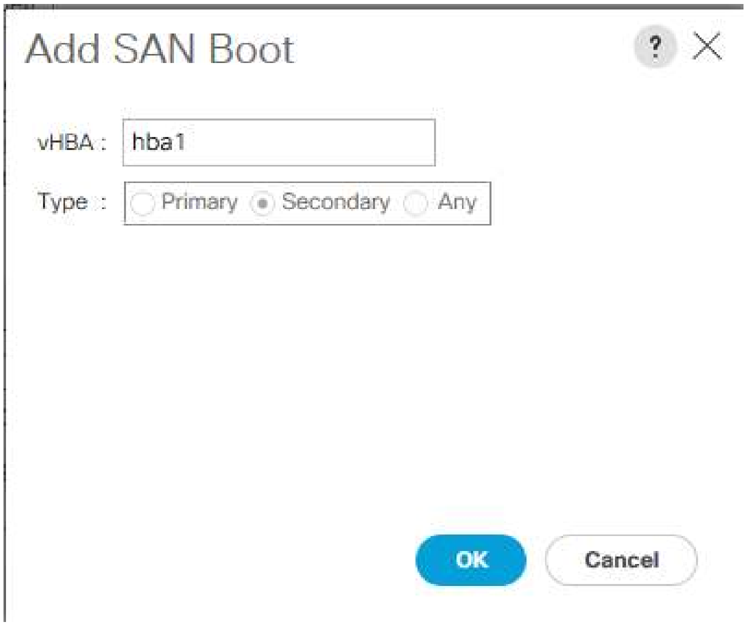

Add SAN Boot

vHBA : hba1

Type : ○ Primary ● Secondary ○ Any

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

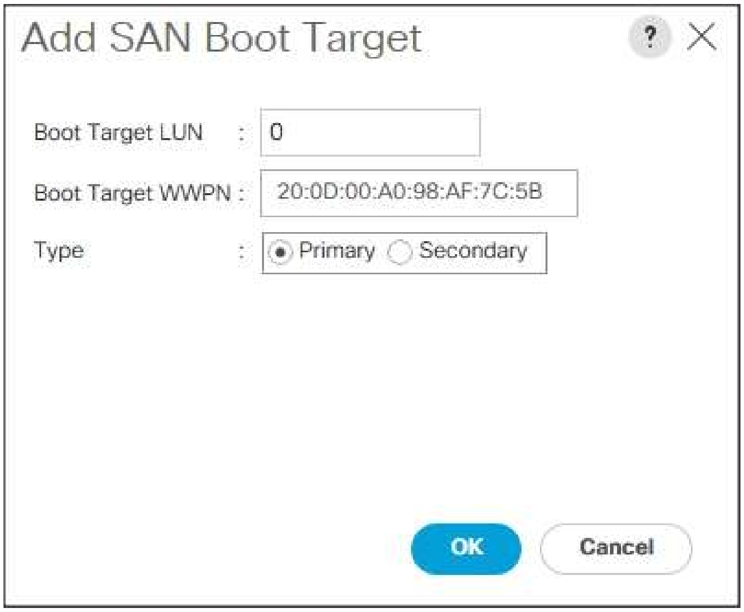

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:0D:00:A0:98:AF:7C:5B

Type : ● Primary ○ Secondary

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

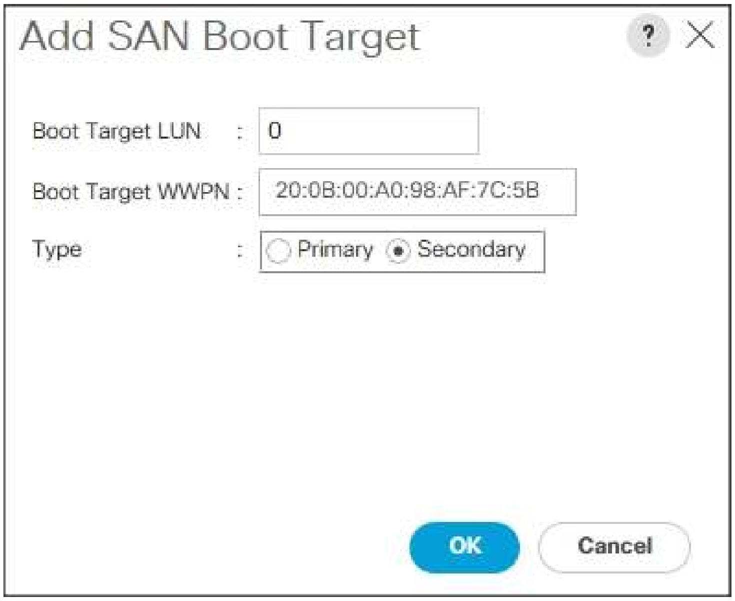

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20:0B:00:A0:98:AF:7C:5B

Type : ○ Primary ● Secondary

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

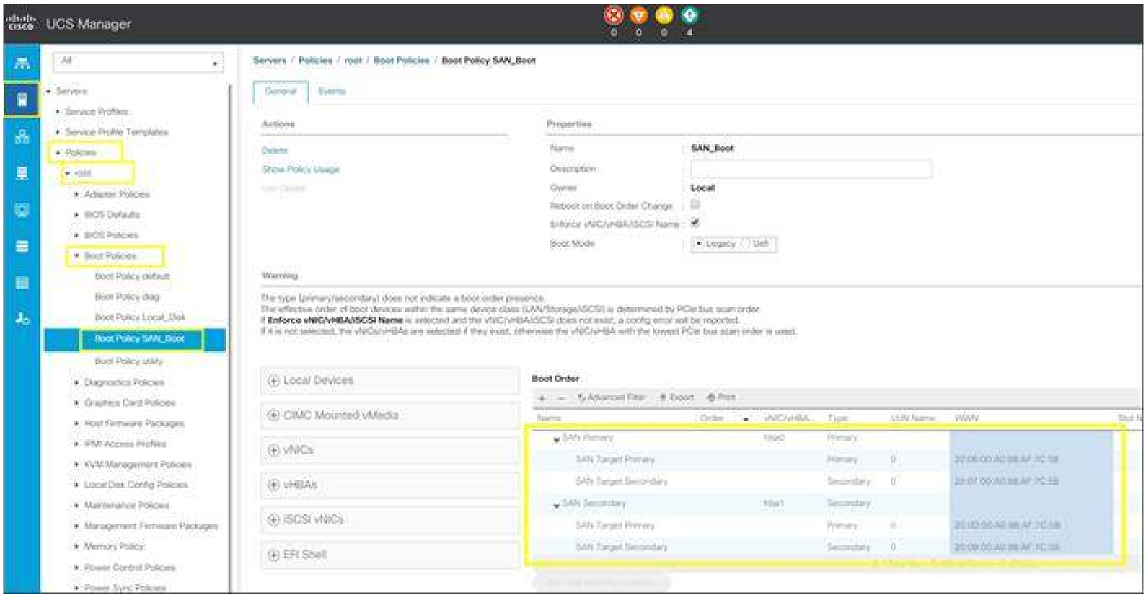

Servers / Policies / root / Boot Policies / Boot Policy SAN_Boot

Properties

Name: SAN_Boot
Description:
Owner: Local
Reboot on Boot Order Change: □
Enforce vNIC/vHBA/iSCSI Name: ☑
Boot Mode: ● Legacy ○ Uefi

Warning

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	Order	vNIC/vHBA...	Type	LUN Name	WWN
SAN Primary		hba0	Primary		
SAN Target Primary			Primary	0	20:06:00:A0:98:AF:7C:58
SAN Target Secondary			Secondary	0	20:01:00:A0:98:AF:7C:58
SAN Secondary		hba1	Secondary		
SAN Target Primary			Primary	0	20:02:00:A0:98:AF:7C:58
SAN Target Secondary			Secondary	0	20:08:00:A0:98:AF:7C:58

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

Configure and Create a Service Profile Template

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

The Cisco UCS service profiles with SAN boot policy provides the following benefits:

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

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

The following sections detail how to create ORA_FLEXPOD.

Create Service Profile Template

To create a service profile template, follow these steps:

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

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

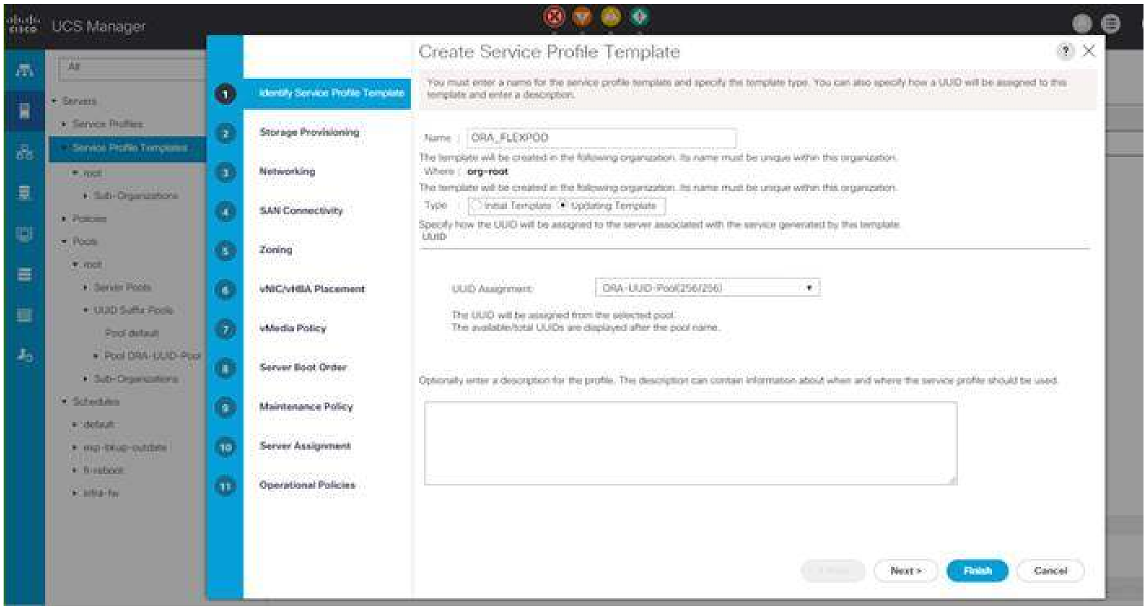

Create Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name : ORA_FLEXPOD

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root**

The template will be created in the following organization. Its name must be unique within this organization.

Type : ○ Initial Template ● Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

UUID

UUID Assignment: ORA-UUID-Pool(256/256)

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

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

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

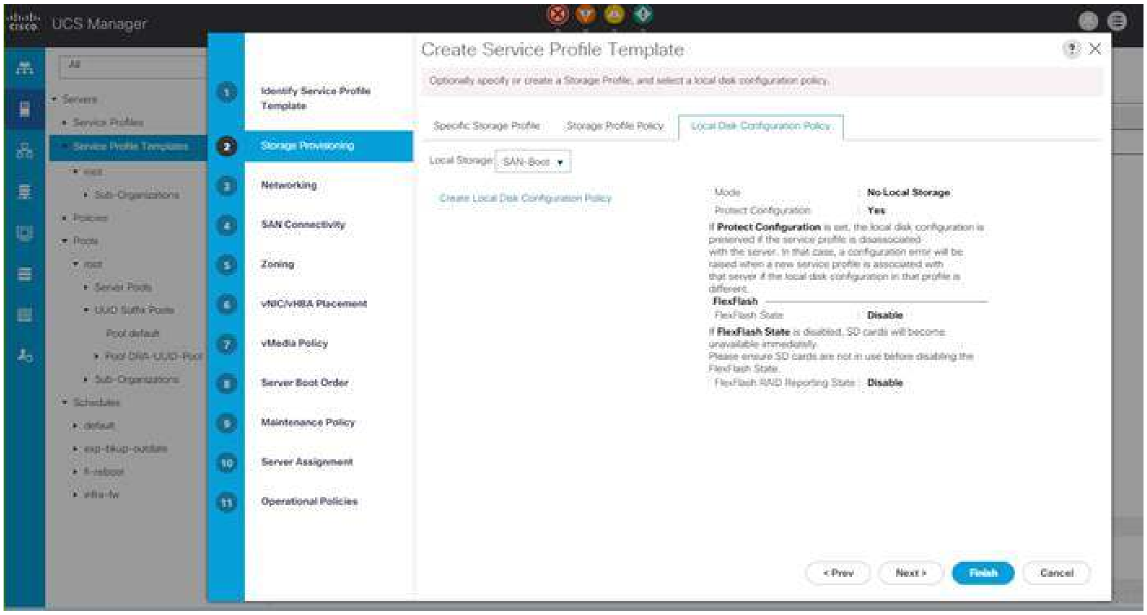

Create Service Profile Template

Optionally specify or create a Storage Profile, and select a local disk configuration policy.

Local Storage: SAN-Boot

Mode : **No Local Storage**

Protect Configuration : **Yes**

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State : **Disable**

If **FlexFlash State** is disabled, SD cards will become unavailable immediately.
Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : **Disable**

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

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

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

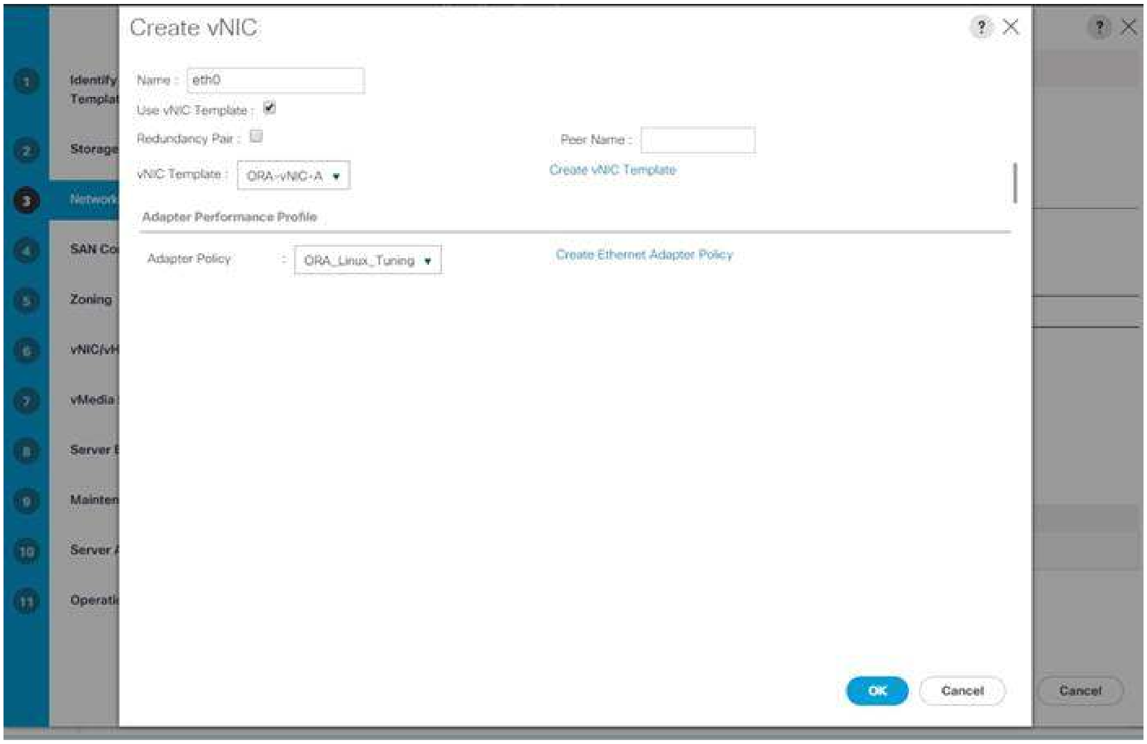

Create vNIC

Name : eth0

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : ORA-vNIC-A

Create vNIC Template

Adapter Performance Profile

Adapter Policy : ORA_Linux_Tuning

Create Ethernet Adapter Policy

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

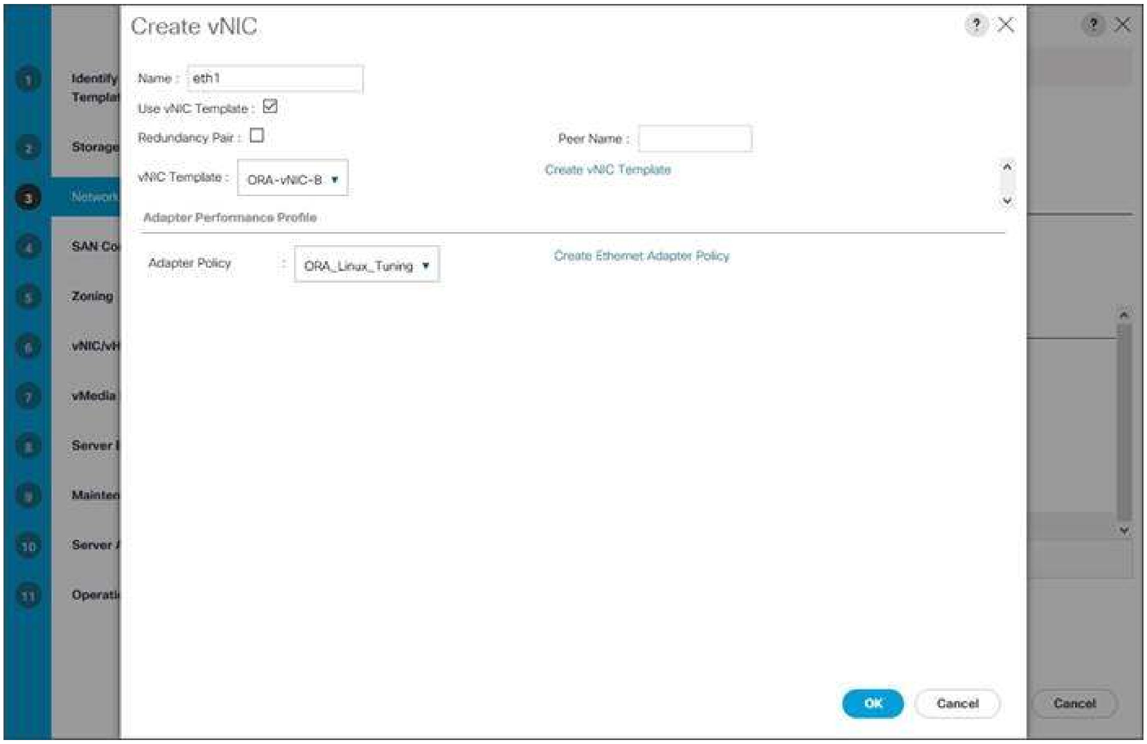

Create vNIC

Name : eth1

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : ORA-vNIC-B

Create vNIC Template

Adapter Performance Profile

Adapter Policy : ORA_Linux_Tuning

Create Ethernet Adapter Policy

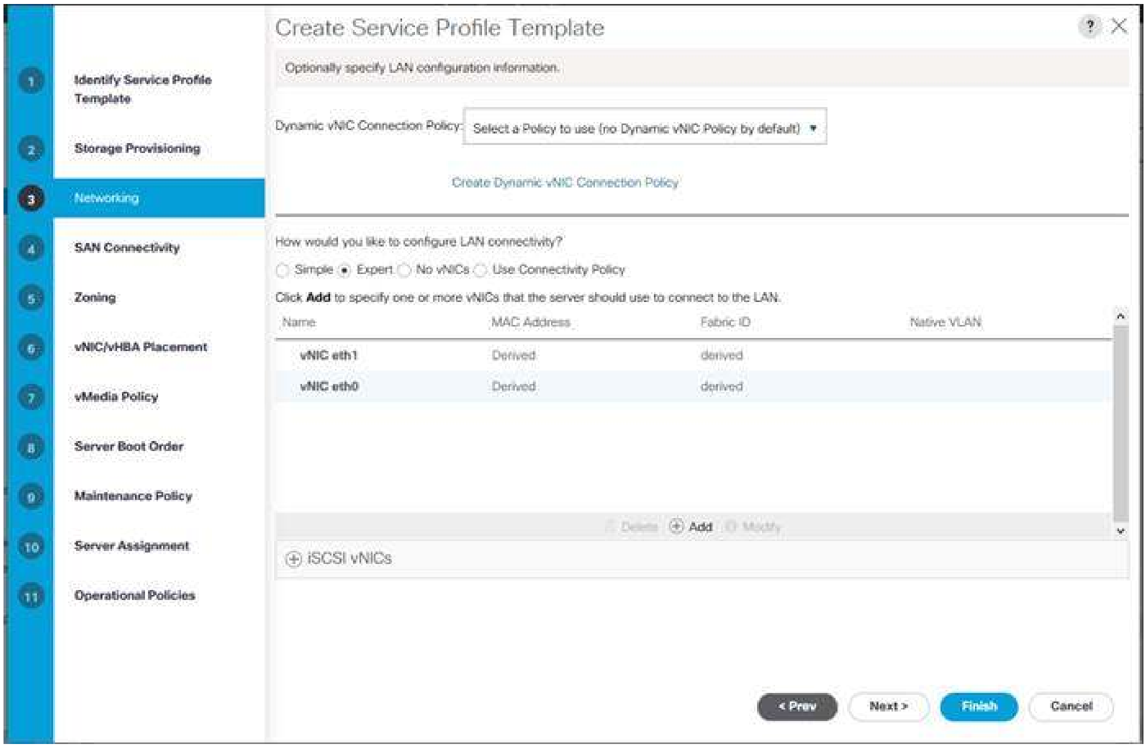

Create Service Profile Template

Optionally specify LAN configuration information.

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

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	

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

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

a. hba0 using vHBA Template Oracle-HBA-A

b. hba1 using vHBA Template Oracle-HBA-B

c. hba2 using vHBA Template Oracle-HBA-A

d. hba3 using vHBA Template Oracle-HBA-B

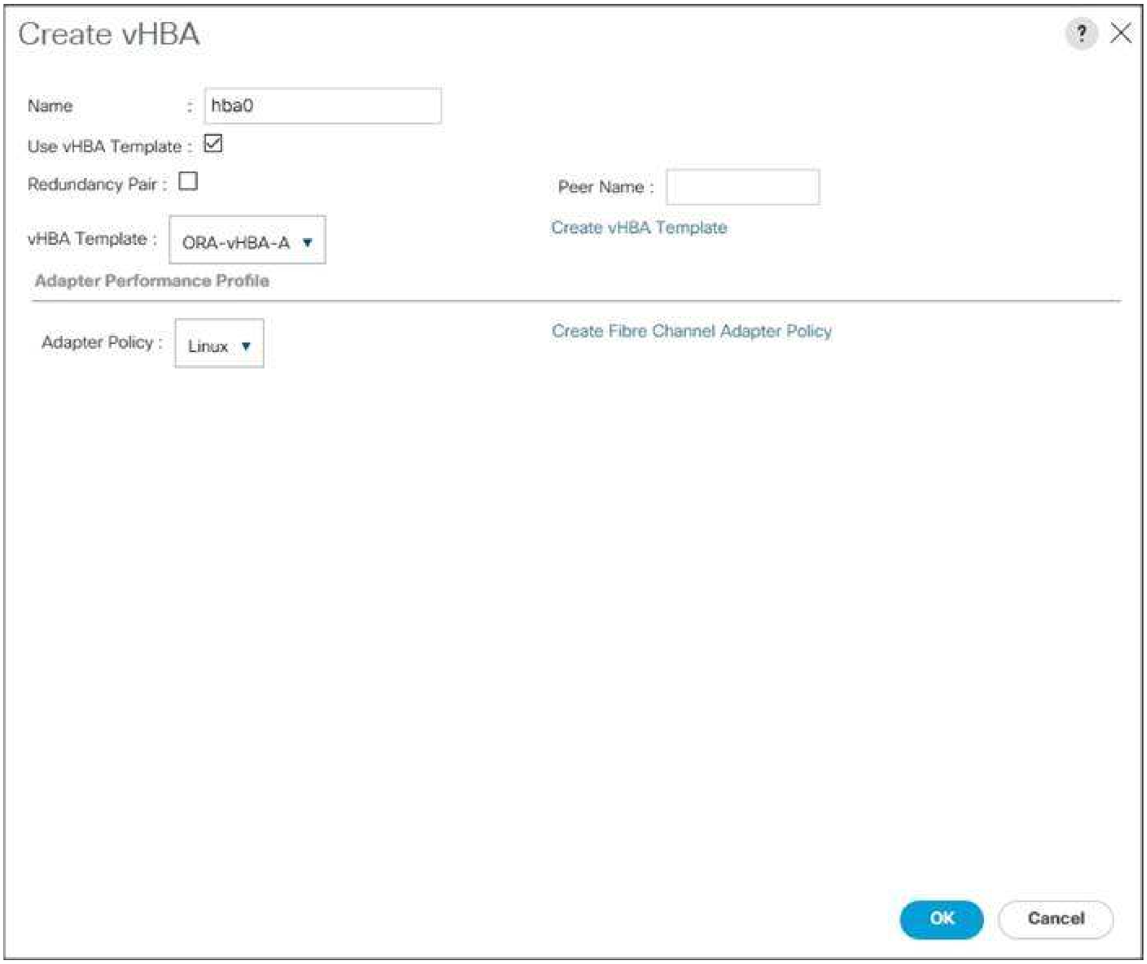

Create vHBA

Name : hba0

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : ORA-vHBA-A

Adapter Performance Profile

Adapter Policy : Linux

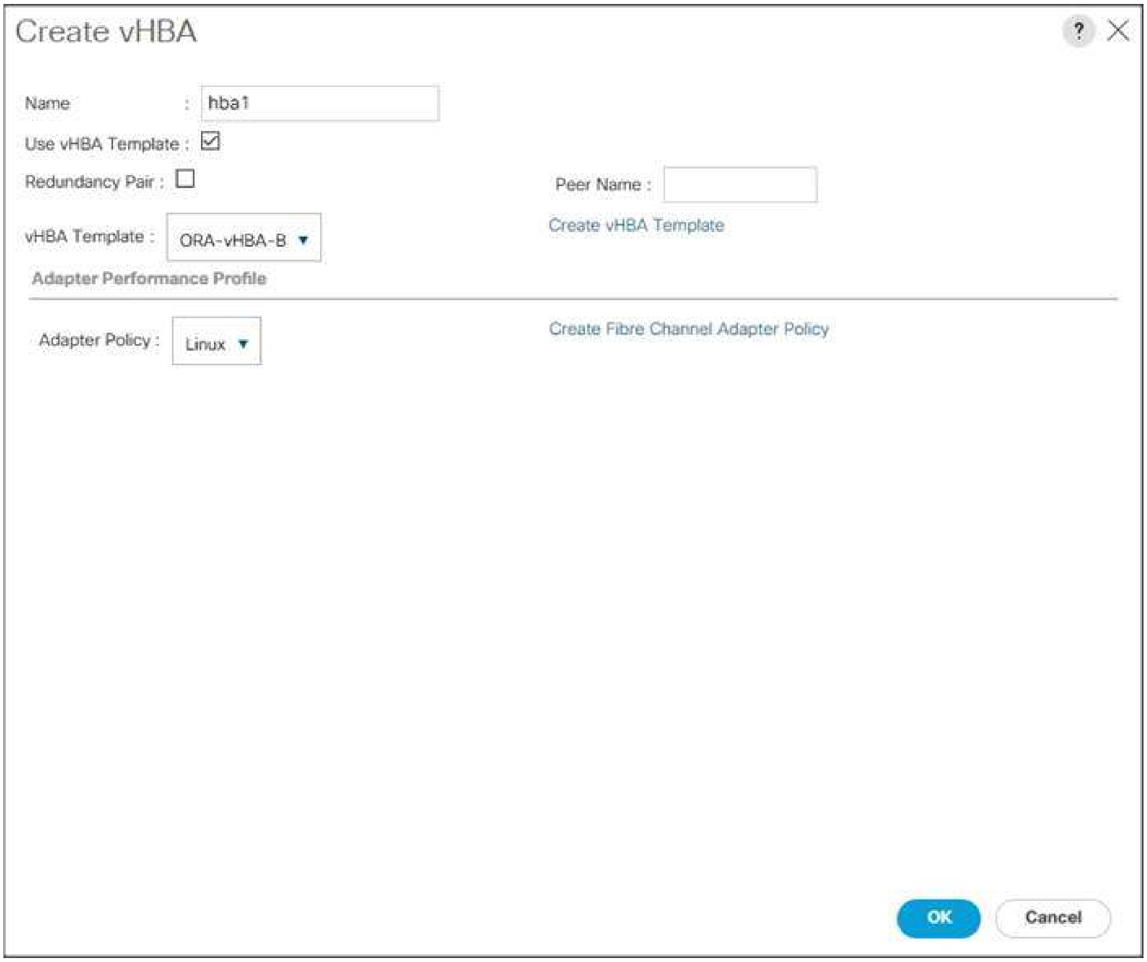

Create vHBA

Name : hba1

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : ORA-vHBA-B ▾

Adapter Performance Profile

Adapter Policy : Linux ▾

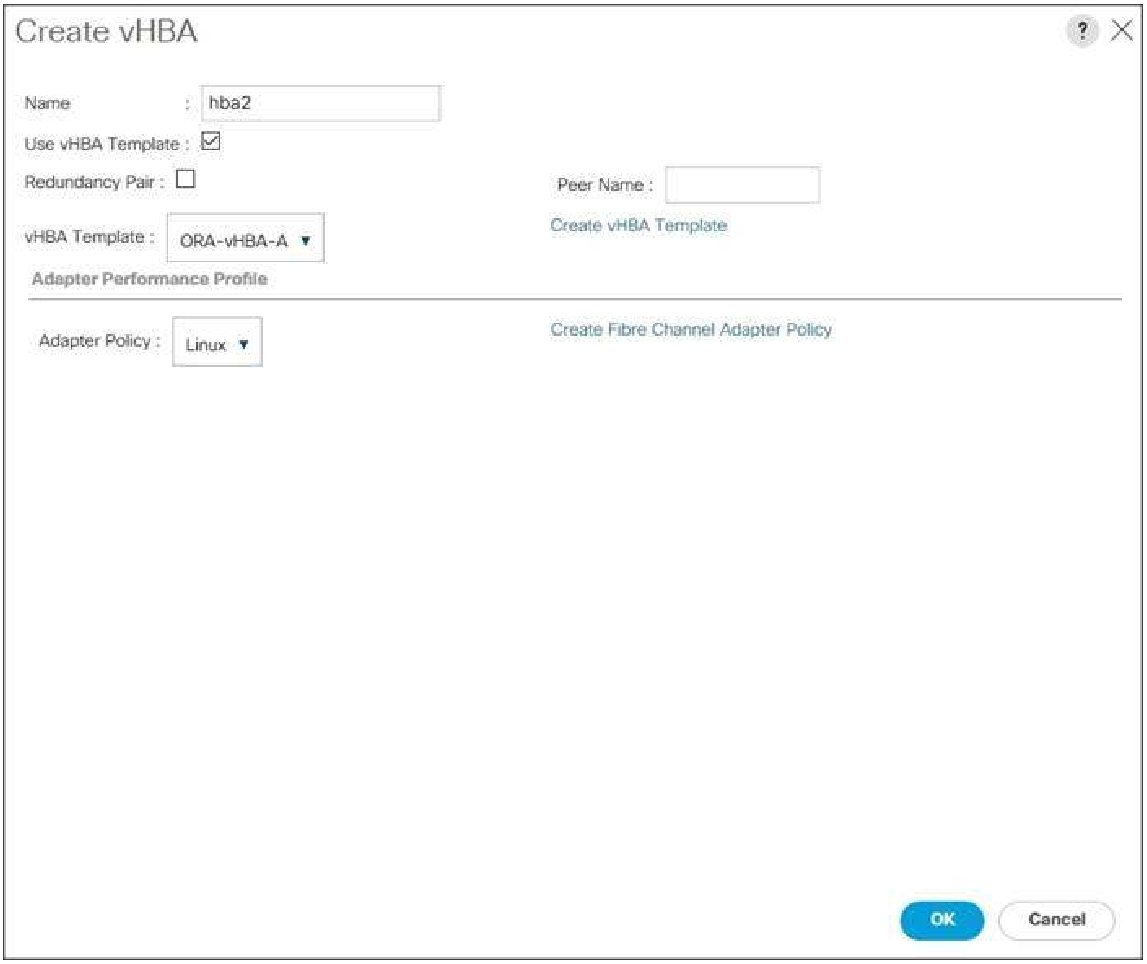

Create vHBA

Name : hba2

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : ORA-vHBA-A ▼

Adapter Performance Profile

Adapter Policy : Linux ▼

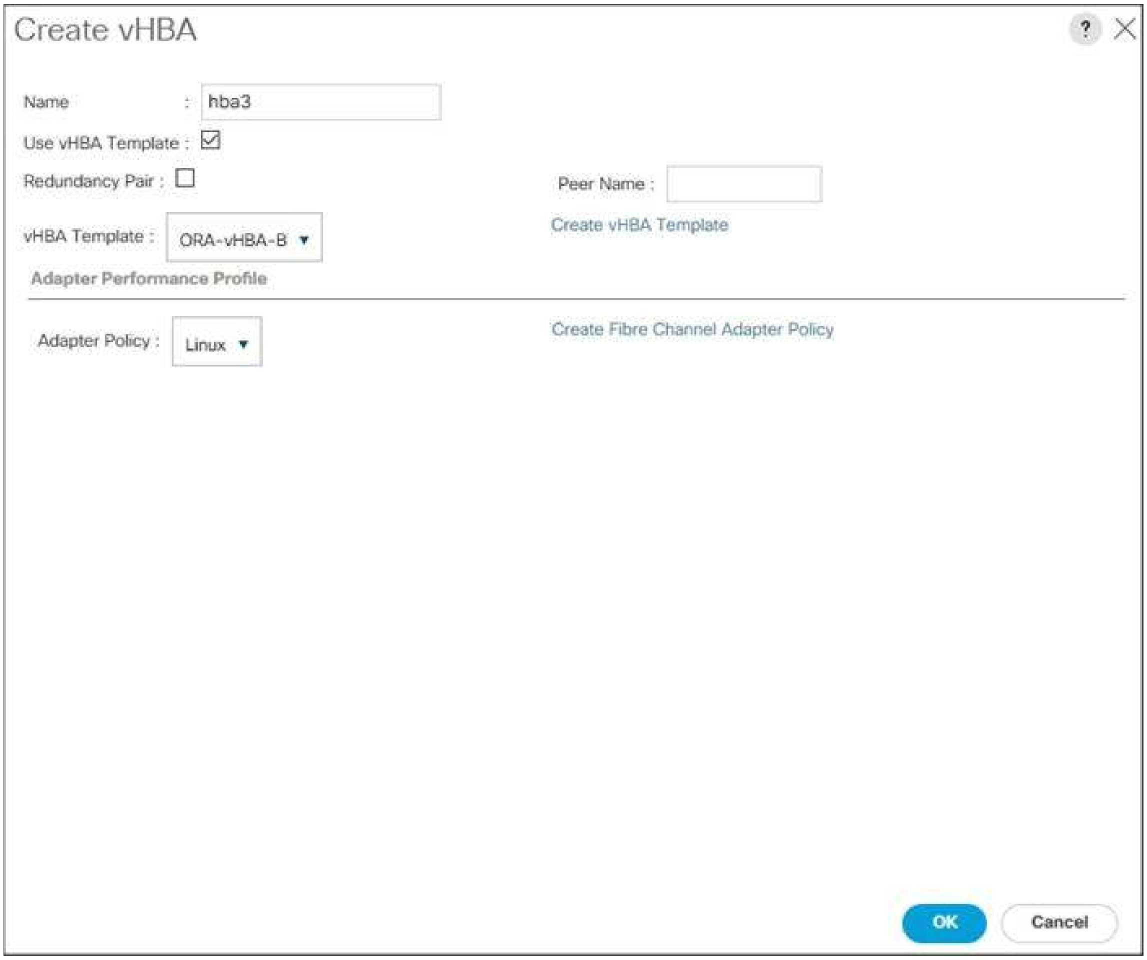

Create vHBA

Name : hba3

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : ORA-vHBA-B

Adapter Performance Profile

Adapter Policy : Linux

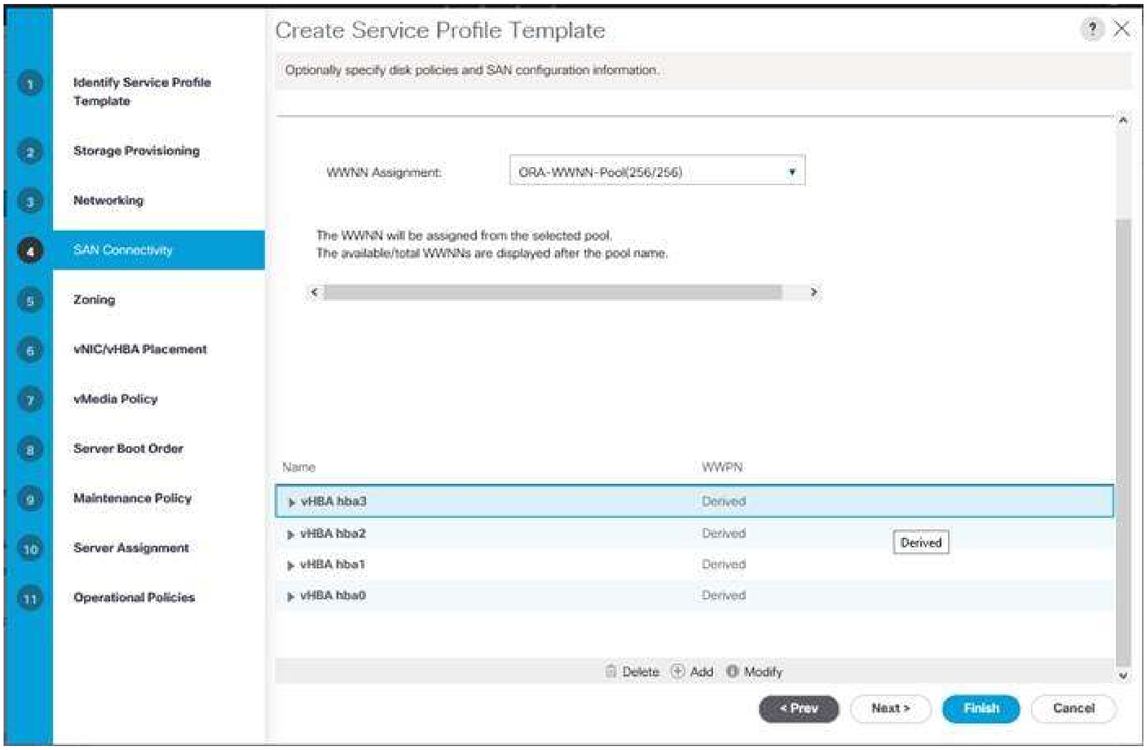

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

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

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

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

9. In vNIC/vHBA Placement menu, keep option as Let System Perform Placement.
10. For this solution, do not configure any vMedia Policy. Click Next.
11. In the Server Boot Order, select "SAN_Boot" as Boot Policy created previously.

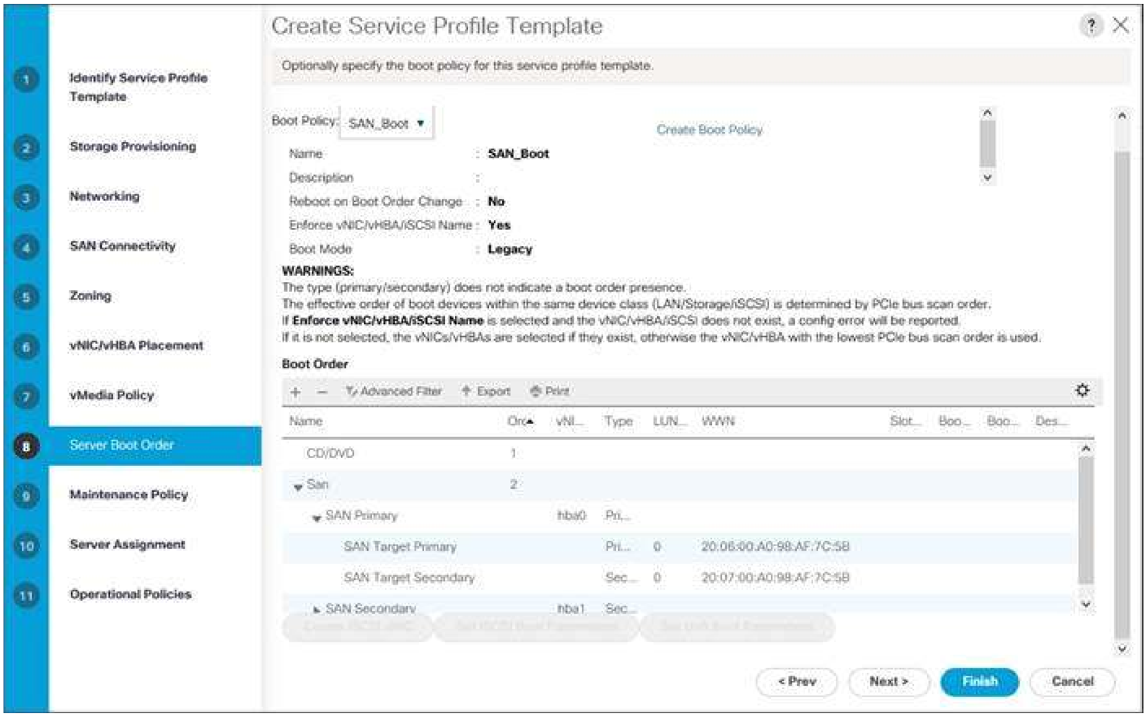

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

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	Ord	vNI...	Type	LUN...	WWN	Slot...	Boo...	Boo...	Des...
CD/DVD	1								
San	2								
SAN Primary		hba0	Pri...						
SAN Target Primary			Pri...	0	20:06:00:A0:98:AF:7C:58				
SAN Target Secondary			Sec...	0	20:07:00:A0:98:AF:7C:58				
SAN Secondary		hba1	Sec...						

12. The maintenance policy and server assignment options were left as default in the configuration.
13. In operational policies menu, select OLTP_BIOS as BIOS Policy, created previously.

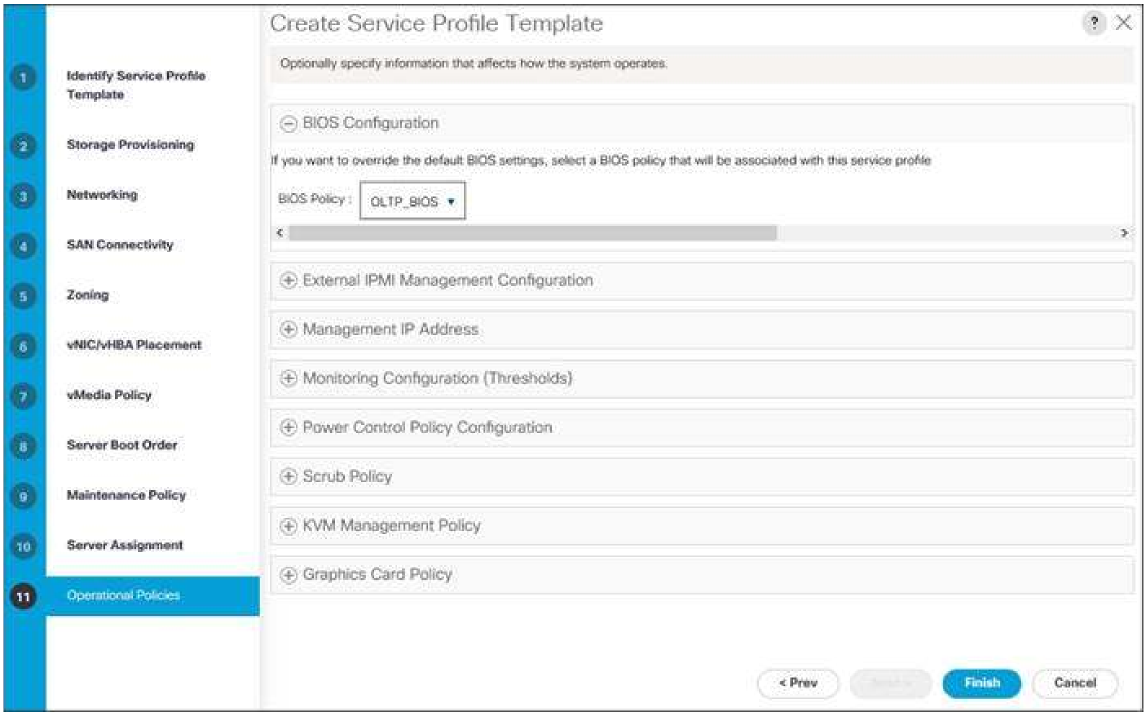

Create Service Profile Template

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy : OLTP_BIOS ▾

⊕ External IPMI Management Configuration
⊕ Management IP Address
⊕ Monitoring Configuration (Thresholds)
⊕ Power Control Policy Configuration
⊕ Scrub Policy
⊕ KVM Management Policy
⊕ Graphics Card Policy

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

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

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

Create Service Profiles from Template and Associate to Servers

Create Service Profiles from Template

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

To create Service Profiles from Template, follow these steps:

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

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

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

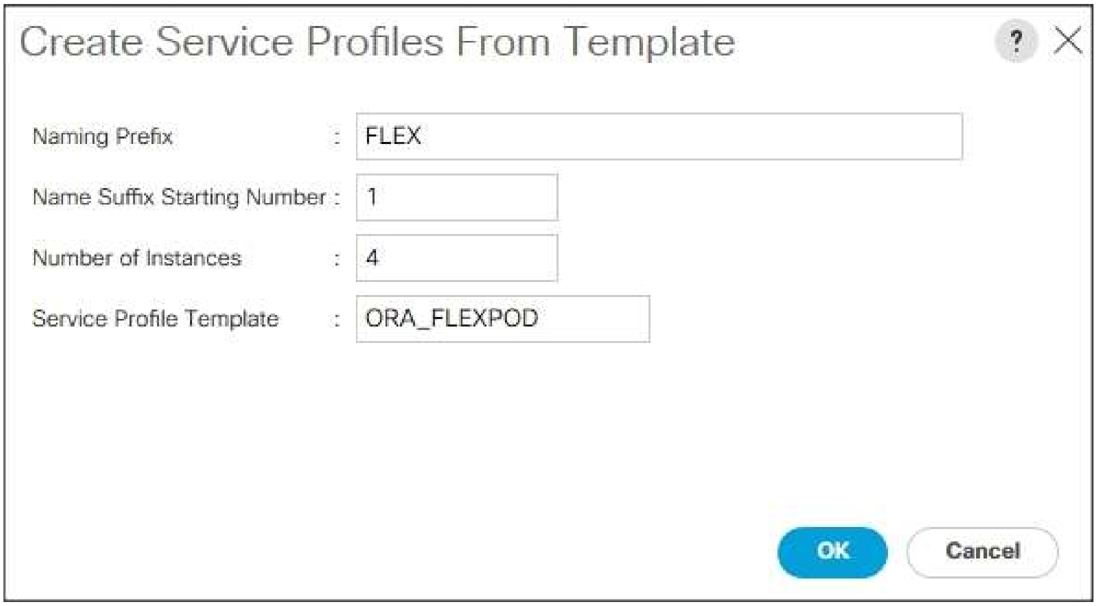

Create Service Profiles From Template

Naming Prefix : FLEX

Name Suffix Starting Number : 1

Number of Instances : 4

Service Profile Template : ORA_FLEXPOD

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

Associate Service Profiles to the Servers

To associate service profiles to the servers, follow these steps:

1. Under the Servers tab, select the desired service profile, and select change service profile association.

2. Right-click the name of service profile you want to associate with the server and select the option " Change Service Profile Association" .

3. In the Change Service Profile Association page, from the Server Assignment drop-down list, select the existing server that you want to assign and click OK.

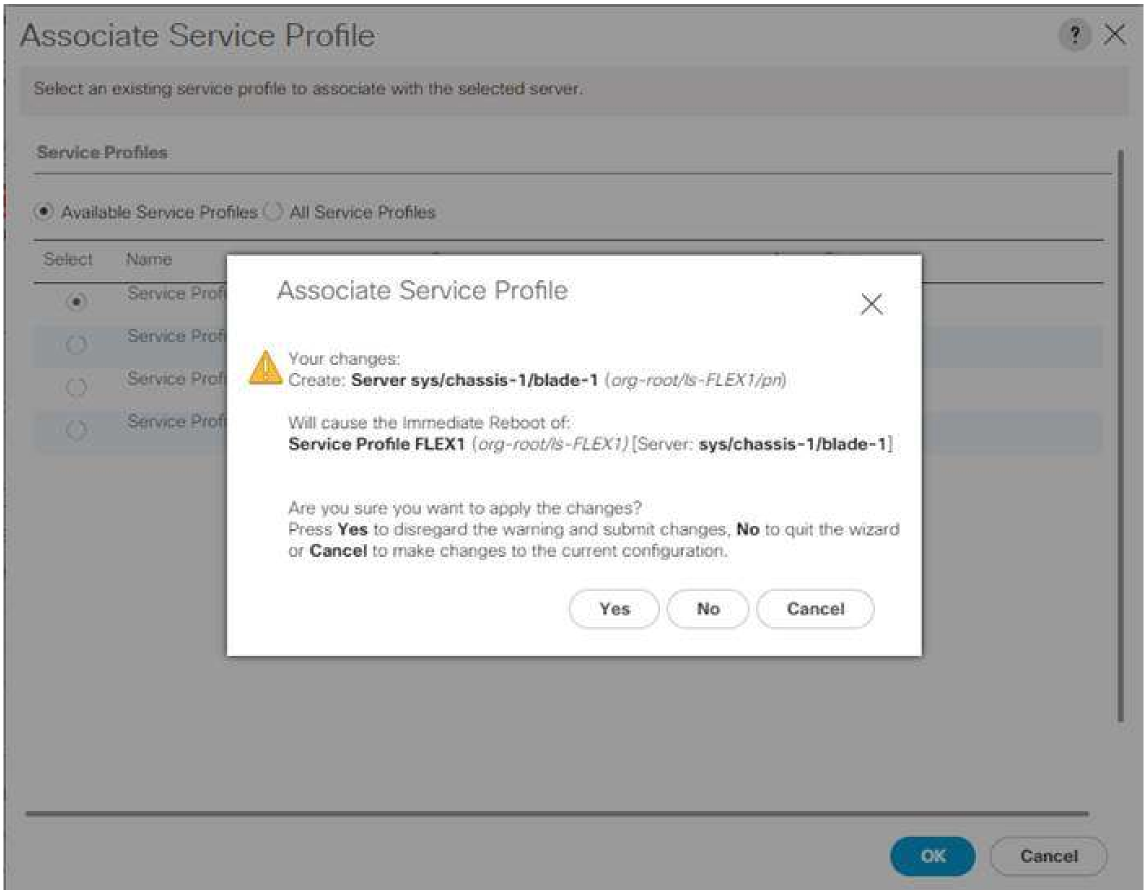

Associate Service Profile

Select an existing service profile to associate with the selected server.

Service Profiles

● Available Service Profiles ○ All Service Profiles

Select	Name
●	Service Prof
○	Service Prof
○	Service Prof
○	Service Prof

Associate Service Profile

Your changes:
Create: **Server sys/chassis-1/blade-1** (*org-root/ls-FLEX1/pn*)

Will cause the Immediate Reboot of:
Service Profile FLEX1 (*org-root/ls-FLEX1*) [Server: **sys/chassis-1/blade-1**]

Are you sure you want to apply the changes?
Press **Yes** to disregard the warning and submit changes, **No** to quit the wizard or **Cancel** to make changes to the current configuration.

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

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

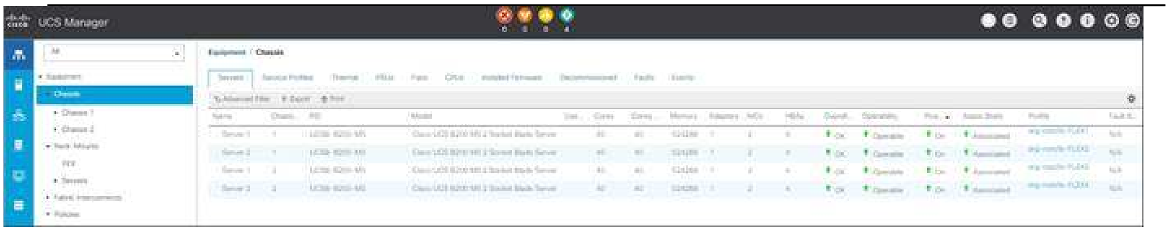

Name	Chassis	PID	Model	User...	Cores	Cores...	Memory	Adapters	NICs	HBAs	Overall...	Operability	Pow...	Assoc State	Profile	Fault S...
Server 1	1	UCSB-B200-M5	Cisco UCS B200 M5 2 Socket Blade Server		40	40	524288	1	2	4	OK	Operable	On	Associated	org-root/ls-FLEX1	N/A
Server 2	1	UCSB-B200-M5	Cisco UCS B200 M5 2 Socket Blade Server		40	40	524288	1	2	4	OK	Operable	On	Associated	org-root/ls-FLEX2	N/A
Server 1	2	UCSB-B200-M5	Cisco UCS B200 M5 2 Socket Blade Server		40	40	524288	1	2	4	OK	Operable	On	Associated	org-root/ls-FLEX3	N/A
Server 2	2	UCSB-B200-M5	Cisco UCS B200 M5 2 Socket Blade Server		40	40	524288	1	2	4	OK	Operable	On	Associated	org-root/ls-FLEX4	N/A

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

This completes the configuration required for Cisco UCS Manager Setup.

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

Configure Cisco MDS 9148S Switch

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

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

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

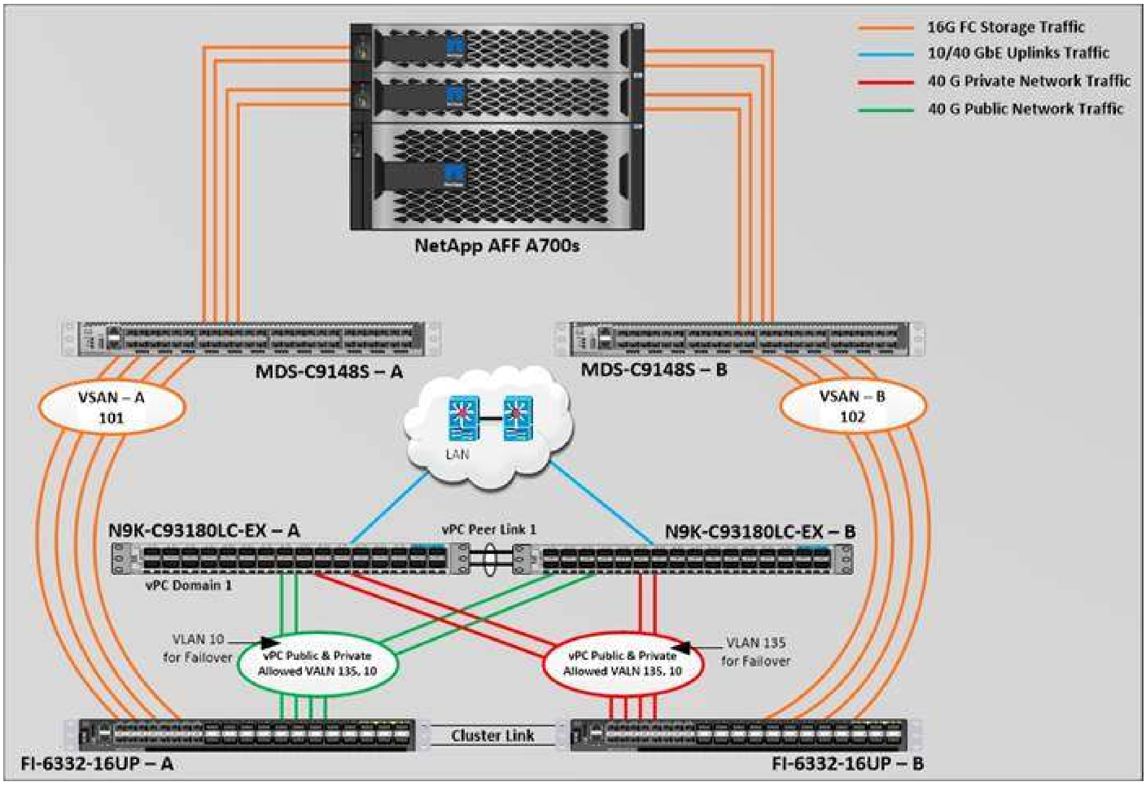

For this solution, we connected four ports (ports 1 to 4) of MDS Switch A to Fabric Interconnect A (ports 1-4). We also connected four ports (ports 1 to 4) of MDS Switch B to Fabric Interconnect B (ports 1-4). All ports carry 16 Gb/s FC Traffic. Table 8 lists port connectivity of Cisco MDS Switches to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

Table 8 MDS Switch Connectivity to the Fabric Interconnects

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

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

Table 9 MDS Switch Connectivity to the NetApp AFF A700s Storage

MDS Switch	MDS Switch Port	NetApp Storage Controller	NetApp Controller Ports	NetApp Port Description
MDS Switch	FC Port 1/9	NetApp AFF A700s	FC Port 0 - Slot	NetApp-A700s-

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

Configure feature for MDS Switch A and MDS Switch B

To configure feature on MDS Switches, follow these steps:

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

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

2. Login as admin user into MDS Switch B:

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

Configure VSANs for MDS Switch A and MDS Switch B

To create VSANs, follow these steps:

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

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

FLEXPOD-MDS-A(config)# interface fc 1/1-12
FLEXPOD-MDS-A(config-if)# switchport trunk allowed vsan 101
```

```
FLEXPOD-MDS-A(config-if)# switchport trunk mode off
FLEXPOD-MDS-A(config-if)# port-license acquire
FLEXPOD-MDS-A(config-if)# no shutdown
FLEXPOD-MDS-A(config-if)# exit
FLEXPOD-MDS-A(config)# copy running-config startup-config
```

3. Login as admin user into MDS Switch B.

4. Create VSAN 102 for Storage Traffic:

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

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

Create and Configure Fibre Channel Zoning

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

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

To create and configure the Fibre channel zoning, follow these steps:

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

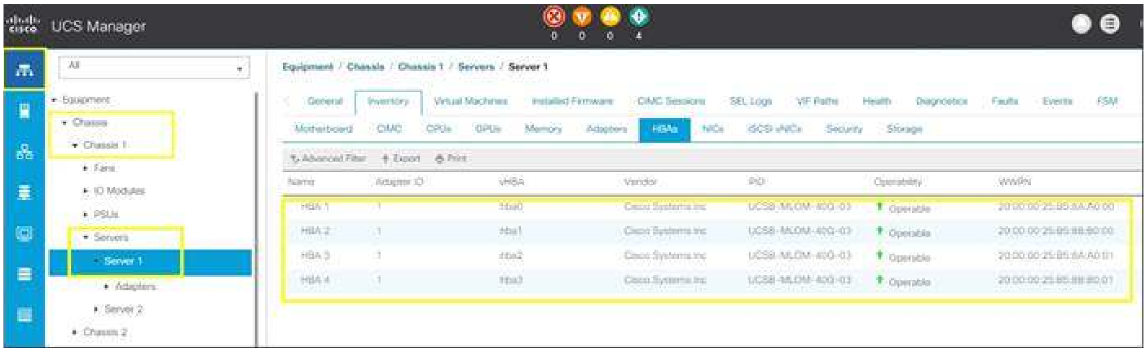

Name	Adapter ID	vHBA	Vendor	PID	Operability	WWPN
HBA 1	1	hba0	Cisco Systems Inc	UCSB-MLOM-40G-03	Operable	20:00:00:25:B5:AA:A0:00
HBA 2	1	hba1	Cisco Systems Inc	UCSB-MLOM-40G-03	Operable	20:00:00:25:B5:BB:B0:00
HBA 3	1	hba2	Cisco Systems Inc	UCSB-MLOM-40G-03	Operable	20:00:00:25:B5:AA:A0:01
HBA 4	1	hba3	Cisco Systems Inc	UCSB-MLOM-40G-03	Operable	20:00:00:25:B5:BB:B0:01

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

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

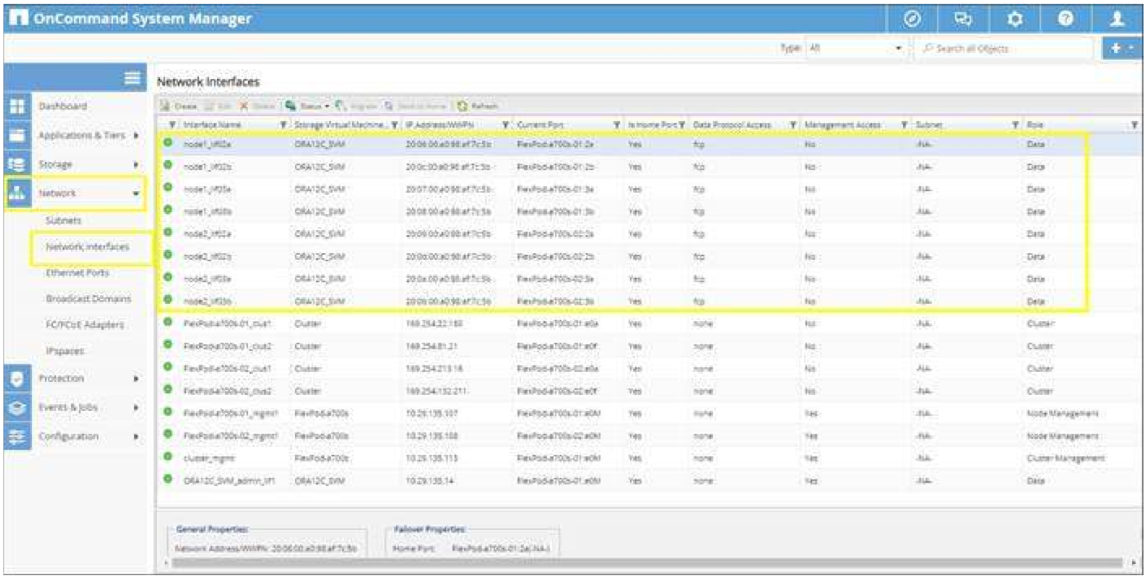

Network Interfaces

Interface Name	Storage Virtual Machine...	IP Address/WWPN	Current Port	Is Home Port	Data Protocol Access	Management Access	Subnet	Role
node1_lif02a	ORA12C_SVM	20:06:00:a0:98:af:7c:5b	FlexPod-a700s-01:2a	Yes	fcp	No	-NA-	Data
node1_lif02b	ORA12C_SVM	20:0c:00:a0:98:af:7c:5b	FlexPod-a700s-01:2b	Yes	fcp	No	-NA-	Data
node1_lif03a	ORA12C_SVM	20:07:00:a0:98:af:7c:5b	FlexPod-a700s-01:3a	Yes	fcp	No	-NA-	Data
node1_lif03b	ORA12C_SVM	20:08:00:a0:98:af:7c:5b	FlexPod-a700s-01:3b	Yes	fcp	No	-NA-	Data
node2_lif02a	ORA12C_SVM	20:09:00:a0:98:af:7c:5b	FlexPod-a700s-02:2a	Yes	fcp	No	-NA-	Data
node2_lif02b	ORA12C_SVM	20:0d:00:a0:98:af:7c:5b	FlexPod-a700s-02:2b	Yes	fcp	No	-NA-	Data
node2_lif03a	ORA12C_SVM	20:0a:00:a0:98:af:7c:5b	FlexPod-a700s-02:3a	Yes	fcp	No	-NA-	Data
node2_lif03b	ORA12C_SVM	20:0b:00:a0:98:af:7c:5b	FlexPod-a700s-02:3b	Yes	fcp	No	-NA-	Data
FlexPod-a700s-01_clus1	Cluster	169.254.22.180	FlexPod-a700s-01:e0a	Yes	none	No	-NA-	Cluster
FlexPod-a700s-01_clus2	Cluster	169.254.81.21	FlexPod-a700s-01:e0f	Yes	none	No	-NA-	Cluster
FlexPod-a700s-02_clus1	Cluster	169.254.213.16	FlexPod-a700s-02:e0a	Yes	none	No	-NA-	Cluster
FlexPod-a700s-02_clus2	Cluster	169.254.132.211	FlexPod-a700s-02:e0f	Yes	none	No	-NA-	Cluster
FlexPod-a700s-01_mgmt1	FlexPod-a700s	10.29.135.107	FlexPod-a700s-01:e0M	Yes	none	Yes	-NA-	Node Management
FlexPod-a700s-02_mgmt1	FlexPod-a700s	10.29.135.108	FlexPod-a700s-02:e0M	Yes	none	Yes	-NA-	Node Management
cluster_mgmt	FlexPod-a700s	10.29.135.113	FlexPod-a700s-01:e0M	Yes	none	Yes	-NA-	Cluster Management
ORA12C_SVM_admin_lif1	ORA12C_SVM	10.29.135.14	FlexPod-a700s-01:e0M	Yes	none	Yes	-NA-	Data

General Properties:
Network Address/WWPN: 20:06:00:a0:98:af:7c:5b

Failover Properties:
Home Port: FlexPod-a700s-01:2a(-NA-)

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

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

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

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

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

```
configure terminal
device-alias database
 device-alias name flex1-hba0 pwwn 20:00:00:25:b5:8a:a0:00
 device-alias name flex1-hba2 pwwn 20:00:00:25:b5:8a:a0:01
 device-alias name flex2-hba0 pwwn 20:00:00:25:b5:8a:a0:02
 device-alias name flex2-hba2 pwwn 20:00:00:25:b5:8a:a0:03
 device-alias name flex3-hba0 pwwn 20:00:00:25:b5:8a:a0:04
 device-alias name flex3-hba2 pwwn 20:00:00:25:b5:8a:a0:05
 device-alias name flex4-hba0 pwwn 20:00:00:25:b5:8a:a0:06
 device-alias name flex4-hba2 pwwn 20:00:00:25:b5:8a:a0:07
 device-alias name NetApp-A700s-01-2A pwwn 20:06:00:a0:98:af:7c:5b
 device-alias name NetApp-A700s-01-3A pwwn 20:07:00:a0:98:af:7c:5b
```

```
  device-alias name NetApp-A700s-02-2A pwwn 20:09:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-3A pwwn 20:0a:00:a0:98:af:7c:5b
  device-alias commit
copy run start
```

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

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

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

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

Create Zoning

Cisco MDS Switch A

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

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

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

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

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

zone name flex4 vsan 101
    member pwwn 20:00:00:25:b5:8a:a0:06
    member pwwn 20:00:00:25:b5:8a:a0:07
    member pwwn 20:06:00:a0:98:af:7c:5b
    member pwwn 20:07:00:a0:98:af:7c:5b
    member pwwn 20:09:00:a0:98:af:7c:5b
    member pwwn 20:0a:00:a0:98:af:7c:5b
```

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

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

4. Activate the zone set by running following commands

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

Cisco MDS Switch B

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

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

```
configure terminal
zone name flex1 vsan 102
```

```
    member pwwn 20:00:00:25:b5:8b:b0:00
    member pwwn 20:00:00:25:b5:8b:b0:01
    member pwwn 20:0c:00:a0:98:af:7c:5b
    member pwwn 20:08:00:a0:98:af:7c:5b
    member pwwn 20:0d:00:a0:98:af:7c:5b
    member pwwn 20:0b:00:a0:98:af:7c:5b

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

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

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

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

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

4. Activate the zone set by running following commands:

```
zoneset activate name flex vsan 102
```

copy run start

Verify FC Ports on MDS Switch

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

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

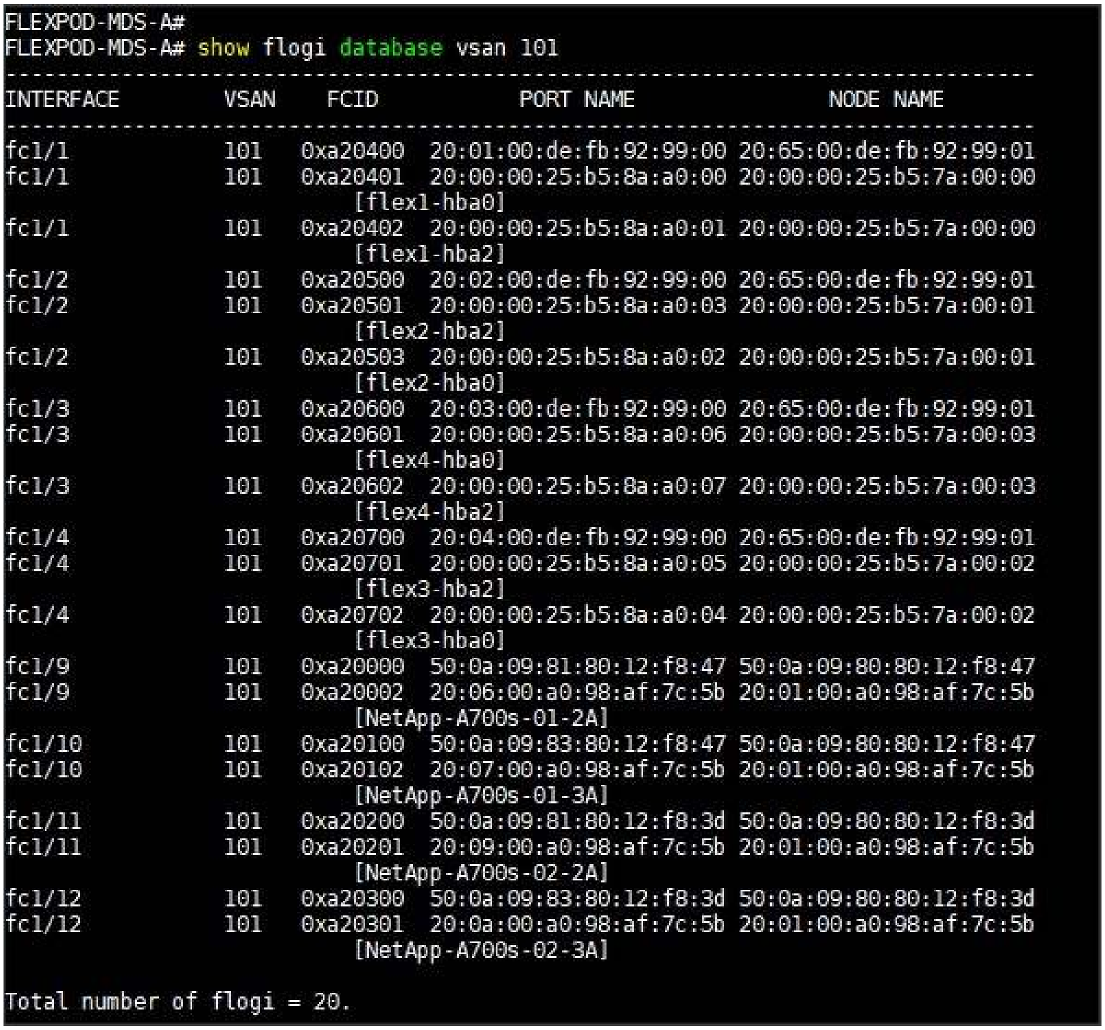

```
FLEXPOD-MDS-A#
FLEXPOD-MDS-A# show flogi database vsan 101
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/1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

Total number of flogi = 20.
```

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

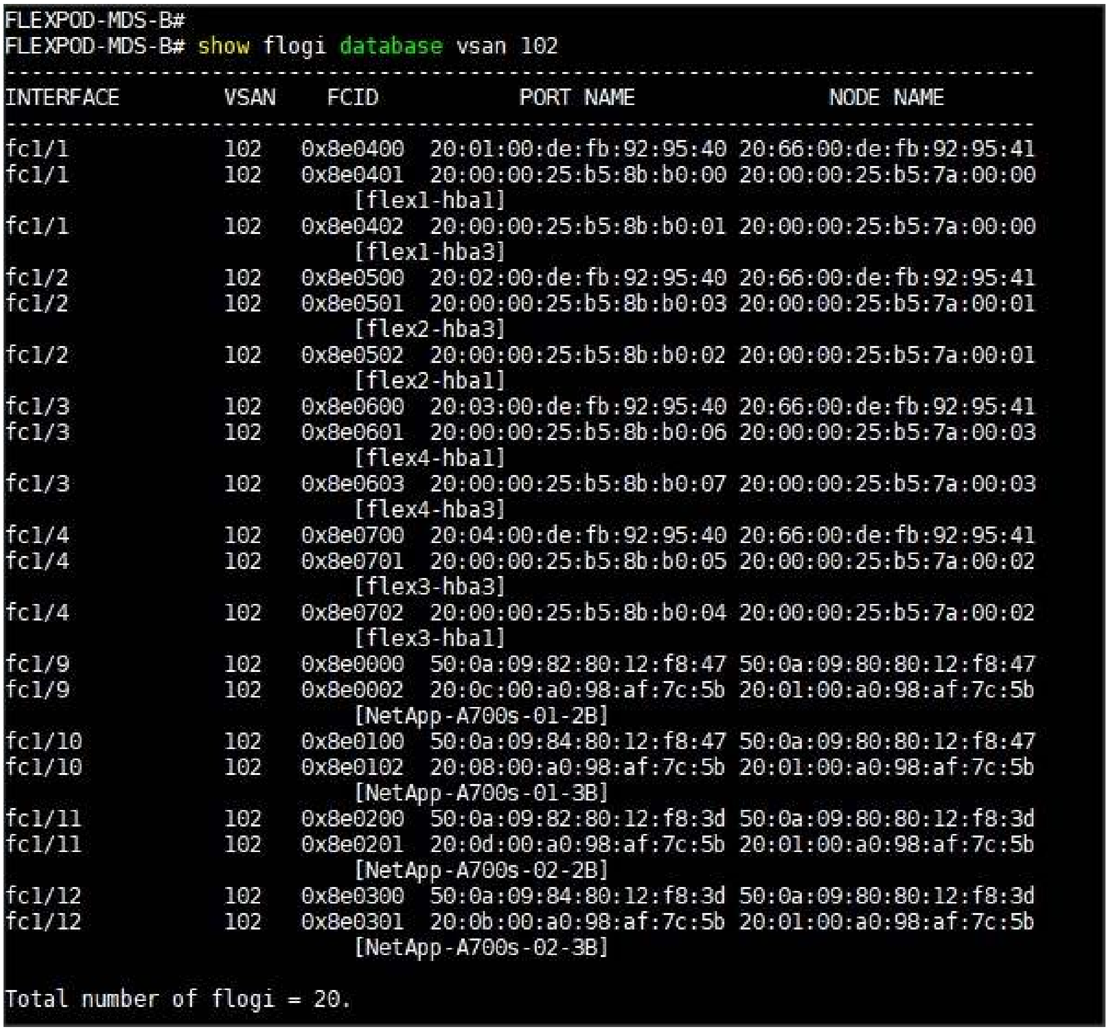

```
FLEXPOD-MDS-B#
FLEXPOD-MDS-B# show flogi database vsan 102
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/1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

Total number of flogi = 20.
```

Configure NetApp AFF A700s Storage

NetApp Storage Connectivity

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

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

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

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

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

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

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

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

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

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

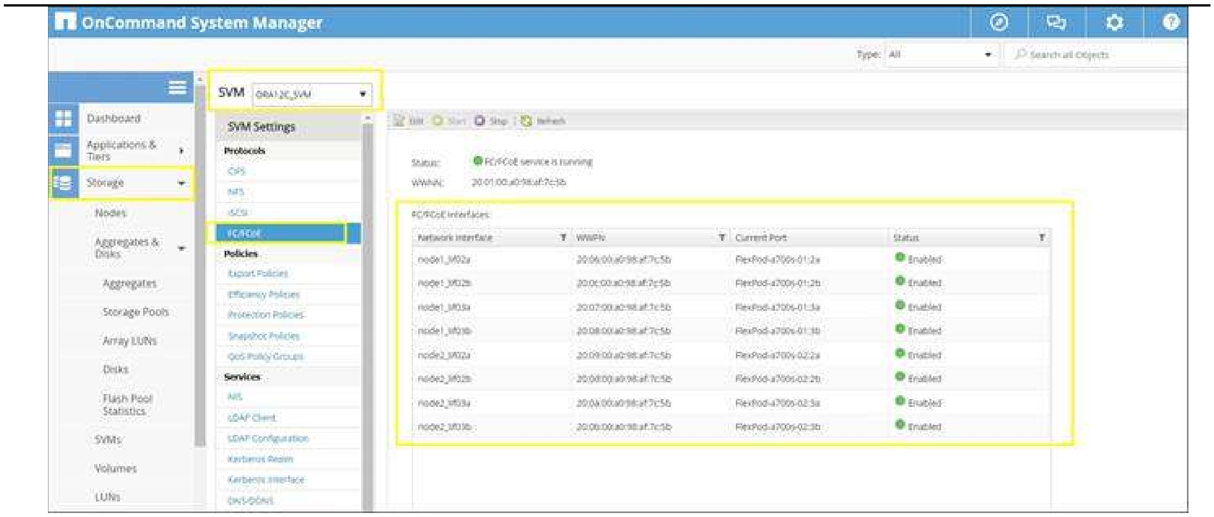

Status: FC/FCoE service is running
WWNN: 20:01:00:a0:98:af:7c:5b

FC/FCoE Interfaces:

Network Interface	WWPN	Current Port	Status
node1_lif02a	20:06:00:a0:98:af:7c:5b	FlexPod-a700s-01:2a	Enabled
node1_lif02b	20:0c:00:a0:98:af:7c:5b	FlexPod-a700s-01:2b	Enabled
node1_lif03a	20:07:00:a0:98:af:7c:5b	FlexPod-a700s-01:3a	Enabled
node1_lif03b	20:08:00:a0:98:af:7c:5b	FlexPod-a700s-01:3b	Enabled
node2_lif02a	20:09:00:a0:98:af:7c:5b	FlexPod-a700s-02:2a	Enabled
node2_lif02b	20:0d:00:a0:98:af:7c:5b	FlexPod-a700s-02:2b	Enabled
node2_lif03a	20:0a:00:a0:98:af:7c:5b	FlexPod-a700s-02:3a	Enabled
node2_lif03b	20:0b:00:a0:98:af:7c:5b	FlexPod-a700s-02:3b	Enabled

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

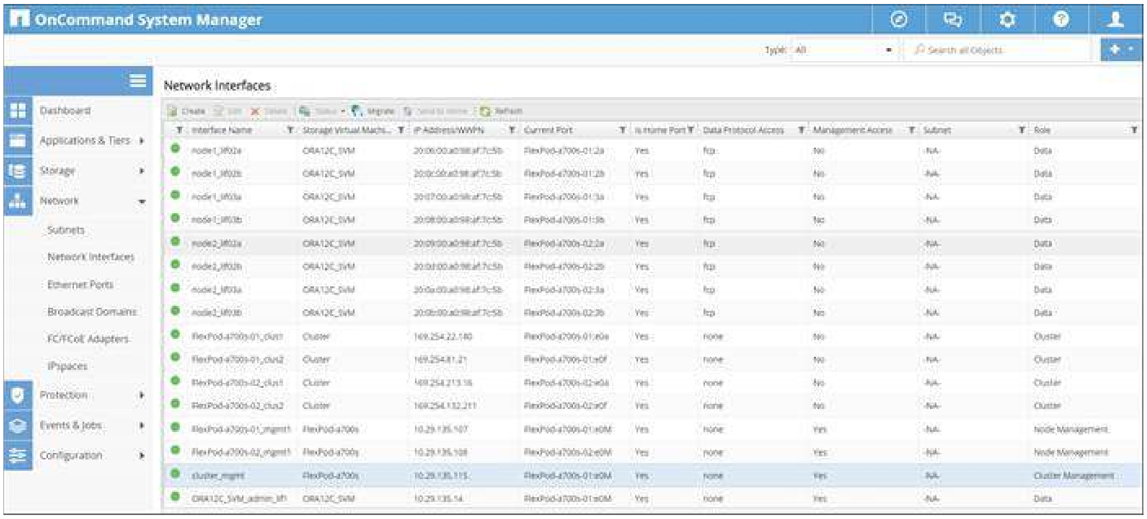

Network Interfaces

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

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

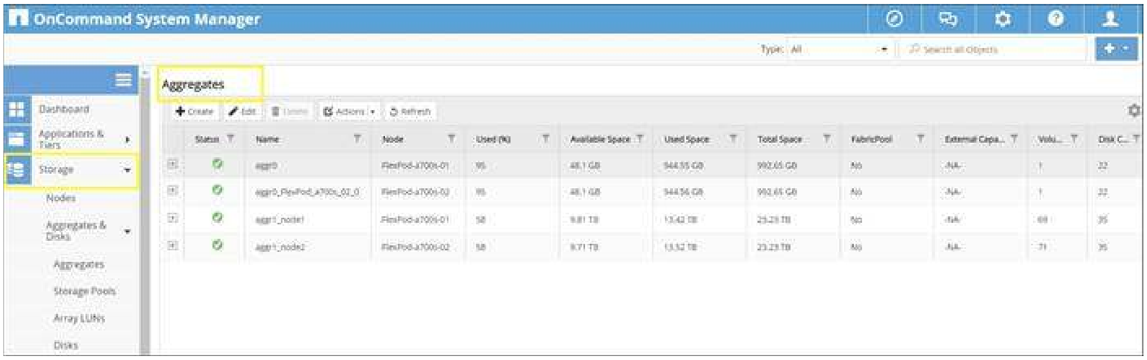

Aggregates

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

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

Operating System and Database Deployment

The design goal of the reference architecture was to best represent a real-world environment as closely as possible. The approach included features of Cisco UCS to rapidly deploy stateless servers and use NetApp AFF Storage A700s boot LUNs to provision the O.S on top it. Zoning was performed on the Cisco MDS 9148S switches to enable the initiators discover the targets during boot process.

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

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

Operating System Configuration

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

1. Download Oracle Linux 7.5 OS image from https://edelivery.oracle.com/linux.

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

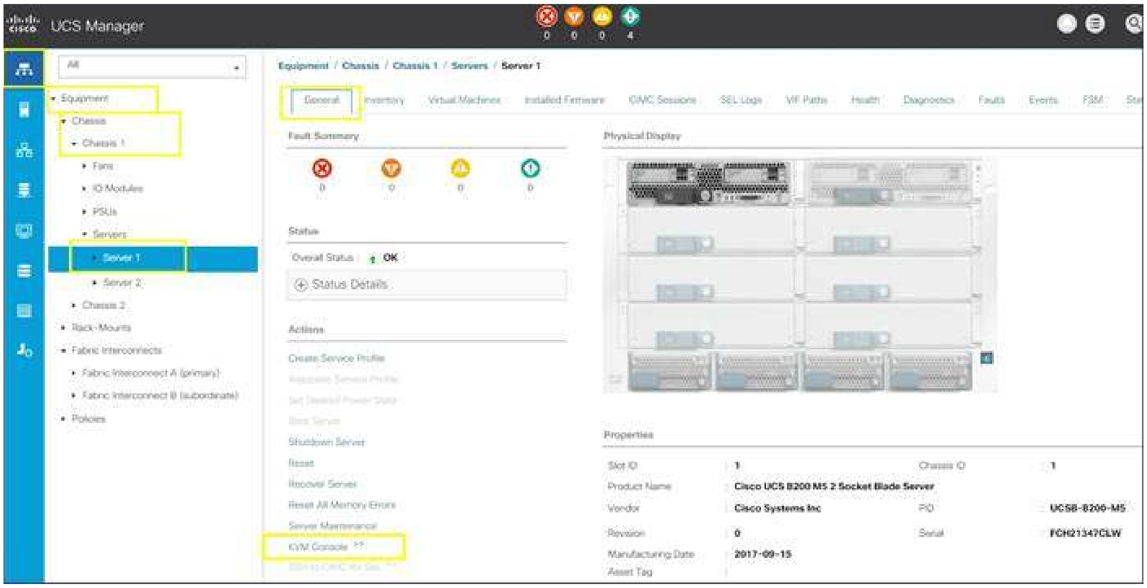

Overall Status : OK

Slot ID : 1
Chassis ID : 1
Product Name : Cisco UCS B200 M5 2 Socket Blade Server
Vendor : Cisco Systems Inc
PID : UCSB-B200-M5
Revision : 0
Serial : FCH21347CLW
Manufacturing Date : 2017-09-15
Asset Tag :

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

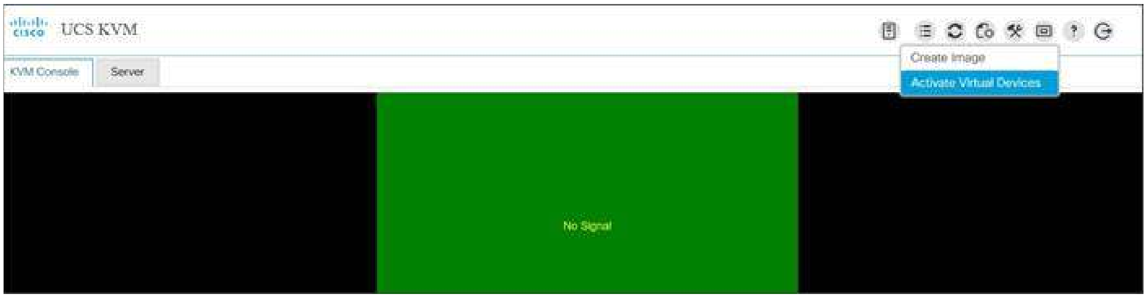

4. When the Server starts booting, it will detect the NetApp Storage active FC paths as shown below. If you see the following message in the KVM console while the server is rebooting along with the target WWPNs, it confirms the setup is done correctly and boot from SAN will be successful.

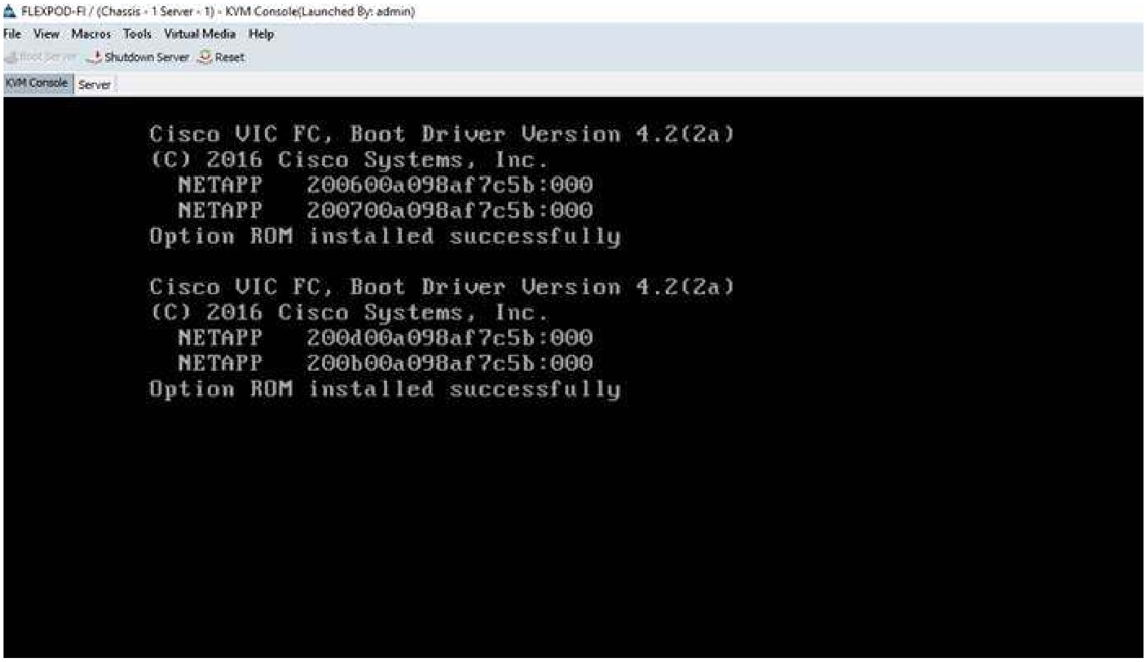

```
Cisco VIC FC, Boot Driver Version 4.2(2a)
(C) 2016 Cisco Systems, Inc.
  NETAPP   200600a098af7c5b:000
  NETAPP   200700a098af7c5b:000
Option ROM installed successfully

Cisco VIC FC, Boot Driver Version 4.2(2a)
(C) 2016 Cisco Systems, Inc.
  NETAPP   200d00a098af7c5b:000
  NETAPP   200b00a098af7c5b:000
Option ROM installed successfully
```

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

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

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

Operating System Prerequisites **for Oracle Software Installation**

Configure BIOS

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

Configure BIOS for OLTP Workloads

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

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

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

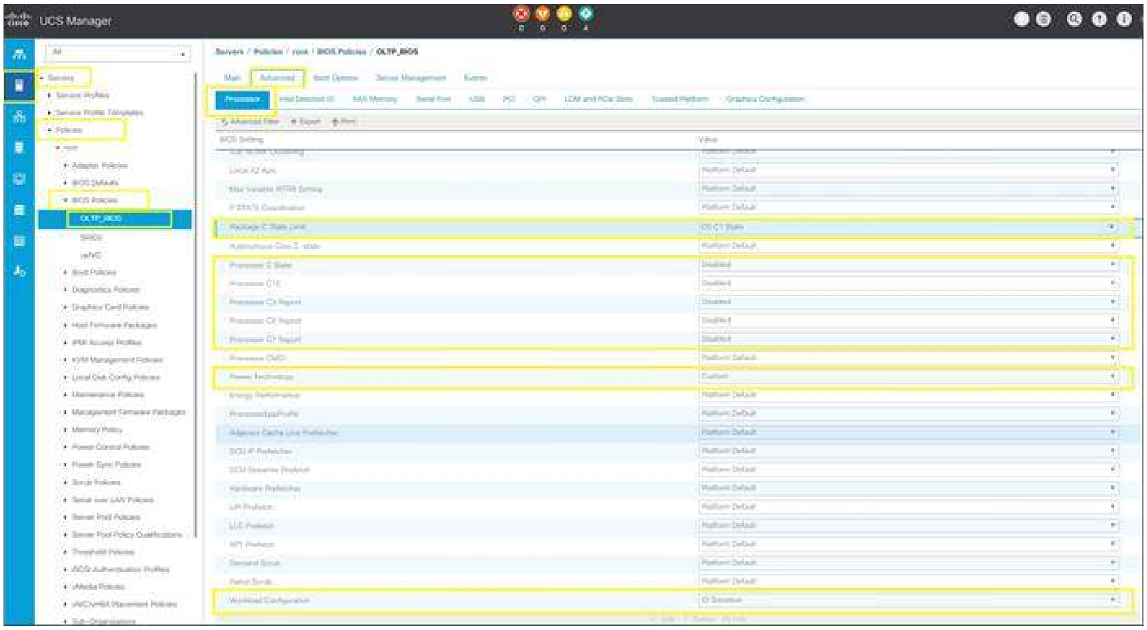

BIOS Setting	Value
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	C0 C1 State
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Custom
Energy Performance	Platform Default
ProcessorEppProfile	Platform Default
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	IO Sensitive

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

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

```
[root@flex1 ~]# cat /etc/default/grub

GRUB_TIMEOUT=5

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

GRUB_DEFAULT=saved

GRUB_DISABLE_SUBMENU=true

GRUB_TERMINAL_OUTPUT="console"

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

GRUB_DISABLE_RECOVERY="true"
```

Prerequisites Automatic Installation

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

To configure operating system pre-requisite for Oracle 12cR2 software on all four nodes, follow these steps:

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

To configure the prerequisites on all the four nodes, follow these steps:

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

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

```
[root@flex1 ~]# yum install oracle-database-server-12cR2-preinstall
```

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

Additional Prerequisites Setup

After configuring automatic or manual prerequisites steps, you have to configure a few additional steps to complete the prerequisites for the Oracle database software installations on all the four nodes as described in the following sections.

Disable SELinux

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

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

```
SELINUX= permissive
```

Disable Firewall

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

```
systemctl status firewalld.service
```

```
systemctl stop firewalld.service
```

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

```
systemctl disable firewalld.service
```

Set the User Passwords

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

```
passwd oracle
```

```
passwd grid
```

Configure Multipath Setup

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

---

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

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

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

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

```
[root@oraracx1 ~]# cat /etc/multipath.conf
defaults {
    find_multipaths yes
    user_friendly_names no
}
multipaths {
    multipath {
        wwid          3600a098038304173475d4c766a49744e
        alias         node1_os
    }
    multipath {
        wwid          3600a098038304173475d4c766a49754a
        alias         ocrvote_1
    }
}
```

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

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

Configure UDEV Rules

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

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

```
 [root@flex1 ~]# cat /etc/udev/rules.d/99-oracle-asmdevices.rules

#All LUNs which starts with dg_orarac_* #

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

#All LUNs which starts with dg_oradata_* #

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

#All LUNs which starts with dg_oraredo_* #

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

#All LUNs which starts with dg_oraarchive_* #

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

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

Configure Public and Private NICs on Each RAC Node

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

Log in as a root user on each node and go to "`/etc/sysconfig/network-scripts`" and configure Public network and Private network IP Address. Configure the private and public NICs with the appropriate IP addresses across all the Oracle RAC nodes.

Configure "/etc/hosts" on Each RAC Node

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

```
[root@flex1 ~]# cat /etc/hosts

127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4


##### Public IP #####

10.29.135.121   flex1.cisco.com flex1

10.29.135.122   flex2.cisco.com flex2

10.29.135.123   flex3.cisco.com flex3

10.29.135.124   flex4.cisco.com flex4

##### Virtual IP #####

10.29.135.125   flex1-vip       flex1-vip.cisco.com

10.29.135.126   flex2-vip       flex2-vip.cisco.com

10.29.135.127   flex3-vip       flex3-vip.cisco.com
```

```
10.29.135.128   flex4-vip      flex4-vip.cisco.com
###### Private IP ######
10.10.10.121  flex1-priv      flex1-priv.cisco.com
10.10.10.122  flex2-priv      flex2-priv.cisco.com
10.10.10.123  flex3-priv      flex3-priv.cisco.com
10.10.10.124  flex4-priv      flex4-priv.cisco.com
#SCAN IP
10.29.135.129   flex-cluster   flex-cluster.cisco.com
10.29.135.130   flex-cluster   flex-cluster.cisco.com
10.29.135.131   flex-cluster   flex-cluster.cisco.com
```

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

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

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

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

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

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

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

Oracle Database 12c GRID Infrastructure Deployment

For this Oracle Database FlexPod solution, you will install the Oracle Grid and Database software on all four nodes (flex1, flex2, flex3 and flex4). The installation guides you through the gathering of all node information and configuring the ASM devices and all the prerequisite validations for GI.

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

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

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

Create Directory Structure

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

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

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

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

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

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

```
cd /u01/app/12.2.0/grid
unzip -q download_location/linuxx64_12201_grid_home
```

Run Cluster Verification Utility

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

```
./runcluvfy.sh stage -pre crsinst -n flex1,flex2,flex3,flex4 –verbose
```

Configure HugePages

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

Advantages of HugePages:

- HugePages are not swappable so there is no page-in/page-out mechanism overhead.
- HugePages uses fewer pages to cover the physical address space, so the size of " book keeping" (mapping from the virtual to the physical address) decreases, so it requiring fewer entries in the TLB and so TLB hit ratio improves.
- HugePages reduces page table overhead. Also, HugePages eliminated page table lookup overhead: Since the pages are not subject to replacement, page table lookups are not required.
- Faster overall memory performance: On virtual memory systems each memory operation is actually two abstract memory operations. Since there are fewer pages to work on, the possible bottleneck on page table access is clearly avoided.

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

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

Install and Configure Oracle Database Grid Infrastructure Software

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

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

To install Oracle Database Grid Infrastructure Software, follow these steps:

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

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

```
./gridSetup.sh
```

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

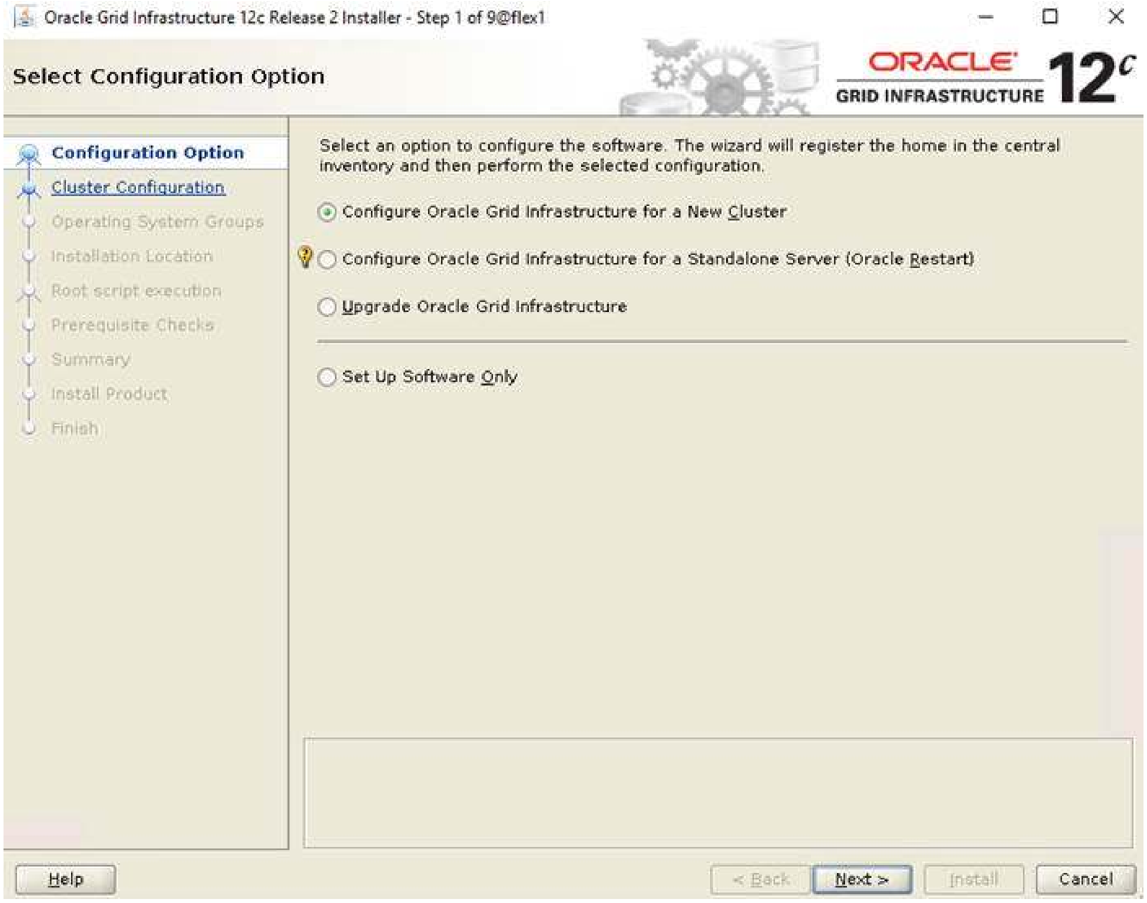

Select Configuration Option

Select an option to configure the software. The wizard will register the home in the central inventory and then perform the selected configuration.

● Configure Oracle Grid Infrastructure for a New Cluster

○ Configure Oracle Grid Infrastructure for a Standalone Server (Oracle Restart)

○ Upgrade Oracle Grid Infrastructure

○ Set Up Software Only

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

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

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

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

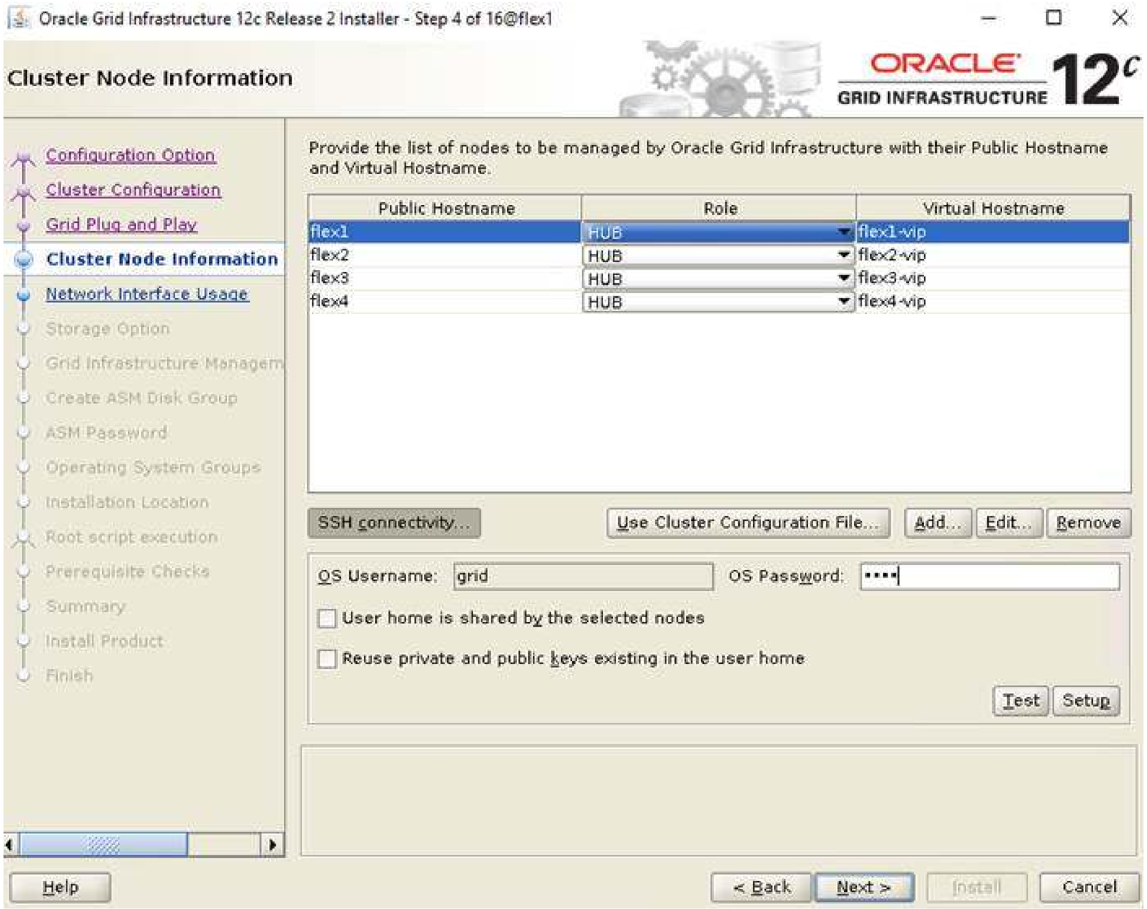

Oracle Grid Infrastructure 12c Release 2 Installer - Step 4 of 16@flex1

Cluster Node Information

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

Public Hostname	Role	Virtual Hostname
flex1	HUB	flex1-vip
flex2	HUB	flex2-vip
flex3	HUB	flex3-vip
flex4	HUB	flex4-vip

OS Username: grid

OS Password: ••••

□ User home is shared by the selected nodes

□ Reuse private and public keys existing in the user home

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

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

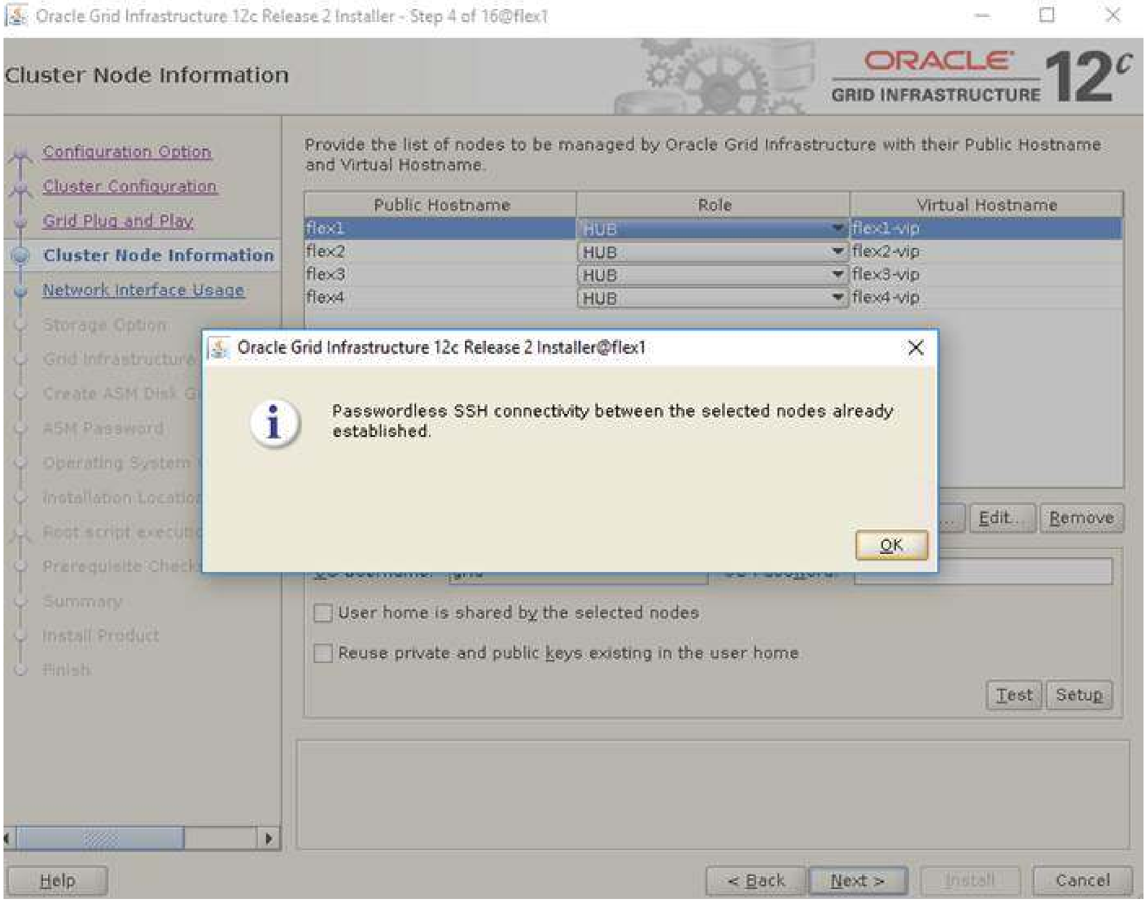

Provide the list of nodes to be managed by Oracle Grid Infrastructure with their Public Hostname and Virtual Hostname.

Public Hostname	Role	Virtual Hostname
flex1	HUB	flex1-vip
flex2	HUB	flex2-vip
flex3	HUB	flex3-vip
flex4	HUB	flex4-vip

□ User home is shared by the selected nodes

□ Reuse private and public keys existing in the user home

Oracle Grid Infrastructure 12c Release 2 Installer@flex1

Passwordless SSH connectivity between the selected nodes already established.

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

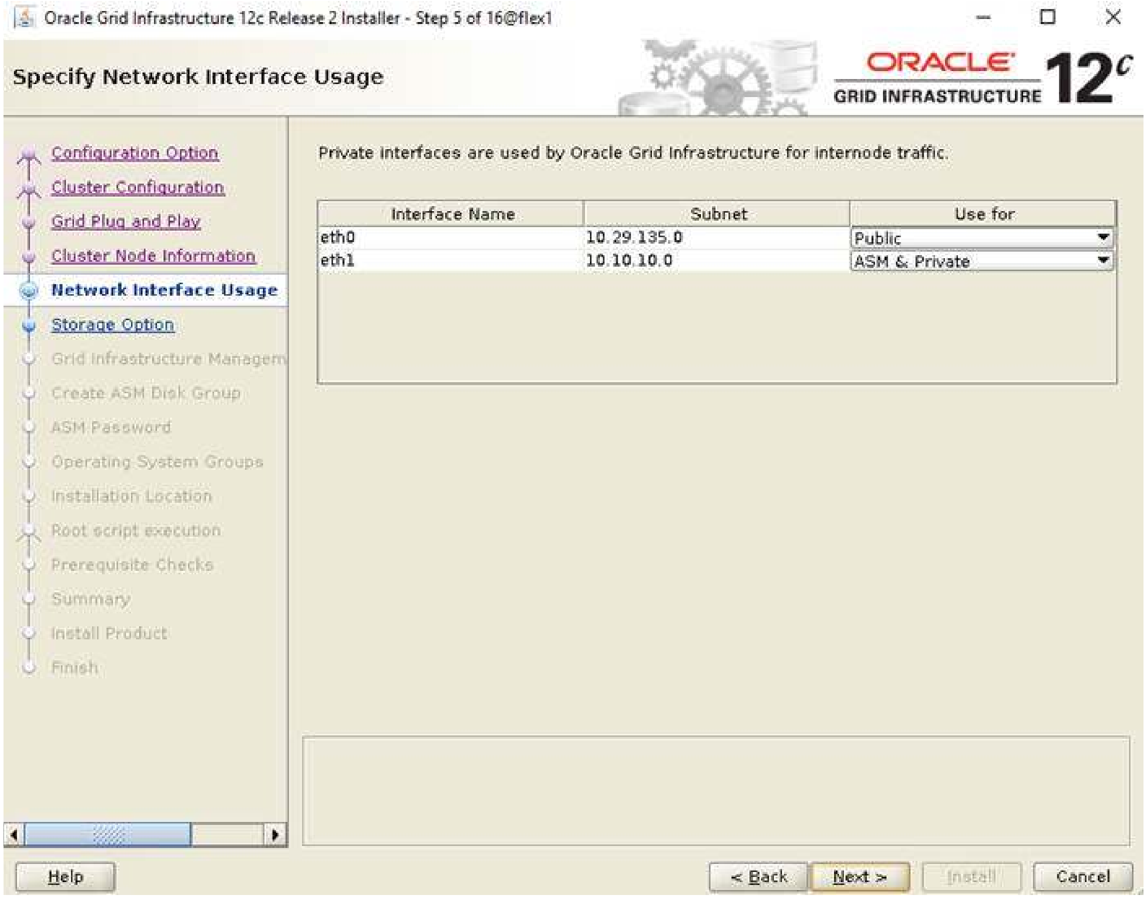

Specify Network Interface Usage

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

Interface Name	Subnet	Use for
eth0	10.29.135.0	Public
eth1	10.10.10.0	ASM & Private

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

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

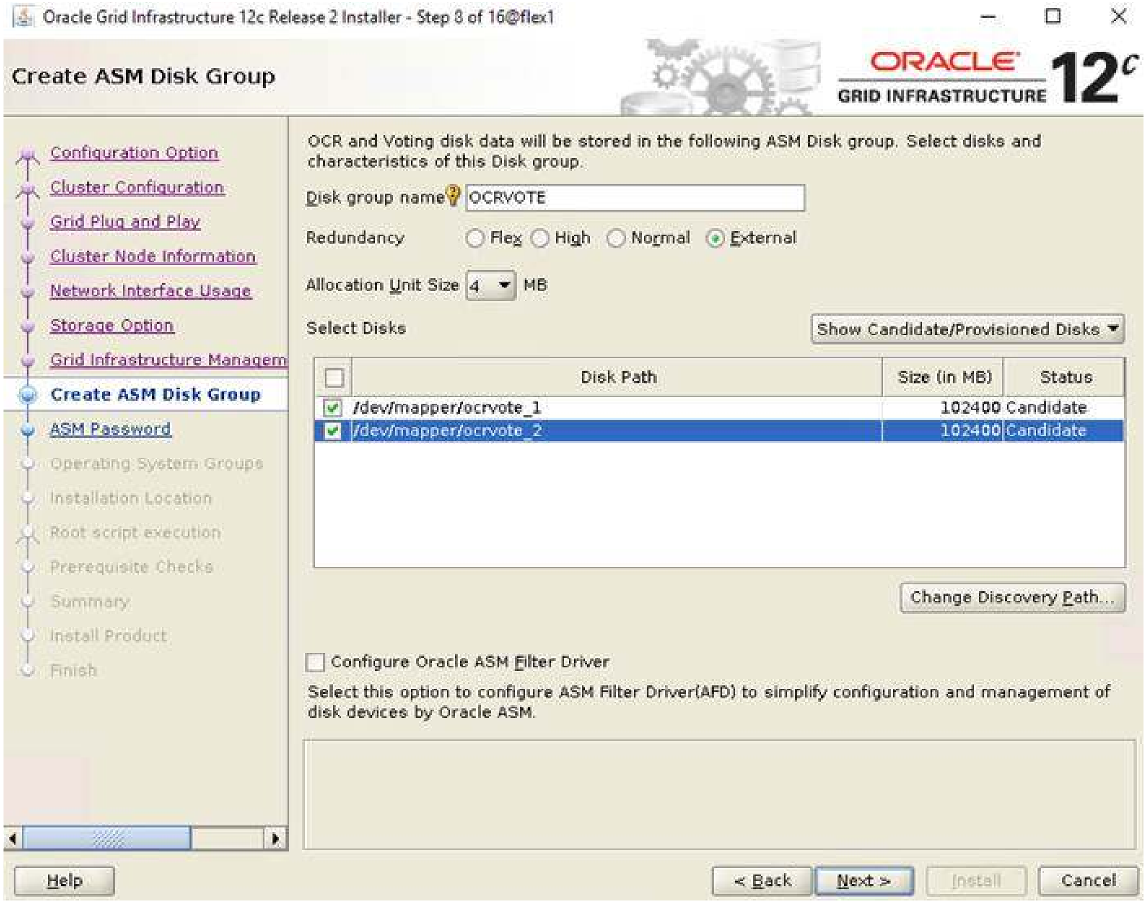

Oracle Grid Infrastructure 12c Release 2 Installer - Step 8 of 16@flex1

Create ASM Disk Group

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: OCRVOTE

Redundancy ○ Flex ○ High ○ Normal ● External

Allocation Unit Size 4 MB

Select Disks — Show Candidate/Provisioned Disks

□	Disk Path	Size (in MB)	Status
☑	/dev/mapper/ocrvote_1	102400	Candidate
☑	/dev/mapper/ocrvote_2	102400	Candidate

□ 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.

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

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

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

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

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

15. Specify the directory to use for the Oracle base for the Oracle Grid Infrastructure installation and then click Next. The Oracle base directory must be different from the Oracle home directory. Click Next and choose Inventory Directory according to your setup.

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

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

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

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

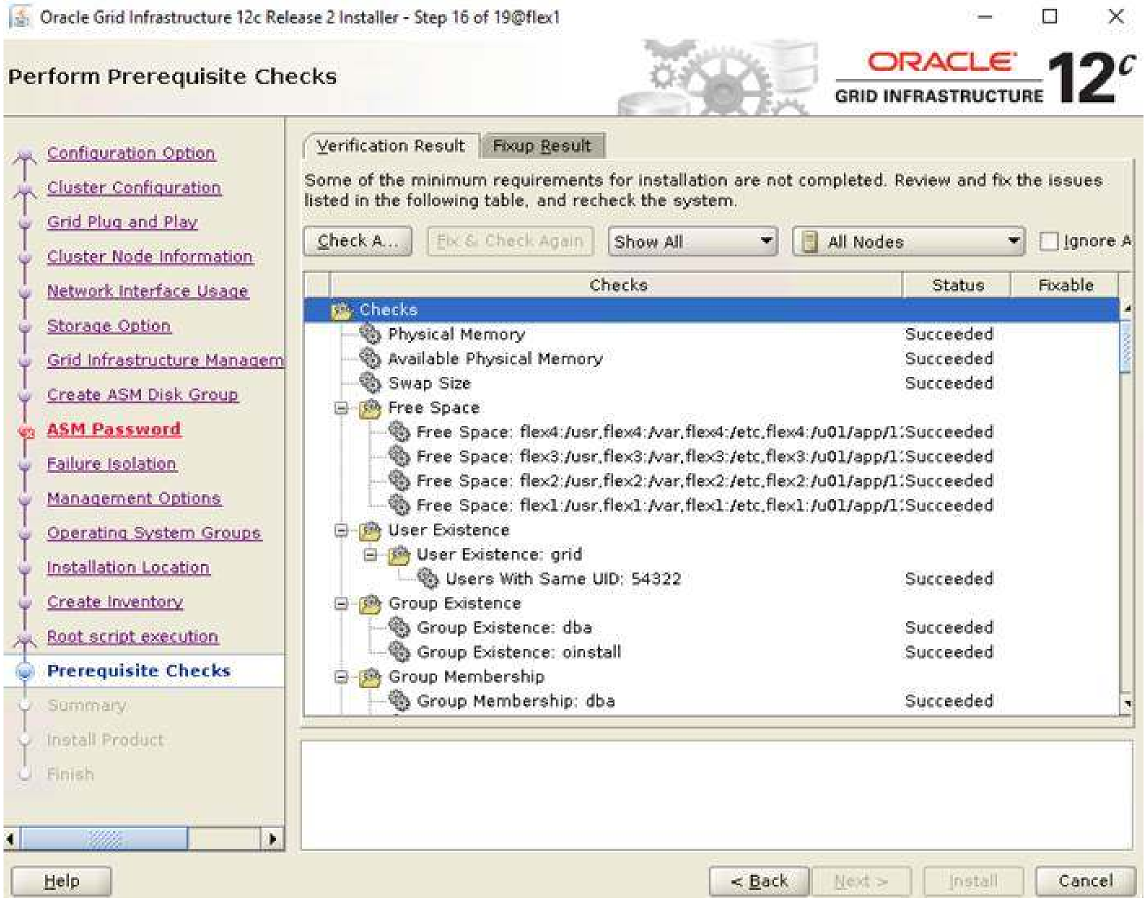

Some of the minimum requirements for installation are not completed. Review and fix the issues listed in the following table, and recheck the system.

Checks	Status	Fixable
Checks		
Physical Memory	Succeeded	
Available Physical Memory	Succeeded	
Swap Size	Succeeded	
Free Space		
Free Space: flex4:/usr,flex4:/var,flex4:/etc,flex4:/u01/app/1:	Succeeded	
Free Space: flex3:/usr,flex3:/var,flex3:/etc,flex3:/u01/app/1:	Succeeded	
Free Space: flex2:/usr,flex2:/var,flex2:/etc,flex2:/u01/app/1:	Succeeded	
Free Space: flex1:/usr,flex1:/var,flex1:/etc,flex1:/u01/app/1:	Succeeded	
User Existence		
User Existence: grid		
Users With Same UID: 54322	Succeeded	
Group Existence		
Group Existence: dba	Succeeded	
Group Existence: oinstall	Succeeded	
Group Membership		
Group Membership: dba	Succeeded	

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

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

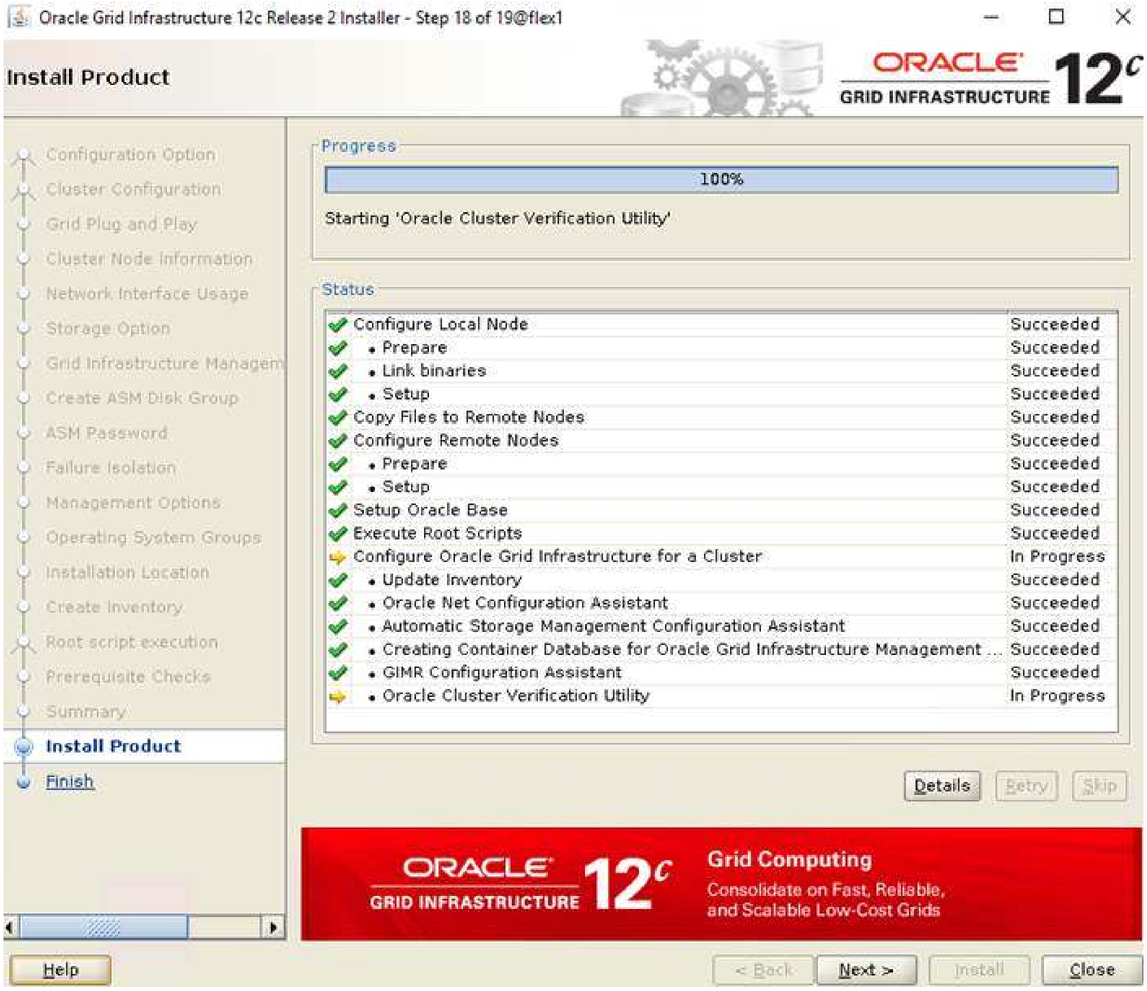

100%

Starting 'Oracle Cluster Verification Utility'

✓ Configure Local Node	Succeeded
✓ • Prepare	Succeeded
✓ • Link binaries	Succeeded
✓ • Setup	Succeeded
✓ Copy Files to Remote Nodes	Succeeded
✓ Configure Remote Nodes	Succeeded
✓ • Prepare	Succeeded
✓ • Setup	Succeeded
✓ Setup Oracle Base	Succeeded
✓ Execute Root Scripts	Succeeded
➜ Configure Oracle Grid Infrastructure for a Cluster	In Progress
✓ • Update Inventory	Succeeded
✓ • Oracle Net Configuration Assistant	Succeeded
✓ • Automatic Storage Management Configuration Assistant	Succeeded
✓ • Creating Container Database for Oracle Grid Infrastructure Management ...	Succeeded
✓ • GIMR Configuration Assistant	Succeeded
➜ • Oracle Cluster Verification Utility	In Progress

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

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

Install Oracle Database Software

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

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

To install Oracle Database Software, follow these steps:

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

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

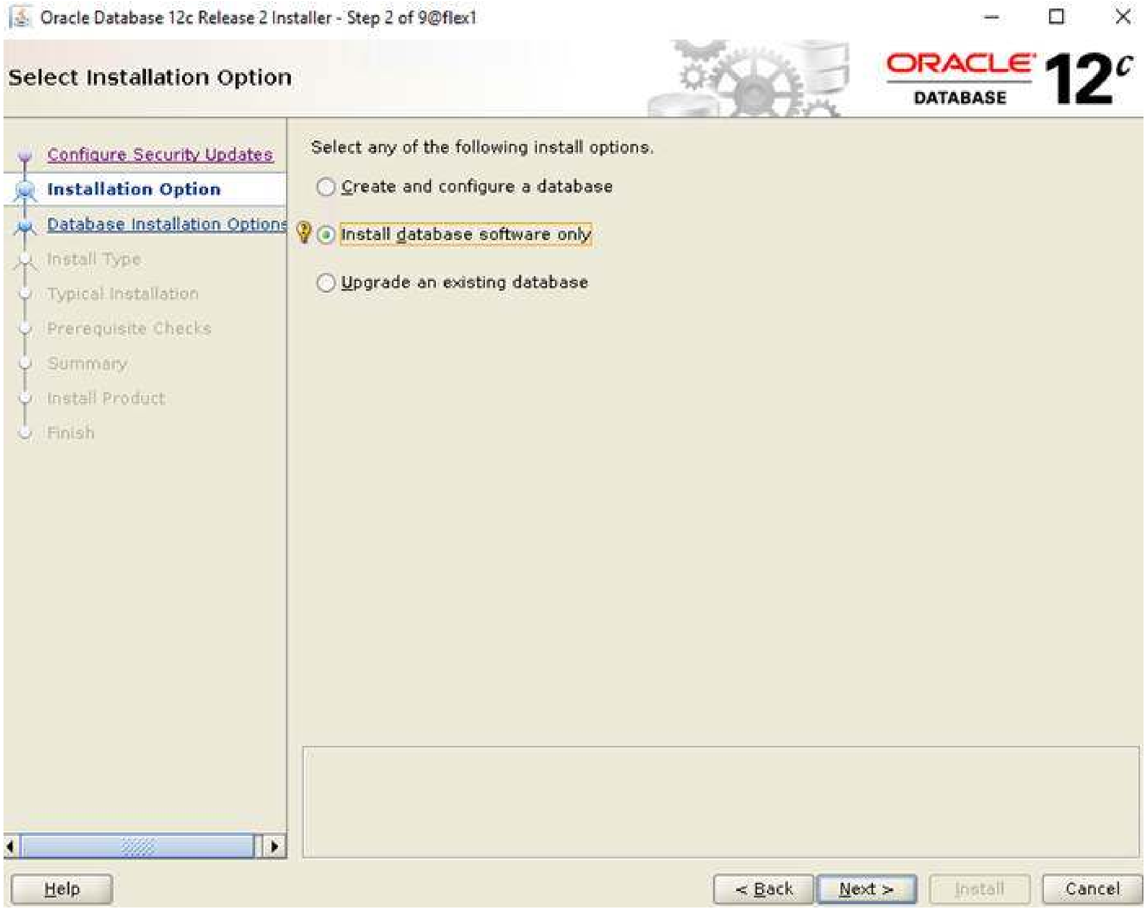

Select Installation Option

Select any of the following install options.

○ Create and configure a database

● Install database software only

○ Upgrade an existing database

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

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

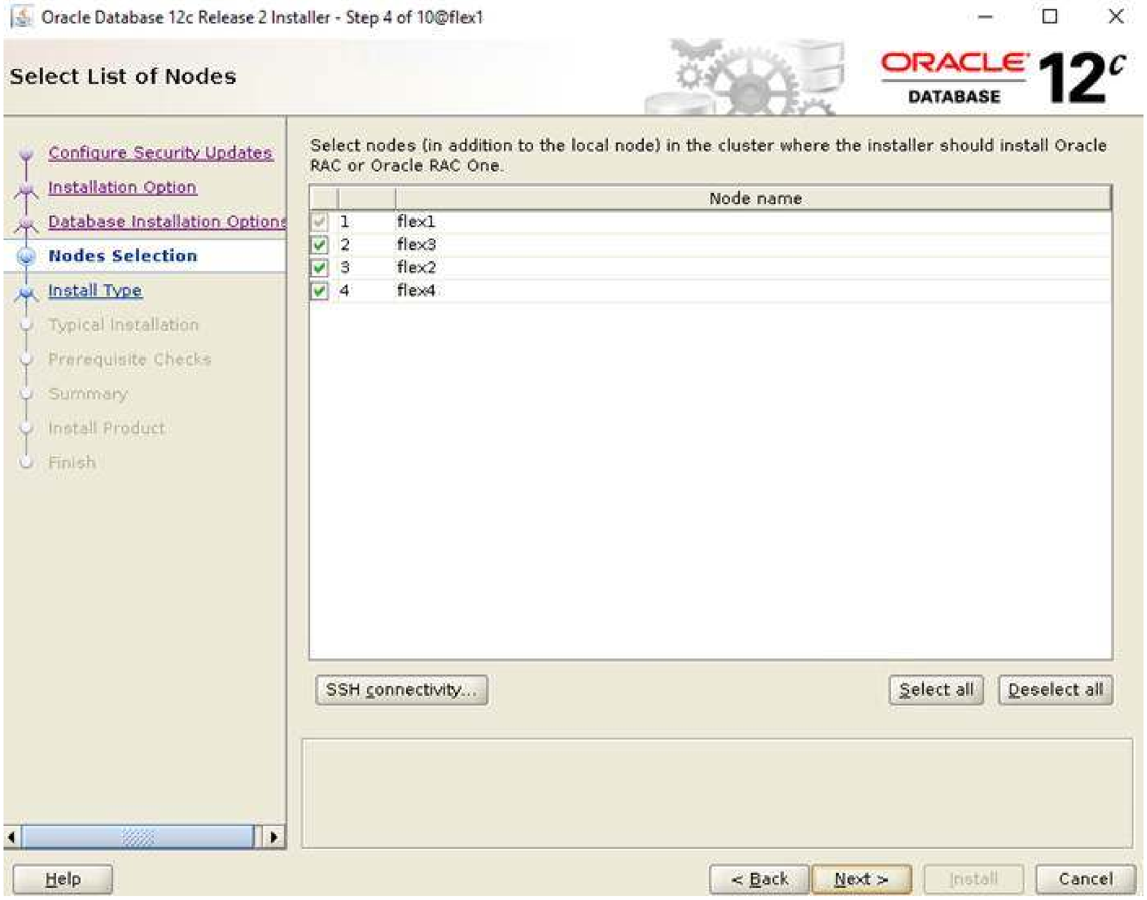

Oracle Database 12c Release 2 Installer - Step 4 of 10@flex1

Select List of Nodes

Select nodes (in addition to the local node) in the cluster where the installer should install Oracle RAC or Oracle RAC One.

		Node name
☑	1	flex1
☑	2	flex3
☑	3	flex2
☑	4	flex4

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

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

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

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

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

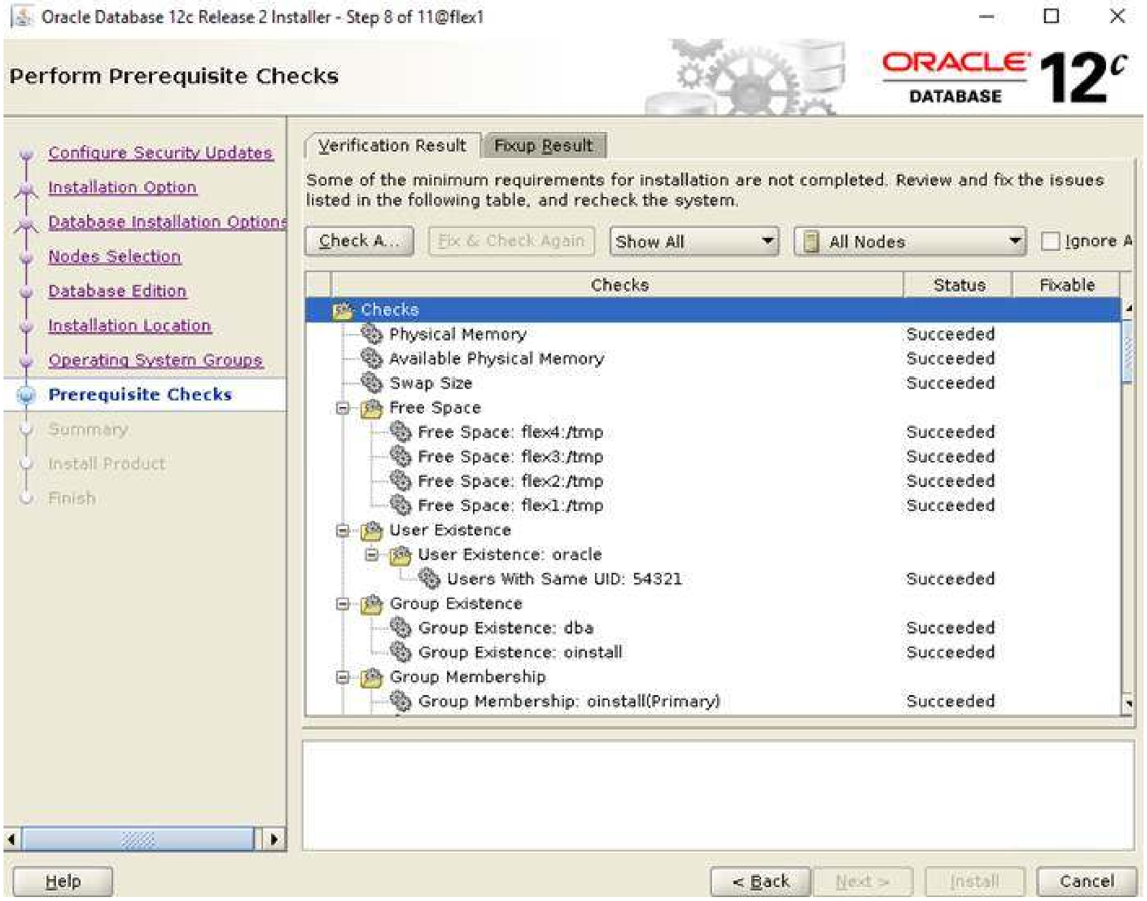

Oracle Database 12c Release 2 Installer - Step 8 of 11@flex1

Perform Prerequisite Checks

Some of the minimum requirements for installation are not completed. Review and fix the issues listed in the following table, and recheck the system.

Checks	Status	Fixable
Checks		
Physical Memory	Succeeded	
Available Physical Memory	Succeeded	
Swap Size	Succeeded	
Free Space		
Free Space: flex4:/tmp	Succeeded	
Free Space: flex3:/tmp	Succeeded	
Free Space: flex2:/tmp	Succeeded	
Free Space: flex1:/tmp	Succeeded	
User Existence		
User Existence: oracle		
Users With Same UID: 54321	Succeeded	
Group Existence		
Group Existence: dba	Succeeded	
Group Existence: oinstall	Succeeded	
Group Membership		
Group Membership: oinstall(Primary)	Succeeded	

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

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

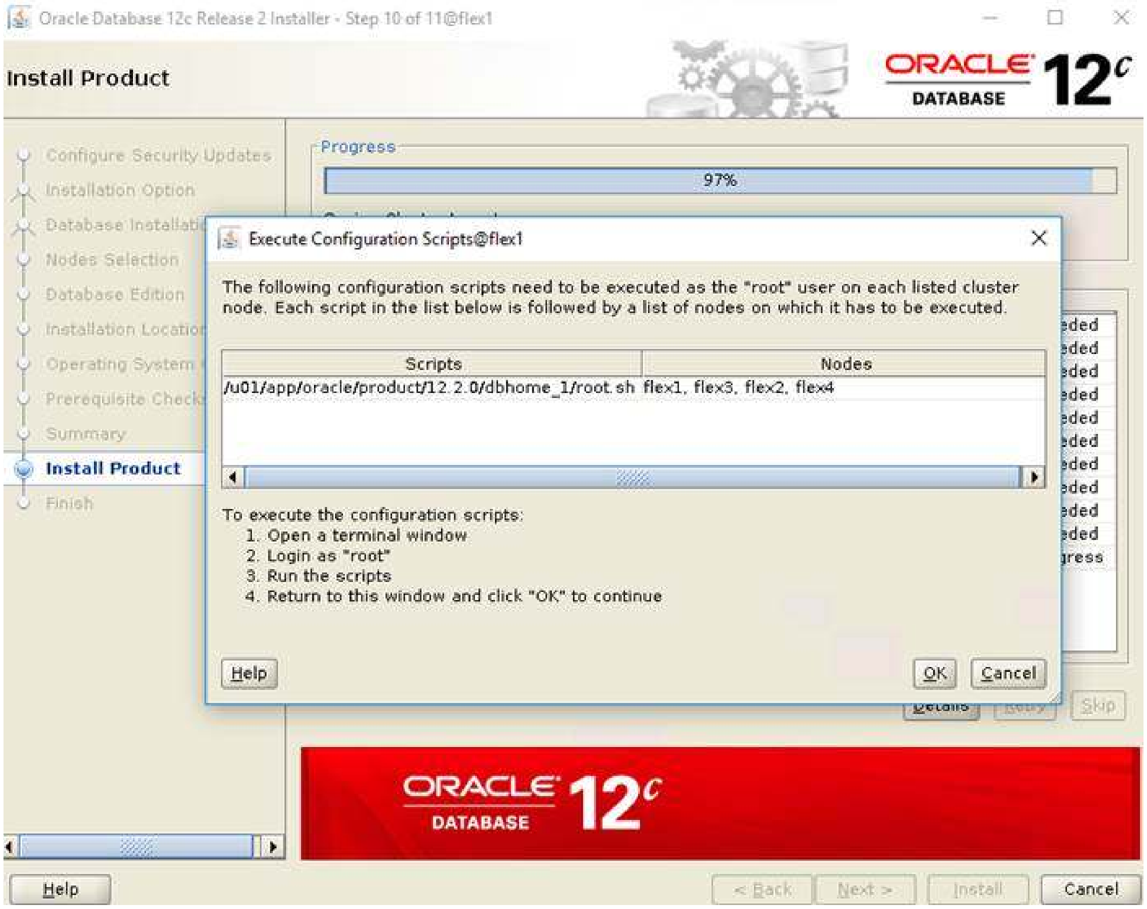

Oracle Database 12c Release 2 Installer - Step 10 of 11@flex1

Install Product

Execute Configuration Scripts@flex1

The following configuration scripts need to be executed as the "root" user on each listed cluster node. Each script in the list below is followed by a list of nodes on which it has to be executed.

Scripts	Nodes
/u01/app/oracle/product/12.2.0/dbhome_1/root.sh	flex1, flex3, flex2, flex4

To execute the configuration scripts:

1. Open a terminal window
2. Login as "root"
3. Run the scripts
4. Return to this window and click "OK" to continue

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

Oracle Flex ASM

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

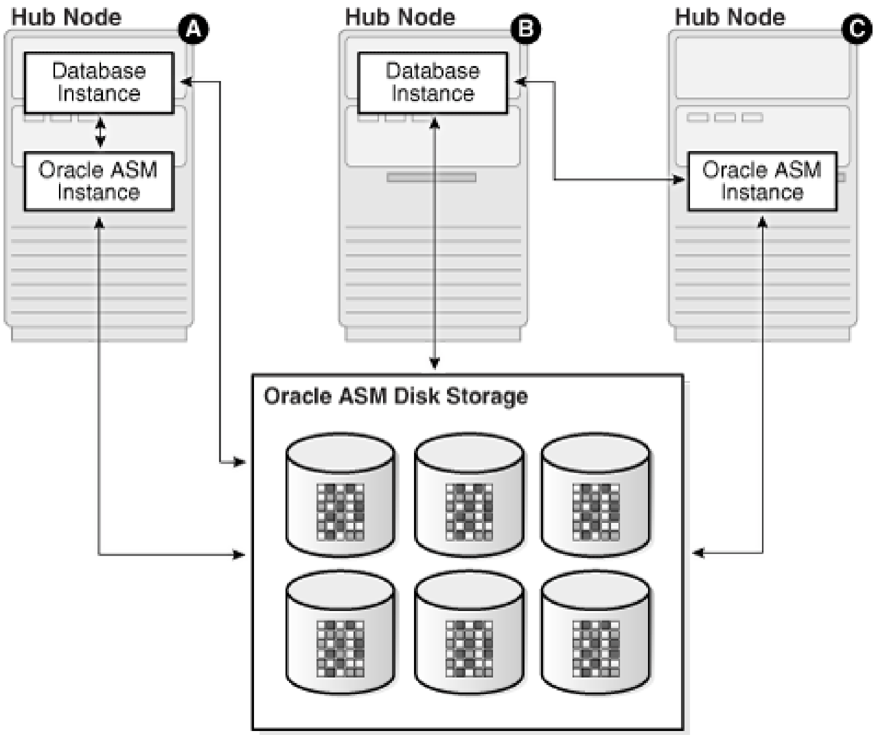

When using Oracle Flex ASM, Oracle ASM clients are configured with direct access to storage. With Oracle Flex ASM, you can consolidate all the storage requirements into a single set of disk groups. All these disk groups are mounted by and managed by a small set of Oracle ASM instances running in a single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

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

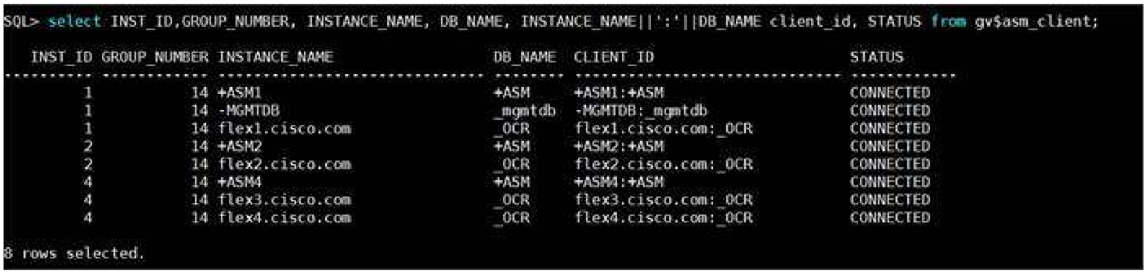
```
SQL> select INST_ID,GROUP_NUMBER, INSTANCE_NAME, DB_NAME, INSTANCE_NAME||':'||DB_NAME client_id, STATUS from gv$asm_client;

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

8 rows selected.
```

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

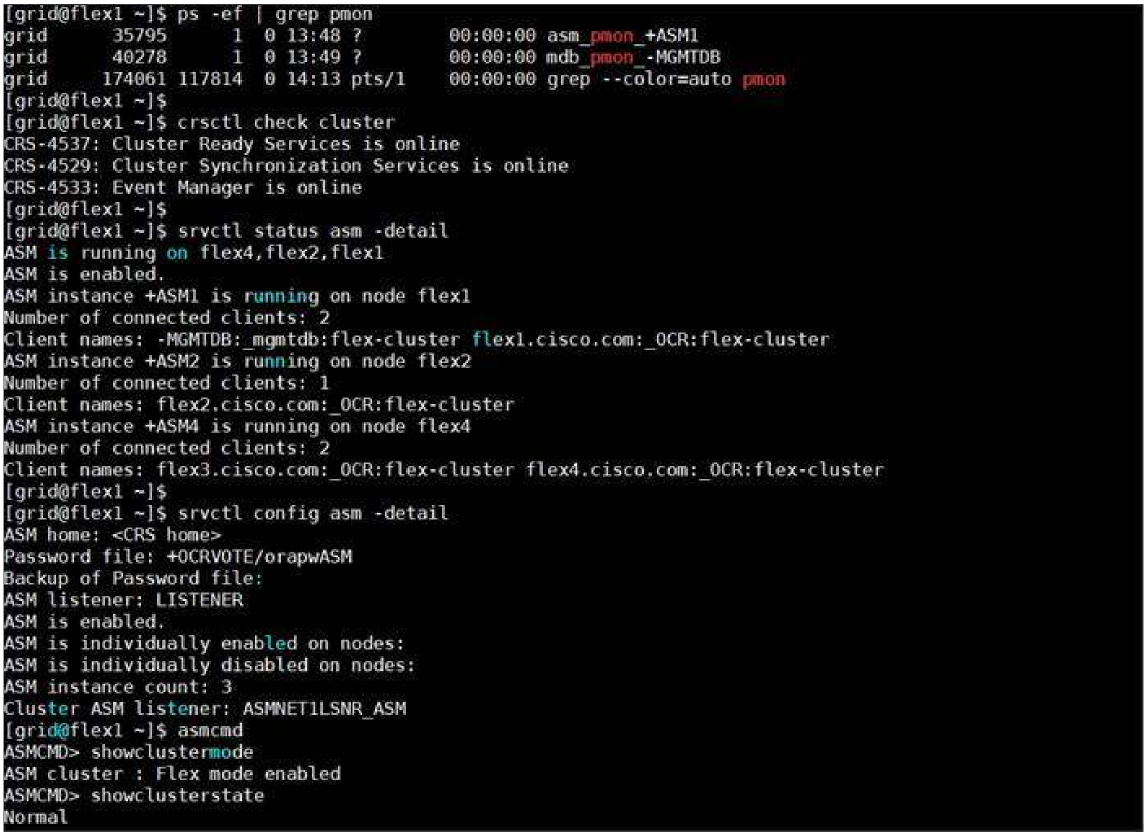

```
[grid@flex1 ~]$ ps -ef | grep pmon
grid      35795      1  0 13:48 ?        00:00:00 asm_pmon_+ASM1
grid      40278      1  0 13:49 ?        00:00:00 mdb_pmon_-MGMTDB
grid     174061 117814  0 14:13 pts/1    00:00:00 grep --color=auto pmon
[grid@flex1 ~]$
[grid@flex1 ~]$ crsctl check cluster
CRS-4537: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

Scalability Test and Results

Before configuring a database for workload tests, it is extremely important to validate that this is indeed a balanced configuration that is capable of delivering expected performance. In this solution, you will test and validate node and user scalability on all 4 node Oracle RAC Databases with various database benchmarking tools as explained below.

Hardware Calibration Test using FIO

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

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

4k and 8k Random Read/Write IOPs Tests

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

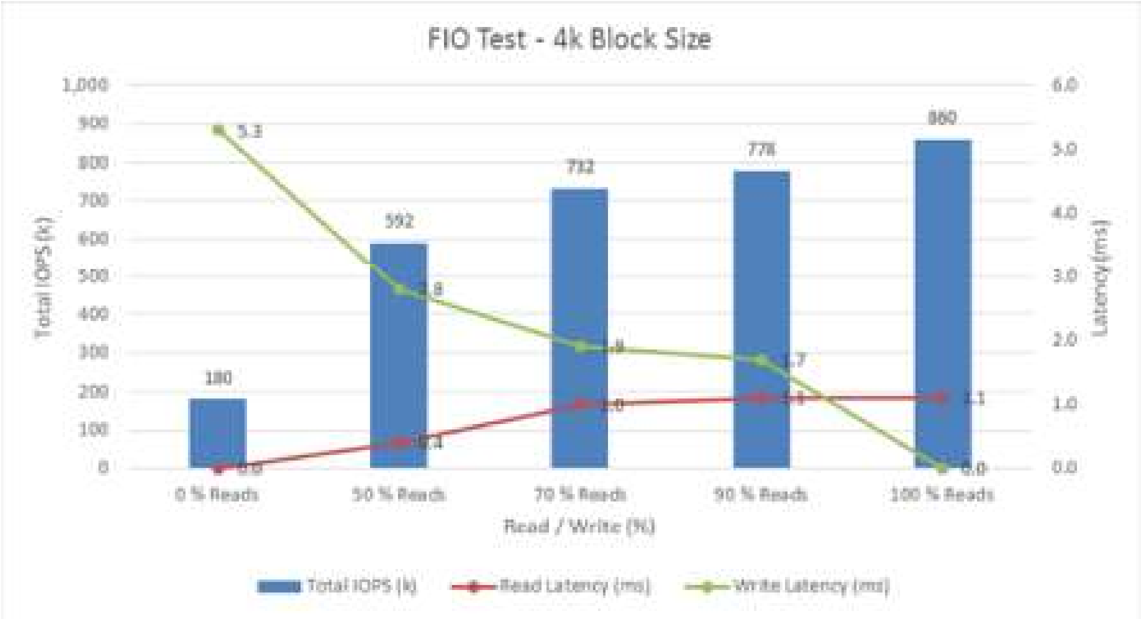

For the 100/0 percent read/write test, we achieved about 860k IOPS with the read latency about 1.1 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 778k IOPS with the read latency about 1.1 millisecond and the write latency about 1.7 millisecond. For the 70/30 percent read/write test, we achieved 732k IOPS with the read latency about 1 millisecond and the write latency about 1.9 millisecond. For the 50/50 percent read/write test, we achieved about 592k IOPS with the read latency about 0.4 millisecond and the write latency about 2.8 millisecond. For the 0/100 percent read/write test, we achieved 180k IOPS with the write latency about 5.3 millisecond.

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

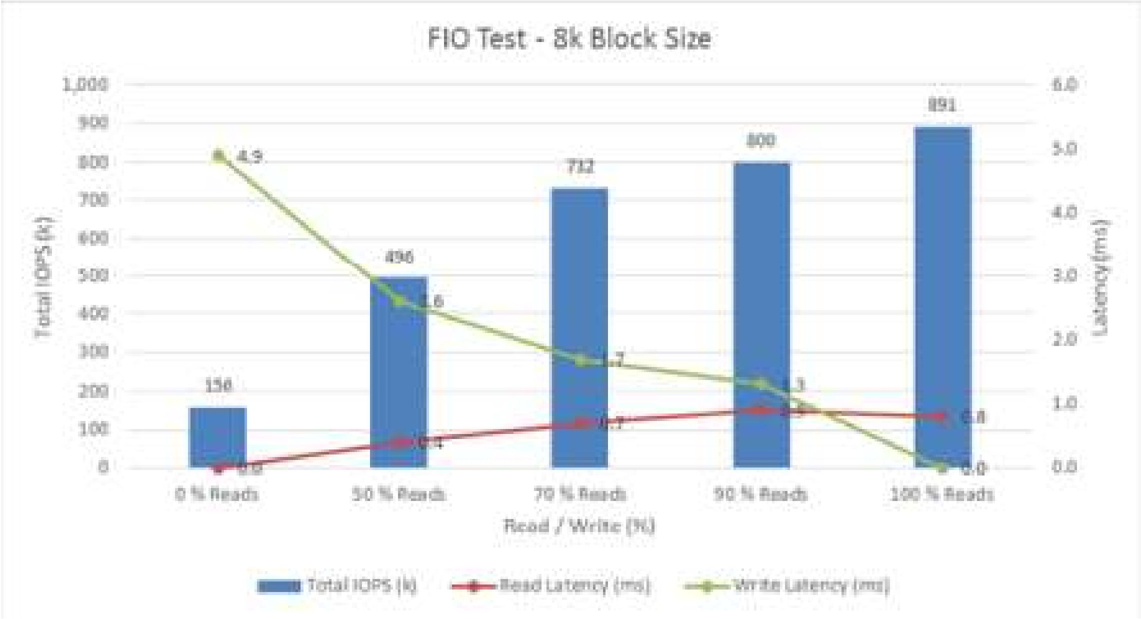

For the 100/0 percent read/write test, we achieved about 891k IOPS with the read latency about 0.8 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 800k IOPS with the read latency about 0.9 millisecond and the write latency about 1.3 millisecond. For the 70/30 percent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

Bandwidth Tests

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

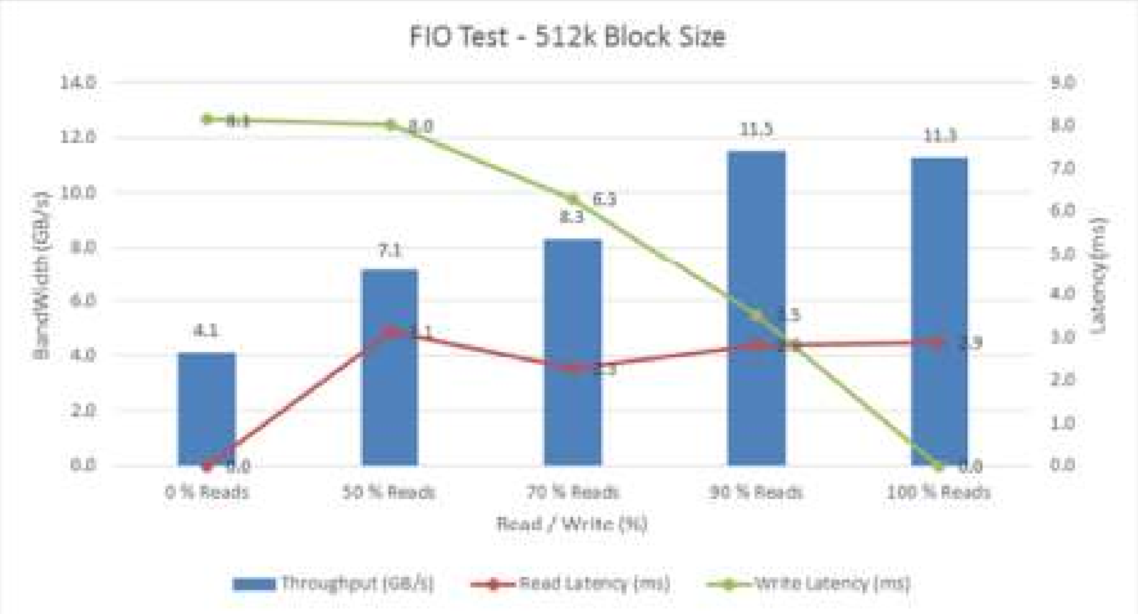

For the 100/0 percent read/write test, we achieved about 11.3 GB/s throughput with the read latency about 2.9 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 11.5 GB/s throughput with the read latency about 2.8 millisecond and the write latency about 3.5 millisecond.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

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

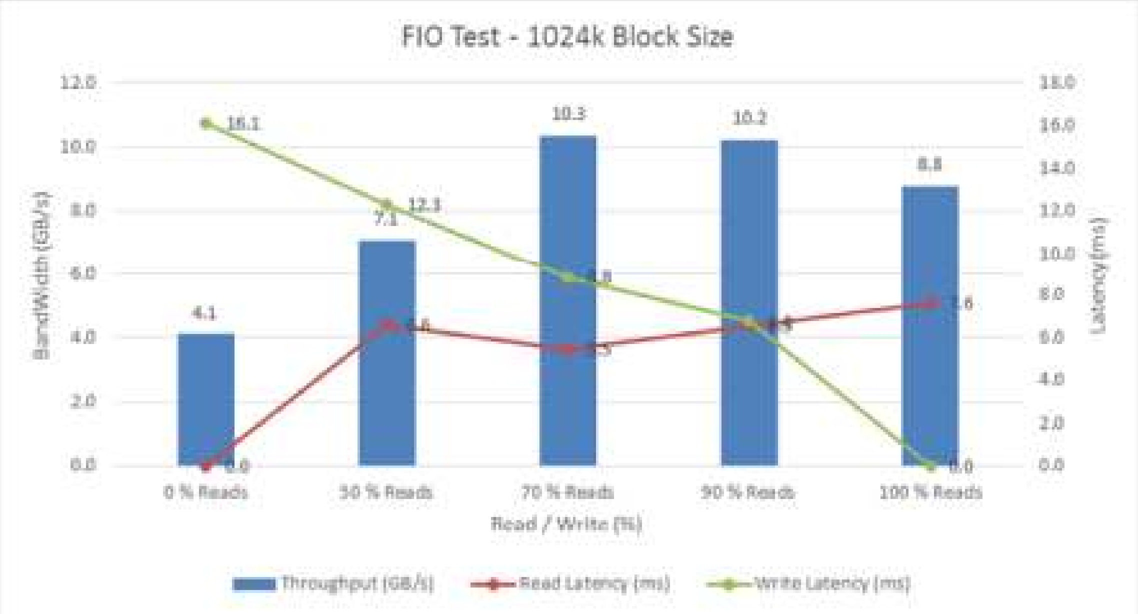

For the 100/0 percent read/write test, we achieved around 8.8 GB/s throughput with the read latency around 7.6 millisecond. Similarly, for the 90/10 percent read/write test, we achieved about 10.2 GB/s throughput with the read latency about 6.5 millisecond and the write latency about 6.7 millisecond.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

millisecond. For the 50/50 percent read/write test, we achieved about 7.1 GB/s throughput with the read latency about 6.6 millisecond and the write latency about 7.1 millisecond. For the 0/100 percent read/write test, we achieved about 4.1 GB/s throughput with the write latency about 16.1 millisecond.

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

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

Database Creation with DBCA

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

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

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

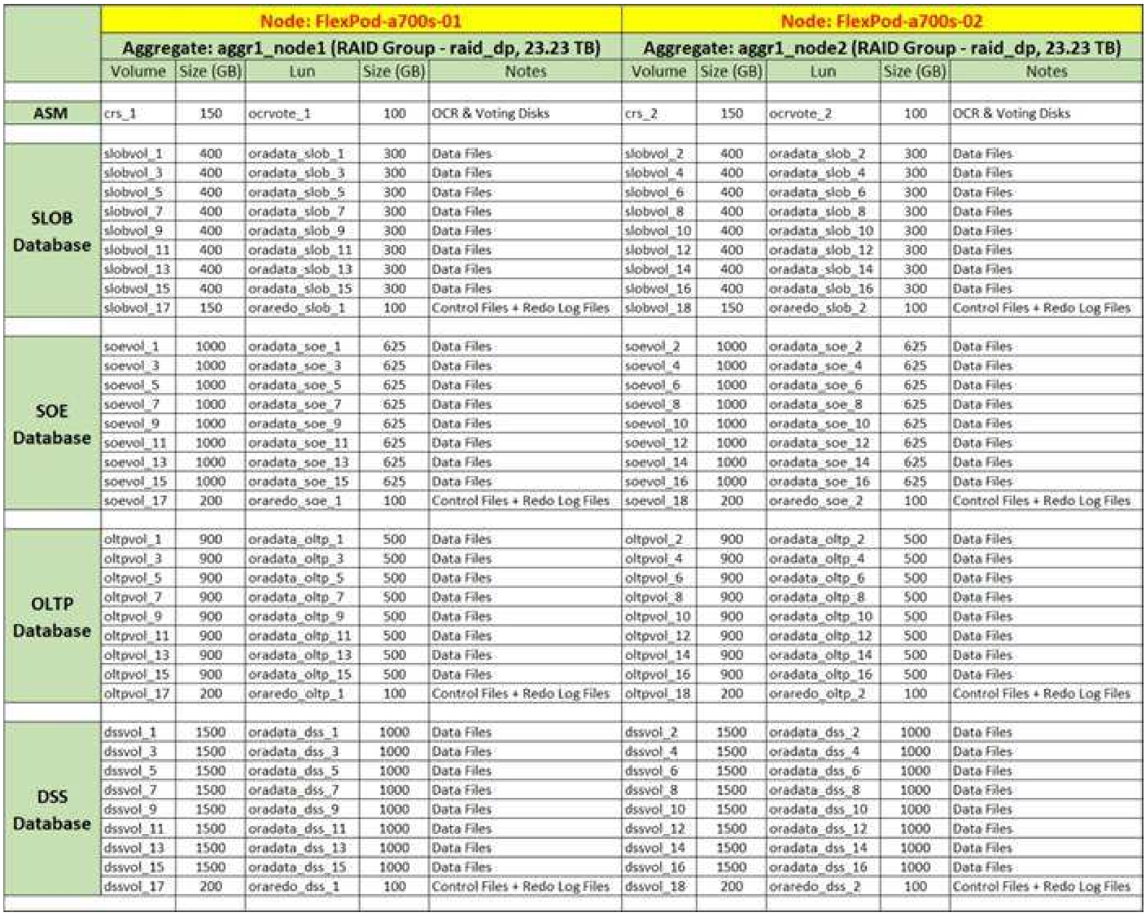

	Node: FlexPod-a700s-01					Node: FlexPod-a700s-02				
	Aggregate: aggr1_node1 (RAID Group - raid_dp, 23.23 TB)					Aggregate: aggr1_node2 (RAID Group - raid_dp, 23.23 TB)				
	Volume	Size (GB)	Lun	Size (GB)	Notes	Volume	Size (GB)	Lun	Size (GB)	Notes
ASM	crs_1	150	ocrvote_1	100	OCR & Voting Disks	crs_2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

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

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

SLOB Test

The Silly Little Oracle Benchmark (SLOB) is a toolkit for generating and testing I/O through an Oracle database. SLOB is very effective to test the I/O subsystem with genuine Oracle SGA-buffered physical I/O. SLOB supports testing physical random single-block reads (db file sequential read) and random single block writes (DBWR flushing capability). SLOB issues single block reads for the read workload that are generally 8K (as the database block size was 8K).

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

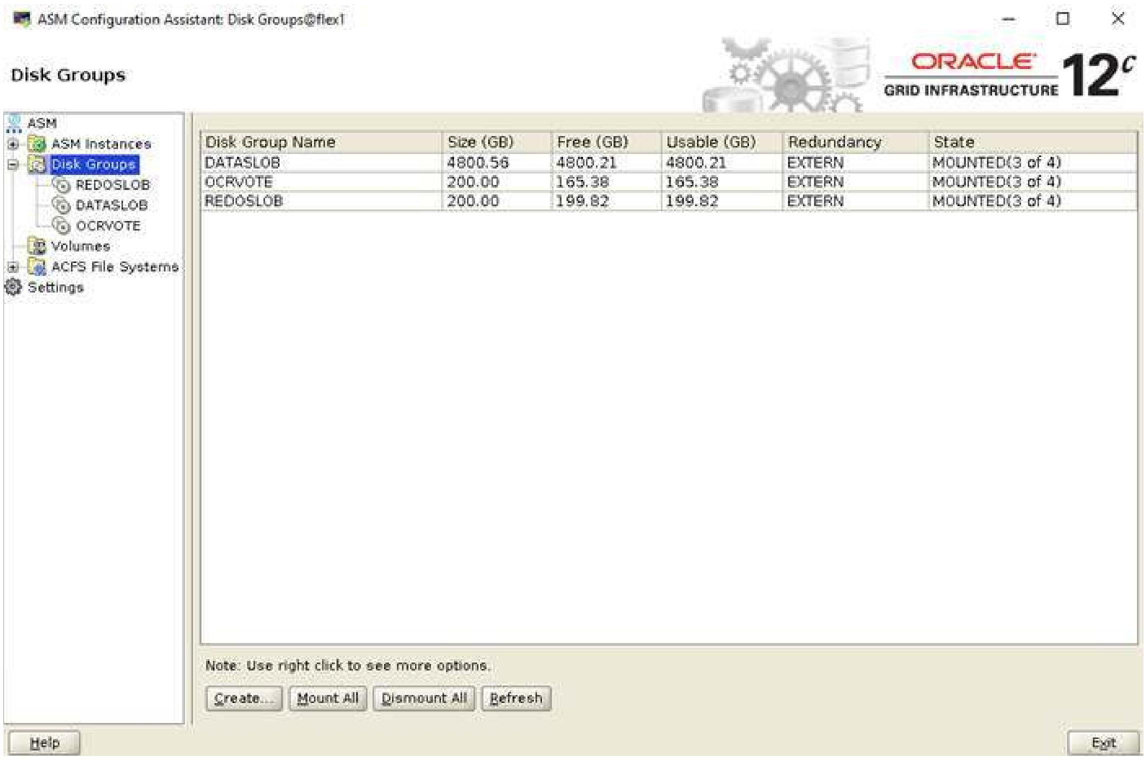

Disk Group Name	Size (GB)	Free (GB)	Usable (GB)	Redundancy	State
DATASLOB	4800.56	4800.21	4800.21	EXTERN	MOUNTED(3 of 4)
OCRVOTE	200.00	165.38	165.38	EXTERN	MOUNTED(3 of 4)
REDOSLOB	200.00	199.82	199.82	EXTERN	MOUNTED(3 of 4)

Note: Use right click to see more options.

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

We used SLOB2 to generate our OLTP workload. Each database server applied the workload to Oracle database, log, and temp files. The following tests were performed and various metrics like IOPS and latency were captured along with Oracle AWR reports for each test scenario.

User Scalability Test

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

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

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

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

Table 10 Total IOPS for SLOB User Scalability Tests

Users	IOPS for Read/Write % (100/0)	IOPS for Read/Write % (90/10)	IOPS for Read/Write % (70/30)	IOPS for Read/Writ % (50/50)
32	186,386	188,906	209,264	220,718
64	316,731	324,739	333,120	343,007
128	461,265	452,060	456,660	469,478

256	541,849	532,473	529,033	532,116
512	603,039	559,831	547,240	555,170

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

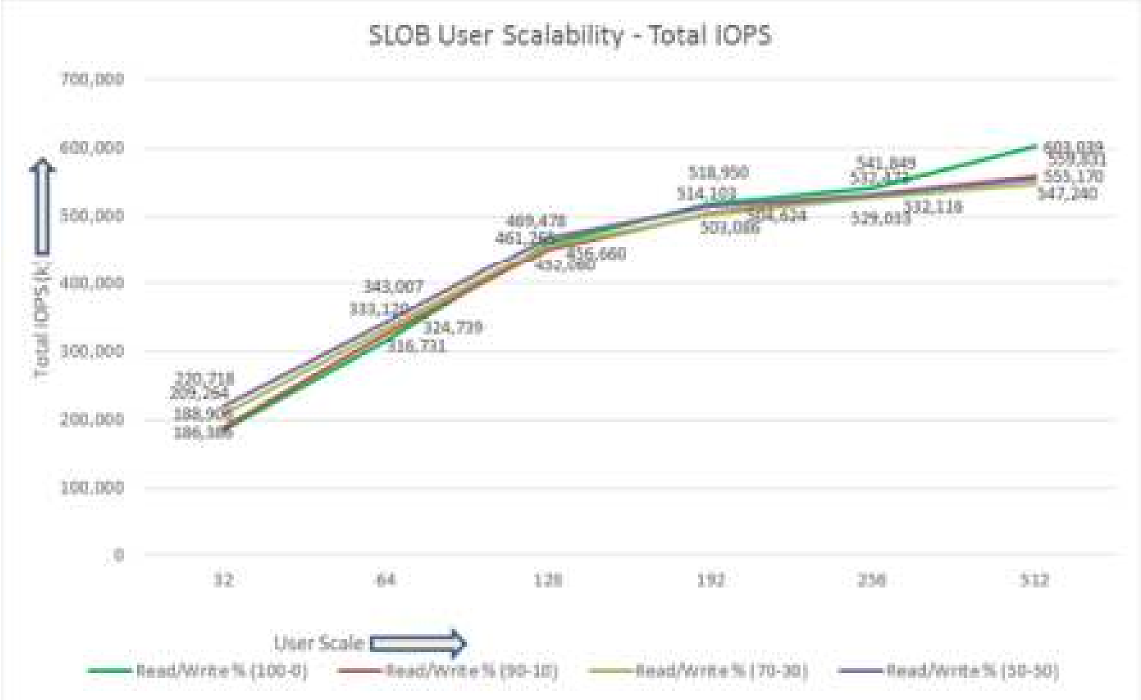

The above graph shows the linear scalability with increased users and similar IOPS from 32 users to 512 users with 100% read, 90% read, 70% read and 50% read.

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

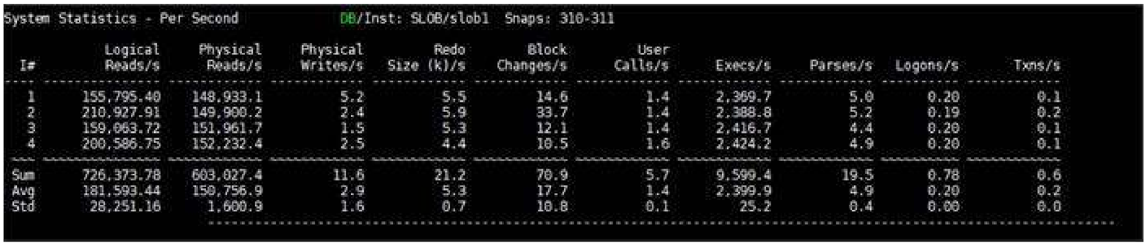

```
System Statistics - Per Second           DB/Inst: SLOB/slob1  Snaps: 310-311

            Logical     Physical     Physical         Redo        Block         User
 I#         Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s      Calls/s      Execs/s     Parses/s   Logons/s        Txns/s
---- -------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- -------------
   1     155,795.40    148,933.1          5.2          5.5         14.6          1.4      2,369.7          5.0       0.20           0.1
   2     210,927.91    149,900.2          2.4          5.9         33.7          1.4      2,388.8          5.2       0.19           0.2
   3     159,063.72    151,961.7          1.5          5.3         12.1          1.4      2,416.7          4.4       0.20           0.1
   4     200,586.75    152,232.4          2.5          4.4         10.5          1.6      2,424.2          4.9       0.20           0.1
 ~~~ ~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~~
 Sum     726,373.78    603,027.4         11.6         21.2         70.9          5.7      9,599.4         19.5       0.78           0.6
 Avg     181,593.44    150,756.9          2.9          5.3         17.7          1.4      2,399.9          4.9       0.20           0.2
 Std      28,251.16      1,600.9          1.6          0.7         10.8          0.1         25.2          0.4       0.00           0.0
                          ------------------------------------------------------------------------------------------------------------------
```

The screenshot above highlights that IO load is distributed across all the cluster nodes performing workload operations. Due to variations in workload randomness, we conducted multiple runs to help ensure consistency in behavior and test results.

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

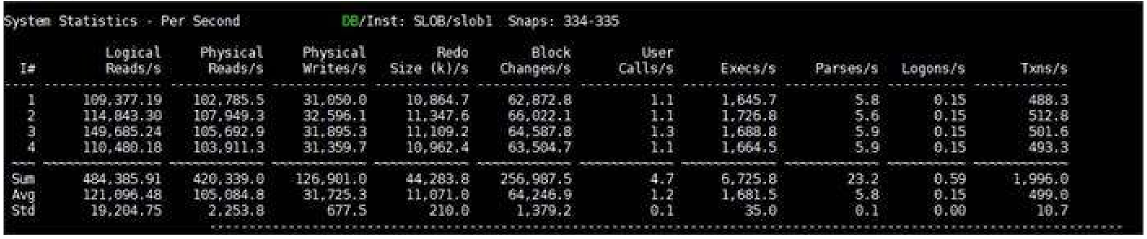

```
System Statistics - Per Second           DB/Inst: SLOB/slob1  Snaps: 334-335

            Logical     Physical     Physical         Redo        Block         User
 I#         Reads/s      Reads/s     Writes/s   Size (k)/s    Changes/s      Calls/s      Execs/s     Parses/s   Logons/s        Txns/s
---- -------------- ------------ ------------ ------------ ------------ ------------ ------------ ------------ ---------- -------------
   1     109,377.19    102,785.5     31,050.0     10,864.7     62,872.8          1.1      1,645.7          5.8       0.15         488.3
   2     114,843.30    107,949.3     32,596.1     11,347.6     66,022.1          1.1      1,726.8          5.6       0.15         512.8
   3     149,685.24    105,692.9     31,895.3     11,109.2     64,587.8          1.3      1,688.8          5.9       0.15         501.6
   4     110,480.18    103,911.3     31,359.7     10,962.4     63,504.7          1.1      1,664.5          5.9       0.15         493.3
 ~~~ ~~~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~~~ ~~~~~~~~~~ ~~~~~~~~~~~~~
 Sum     484,385.91    420,339.0    126,901.0     44,283.8    256,987.5          4.7      6,725.8         23.2       0.59       1,996.0
 Avg     121,096.48    105,084.8     31,725.3     11,071.0     64,246.9          1.2      1,681.5          5.8       0.15         499.0
 Std      19,204.75      2,253.8        677.5        210.0      1,379.2          0.1         35.0          0.1       0.00          10.7
                          ------------------------------------------------------------------------------------------------------------------
```

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

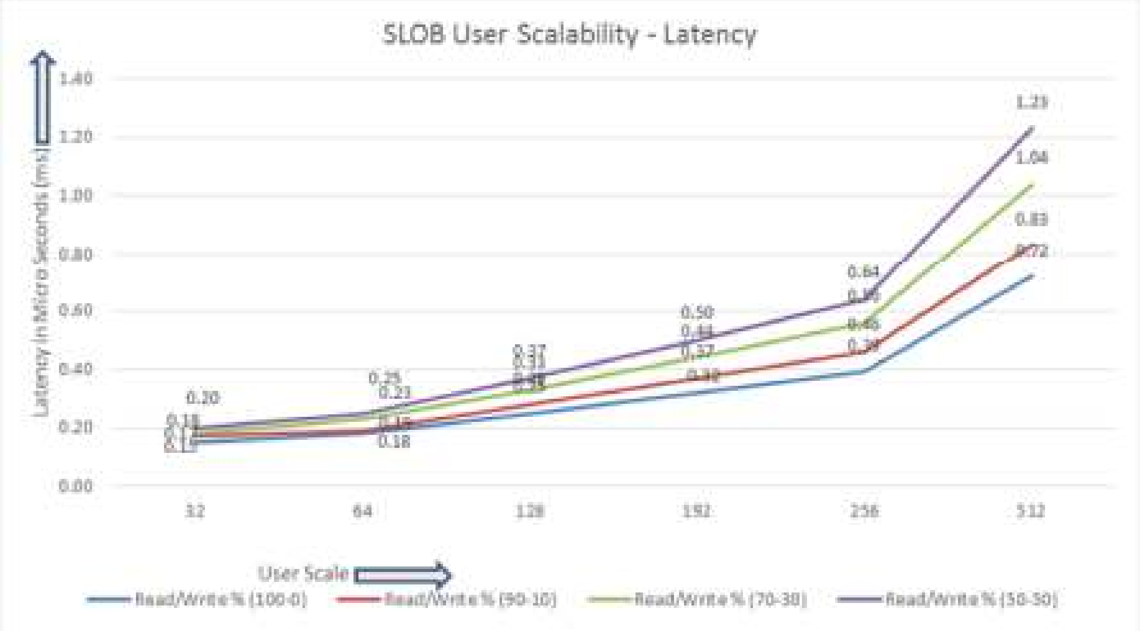

The following screenshot was captured from 100 % Read (0% Update) Test scenario while running SLOB test for 512 users. The screenshot shows a section of AWR report from the run that highlights top timed Events.

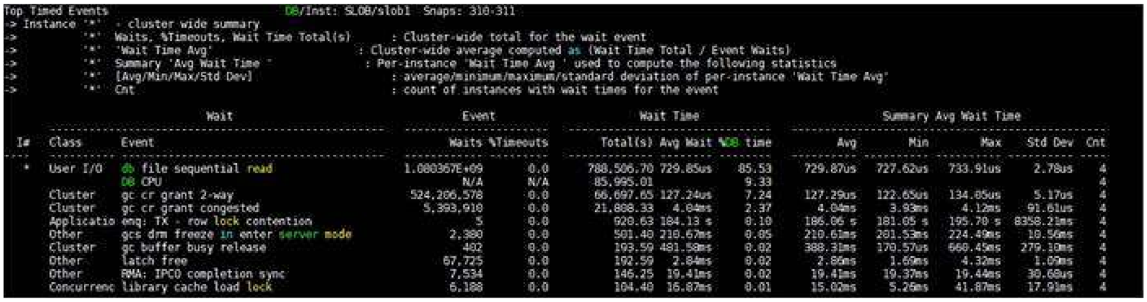

```
Top Timed Events                                   DB/Inst: SLOB/slob1  Snaps: 310-311
-> Instance '*' - cluster wide summary
->          '*'  Waits, %Timeouts, Wait Time Total(s)       : Cluster-wide total for the wait event
->          '*'  'Wait Time Avg'                          : Cluster-wide average computed as (Wait Time Total / Event Waits)
->          '*'  Summary 'Avg Wait Time '                  : Per-instance 'Wait Time Avg ' used to compute the following statistics
->          '*'  [Avg/Min/Max/Std Dev]                         : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
->          '*'  Cnt                                           : count of instances with wait times for the event

                             Wait                                   Event                Wait Time                        Summary Avg Wait Time
        ------------------------------------------------  ------------------------  -------------------------------  ------------------------------------------------
 I#   Class      Event                                       Waits %Timeouts        Total(s) Avg Wait %DB time       Avg        Min        Max    Std Dev  Cnt
----  ---------- ----------------------------------------  ------------ ----------  ------------ --------- -------  ---------- ---------- ---------- ---------- ----
   *  User I/O   db file sequential read                    1.080367E+09       0.0    788,506.70 729.85us    85.53    729.87us   727.62us   733.91us     2.78us    4
                 DB CPU                                              N/A       N/A     85,995.01               9.33                                                4
      Cluster    gc cr grant 2-way                           524,206,578       0.0     66,697.65 127.24us     7.24    127.29us   122.65us   134.05us     5.17us    4
      Cluster    gc cr grant congested                         5,393,910       0.0     21,888.33   4.04ms     2.37      4.04ms     3.93ms     4.12ms    91.61us    4
      Applicatio enq: TX - row lock contention                         5       0.0        920.63 184.13 s     0.10    186.06 s   181.05 s   195.70 s  8358.21ms    4
      Other      gcs drm freeze in enter server mode               2,380       0.0        501.40 210.67ms     0.05    210.61ms   201.53ms   224.49ms    10.56ms    4
      Cluster    gc buffer busy release                              402       0.0        193.59 481.58ms     0.02    388.31ms   170.57us   660.45ms   279.10ms    4
      Other      latch free                                       67,725       0.0        192.59   2.84ms     0.02      2.86ms     1.69ms     4.32ms     1.09ms    4
      Other      RMA: IPC0 completion sync                         7,534       0.0        146.25  19.41ms     0.02     19.41ms    19.37ms    19.44ms    30.68us    4
      Concurrenc library cache load lock                           6,188       0.0        104.40  16.87ms     0.01     15.02ms     5.26ms    41.87ms    17.91ms    4
```

Swingbench Test

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

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

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

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

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

- OLTP database user scalability representing small and random transactions
- DSS database workload representing larger transactions

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

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

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

One (SOE) Database Performance

For one OLTP database workload featuring Order Entry schema, we used one SOE database. For the SOE database (4 TB), we used 128GB size of System Global Area (SGA). We also made sure that HugePages were in use. The OLTP Database scalability test was run for at least 12 hours and made sure that results are consistent for the duration of the full run.

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

User Scalability

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

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

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

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

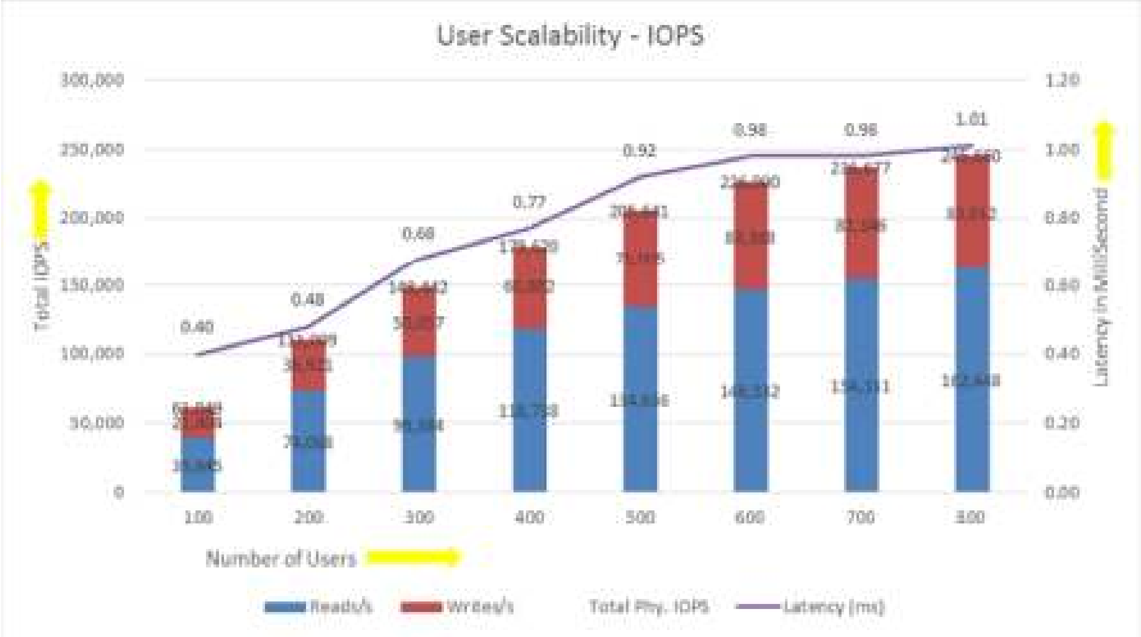

The chart below shows TPM and System Utilization of the system while running swingbench workload for 100 user to 800 users.

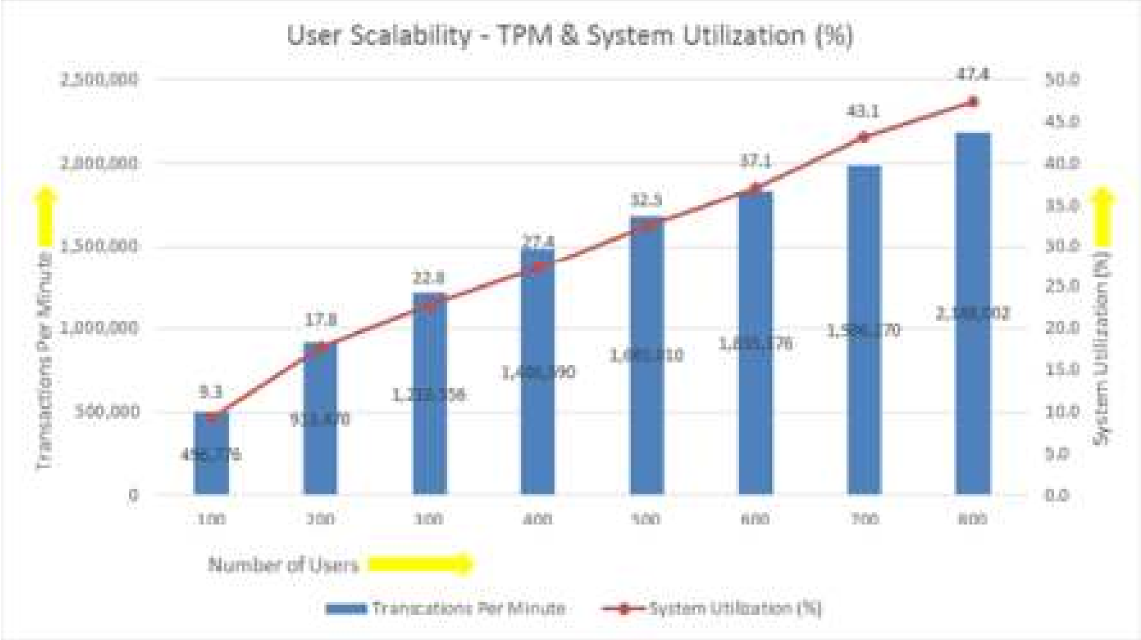

The screenshot below was captured from the 800 User Scale Test scenario while running Swingbench workload on SOE database. The snapshot shows a section from 4-hour window of AWR Global report from the run that highlights Physical Reads/Sec and Physical Writes/Sec for each instance. Notice that IO load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

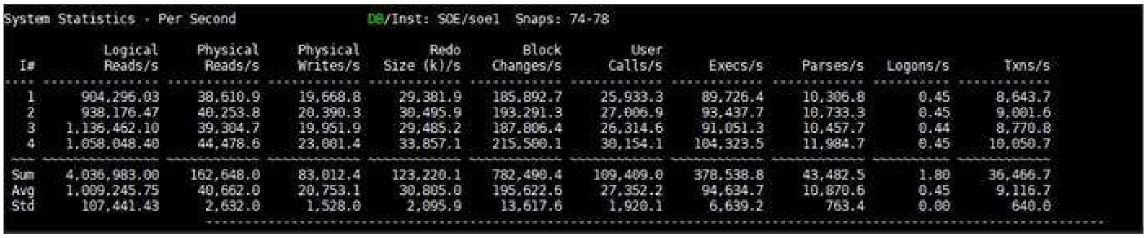

```
System Statistics - Per Second              DB/Inst: SOE/soe1  Snaps: 74-78

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

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

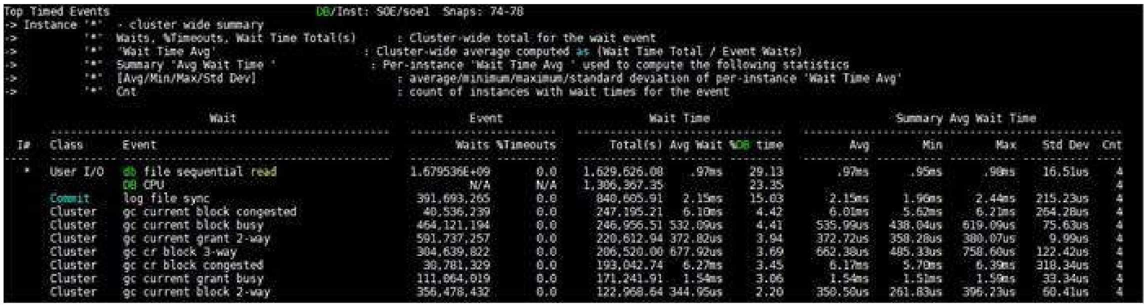

```
Top Timed Events                                    DB/Inst: SOE/soel  Snaps: 74-78
-> Instance '*'   - cluster wide summary
->           '*'   Waits, %Timeouts, Wait Time Total(s)       : Cluster-wide total for the wait event
->           '*'   'Wait Time Avg'                          : Cluster-wide average computed as (Wait Time Total / Event Waits)
->           '*'   Summary 'Avg Wait Time '                  : Per-instance 'Wait Time Avg ' used to compute the following statistics
->           '*'   [Avg/Min/Max/Std Dev]                       : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
->           '*'   Cnt                                         : count of instances with wait times for the event

                              Wait                                   Event                      Wait Time                           Summary Avg Wait Time
          ----------------------------------------------  ---------------------  ---------------------------------  ---------------------------------------------
  I#   Class      Event                                           Waits %Timeouts      Total(s) Avg Wait %DB time        Avg        Min        Max    Std Dev  Cnt
 ----  ---------- --------------------------------------  ------------ ---------  ------------- -------- --------  ---------- ---------- ---------- ---------- ----
   *   User I/O   db file sequential read                 1.679536E+09       0.0   1,629,626.08    .97ms    29.13       .97ms      .95ms      .98ms    16.51us    4
                  DB CPU                                           N/A       N/A   1,306,367.35             23.35                                                4
       Commit     log file sync                            391,693,265       0.0     840,665.91   2.15ms    15.03      2.15ms     1.96ms     2.44ms   215.23us    4
       Cluster    gc current block congested                40,536,239       0.0     247,195.21   6.10ms     4.42      6.01ms     5.62ms     6.21ms   264.28us    4
       Cluster    gc current block busy                    464,121,194       0.0     246,956.51 532.09us     4.41    535.99us   438.04us   619.09us    75.63us    4
       Cluster    gc current grant 2-way                   591,737,257       0.0     220,612.94 372.82us     3.94    372.72us   358.28us   380.07us     9.99us    4
       Cluster    gc cr block 3-way                        304,639,822       0.0     206,520.00 677.92us     3.69    662.38us   485.33us   758.60us   122.42us    4
       Cluster    gc cr block congested                     30,781,329       0.0     193,042.74   6.27ms     3.45      6.17ms     5.70ms     6.39ms   318.34us    4
       Cluster    gc current grant busy                    111,064,019       0.0     171,241.91   1.54ms     3.06      1.54ms     1.51ms     1.59ms    33.34us    4
       Cluster    gc current block 2-way                   356,478,432       0.0     122,968.64 344.95us     2.20    350.50us   261.83us   396.23us    60.41us    4
```

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

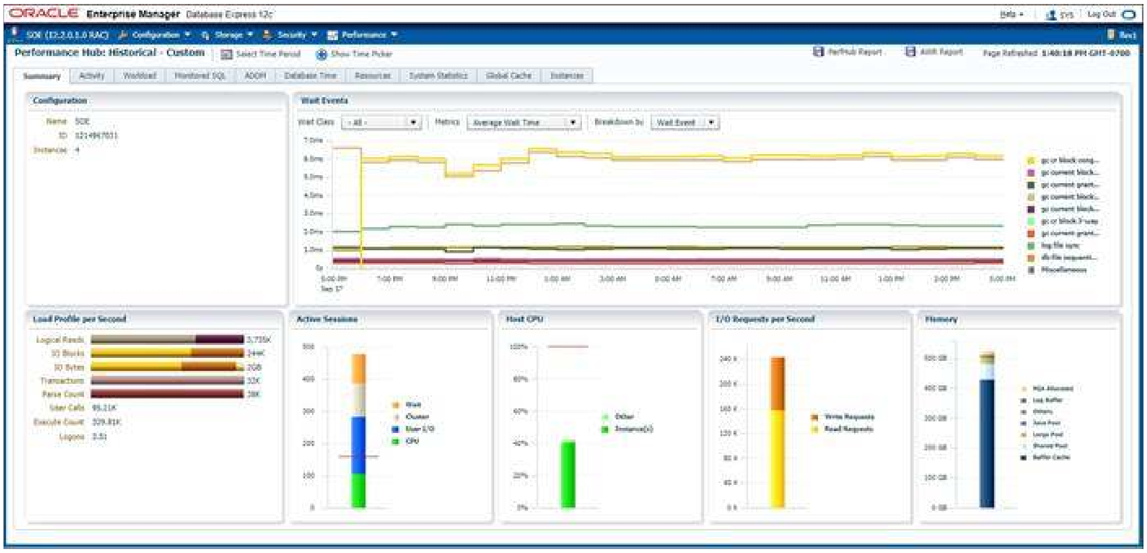

Two (SOE and OLTP) Database Performance

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

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

User Scalability

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

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

Users	IOPS for SOE DB	IOPS for OLTP DB	Total IOPS	System Utilization (%)
200	81,838	27,755	109,593	18.9
300	119,816	26,097	145,913	25.3
400	128,855	44,909	173,764	32.1
500	139,657	61,445	201,102	39.3

600	145,508	75,465	220,973	44.2
700	150,922	89,918	240,840	50.7
800	142,873	96,206	239,079	55.9

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

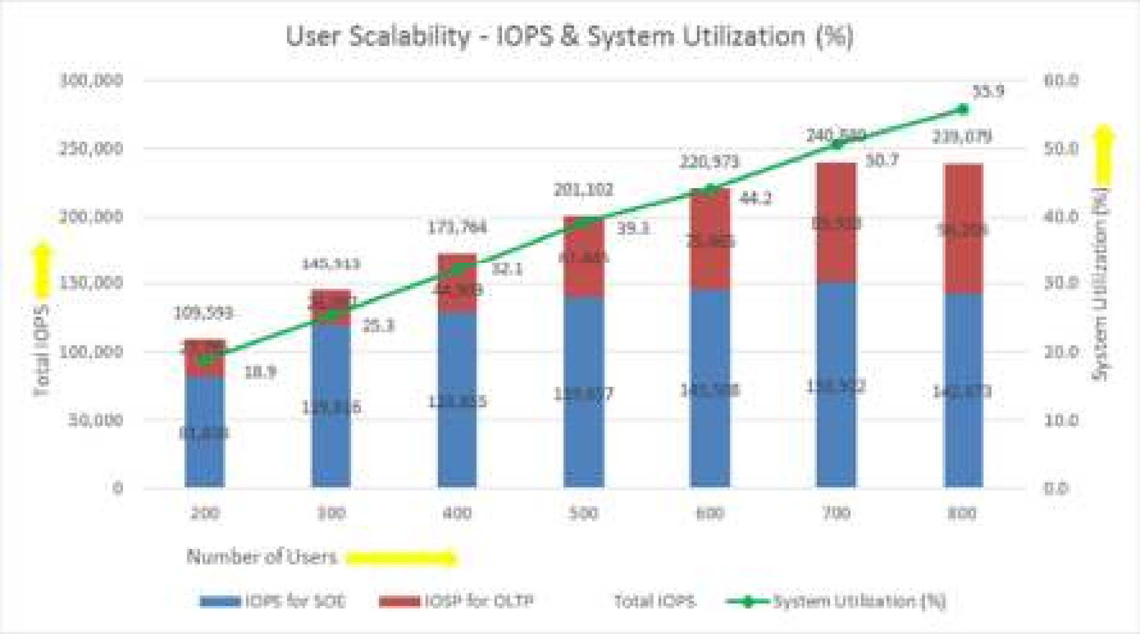

Table 13 lists the TPM and System Utilization for SOE + OLTP databases running Swingbench workloads together from 100 users to 800 users.

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

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

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

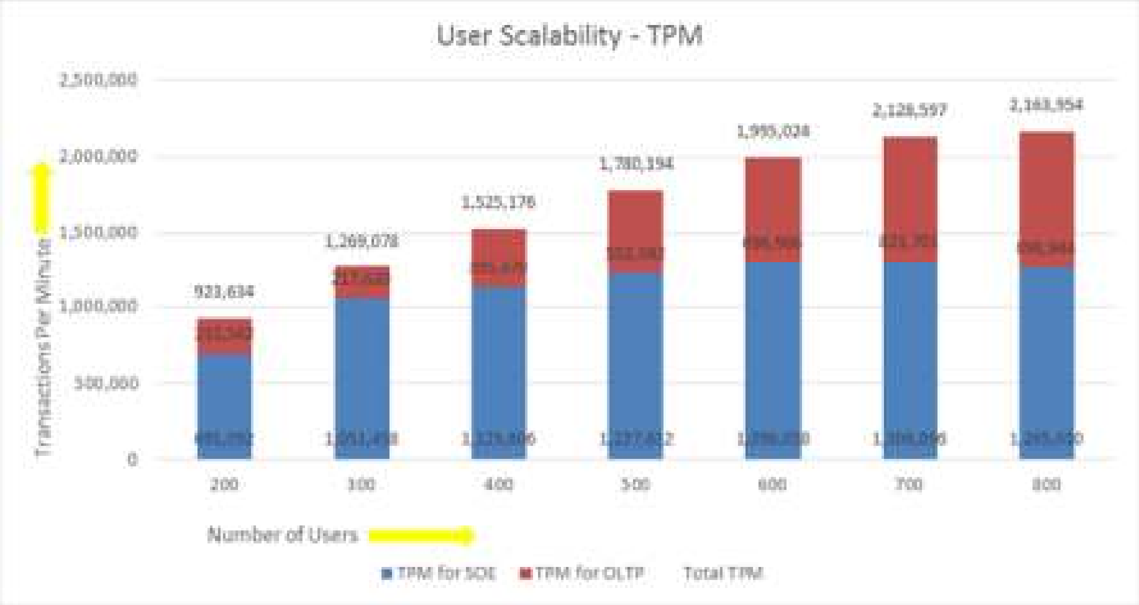

The results were in line with prior assessments where the user scalability was almost linear till 600 users and beyond 600 users the rate of IOPS increase slowed down.

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

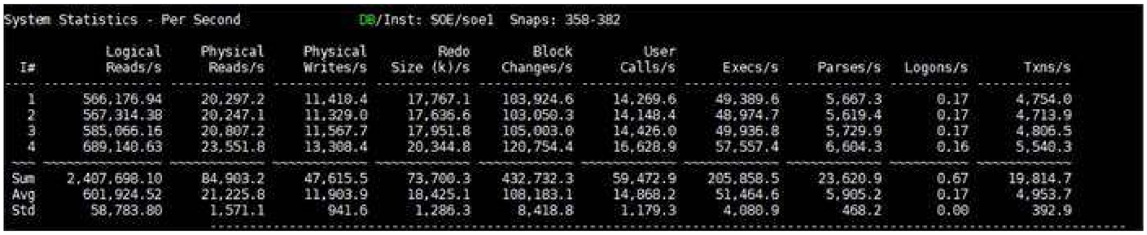

```
System Statistics - Per Second              DB/Inst: SOE/soel  Snaps: 358-382

            Logical    Physical    Physical        Redo       Block        User
  I#        Reads/s     Reads/s    Writes/s  Size (k)/s   Changes/s     Calls/s      Execs/s    Parses/s   Logons/s       Txns/s
---- -------------- ----------- ----------- ----------- ----------- ----------- ----------- ----------- --------- -----------
   1     566,176.94    20,297.2    11,418.4    17,767.1   103,924.6    14,269.6     49,389.6     5,667.3      0.17      4,754.0
   2     567,314.38    20,247.1    11,329.0    17,636.6   103,050.3    14,148.4     48,974.7     5,619.4      0.17      4,713.9
   3     585,066.16    20,807.2    11,567.7    17,951.8   105,003.0    14,426.0     49,936.8     5,729.9      0.17      4,806.5
   4     689,140.63    23,551.8    13,308.4    20,344.8   120,754.4    16,628.9     57,557.4     6,604.3      0.16      5,540.3
 ~~~ ~~~~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~~~ ~~~~~~~~~ ~~~~~~~~~~~
 Sum   2,407,698.10    84,903.2    47,615.5    73,700.3   432,732.3    59,472.9    205,858.5    23,620.9      0.67     19,814.7
 Avg     601,924.52    21,225.8    11,903.9    18,425.1   108,183.1    14,868.2     51,464.6     5,905.2      0.17      4,953.7
 Std      58,783.80     1,571.1       941.6     1,286.3     8,418.8     1,179.3      4,080.9       468.2      0.00        392.9
                  ----------------------------------------------------------------------------------------------------------------
```

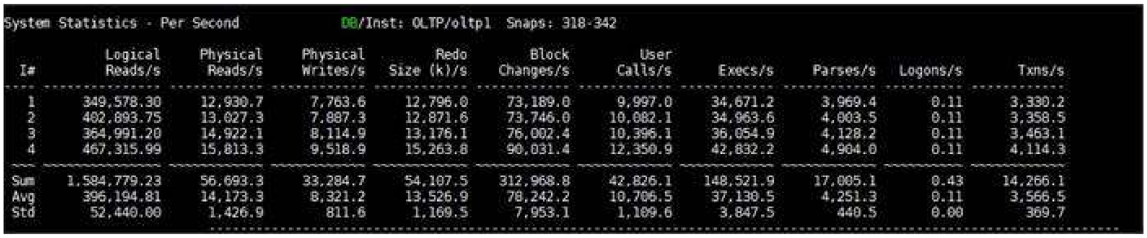

```
System Statistics - Per Second             DB/Inst: OLTP/oltp1  Snaps: 318-342

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

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

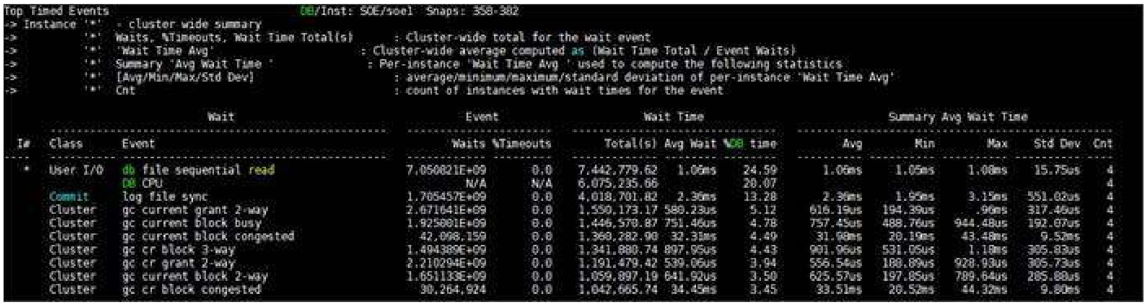

```
Top Timed Events                                    DB/Inst: SOE/soel  Snaps: 358-382
-> Instance '*'  - cluster wide summary
->          '*'  Waits, %Timeouts, Wait Time Total(s)      : Cluster-wide total for the wait event
->          '*'  'Wait Time Avg'                      : Cluster-wide average computed as (Wait Time Total / Event Waits)
->          '*'  Summary 'Avg Wait Time '              : Per-instance 'Wait Time Avg ' used to compute the following statistics
->          '*'  [Avg/Min/Max/Std Dev]                     : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
->          '*'  Cnt                                       : count of instances with wait times for the event

                              Wait                                   Event                    Wait Time                          Summary Avg Wait Time
      ---------------------------------------------------- ----------------------- ------------------------------- -----------------------------------------------
  I#  Class      Event                                            Waits %Timeouts       Total(s) Avg Wait %DB time        Avg        Min        Max    Std Dev  Cnt
 ---- ---------- ---------------------------------------- ------------ --------- ------------- -------- -------- ---------- ---------- ---------- ---------- ----
    *  User I/O   db file sequential read                  7.050821E+09       0.0  7,442,779.62   1.06ms    24.59     1.06ms     1.05ms     1.08ms    15.75us    4
                  DB CPU                                            N/A       N/A  6,075,235.66             20.07                                                 4
       Commit     log file sync                            1.705457E+09       0.0  4,018,701.82   2.36ms    13.28     2.36ms     1.95ms     3.15ms   551.02us    4
       Cluster    gc current grant 2-way                   2.671641E+09       0.0  1,550,173.17 580.23us     5.12   616.19us   194.39us      .96ms   317.46us    4
       Cluster    gc current block busy                    1.925081E+09       0.0  1,446,570.87 751.46us     4.78   757.45us   488.76us   944.48us   192.07us    4
       Cluster    gc current block congested                 42,098,159       0.0  1,360,282.90  32.31ms     4.49    31.98ms    20.19ms    43.48ms     9.52ms    4
       Cluster    gc cr block 3-way                        1.494389E+09       0.0  1,341,880.74 897.95us     4.43   901.96us   531.05us     1.18ms   305.83us    4
       Cluster    gc cr grant 2-way                        2.210294E+09       0.0  1,191,479.42 539.06us     3.94   556.54us   180.89us   928.93us   305.73us    4
       Cluster    gc current block 2-way                   1.651133E+09       0.0  1,059,897.19 641.92us     3.50   625.57us   197.85us   789.64us   285.88us    4
       Cluster    gc cr block congested                      30,264,924       0.0  1,042,665.74  34.45ms     3.45    33.51ms    20.52ms    44.32ms     9.80ms    4
```

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

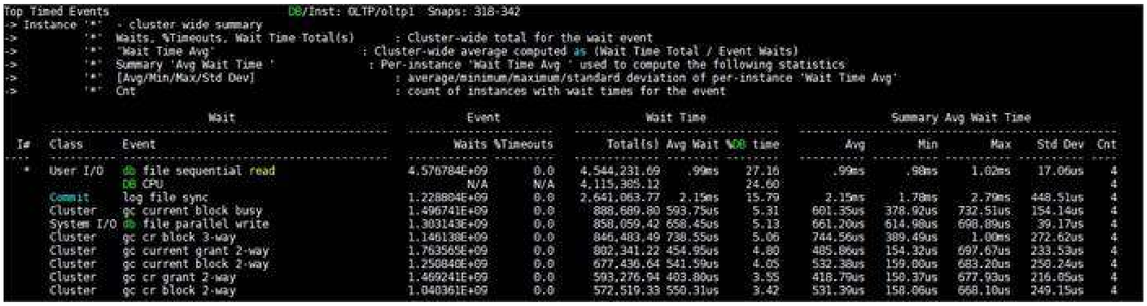

```
Top Timed Events                                   DB/Inst: OLTP/oltp1  Snaps: 318-342
-> Instance '*'  - cluster wide summary
->          '*'  Waits, %Timeouts, Wait Time Total(s)      : Cluster-wide total for the wait event
->          '*'  'Wait Time Avg'                          : Cluster-wide average computed as (Wait Time Total / Event Waits)
->          '*'  Summary 'Avg Wait Time '                  : Per-instance 'Wait Time Avg ' used to compute the following statistics
->          '*'  [Avg/Min/Max/Std Dev]                     : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
->          '*'  Cnt                                       : count of instances with wait times for the event

                          Wait                                       Event                      Wait Time                          Summary Avg Wait Time
    -----------------------------------------------------   ------------------------ ----------------------------- ------------------------------------------------------
 I#  Class      Event                                                 Waits %Timeouts     Total(s) Avg Wait %DB time        Avg        Min        Max    Std Dev  Cnt
---- ---------- -----------------------------------------   ------------ ---------  ------------- -------- ---------  ---------- ---------- ---------- ---------- ----
  *  User I/O   db file sequential read                       4.576784E+09       0.0   4,544,231.69    .99ms     27.16       .99ms      .98ms     1.02ms    17.06us    4
                DB CPU                                                 N/A       N/A   4,115,385.12              24.60                                                  4
     Commit     log file sync                                 1.228804E+09       0.0   2,641,063.77   2.15ms     15.79      2.15ms     1.78ms     2.79ms   448.51us    4
     Cluster    gc current block busy                         1.496741E+09       0.0     888,609.80 593.75us      5.31    601.35us   378.92us   732.51us   154.14us    4
     System I/O db file parallel write                        1.303143E+09       0.0     858,059.42 658.45us      5.13    661.20us   614.98us   698.89us    39.17us    4
     Cluster    gc cr block 3-way                             1.146138E+09       0.0     846,483.49 738.55us      5.06    744.56us   389.49us     1.00ms   272.62us    4
     Cluster    gc current grant 2-way                        1.763565E+09       0.0     802,341.22 454.95us      4.80    485.86us   154.32us   697.67us   233.53us    4
     Cluster    gc current block 2-way                        1.250840E+09       0.0     677,436.64 541.59us      4.05    532.38us   159.00us   683.20us   250.24us    4
     Cluster    gc cr grant 2-way                             1.469241E+09       0.0     593,276.94 403.80us      3.55    418.79us   150.37us   677.93us   216.05us    4
     Cluster    gc cr block 2-way                             1.040361E+09       0.0     572,519.33 550.31us      3.42    531.39us   158.06us   668.10us   249.15us    4
```

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

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

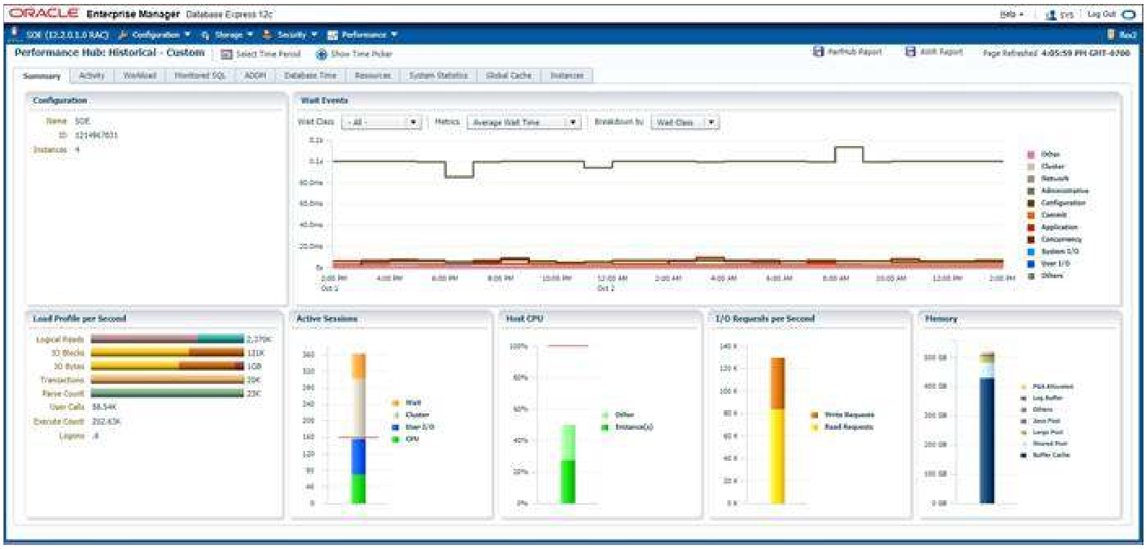

Figure 4 Wait Events, IO Requests per Second and CPU Utilization for the OLTP Database

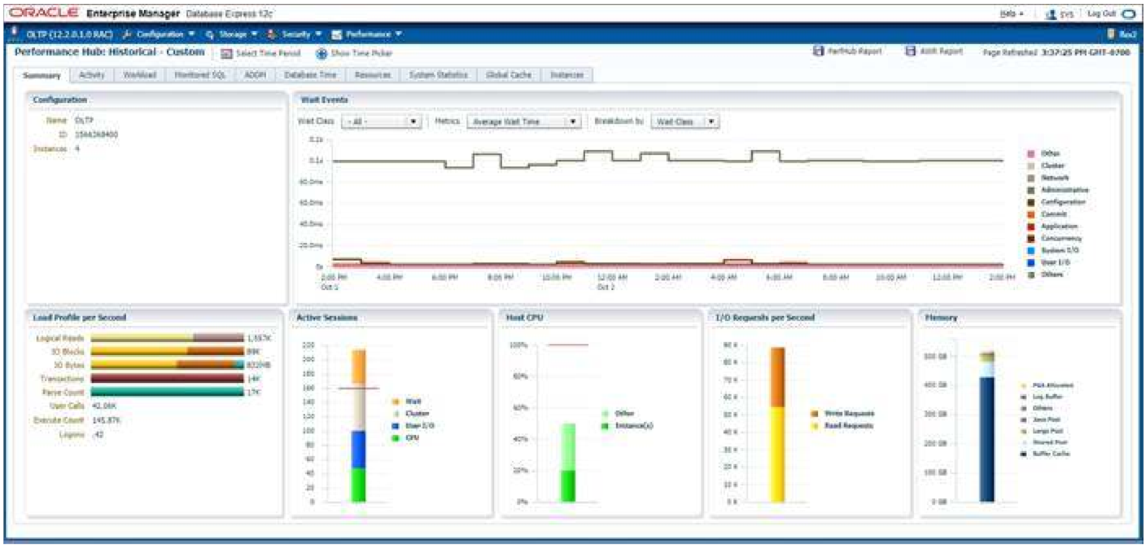

One (DSS) Database Performance

DSS database workloads are generally sequential in nature, read intensive and exercise large IO size. DSS database workload runs a small number of users that typically exercise extremely complex queries that run for hours. We configured 6 TB of DSS database by loading Swingbench sh schema into Datafile Tablespace.

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

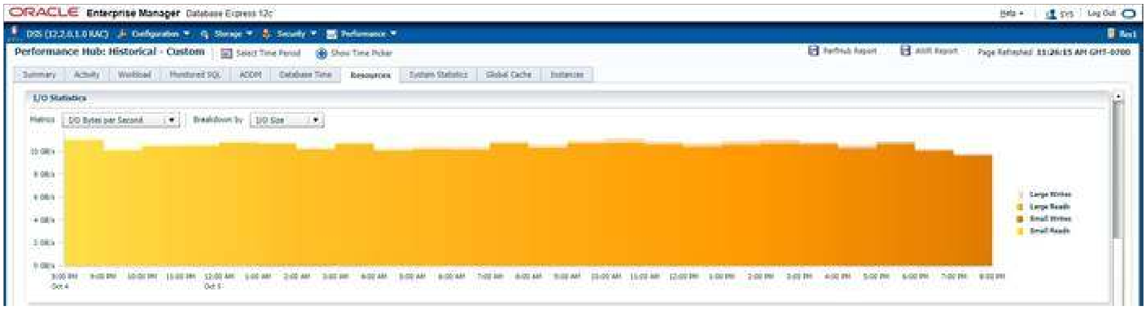

For 24 hour DSS workload test, we observed the total sustained IO bandwidth average was about 10.5 GB/sec after the initial ramp up workload. As shown in above screenshot, the IO was consistent throughout the run and we did not observe any significant dips in performance for complete period of time.

The screenshot shown below was captured from Oracle AWR report while running Swingbench SH workload on DSS database for 24 hours.

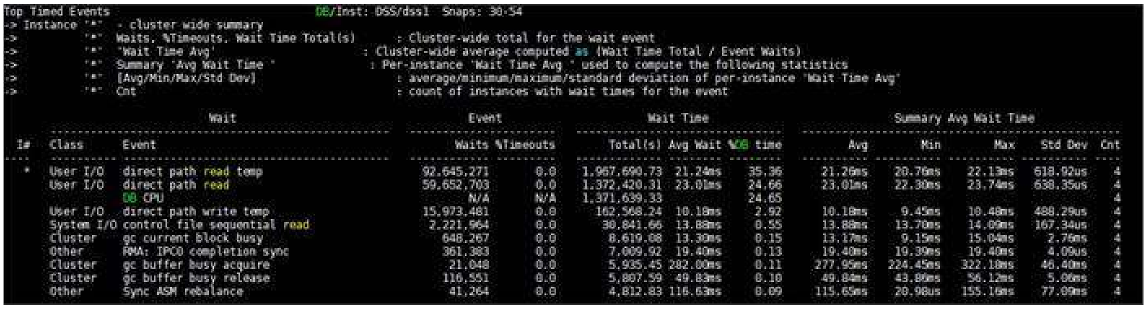

```
Top Timed Events                          DB/Inst: DSS/dss1  Snaps: 30-54
-> Instance '*'  - cluster wide summary
->          '*'  Waits, %Timeouts, Wait Time Total(s)      : Cluster-wide total for the wait event
->          '*'  'Wait Time Avg'                         : Cluster-wide average computed as (Wait Time Total / Event Waits)
->          '*'  Summary 'Avg Wait Time '                  : Per-instance 'Wait Time Avg ' used to compute the following statistics
->          '*'  [Avg/Min/Max/Std Dev]                        : average/minimum/maximum/standard deviation of per-instance 'Wait Time Avg'
->          '*'  Cnt                                      : count of instances with wait times for the event

                              Wait                             Event                    Wait Time                         Summary Avg Wait Time
       ---------------------------------------------   -------------------   ------------------------------   ----------------------------------------------
  I#   Class       Event                                   Waits %Timeouts         Total(s) Avg Wait %DB time         Avg        Min        Max    Std Dev  Cnt
 ----  ----------- --------------------------------   ---------- ---------   -------------- -------- -------   ---------- ---------- ---------- ---------- ----
    *  User I/O    direct path read temp              92,645,271       0.0     1,967,690.73  21.24ms   35.36      21.26ms    20.76ms    22.13ms   618.92us    4
       User I/O    direct path read                   59,652,703       0.0     1,372,420.31  23.01ms   24.66      23.01ms    22.30ms    23.74ms   638.35us    4
                   DB CPU                                    N/A       N/A     1,371,639.33            24.65                                                   4
       User I/O    direct path write temp             15,973,481       0.0       162,568.24  10.18ms    2.92      10.18ms     9.45ms    10.48ms   488.29us    4
       System I/O control file sequential read         2,221,964       0.0        30,841.66  13.88ms    0.55      13.88ms    13.70ms    14.09ms   167.34us    4
       Cluster     gc current block busy                 648,267       0.0         8,619.08  13.30ms    0.15      13.17ms     9.15ms    15.04ms     2.76ms    4
       Other       RMA: IPC0 completion sync             361,383       0.0         7,009.92  19.40ms    0.13      19.40ms    19.39ms    19.40ms     4.09us    4
       Cluster     gc buffer busy acquire                 21,048       0.0         5,935.45 282.00ms    0.11     277.95ms   224.45ms   322.10ms    46.40ms    4
       Cluster     gc buffer busy release                116,551       0.0         5,807.59  49.83ms    0.10      49.84ms    43.86ms    56.12ms     5.06ms    4
       Other       Sync ASM rebalance                     41,264       0.0         4,812.83 116.63ms    0.09     115.65ms    20.98us   155.16ms    77.09ms    4
```

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

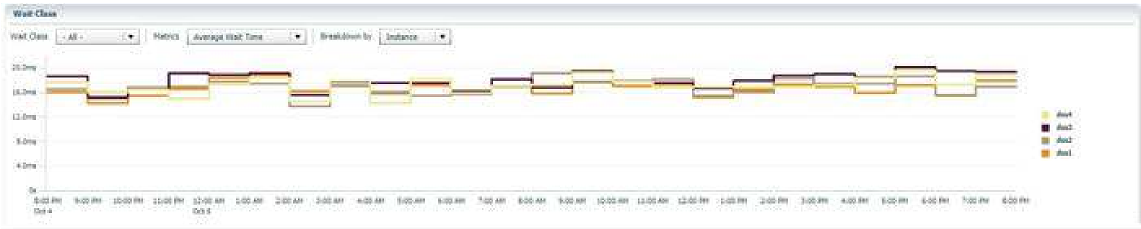

The screenshot below shows IO throughput and CPU utilization for each instance of DSS Database while running Swingbench workload test for 24 hours.

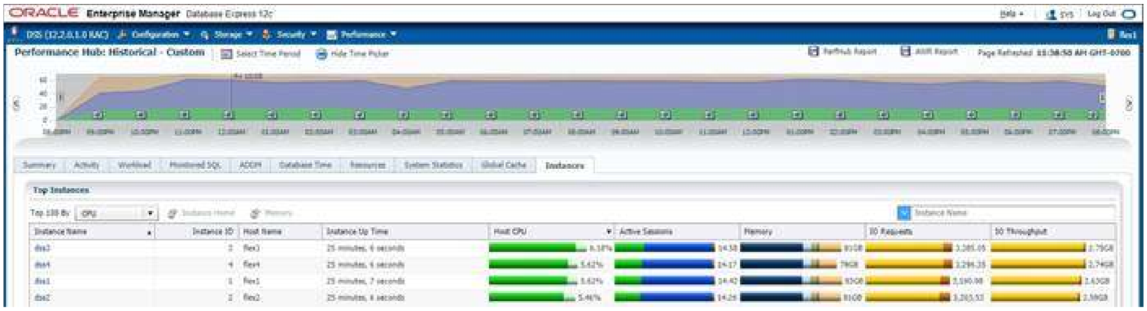

Instance Name	Instance ID	Host Name	Instance Up Time	Host CPU	Active Sessions	Memory	IO Requests	IO Throughput
dss3	3	flex3	25 minutes, 6 seconds	6.18%	14.58	81GB	3,285.05	2.75GB
dss4	4	flex4	25 minutes, 6 seconds	5.62%	14.17	79GB	3,294.35	2.74GB
dss1	1	flex1	25 minutes, 7 seconds	5.62%	14.42	93GB	3,190.98	2.63GB
dss2	2	flex2	25 minutes, 4 seconds	5.46%	14.29	81GB	3,203.53	2.59GB

All Three (SOE, OLTP and DSS) Database Performance

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

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

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

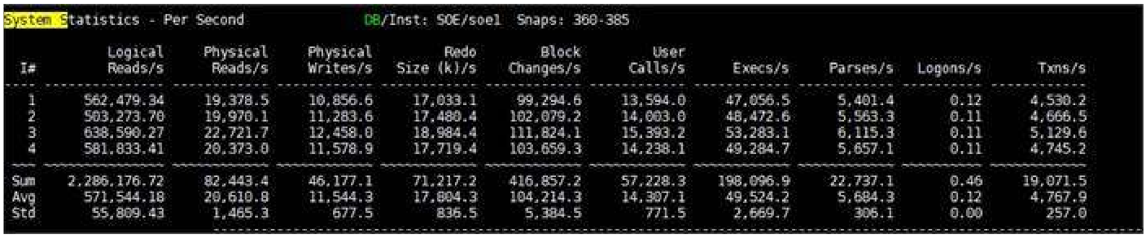

```
System Statistics - Per Second               DB/Inst: SOE/soe1  Snaps: 360-385

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

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

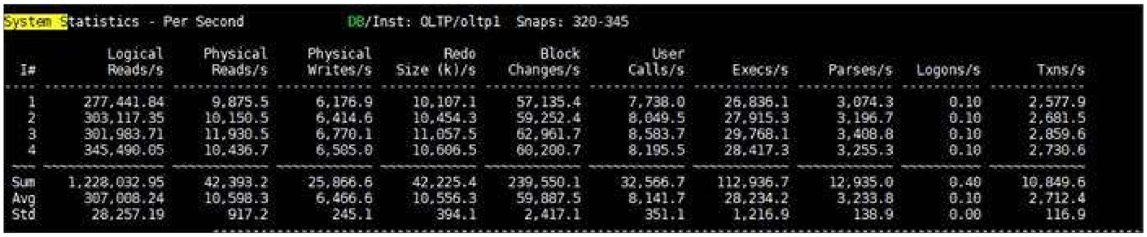

```
System Statistics - Per Second               DB/Inst: OLTP/oltp1  Snaps: 320-345

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

DSS Database AWR snapshot that highlights database MB per Seconds.

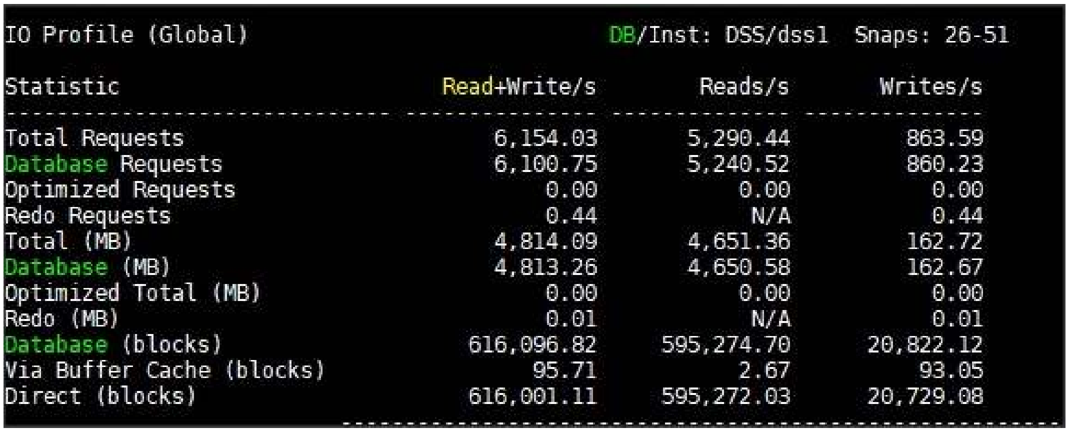

```
IO Profile (Global)                               DB/Inst: DSS/dss1  Snaps: 26-51

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

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

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

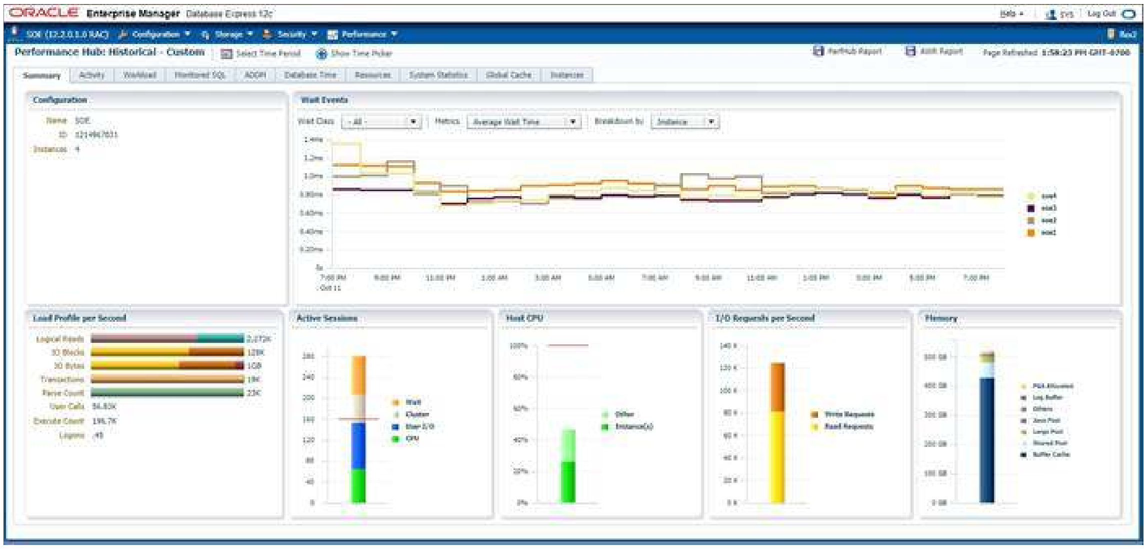

Figure 6 OLTP Database Wait Events per Instance, IO Requests per Second and CPU Utilization

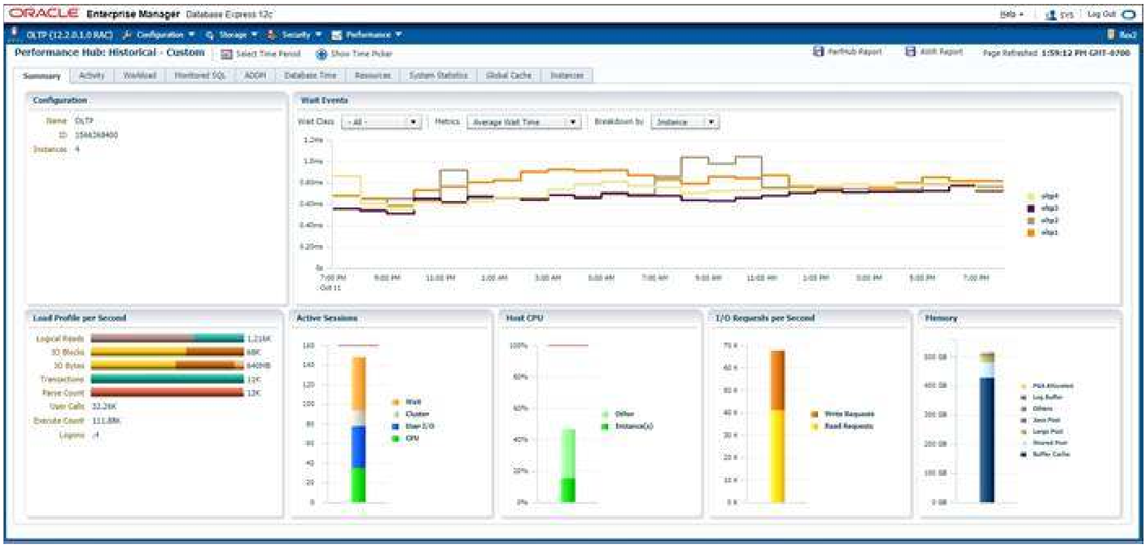

Figure 7 DSS Database Wait Events per Instance, IO Bytes and CPU Utilization

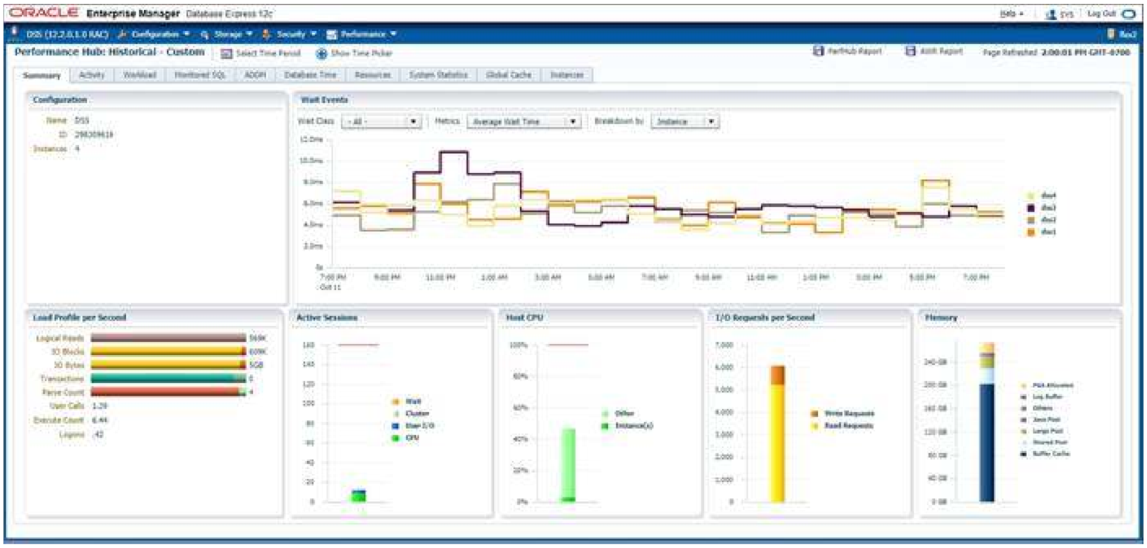

Resiliency and Failure Tests

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

Table 14 Hardware Failover Tests

Scenario	Tests
Test 1 – UCS FI – A Failure	Run the system on Full Database work Load. Power Off Fabric Interconnect – A and check network traffic on Fabr Interconnect – B.
Test 2 – UCS FI – B Failure	Run the system on Full Database work Load.

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

Figure 8 The FlexPod Solution Infrastructure Diagram Under Normal Operating Conditions

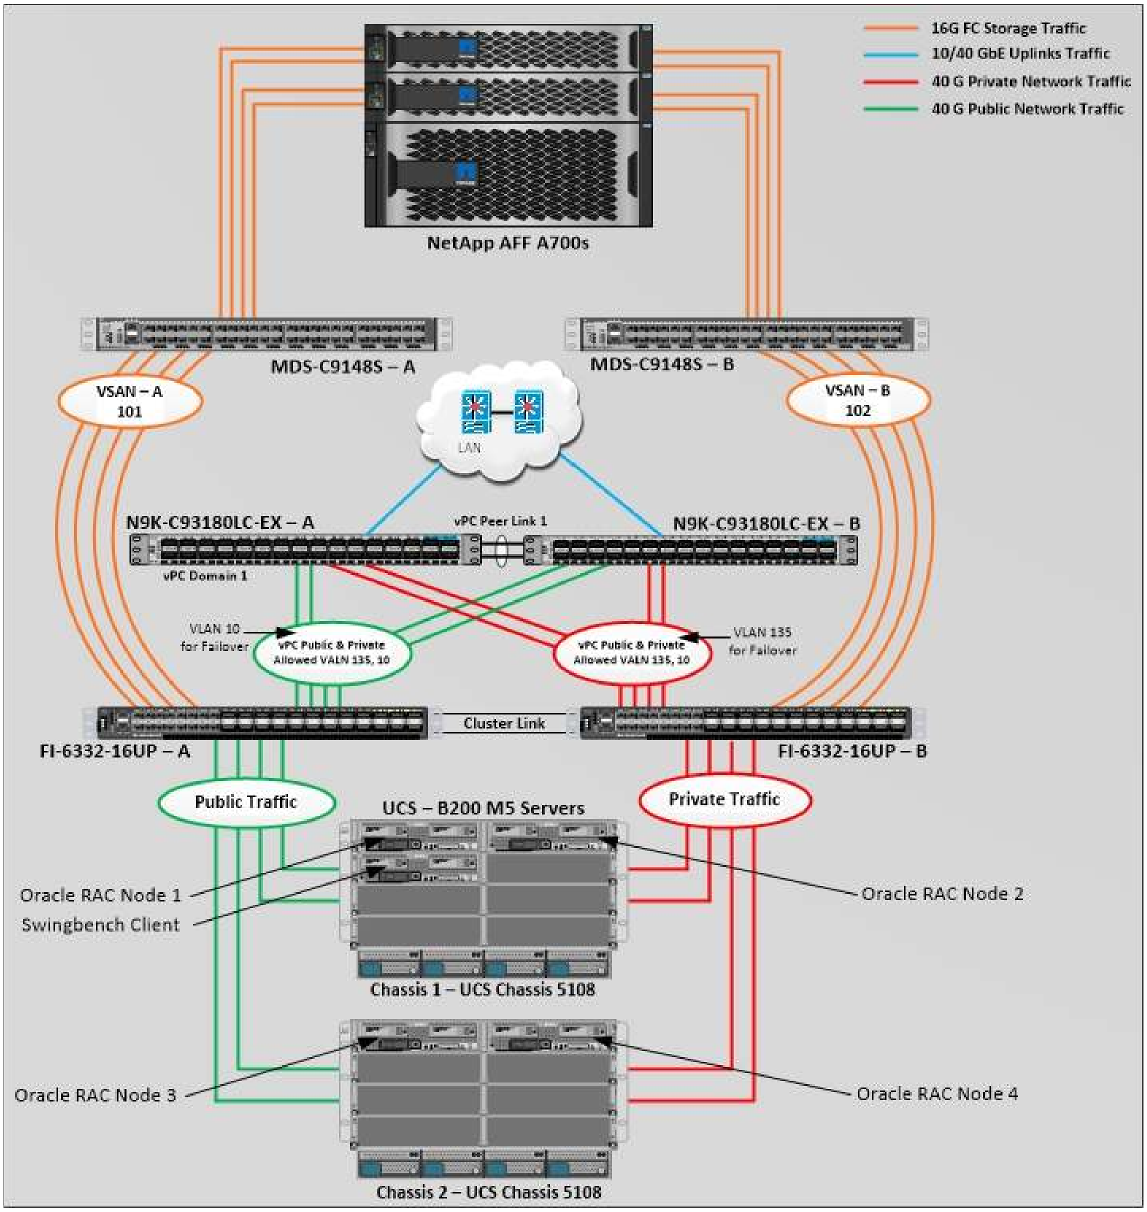

The screenshot below shows a complete infrastructure details of MAC address and VLAN information for Cisco UCS Fabric Interconnect – A switch before failover test.

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

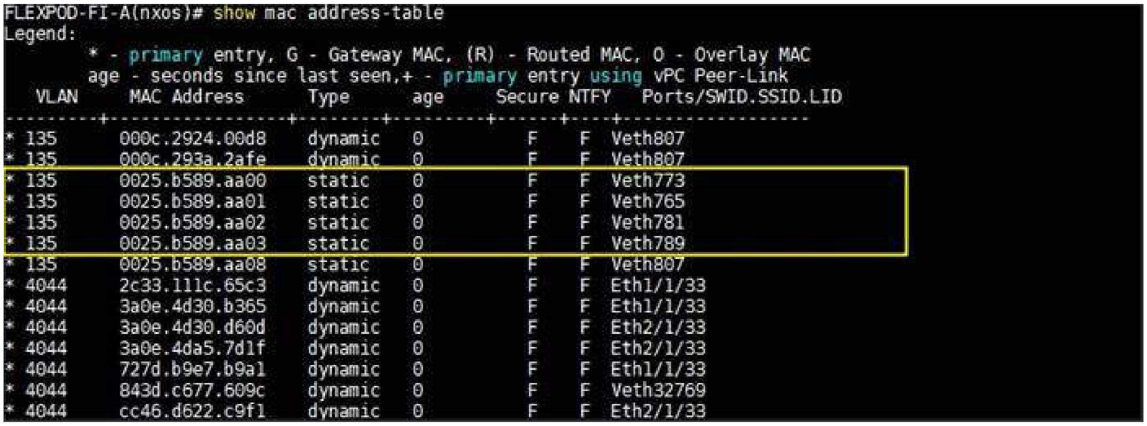

```
FLEXPOD-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8   dynamic   0          F     F  Veth807
* 135      000c.293a.2afe   dynamic   0          F     F  Veth807
* 135      0025.b589.aa00   static    0          F     F  Veth773
* 135      0025.b589.aa01   static    0          F     F  Veth765
* 135      0025.b589.aa02   static    0          F     F  Veth781
* 135      0025.b589.aa03   static    0          F     F  Veth789
* 135      0025.b589.aa08   static    0          F     F  Veth807
* 4044     2c33.111c.65c3   dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.b365   dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.d60d   dynamic   0          F     F  Eth2/1/33
* 4044     3a0e.4da5.7d1f   dynamic   0          F     F  Eth2/1/33
* 4044     727d.b9e7.b9a1   dynamic   0          F     F  Eth1/1/33
* 4044     843d.c677.609c   dynamic   0          F     F  Veth32769
* 4044     cc46.d622.c9f1   dynamic   0          F     F  Eth2/1/33
```

As shown in the screenshot above, Fabric Interconnect – A carry Oracle Public Network traffic on VLAN 135 under normal operating conditions before failover test.

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

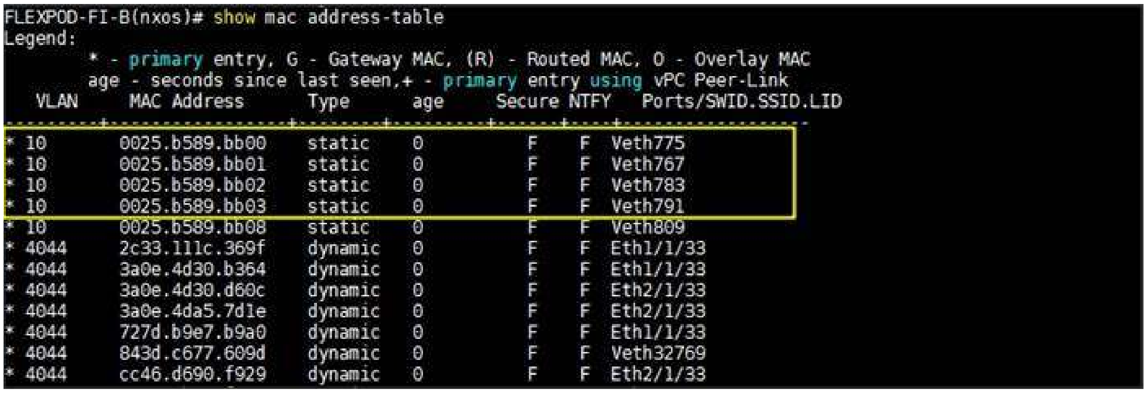

```
FLEXPOD-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b589.bb00    static    0          F     F  Veth775
* 10       0025.b589.bb01    static    0          F     F  Veth767
* 10       0025.b589.bb02    static    0          F     F  Veth783
* 10       0025.b589.bb03    static    0          F     F  Veth791
* 10       0025.b589.bb08    static    0          F     F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F     F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F     F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F     F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F     F  Veth32769
* 4044     cc46.d690.f929    dynamic   0          F     F  Eth2/1/33
```

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

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

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

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

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

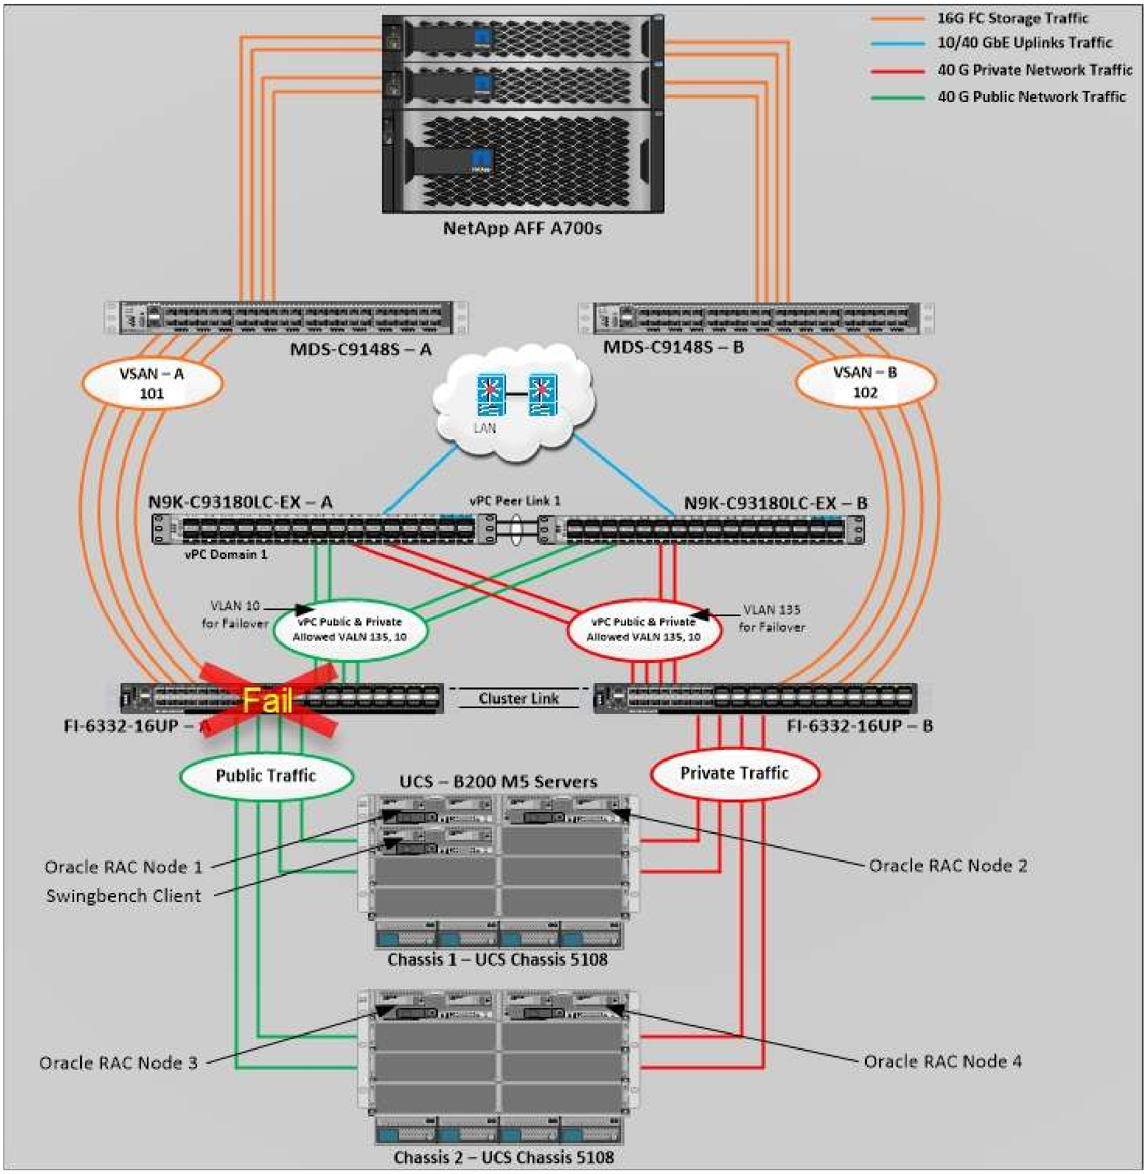

Unplug the power cable from Fabric Interconnect – A, and check the MAC address and VLAN information on Cisco UCS Fabric Interconnect – B.

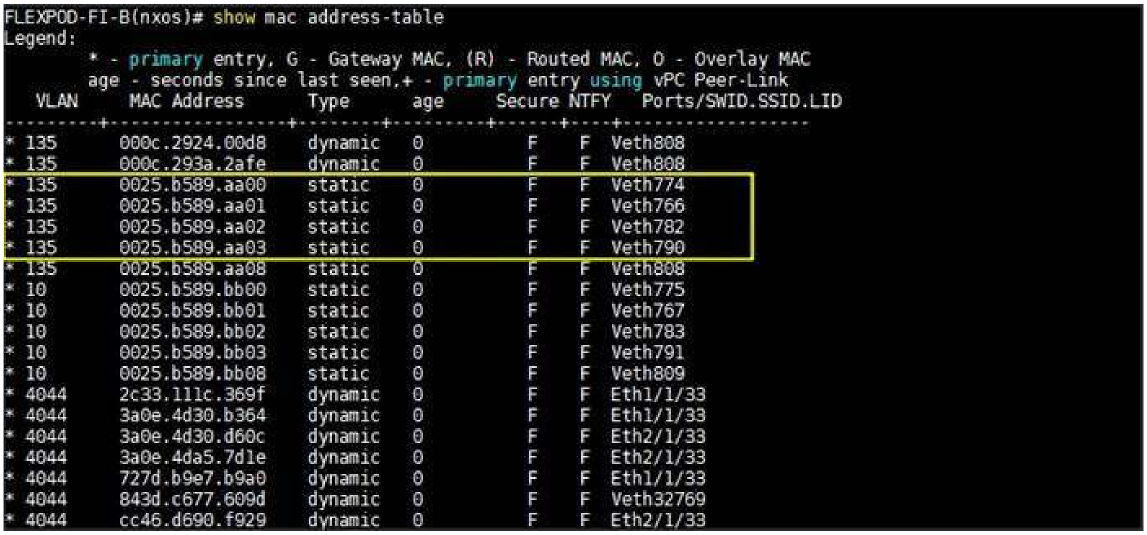

```
FLEXPOD-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F     F  Veth808
* 135      000c.293a.2afe    dynamic   0          F     F  Veth808
* 135      0025.b589.aa00    static    0          F     F  Veth774
* 135      0025.b589.aa01    static    0          F     F  Veth766
* 135      0025.b589.aa02    static    0          F     F  Veth782
* 135      0025.b589.aa03    static    0          F     F  Veth790
* 135      0025.b589.aa08    static    0          F     F  Veth808
* 10       0025.b589.bb00    static    0          F     F  Veth775
* 10       0025.b589.bb01    static    0          F     F  Veth767
* 10       0025.b589.bb02    static    0          F     F  Veth783
* 10       0025.b589.bb03    static    0          F     F  Veth791
* 10       0025.b589.bb08    static    0          F     F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F     F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F     F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F     F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F     F  Veth32769
* 4044     cc46.d690.f929    dynamic   0          F     F  Eth2/1/33
```

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

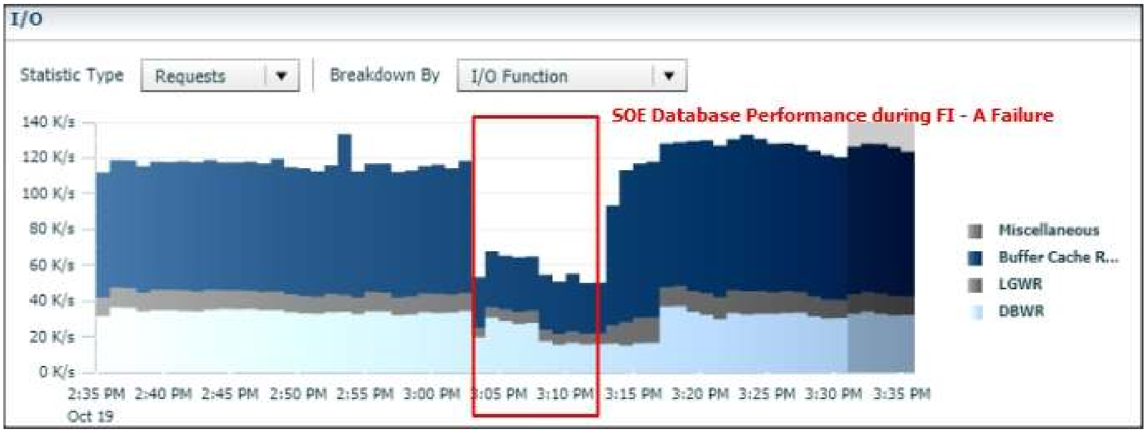
Statistic Type: Requests

Breakdown By: I/O Function

SOE Database Performance during FI - A Failure

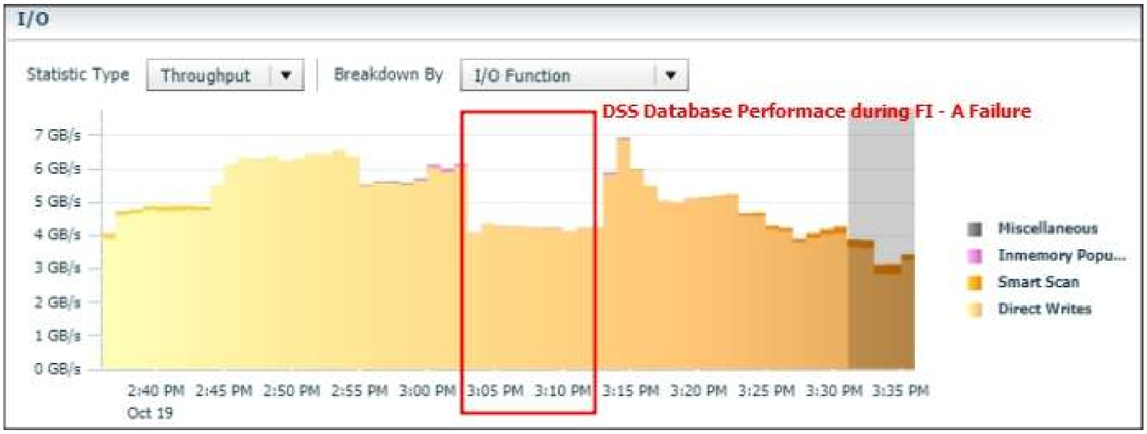
Statistic Type: Throughput

Breakdown By: I/O Function

DSS Database Performace during FI - A Failure

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

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

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

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

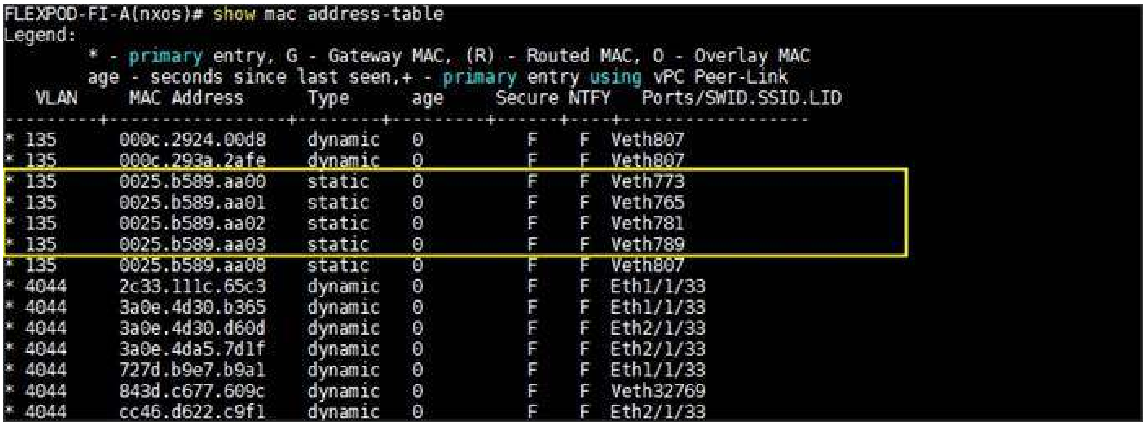

```
FLEXPOD-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F     F   Veth807
* 135      000c.293a.2afe    dynamic   0          F     F   Veth807
* 135      0025.b589.aa00    static    0          F     F   Veth773
* 135      0025.b589.aa01    static    0          F     F   Veth765
* 135      0025.b589.aa02    static    0          F     F   Veth781
* 135      0025.b589.aa03    static    0          F     F   Veth789
* 135      0025.b589.aa08    static    0          F     F   Veth807
* 4044     2c33.111c.65c3    dynamic   0          F     F   Eth1/1/33
* 4044     3a0e.4d30.b365    dynamic   0          F     F   Eth1/1/33
* 4044     3a0e.4d30.d60d    dynamic   0          F     F   Eth2/1/33
* 4044     3a0e.4da5.7d1f    dynamic   0          F     F   Eth2/1/33
* 4044     727d.b9e7.b9a1    dynamic   0          F     F   Eth1/1/33
* 4044     843d.c677.609c    dynamic   0          F     F   Veth32769
* 4044     cc46.d622.c9f1    dynamic   0          F     F   Eth2/1/33
```

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

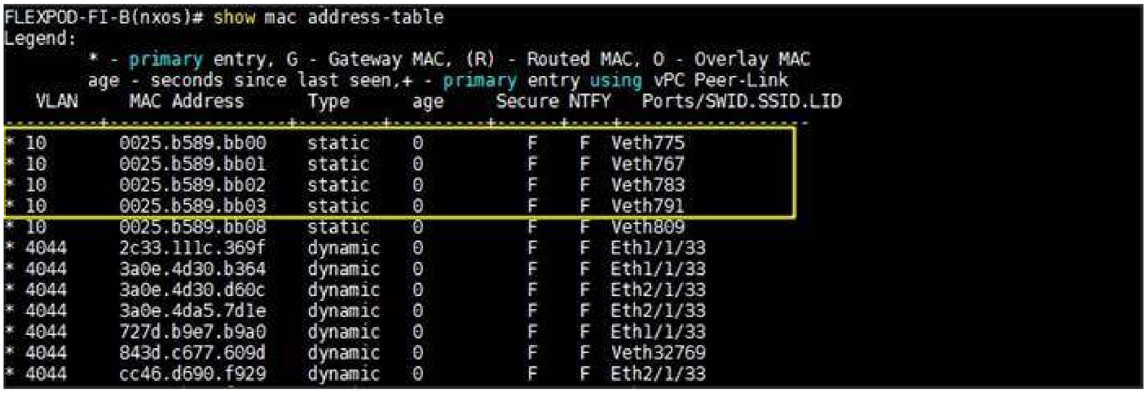

```
FLEXPOD-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b589.bb00    static    0          F     F   Veth775
* 10       0025.b589.bb01    static    0          F     F   Veth767
* 10       0025.b589.bb02    static    0          F     F   Veth783
* 10       0025.b589.bb03    static    0          F     F   Veth791
* 10       0025.b589.bb08    static    0          F     F   Veth809
* 4044     2c33.111c.369f    dynamic   0          F     F   Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F     F   Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F     F   Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F     F   Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F     F   Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F     F   Veth32769
* 4044     cc46.d690.f929    dynamic   0          F     F   Eth2/1/33
```

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

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

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

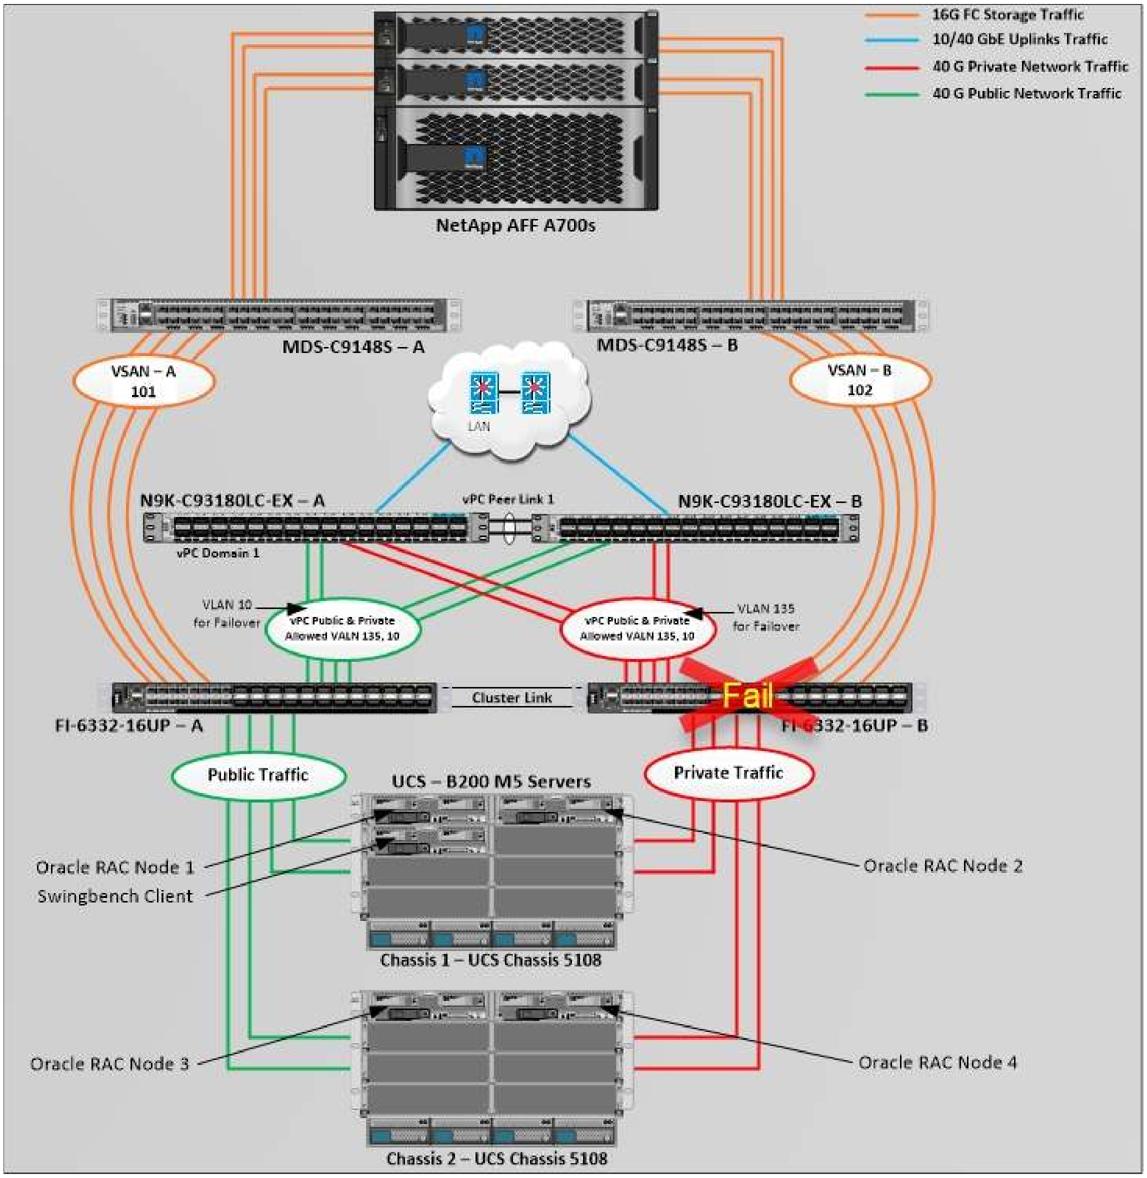

Unplug power cable from Fabric Interconnect – B, and check the MAC address and VLAN information on Cisco UCS Fabric Interconnect – A.

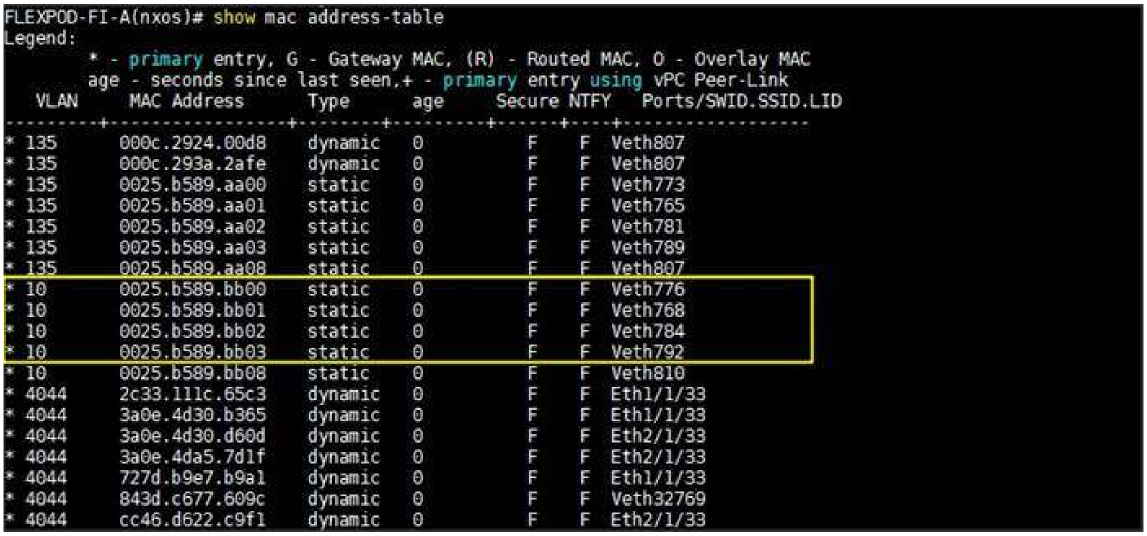

```
FLEXPOD-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8   dynamic   0          F    F   Veth807
* 135      000c.293a.2afe   dynamic   0          F    F   Veth807
* 135      0025.b589.aa00   static    0          F    F   Veth773
* 135      0025.b589.aa01   static    0          F    F   Veth765
* 135      0025.b589.aa02   static    0          F    F   Veth781
* 135      0025.b589.aa03   static    0          F    F   Veth789
* 135      0025.b589.aa08   static    0          F    F   Veth807
* 10       0025.b589.bb00   static    0          F    F   Veth776
* 10       0025.b589.bb01   static    0          F    F   Veth768
* 10       0025.b589.bb02   static    0          F    F   Veth784
* 10       0025.b589.bb03   static    0          F    F   Veth792
* 10       0025.b589.bb08   static    0          F    F   Veth810
* 4044     2c33.111c.65c3   dynamic   0          F    F   Eth1/1/33
* 4044     3a0e.4d30.b365   dynamic   0          F    F   Eth1/1/33
* 4044     3a0e.4d30.d60d   dynamic   0          F    F   Eth2/1/33
* 4044     3a0e.4da5.7d1f   dynamic   0          F    F   Eth2/1/33
* 4044     727d.b9e7.b9a1   dynamic   0          F    F   Eth1/1/33
* 4044     843d.c677.609c   dynamic   0          F    F   Veth32769
* 4044     cc46.d622.c9f1   dynamic   0          F    F   Eth2/1/33
```

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

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

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

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

Test 3 and 4 – Cisco Nexus Switch Failure Test

We conducted a hardware failure test on Nexus Switch – A by disconnecting power cable to the switch and checking the MAC address and VLAN information on Cisco UCS Nexus Switch – B. During Nexus Switch – A failure, it routed all the Private Network and Public Network Traffic of VLAN 10 and VLAN 134 to Nexus Switch – B. So, Nexus Switch – A Failover did not cause any disruption to Private, Public and Storage Network Traffic.

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

Test 5 and 6 – Cisco MDS Switch Failure Test

We conducted hardware failure test on MDS Switch – A by disconnecting power cable to the Switch and checking the storage network traffic on MDS Switch – B. As we expected, we observed some impact on all three databases performance as we lost half of the VSAN (VSAN-A 101) traffic. While VSAN-A (101) stays locally into the switch and only carry storage traffic through the MDS switch A, VSAN-A doesn't failover to MDS Switch B therefore we reduced server to storage connectivity into half during MDS Switch A failure.

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

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

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

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

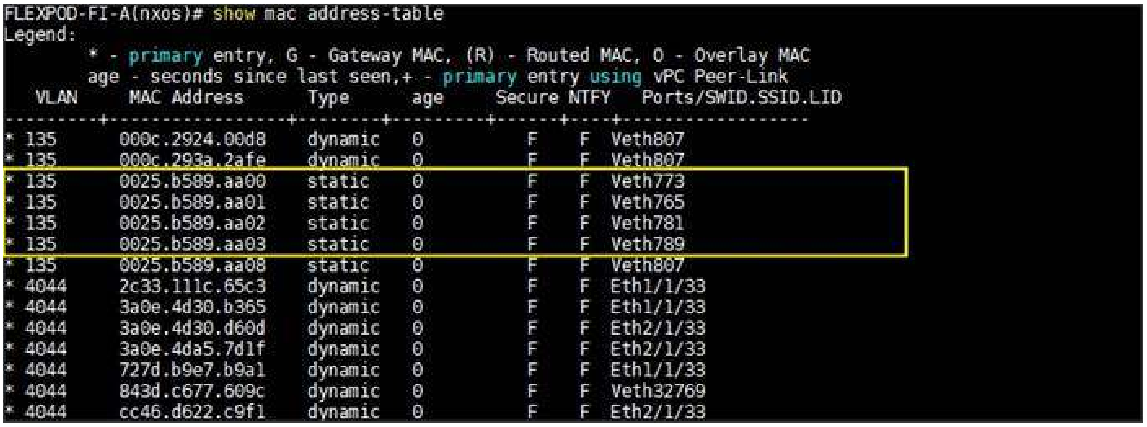

```
FLEXPOD-FI-A(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 135      000c.2924.00d8    dynamic   0          F     F  Veth807
* 135      000c.293a.2afe    dynamic   0          F     F  Veth807
* 135      0025.b589.aa00    static    0          F     F  Veth773
* 135      0025.b589.aa01    static    0          F     F  Veth765
* 135      0025.b589.aa02    static    0          F     F  Veth781
* 135      0025.b589.aa03    static    0          F     F  Veth789
* 135      0025.b589.aa08    static    0          F     F  Veth807
* 4044     2c33.111c.65c3    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.b365    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.d60d    dynamic   0          F     F  Eth2/1/33
* 4044     3a0e.4da5.7d1f    dynamic   0          F     F  Eth2/1/33
* 4044     727d.b9e7.b9a1    dynamic   0          F     F  Eth1/1/33
* 4044     843d.c677.609c    dynamic   0          F     F  Veth32769
* 4044     cc46.d622.c9f1    dynamic   0          F     F  Eth2/1/33
```

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

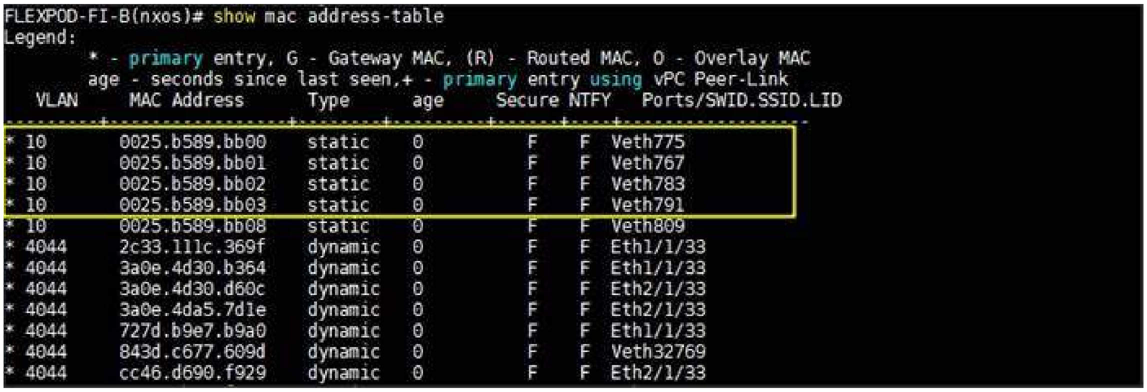

```
FLEXPOD-FI-B(nxos)# show mac address-table
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY   Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0025.b589.bb00    static    0          F     F  Veth775
* 10       0025.b589.bb01    static    0          F     F  Veth767
* 10       0025.b589.bb02    static    0          F     F  Veth783
* 10       0025.b589.bb03    static    0          F     F  Veth791
* 10       0025.b589.bb08    static    0          F     F  Veth809
* 4044     2c33.111c.369f    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.b364    dynamic   0          F     F  Eth1/1/33
* 4044     3a0e.4d30.d60c    dynamic   0          F     F  Eth2/1/33
* 4044     3a0e.4da5.7d1e    dynamic   0          F     F  Eth2/1/33
* 4044     727d.b9e7.b9a0    dynamic   0          F     F  Eth1/1/33
* 4044     843d.c677.609d    dynamic   0          F     F  Veth32769
* 4044     cc46.d690.f929    dynamic   0          F     F  Eth2/1/33
```

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

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

Summary

The Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, network, storage access, and virtualization into a single cohesive system. Cisco UCS is an ideal platform for the architecture of mission critical database workloads such as Oracle RAC. The FlexPod Datacenter with NetApp All Flash AFF system is a converged infrastructure platform that combines best-of breed technologies from Cisco and NetApp into a powerful converged platform for enterprise applications. The pre-validated FlexPod architecture delivers proven value, agility, and performance that drive higher productivity, faster decision making, and greater opportunities for growth.

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

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

- Cisco UCS stateless computing architecture provided by the Service Profile capability of Cisco UCS allows fast, non-disruptive workload changes to be executed simply and seamlessly across the integrated UCS infrastructure and Cisco x86 servers.
- A single platform built from unified compute, fabric, and storage technologies, allowing you to scale to large-scale data centers without architectural changes.
- Enabling faster deployments, greater flexibility of choice, efficiency, high availability and lower risk.

Appendix

Cisco Nexus 93180 LC-EX Configuration

```
FLEXPOD-NEXUS-A# show running-config
!Command: show running-config
!Time: Tue Nov 20 21:01:25 2018
version 7.0(3)I6(1)
switchname FLEXPOD-NEXUS-A
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
vdc FLEXPOD-NEXUS-A id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature vpc

no password strength-check
username admin password 5 $5
$FWM47Q8G$RxxMd920jZrjYX5GXLRCkfWgDtka29dQV1TiP/k4VmD  role network-admin
ip domain-lookup
system default switchport
system qos
  service-policy type network-qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0x54fc9b6e68ba6a06acf4165853e18078 priv
0x54fc9b6e68ba6a06acf4165853e18078 localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
```

```
ntp server 72.163.32.44 use-vrf default

vlan 1,10,135
vlan 10
  name Oracle_RAC_Private_Traffic
vlan 135
  name Oracle_RAC_Public_Traffic

spanning-tree port type edge bpduguard default
spanning-tree port type network default
vrf context management
  ip route 0.0.0.0/0 10.29.135.1
port-channel load-balance src-dst l4port
vpc domain 1
  peer-keepalive destination 10.29.135.104 source 10.29.135.103
  auto-recovery

interface Vlan1

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

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

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

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

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

interface Ethernet1/3
...........
...........
interface Ethernet1/20

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

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

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

interface Ethernet1/24
  description Fabric-Interconnect-B-32
  switchport mode trunk
  switchport trunk allowed vlan 1,10,135
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 52 mode active

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

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

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

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

interface mgmt0
  vrf member management
  ip address 10.29.135.103/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I6.1.bin
no system default switchport shutdown
```

MDS 9148S FC Switch Configuration

```
FLEXPOD-MDS-A# show running-config
!Command: show running-config
!Time: Tue Nov 20 21:19:36 2018
version 7.3(0)D1(1)
power redundancy-mode redundant
feature npiv
no feature http-server
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $5$lC5R4cI9$rfRLLbkccOM1lwxALFaHFjx1XtCse5pPjAk3i/HJi0/  role
network-admin
no password strength-check
ip domain-lookup
ip host FLEXPOD-MDS-A  10.29.135.105
aaa group server radius radius
snmp-server user admin network-admin auth md5 0xe2b9f6bd292aa43dc651e566873540cf priv
0xe2b9f6bd292aa43dc651e566873540cf localizedkey
rmon event 1 log description FATAL(1) owner PMON@FATAL
rmon event 2 log description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log description ERROR(3) owner PMON@ERROR
rmon event 4 log description WARNING(4) owner PMON@WARNING
rmon event 5 log description INFORMATION(5) owner PMON@INFO
ntp server 72.163.32.44
vsan database
  vsan 101
device-alias database
  device-alias name flex1-hba0 pwwn 20:00:00:25:b5:8a:a0:00
  device-alias name flex1-hba2 pwwn 20:00:00:25:b5:8a:a0:01
  device-alias name flex2-hba0 pwwn 20:00:00:25:b5:8a:a0:02
  device-alias name flex2-hba2 pwwn 20:00:00:25:b5:8a:a0:03
  device-alias name flex3-hba0 pwwn 20:00:00:25:b5:8a:a0:04
  device-alias name flex3-hba2 pwwn 20:00:00:25:b5:8a:a0:05
  device-alias name flex4-hba0 pwwn 20:00:00:25:b5:8a:a0:06
  device-alias name flex4-hba2 pwwn 20:00:00:25:b5:8a:a0:07
  device-alias name NetApp-A700s-01-2A pwwn 20:06:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-01-3A pwwn 20:07:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-2A pwwn 20:09:00:a0:98:af:7c:5b
  device-alias name NetApp-A700s-02-3A pwwn 20:0a:00:a0:98:af:7c:5b
device-alias commit
fcdomain fcid database
  vsan 1 wwn 20:06:8c:60:4f:bd:31:80 fcid 0x940000 dynamic
  vsan 1 wwn 20:05:8c:60:4f:bd:31:80 fcid 0x940100 dynamic
  vsan 1 wwn 20:01:00:de:fb:92:99:00 fcid 0x940300 dynamic
  vsan 1 wwn 20:02:00:de:fb:92:99:00 fcid 0x940400 dynamic
  vsan 1 wwn 20:03:00:de:fb:92:99:00 fcid 0x940500 dynamic
  vsan 1 wwn 20:04:00:de:fb:92:99:00 fcid 0x940600 dynamic
  vsan 1 wwn 50:0a:09:82:80:12:f8:47 fcid 0x940700 dynamic
  vsan 1 wwn 50:0a:09:84:80:12:f8:47 fcid 0x940800 dynamic
  vsan 1 wwn 50:0a:09:81:80:12:f8:47 fcid 0x940900 dynamic
  vsan 1 wwn 50:0a:09:83:80:12:f8:47 fcid 0x940a00 dynamic
  vsan 1 wwn 50:0a:09:81:80:12:f8:3d fcid 0x940b00 dynamic
  vsan 1 wwn 50:0a:09:83:80:12:f8:3d fcid 0x940c00 dynamic
  vsan 101 wwn 50:0a:09:81:80:12:f8:47 fcid 0xa20000 dynamic
  vsan 101 wwn 50:0a:09:83:80:12:f8:47 fcid 0xa20100 dynamic
  vsan 101 wwn 50:0a:09:81:80:12:f8:3d fcid 0xa20200 dynamic
  vsan 101 wwn 50:0a:09:83:80:12:f8:3d fcid 0xa20300 dynamic
  vsan 101 wwn 20:01:00:de:fb:92:99:00 fcid 0xa20400 dynamic
  vsan 101 wwn 20:00:00:25:b5:8a:a0:00 fcid 0xa20702 dynamic
!           [flex1-hba0]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:04 fcid 0xa20501 dynamic
!           [flex3-hba0]
  vsan 1 wwn 20:4d:54:7f:ee:76:cd:80 fcid 0x940200 dynamic
  vsan 101 wwn 20:00:00:25:b5:8a:a0:06 fcid 0xa20502 dynamic
!           [flex4-hba0]
  vsan 101 wwn 20:00:00:25:b5:8a:a0:02 fcid 0xa20701 dynamic
!           [flex2-hba0]
  vsan 101 wwn 20:02:00:de:fb:92:99:00 fcid 0xa20500 dynamic
  vsan 101 wwn 20:03:00:de:fb:92:99:00 fcid 0xa20600 dynamic
  vsan 101 wwn 20:04:00:de:fb:92:99:00 fcid 0xa20700 dynamic
  vsan 101 wwn 20:02:00:a0:98:af:7c:5b fcid 0xa20001 dynamic
  vsan 101 wwn 20:04:00:a0:98:af:7c:5b fcid 0xa20101 dynamic
  vsan 101 wwn 20:06:00:a0:98:af:7c:5b fcid 0xa20002 dynamic
!           [NetApp-A700s-01-2A]
  vsan 101 wwn 20:07:00:a0:98:af:7c:5b fcid 0xa20102 dynamic
```

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

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

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

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

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

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

zoneset activate name flex vsan 101
do clear zone database vsan 101
```

```
!Full Zone Database Section for vsan 101
zone name flex1 vsan 101
  member pwwn 20:00:00:25:b5:8a:a0:00
!           [flex1-hba0]
  member pwwn 20:00:00:25:b5:8a:a0:01
!           [flex1-hba2]
  member pwwn 20:06:00:a0:98:af:7c:5b
!           [NetApp-A700s-01-2A]
  member pwwn 20:07:00:a0:98:af:7c:5b
!           [NetApp-A700s-01-3A]
  member pwwn 20:09:00:a0:98:af:7c:5b
!           [NetApp-A700s-02-2A]
  member pwwn 20:0a:00:a0:98:af:7c:5b
!           [NetApp-A700s-02-3A]

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

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

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

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

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

interface mgmt0
  ip address 10.29.135.105 255.255.255.0
vsan database
  vsan 101 interface fc1/1
  vsan 101 interface fc1/2
  vsan 101 interface fc1/3
  vsan 101 interface fc1/4
  vsan 101 interface fc1/5
  vsan 101 interface fc1/6
  vsan 101 interface fc1/7
  vsan 101 interface fc1/8
  vsan 101 interface fc1/9
  vsan 101 interface fc1/10
  vsan 101 interface fc1/11
  vsan 101 interface fc1/12
```

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

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

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

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

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

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

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

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

interface fc1/9
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

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

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

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

interface fc1/13
  port-license acquire
...............
...............
interface fc1/48
  port-license acquire

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

Multipath Configuration “/etc/multipath.conf”

```
defaults {
        find_multipaths yes
        user_friendly_names no
}
multipaths {
        multipath {
                wwid           3600a098038304173475d4c766a49744e
                alias         node1_os
        }
        multipath {
                wwid           3600a098038304173475d4c766a49754a
                alias         ocrvote_1
        }
        multipath {
                wwid           3600a098038304173452b4c7959614634
                alias         ocrvote_2
        }
        multipath {
                wwid           3600a098038304173475d4c766a49754b
                alias         oradata_slob_1
        }
        multipath {
                wwid           3600a098038304173452b4c7959614635
                alias         oradata_slob_2
        }
        multipath {
                wwid           3600a098038304173475d4c766a49754c
                alias         oradata_slob_3
        }
        multipath {
                wwid           3600a098038304173452b4c7959614636
                alias         oradata_slob_4
        }
        multipath {
                wwid           3600a098038304173475d4c766a49754d
                alias         oradata_slob_5
        }
        multipath {
```

```
        wwid          3600a098038304173452b4c7959614637
        alias         oradata_slob_6
    }
    multipath {
        wwid          3600a098038304173475d4c766a49754e
        alias         oradata_slob_7
    }
    multipath {
        wwid          3600a098038304173452b4c7959614638
        alias         oradata_slob_8
    }
    multipath {
        wwid          3600a098038304173475d4c766a49754f
        alias         oradata_slob_9
    }
    multipath {
        wwid          3600a098038304173452b4c7959614639
        alias         oradata_slob_10
    }
    multipath {
        wwid          3600a098038304173475d4c766a497550
        alias         oradata_slob_11
    }
    multipath {
        wwid          3600a098038304173452b4c795961462d
        alias         oradata_slob_12
    }
    multipath {
        wwid          3600a098038304173475d4c766a497551
        alias         oradata_slob_13
    }
    multipath {
        wwid          3600a098038304173452b4c7959614641
        alias         oradata_slob_14
    }
    multipath {
        wwid          3600a098038304173475d4c766a497552
        alias         oradata_slob_15
    }
```

```
 multipath {
        wwid          3600a098038304173452b4c7959614642
        alias        oradata_slob_16
 }
 multipath {
        wwid          3600a098038304173475d4c766a497553
        alias        oraredo_slob_1
 }
 multipath {
        wwid          3600a098038304173452b4c7959614643
        alias        oraredo_slob_2
 }
 multipath {
        wwid          3600a098038304173475d4c766a497563
        alias        oradata_oltp_1
 }
multipath {
        wwid          3600a098038304173452b4c795961464e
        alias        oradata_oltp_2
 }
multipath {
        wwid          3600a098038304173475d4c766a497564
        alias        oradata_oltp_3
 }
multipath {
        wwid          3600a098038304173452b4c795961464f
        alias        oradata_oltp_4
 }
multipath {
        wwid          3600a098038304173475d4c766a497565
        alias        oradata_oltp_5
 }
multipath {
        wwid          3600a098038304173452b4c7959614650
        alias        oradata_oltp_6
 }
multipath {
        wwid          3600a098038304173475d4c766a497566
        alias        oradata_oltp_7
```

```
         }
        multipath {
                wwid            3600a098038304173452b4c7959614651
                alias         oradata_oltp_8
         }
        multipath {
                wwid            3600a098038304173475d4c766a497567
                alias         oradata_oltp_9
         }
        multipath {
                wwid            3600a098038304173452b4c7959614652
                alias         oradata_oltp_10
         }
        multipath {
                wwid            3600a098038304173475d4c766a497568
                alias         oradata_oltp_11
         }
        multipath {
                wwid            3600a098038304173452b4c7959614653
                alias         oradata_oltp_12
         }
        multipath {
                wwid            3600a098038304173475d4c766a497569
                alias         oradata_oltp_13
         }
        multipath {
                wwid            3600a098038304173452b4c7959614654
                alias         oradata_oltp_14
         }
        multipath {
                wwid            3600a098038304173475d4c766a49756a
                alias         oradata_oltp_15
         }
        multipath {
                wwid            3600a098038304173452b4c7959614655
                alias         oradata_oltp_16
         }
         multipath {
                wwid            3600a098038304173475d4c766a49756b
        alias          oraredo_oltp_1
}
multipath {
        wwid            3600a098038304173452b4c7959614656
        alias          oraredo_oltp_2
}
multipath {
        wwid            3600a098038304173475d4c766a49756c
        alias          oradata_soe_1
}
multipath {
        wwid            3600a098038304173452b4c7959614657
        alias          oradata_soe_2
}
multipath {
        wwid            3600a098038304173475d4c766a49756d
        alias          oradata_soe_3
}
multipath {
        wwid            3600a098038304173452b4c7959614658
        alias          oradata_soe_4
}
multipath {
        wwid            3600a098038304173475d4c766a49756e
        alias          oradata_soe_5
}
multipath {
        wwid            3600a098038304173452b4c7959614659
        alias          oradata_soe_6
}
multipath {
        wwid            3600a098038304173475d4c766a49756f
        alias          oradata_soe_7
}
multipath {
        wwid            3600a098038304173452b4c795961465a
        alias          oradata_soe_8
}
multipath {
        wwid          3600a098038304173475d4c766a497570
        alias        oradata_soe_9
    }
    multipath {
        wwid          3600a098038304173452b4c795961462f
        alias        oradata_soe_10
    }
    multipath {
        wwid          3600a098038304173475d4c766a497571
        alias        oradata_soe_11
    }
    multipath {
        wwid          3600a098038304173452b4c7959614661
        alias        oradata_soe_12
    }
    multipath {
        wwid          3600a098038304173475d4c766a497572
        alias        oradata_soe_13
    }
    multipath {
        wwid          3600a098038304173452b4c7959614662
        alias        oradata_soe_14
    }
    multipath {
        wwid          3600a098038304173475d4c766a497573
        alias        oradata_soe_15
    }
    multipath {
        wwid          3600a098038304173452b4c7959614663
        alias        oradata_soe_16
    }
    multipath {
        wwid          3600a098038304173475d4c766a497574
        alias        oraredo_soe_1
    }
    multipath {
        wwid          3600a098038304173452b4c7959614664
        alias        oraredo_soe_2
    }
```

```
multipath {
        wwid           3600a098038304173475d4c766a497575
        alias         oradata_dss_1
}
multipath {
        wwid           3600a098038304173452b4c7959614665
        alias         oradata_dss_2
}
multipath {
        wwid           3600a098038304173475d4c766a497576
        alias         oradata_dss_3
}
multipath {
        wwid           3600a098038304173452b4c7959614666
        alias         oradata_dss_4
}
multipath {
        wwid           3600a098038304173475d4c766a497577
        alias         oradata_dss_5
}
multipath {
        wwid           3600a098038304173452b4c7959614667
        alias         oradata_dss_6
}
multipath {
        wwid           3600a098038304173475d4c766a497578
        alias         oradata_dss_7
}
multipath {
        wwid           3600a098038304173452b4c7959614668
        alias         oradata_dss_8
}
multipath {
        wwid           3600a098038304173475d4c766a497579
        alias         oradata_dss_9
}
multipath {
        wwid           3600a098038304173452b4c7959614669
        alias         oradata_dss_10
        }
        multipath {
                wwid             3600a098038304173475d4c766a49757a
                alias           oradata_dss_11
        }
        multipath {
                wwid             3600a098038304173452b4c795961466a
                alias           oradata_dss_12
        }
        multipath {
                wwid             3600a098038304173475d4c766a497630
                alias           oradata_dss_13
        }
        multipath {
                wwid             3600a098038304173452b4c795961466b
                alias           oradata_dss_14
        }
        multipath {
                wwid             3600a098038304173475d4c766a497631
                alias           oradata_dss_15
        }
        multipath {
                wwid             3600a098038304173452b4c795961466c
                alias           oradata_dss_16
        }
        multipath {
                wwid             3600a098038304173475d4c766a497632
                alias           oraredo_dss_1
        }
        multipath {
                wwid             3600a098038304173452b4c795961466d
                alias           oraredo_dss_2
        }
}
```

Configuration of “/etc/sysctl.conf”

```
# oracle-database-server-12cR2-preinstall setting for fs.file-max is 6815744
fs.file-max = 6815744
# oracle-database-server-12cR2-preinstall setting for kernel.sem is '250 32000 100 128'
kernel.sem = 250 32000 100 128
# oracle-database-server-12cR2-preinstall setting for kernel.shmmni is 4096
kernel.shmmni = 4096
# oracle-database-server-12cR2-preinstall setting for kernel.shmall is 1073741824 on x86_64
kernel.shmall = 1073741824
# oracle-database-server-12cR2-preinstall setting for kernel.shmmax is 4398046511104 on x86_64
kernel.shmmax = 4398046511104
# oracle-database-server-12cR2-preinstall setting for kernel.panic_on_oops is 1 per Orabug 19212317
kernel.panic_on_oops = 1
# oracle-database-server-12cR2-preinstall setting for net.core.rmem_default is 262144
net.core.rmem_default = 262144
# oracle-database-server-12cR2-preinstall setting for net.core.rmem_max is 4194304
net.core.rmem_max = 4194304
# oracle-database-server-12cR2-preinstall setting for net.core.wmem_default is 262144
net.core.wmem_default = 262144
# oracle-database-server-12cR2-preinstall setting for net.core.wmem_max is 1048576
net.core.wmem_max = 1048576
# oracle-database-server-12cR2-preinstall setting for net.ipv4.conf.all.rp_filter is 2
net.ipv4.conf.all.rp_filter = 2
# oracle-database-server-12cR2-preinstall setting for net.ipv4.conf.default.rp_filter is 2
net.ipv4.conf.default.rp_filter = 2
# oracle-database-server-12cR2-preinstall setting for fs.aio-max-nr is 1048576
fs.aio-max-nr = 1048576
# oracle-database-server-12cR2-preinstall setting for net.ipv4.ip_local_port_range is 9000 65500
net.ipv4.ip_local_port_range = 9000 65500
vm.nr_hugepages=180000
```

Configuration of “/etc/security/limits.d/oracle-database-server-12cR2-preinstall.conf”

```
# oracle-database-server-12cR2-preinstall setting for nofile soft limit is 1024
oracle   soft   nofile    1024
# oracle-database-server-12cR2-preinstall setting for nofile hard limit is 65536
oracle   hard   nofile    65536
# oracle-database-server-12cR2-preinstall setting for nproc soft limit is 16384
# refer orabug15971421 for more info.
oracle   soft   nproc    16384
# oracle-database-server-12cR2-preinstall setting for nproc hard limit is 16384
oracle   hard   nproc    16384
# oracle-database-server-12cR2-preinstall setting for stack soft limit is 10240KB
oracle   soft   stack    10240
# oracle-database-server-12cR2-preinstall setting for stack hard limit is 32768KB
oracle   hard   stack    32768
# oracle-database-server-12cR2-preinstall setting for memlock hard limit is maximum of 128GB on x86_64 or 3GB on x86 OR 90 % of RAM
oracle   hard   memlock    473985890
# oracle-database-server-12cR2-preinstall setting for memlock soft limit is maximum of 128GB on x86_64 or 3GB on x86 OR 90% of RAM
oracle   soft   memlock    473985890
```

Configuration of “/etc/udev/rules.d/99-oracle-asmdevices.rules”

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

References

The following references were used in preparing this document.

Cisco Unified Computing System

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

Cisco UCS Data Center Design Guides

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

FlexPod Converged Infrastructure

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

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

Cisco UCS B200 M5 Servers

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

Oracle Database 12c Release 2

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

NetApp AFF A-Series All Flash Storage

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

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

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

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

NetApp Support

https://mysupport.netapp.com/

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

About the Authors

Tushar Patel, Principal Engineer, CSPG UCS Product Management and Data Center Solutions Engineering Group, Cisco Systems, Inc.

Tushar Patel is a Principal Engineer in the Cisco Systems CSPG UCS Product Management and Data Center Solutions Engineering Group and a specialist in flash storage technologies and Oracle RAC RDBMS. Tushar has over 23 years of experience in Flash Storage architecture, database architecture, design and performance. Tushar also has strong background in Intel X86 architecture, hyper converged systems, storage technologies and virtualization. He has worked with large number of enterprise customers, evaluate and deploy mission critical database solutions. Tushar has presented to both internal and external audiences at various conferences and customer events.

Hardikkumar Vyas, Solution Engineer, CSPG UCS Product Management and Data Center Solutions Engineering Group, Cisco Systems, Inc.

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

Acknowledgements

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

James Bradshaw, NetApp

Solution Brief

FlexPod®

Flexible Data Center Infrastructure Simplifies IT and Accelerates Applications

The Challenge

The need for agility and simplicity in deploying and managing infrastructure

Supporting business objectives from an IT perspective is becoming more complex and elusive, as users require on-demand resources that are secure, reliable, and easy to use. IT organizations need to adapt with simpler approaches to deploying and managing infrastructure. IT must enable greater agility by optimizing costs and existing staff resources, while also maintaining data security and adherence to compliance requirements. With growing businesses' requests such as expansion to new markets and support for mobile users—as well as opportunities in analytics and hybrid cloud—IT must anticipate new requests by modernizing the data center to become an agile, innovative service provider.

The Solution

The FlexPod converged infrastructure platform

Built on groundbreaking technology from NetApp and Cisco, the FlexPod® converged infrastructure platform meets and exceeds these challenges. FlexPod is trusted by thousands of customers across the globe. Composed of prevalidated storage, networking, and server technologies, FlexPod is designed to increase IT responsiveness to business needs and reduce the cost of computing with maximum uptime and minimal risk.

The prevalidated FlexPod architecture delivers unmatched application performance that drives higher productivity, faster decision making, and greater opportunities for growth. FlexPod simplifies and modernizes IT with continuous innovation, broad support for any cloud strategy, and improved operational efficiency to accelerate data center transformation and business evolution.

Key Features

Optimize your applications

Deliver increased application performance that drives greater productivity, speeds decision making, and provides opportunities for growth.

Embrace hybrid cloud

Transcend data center constraints with a secure and scalable foundation for cloud and as-a-service initiatives.

Simplify your IT

Reduce complexities, making room for better IT services, increased productivity, and continued innovation across the enterprise.

Solution Differentiators

- Flexible design with a broad range of reference architectures for popular business applications
- Elimination of costly, disruptive downtime with clustered Data ONTAP® nondisruptive operations
- Cisco ACI centralized, policy-driven automation that accelerates application deployments
- Multiprotocol storage platform that unifies application silos, allowing NAS or SAN, file or block, on one converged platform
- Support for private, public, or hybrid cloud strategies with a consistent set of data management tools from flash to disk to cloud

Application Optimization Starts with FlexPod

- Host multiple instances of mixed applications, consolidated on a shared infrastructure with centralized, simplified management
- Speed database application performance by up to 20x with All Flash FAS
- Scale out and scale up as workloads increase, adding storage and compute layers as your business grows

Hybrid Cloud Flexibility

- Operate across hybrid cloud resources with the software-defined capabilities of NetApp® Data Fabric, while maintaining security, control and workload portability with Cisco Intercloud Fabric
- Manage data from flash to disk to cloud with the simplicity of a single set of tools
- Optionally use OpenStack software on FlexPod to create a private or hybrid cloud

Simplified IT Infrastructure

- Use Cisco Unified Computing System (Cisco UCS) Director for end-to-end, single-view automation and orchestration, freeing IT staff to focus on new services
- Validated designs help you deploy FlexPod platforms in a wide range of operating environments with less risk and accelerated ROI
- Cooperative support is designed to simplify and streamline support for your FlexPod converged infrastructure

The FlexPod Family

FlexPod is offered in three solution categories that are designed to meet your specific capacity and performance requirements:

> *"Among the biggest benefits of FlexPod are integrated components that help enable us to centrally manage all our data center requirements."*
>
> **Darrell Williams**
> **Director of Information Systems, Katz, Sapper & Miller**

- **FlexPod Express** is ideal for midsized businesses and branch offices. It can be used as a cost-effective starting point for infrastructure consolidation and virtualization solutions.
- **FlexPod Datacenter** is suited for large enterprises and cloud service providers that have mature IT processes and rapid growth expectations and want to deploy a highly scalable shared infrastructure for multiple business-critical applications.
- **FlexPod Select** supports high-performance computing or very large data capacity environments such as big data analytics, scientific computing, and dedicated application optimization.

Any of these FlexPod solutions can be scaled up or out and duplicated in modular fashion to accommodate your future growth. They can also scale to a larger FlexPod configuration with a clearly defined upgrade path that leverages all existing components and management processes.

Proven Across a Broad Range of Environments

FlexPod has been pretested and jointly validated with popular hypervisors, operating systems, applications, and infrastructure software, including:

- VMware vSphere
- VMware View
- Citrix XenDesktop
- Red Hat Enterprise Linux
- Red Hat Enterprise Linux OpenStack Platform
- Oracle (RAC, JD Edwards, Oracle Linux, Oracle VM Server)
- SAP
- Microsoft Exchange, SQL Server, and SharePoint
- Microsoft Private Cloud
- Hortonworks Data Platform
- Cloudera's Distribution, including Apache Hadoop
- NetApp SnapProtect® technology
- Cisco Nexus data center switches

Reference Architectures

NetApp and Cisco have jointly developed numerous reference architectures to help you integrate and flex the solution to meet your specific requirements for the following critical environments:

Workload consolidation

FlexPod helps you consolidate and virtualize your business applications onto less hardware. Along with improved hardware utilization, this approach frees up data center space and reduces power and cooling requirements, enabling you to slash your infrastructure costs by up to 50%.

Virtual desktop infrastructure

FlexPod is a self-contained virtual desktop solution in a rack. Its modular design facilitates rapid, repeatable deployment of thousands of virtual desktops. FlexPod is optimized for VMware View and Citrix XenDesktop. You gain efficiencies by deduplicating up to 90% of redundant user and OS

> *““Our executives loved the simplicity and power of the integrated stack in FlexPod. And for IT, the prevalidated architecture with prescriptive sizing and design guides reduced our risk.”*
>
> **Wojciech Biernacki**
> **IT Systems Administrator, University of Tennessee**

data, and I/O performance can be accelerated by up to 50% with the NetApp Virtual Storage Tier. The extended memory technology of Cisco Unified Computing System provides the industry’s greatest number of VMs per core density.

Development and test

FlexPod enables rapid provisioning and deprovisioning of virtual resources, making it ideal for development and test environments. NetApp FlexClone® software facilitates rapid development/test setup with cloning technology that lets you deploy thousands of space-efficient VMs for new projects in minutes, accelerating time to market. Clones can also be redeployed to secondary sites, reducing preparation time for initiatives such as disaster recovery testing.

Business and disaster recovery

FlexPod can be configured with integrated data protection software to provide fast recovery from system, site, and regional outages for business continuity. The combination of NetApp MetroCluster™ and SnapMirror® technologies with Cisco UCS Manager offers automated monitoring and failover, as well as cost-effective replication to a secondary site for continuous protection against unplanned downtime. FlexPod lets you move virtual server and storage resources and data nondisruptively across hardware to eliminate planned downtime.

Secure multi-tenancy and secure separation

FlexPod leverages Cisco and NetApp technologies to deliver secure multi-tenancy with solutions such as Cisco Secure Enclaves. Resources and data for each tenant—application, business unit, or customer—are securely isolated within the FlexPod environment. This combines the data separation and service-level guarantees offered by application silos with the efficiencies of a converged, virtualized infrastructure.

Best-in-Class Components for Enhanced Data Center Efficiency

FlexPod components are integrated in a standardized configuration that scales from entry-level designs for hundreds of users up to high-performance big data workloads for thousands of users. This integrated approach can significantly reduce your capital and operating expenses through end-to-end virtualization and higher efficiencies at each layer.

Cisco Unified Computing System

Cisco UCS is a data center platform that eliminates time-consuming manual configuration, reduces TCO, and increases business agility. It combines compute and network resources, storage access, and virtualization into a scalable, modular system that is easily managed as a single entity by Cisco UCS Manager. Service profile templates enable automatic, policy-based hardware configuration and deployment for large, stateless computing environments. Highly efficient Cisco UCS extended memory technology reduces memory requirements by up to 60%.

Cisco Nexus data center switches

Cisco Nexus switches use award-winning unified fabric technology to identify and consolidate all network traffic onto a single simplified, cost-effective architecture based on Fibre Channel over Ethernet. The switches offer “zero-touch” installation, automatic configuration, enterprise-class scalability, and nondisruptive in-service upgrades. A single point of policy management also increases efficiency, availability, and security.

The option of Cisco Nexus 7000 Series switches provides even greater networking scale, throughput, availability, and advanced features for data center interconnect requirements. Cisco Nexus 9000 switches lay the foundation for software-defined innovations such as Cisco Application Centric Infrastructure, allowing intelligent software to automate hardware resources across next-generation data centers.

NetApp FAS storage

NetApp FAS storage systems reduce the cost and complexity for virtualized infrastructures by meeting all of your storage requirements with a single, highly scalable solution. NetApp’s unified storage platform supports all protocols, so you no longer need to purchase separate systems to accommodate different storage needs. You can slash capacity use by up to 50% with built-in deduplication, thin provisioning, and space-efficient backup and cloning.

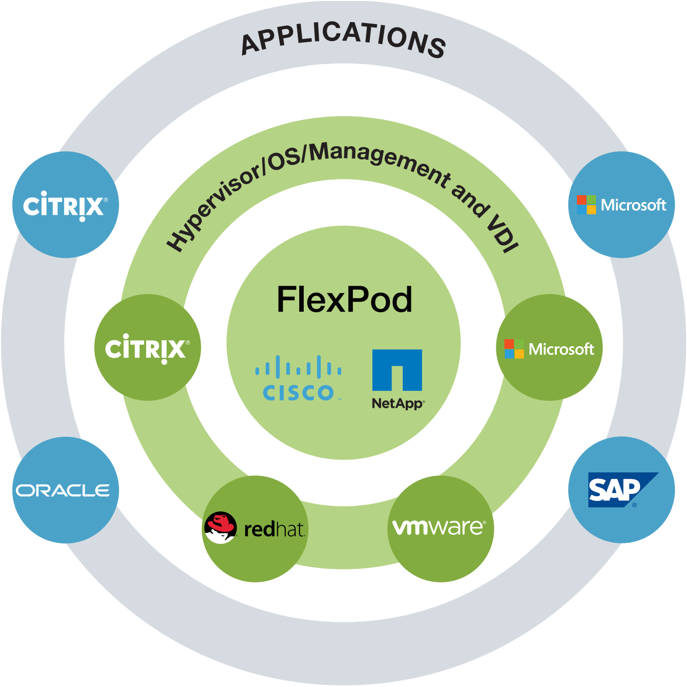

Figure 1) FlexPod Cooperative Support model: an ecosystem of a multivendor engagement.

NetApp systems enhance operational efficiency with automated storage management, data protection, and security. The clustered Data ONTAP operating system brings a new level of nondisruptive operations, scalability, and efficiency to enterprise storage. Performance is optimized with innovative flash technologies and 10GbE and FCoE support. With NetApp storage, you can deploy the exact proportion of flash to spinning media for your particular environment. And for extreme performance for dedicated workloads, NetApp All Flash FAS systems increase database application performance by up to 20x.

Cooperative Support Speeds Problem Resolution

FlexPod Cooperative Support is a partnership between NetApp; Cisco; and our technology partners Microsoft, VMware, Citrix, and Red Hat. Your IT staff chooses which vendor to call based on your initial assessment of the problem’s origin. Multivendor engineers work to resolve your issue quickly using shared communications, expertise gained through ongoing joint training, and a formal escalation process. The result is a rapid, coordinated resolution to your technical issues.

Choice of Management Solutions

NetApp and Cisco work with trusted partners to provide you with a choice of management solutions. The FlexPod architecture provides APIs at each layer so that it can easily integrate with a broad range of software solutions for end-to-end management. Validated FlexPod management solutions have been tested to verify that they deliver essential functionality. Together with partners, we provide a variety of capabilities, including automation and orchestration, monitoring and analytics, and configuration management.

Global Service Delivery Ecosystem

You can choose from a global network of FlexPod Premium Partners and other highly qualified solution delivery partners to implement your FlexPod solution. These partners understand your business requirements and are certified on NetApp, Cisco, and complementary technologies to deliver a complete cloud solution that fits your business needs.

Getting Started

To learn how the FlexPod platform enables you to build a flexible and efficient converged virtualized infrastructure to modernize your data center, contact your local NetApp or Cisco representative or partner. Learn more at www.netapp.com/flexpod.

FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1

Updated: August 26, 2016

FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect and VMware vSphere 6.0 U1

Deployment Guide for FlexPod Datacenter with Cisco UCS Manager 3.1 and VMware vSphere 6.0 U1

Last Updated: August 26, 2016

About Cisco Validated Designs

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

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

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

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

Table of Contents

Executive Summary

Cisco Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

This document describes the Cisco and NetApp® FlexPod Datacenter with Cisco UCS Manager unified software release 3.1 and VMware vSphere 6.0 U1. Cisco UCS Manager (UCSM) 3.1 provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with Cisco UCS unified software release and VMware vSphere 6.0 U1 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, and NetApp AFF.

Solution Overview

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod, which uses best of breed storage, server and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides a step by step configuration and implementation guide for the FlexPod Datacenter with Cisco UCS 6300 Fabric Interconnect, NetApp AFF, and Cisco Nexus 9000 solution. For the design decisions and technology discussion of the solution, please refer to FlexPod Datacenter with Cisco Unified Software Release and VMware vSphere 6 Design Guide:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi60_n9k_design.html

What’s New?

The following design elements distinguish this version of FlexPod from previous FlexPod models:

- Validation of Cisco UCS 6300 Fabric Interconnects
- Support for the Cisco UCS 3.1(1h) unified software release, Cisco UCS B200-M4 servers, and Cisco UCS C220-M4 servers
- Support for the latest release of NetApp Data ONTAP® 8.3.2
- An IP-based storage design, supplemented with direct attached fibre channel connectivity, supporting both NAS datastores, and FC and iSCSI based SAN LUNs
- Validation of VMware vSphere 6.0 U1b
- HTML-based Cisco UCS Manager

Solution Design

Architecture

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere® built on FlexPod includes NetApp All Flash FAS storage, Cisco Nexus® networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or " flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

Figure 1 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with IP-based storage. This design uses the Cisco UCS Fabric Interconnect, Cisco Nexus 9000, and Cisco UCS C-Series and B-Series servers and the NetApp AFF family of storage controllers connected in a highly available modular design. This infrastructure is deployed to provide either FC or iSCSI-booted hosts with file-level and block-level access to shared storage. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Physical Topology

Figure 1 illustrates the physical architectures.

Figure 1 FlexPod Design with Cisco Nexus 9000 and NetApp Data ONTAP

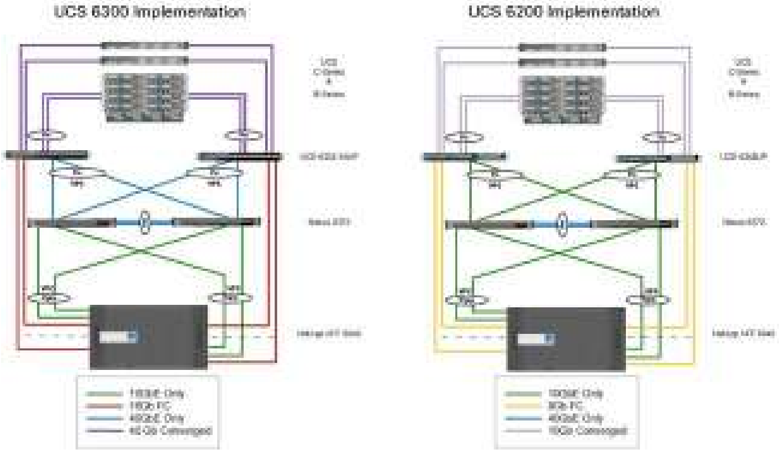

The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches
- Two Cisco UCS 6332-16UP or Two Cisco UCS 6248UP fabric interconnects
- One NetApp AFF8040 (HA pair) running clustered Data ONTAP with Disk shelves and Solid State Drives (SSD)

For server virtualization, the deployment includes VMware vSphere 6.0 U1. Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps for deploying the base architecture, as shown in Figure 1. These procedures cover everything from physical cabling to network, compute and storage device configurations.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6300 Series, UCS B-200 M4, UCS C-220 M4	3.1(1h)	Includes the Cisco UCS-IOM 2304, Cisco UCS Manager, Cisco UCS VIC 1340 and Cisco UCS VIC 1385
	Cisco eNIC	2.3.0.7	
	Cisco fNIC	1.6.0.25	
Network	Cisco Nexus 9000 NX-OS	7.0(3)I1(3)	
	Cisco Nexus 1000V	5.2(1)SV3 (1.5b)	
	Cisco Nexus 1110-X	5.2(1)SP1(7.3)	
Storage	NetApp AFF 8040	Data ONTAP 8.3.2	
Software	Cisco UCS Manager	3.1(1h)	
	Cisco UCS Performance Manager	2.0	
	VMware vSphere ESXi	6.0 U1b	
	VMware vCenter	6.0 U1b	
	NetApp Virtual Storage Console (VSC)	6.2	
	OnCommand Performance Manager	2.0	

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-01, VM-Host-Prod-02 to represent infrastructure and production hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?
 [-node] <nodename>                Node
 { [-vlan-name] {<netport>|<ifgrp>}  VLAN Name
 |  -port {<netport>|<ifgrp>}       Associated Network Port
 [-vlan-id] <integer> }            Network Switch VLAN Identifier
```

Example:

```
network port vlan -node <node01> -vlan-name i0a-<vlan id>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 3 lists the virtual machines (VMs) necessary for deployment as outlined in this guide. Table 2 describes the VLANs necessary for deployment as outlined in this guide.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out of Band Mgmt	VLAN for out-of-band management interfaces	13
In-Band Mgmt	VLAN for in-band management interfaces	113
Native	VLAN to which untagged frames are assigned	2
NFS	VLAN for Infrastructure NFS traffic	3170
vMotion	VLAN for VMware vMotion	3173
VM-Traffic	VLAN for Production VM Interfaces	3174
iSCSI-A	VLAN for Fabric A iSCSI	901
iSCSI-B	VLAN for Fabric B iSCSI	902
Packet-Ctrl	VLAN Nexus 1110-X Packet and Control	3176

Table 3 lists the VMs necessary for deployment as outlined in this document.

Table 3 Virtual Machines

Virtual Machine Description	Host Name
Active Directory	
vCenter Server	
NetApp VSC	
NetApp OnCommand Unified Manager	
OnCommand Performance Manager	

Table 4 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps.

Table 4 Configuration Variables

Variable	Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01 (Example: 192.168.156.21)
<<var_node01_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_url_boot_software>>	Data ONTAP 8.3.2 URL (Example: http://192.168.156.9/832_q_image.tgz)
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02 (Example: 192.168.156.22)
<<var_node02_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_clustername>>	Storage cluster host name (Example: clus)
<<var_cluster_base_license_key>>	Cluster base license key (Example: 1234567890ABCD1234567890ABCD)
<<var_nfs_license>>	NFS license key (Example: 1234567890ABCD1234567890ABCD)
<<var_fc_license>>	Fiber Channel license key (if using Fiber Channel) (Example: 1234567890ABCD1234567890ABCD)
<<var_iscsi_license>>	iSCSI license key (if using iSCSI) (Example: 1234567890ABCD1234567890ABCD)
<<var_password>>	Global default administrative password (Example: Fl3xP0d9)

Variable	Value
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster (Example: 192.168.157.20)
<<var_clustermgmt_mask>>	In-band management network netmask (Example: 255.255.255.0)
<<var_clustermgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.157.1)
<<var_dns_domain_name>>	DNS domain name (Example: flexpod.com)
<<var_nameserver_ip>>	DNS server IP(s) (Example: 192.168.156.9)
<<var_node_location>>	Node location string for each node (Example: RTP9-D04)
<<var_node01_sp_ip>>	Out-of-band cluster node 01 service processor management IP (Example: 192.168.156.18)
<<var_node01_sp_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node01_sp_gateway>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP (Example: 192.168.156.19)
<<var_node02_sp_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_node02_sp_gateway>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_node01>>	Cluster node 01 hostname (Example: clus-01)
<<var_node02>>	Cluster node 02 hostname (Example: clus-02)
<<var_num_disks>>	Number of disks to assign to each storage controller (Example: 5)
<<var_nfs_vlan_id>>	Infrastructure NFS VLAN ID for LIF (Example: 3170)
<<var_iscsi_vlan_A_id>>	Infrastructure iSCSI-A VLAN ID for LIF (Example: 901)
<<var_iscsi_vlan_B_id>>	Infrastructure iSCSI-B VLAN ID for LIF (Example: 902)
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID (Example: 113)
<<var_oob_mgmt_vlan_id>>	Out-of-band management network VLAN ID (Example: 13)
<<var_timezone>>	FlexPod time zone (Example: America/New_York)
<<var_global_ntp_server_ip>>	NTP server IP address for out-of-band mgmt. (Example: 192.168.156.1)
<<var_switch_a_ntp_ip>>	NTP server IP address for Nexus 9372 Switch A (Example: 192.168.156.1)
<<var_switch_b_ntp_ip>>	NTP server IP address for Nexus 9372 Switch B (Example: 192.168.156.1)
<<var_ib-mgmt_vlan_netmask_length>>	Length of IB-MGMT-VLAN Netmask (Example: /24)
<<var_snmp_contact>>	Administrator e-mail address (Example: admin@flexpod.com)
<<var_snmp_location>>	Cluster location string (Example: RTP9-D04)
<<var_cert_common_name>>	Common name string for certificate (Example: “clus.flexpod.com”)
<<var_cert_country>>	Country for certificate (Example: “USA”)
<<var_cert_state>>	State for certificate (Example: “NC”)
<<var_cert_locality>>	Locality for certificate (Example: “RTP”)
<<var_cert_org>>	Organization for certificate (Example: “FlexPod”)
<<var_cert_unit>>	Organizational Unit for certificate (Example: “Dev”)
<<var_cert_email>>	E-mail address for certificate (Example: “admin@flexpod.com”)
<<var_cert_days>>	Days until certificate expiration (Example: 365)
<<var_oncommand_server_fqdn>>	VSC or OnCommand VM fully qualified domain name (FQDN) (Example: ocum.flexpod.com)
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name (Example: fl3xp0d)
<<var_mailhost>>	Mail server host name (Example: smtp.flexpod.com)
<<var_storage_admin_email>>	Administrator e-mail address (Example: storage@flexpod.com)
<<var_node01_nfs_lif_infra_swap_ip>>	IP address of Infra Swap (Example: 192.168.170.21)
<<var_node01_nfs_lif_infra_swap_mask>>	Subnet Mask of Infra Swap (Example: 255.255.255.0)
<<var_node02_nfs_lif_infra_datastore_1_ip>>	IP address of Datastore 1 (Example: 192.168.170.22)
<<var_node02_nfs_lif_infra_datastore_1_mask>>	Subnet mask of Datastore 1 (Example: 255.255.255.0)
<<var_vserver_mgmt_ip>>	Management IP address for Vserver (Example: 192.168.156.23)
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver (Example: 255.255.255.0)

Variable	Value
<<var_vsadmin_password>>	Password for VS admin account (Example: Fl3xP0d)
<<var_ucs_6248_clustername>>	Cisco UCS Manager cluster host name (Example: ucs-6248)
<<var_ucs_6332_clustername>>	Cisco UCS Manager cluster host name (Example: ucs-6332)
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address (Example: 192.168.156.51
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway (Example: 192.168.156.1)
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address (Example: 192.168.156.52)
<<var_vm_host_infra_01_iqn>>	Cisco UCS Service Profile generated IQN of Infra 02 (Example: iqn.1992-08.com.cisco:ucs-6248-host:1
<<var_vm_host_prod_02_iqn>>	Cisco UCS Service Profile generated IQN of Infra 01 (Example: iqn.1992-08.com.cisco:ucs-6332-host:1
<<var_vm_host_infra_01_ip>>	VMware ESXi host 02 out-of-band management IP (Example: 10.1.156.25)
<<var_vm_host_prod_02_ip>>	VMware ESXi host 01 out-of-band management IP (Example: 10.1.156.28)
<<var_nfs_vlan_ip_host_01>>	ESXi host 1, NFS VLAN IP (Example: 192.168.170.25)
<<var_nfs_vlan_ip_mask_host_01>>	ESXi host1, NFS VLAN subnet mask (Example: 255.255.255.0)
<<var_nfs_vlan_ip_host_02>>	ESXi host 2, NFS VLAN IP (Example: 192.168.170.28)
<<var_nfs_vlan_ip_mask_host_02>>	ESXi host2, NFS VLAN subnet mask (Example: 255.255.255.0)
<<var_vcenter_server_ip>>	IP address of the vCenter Server (Example: 10.1.156.100)
<<var_svm_mgmt_vlan_id>>	Infrastructure Vserver management VLAN ID (Example: 13)
<<var_node01_iscsi_lif01a_ip>>	iSCSI LIF 01a IP address (Example: 192.168.91.21)
<<var_node01_iscsi_lif01a_mask>>	iSCSI LIF 01a subnet mask (Example: 255.255.255.0)
<<var_node01_iscsi_lif01b_ip>>	iSCSI LIF 01b IP address (Example: 192.168.92.21)
<<var_node01_iscsi_lif01b_mask>>	iSCSI LIF 01b subnet mask (Example: 255.255.255.0)
<<var_node01_iscsi_lif02a_ip>>	iSCSI LIF 02a IP address (Example: 192.168.91.22)
<<var_node01_iscsi_lif02a_mask>>	iSCSI LIF 02a subnet mask (Example: 255.255.255.0)
<<var_node01_iscsi_lif02b_ip>>	iSCSI LIF 02b IP address (Example: 192.168.92.22)
<<var_node01_iscsi_lif02b_mask>>	iSCSI LIF 02b subnet mask (Example: 255.255.255.0)
<<var_node01_fcp_p0rt1 >>	Node 1 FC port 1 (Example: 0a)
<<var_node01_fcp_p0rt2 >>	Node 1 FC port 2 (Example: 0b)
<<var_node02_fcp_p0rt1 >>	Node 2 FC port 1 (Example: 0a)
<<var_node02_fcp_p0rt2 >>	Node 2 FC port 2 (Example: 0b)
<<var_vserver_mgmt_ip>>	Management IP address for Infrastructure Vserver (Example: 192.168.156.23)
<<var_vserver_mgmt_mask>>	Management subnet mask for Infrastructure Vserver (Example: 255.255.255.0)
<<var_oncommand_server_ip>>	IP address of the OnCommand Unified Manager (Example: 10.1.156.10)
<<var_rule_index>>	Rule index number (Example: 1)

Physical Infrastructure

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain the details for the prescribed and supported configuration of the NetApp AFF8040 running clustered Data ONTAP 8.3.2.

For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps

Be sure to follow the cabling directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.

Figure 2 shows a cabling diagram for a FlexPod configuration using the Cisco Nexus 9000 and NetApp storage systems with clustered Data ONTAP connected to the UCS 6332-16UP Fabric Interconnect, and Figure 3 shows that same configuration utilizing the UCS 6248UP Fabric Interconnect. Cabling of both the 6332-16UP and the 6248UP in adjacency are done as an example of interchangeable viability between the two fabric interconnect models, and not intended to imply a requirement of simultaneous deployment.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 2 FlexPod Cabling Diagram (6332-16UP)

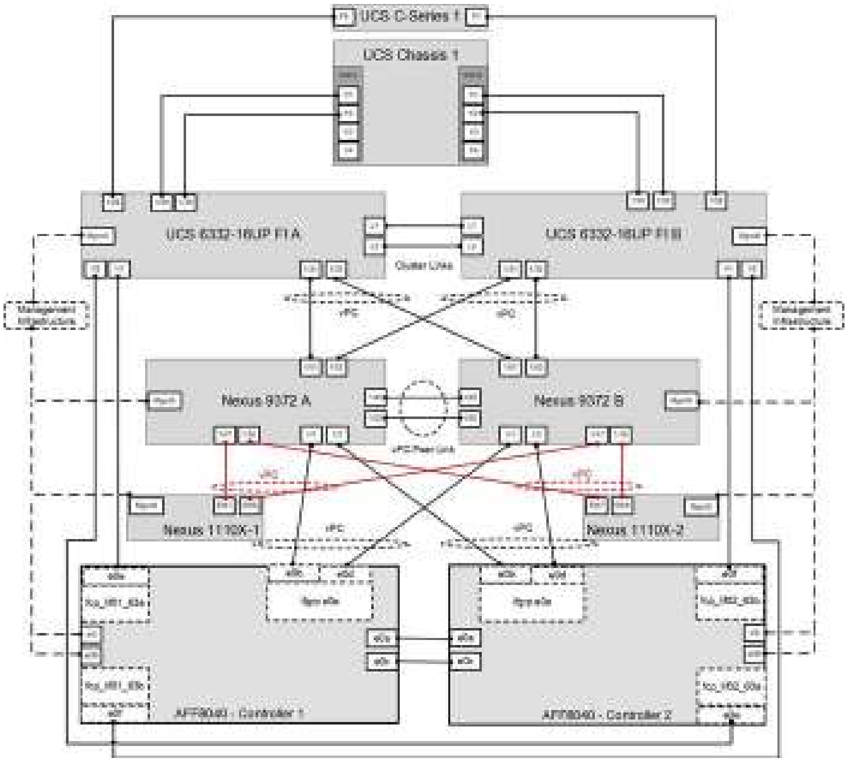

Figure 3 FlexPod Cabling Diagram (6248UP)

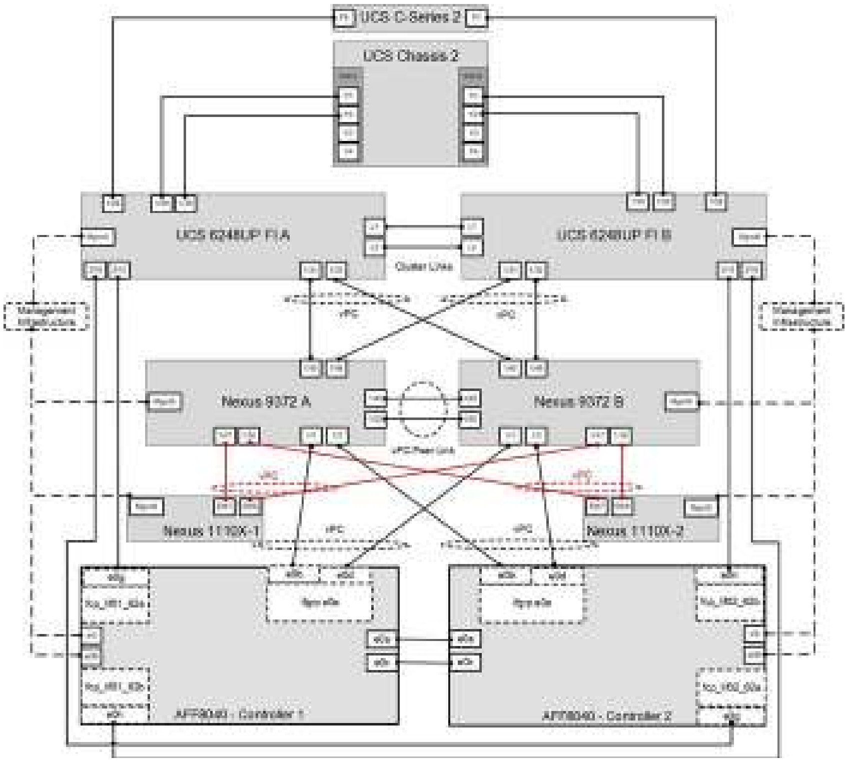

Table 5 through Table 13 provide the details of all the connections in use.

Table 5 Cisco Nexus 9372-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 A	Eth1/1	10GbE	NetApp Controller 1	e0b
	Eth1/2	10GbE	NetApp Controller 2	e0b
	Eth1/17	10GbE	Nexus 1110-X 1	Eth7
	Eth1/18	10GbE	Nexus 1110-X 2	Eth7
	Eth1/45	10GbE	Cisco UCS 6248UP FI A	Eth1/31
	Eth1/46	10GbE	Cisco UCS 6248UP FI B	Eth1/31
	Eth1/49	40GbE	Cisco Nexus 9372 B	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 B	Eth1/50
	Eth1/51	40GbE	Cisco UCS 6332-16UP FI A	Eth1/31
	Eth1/52	40GbE	Cisco UCS 6332-16UP FI B	Eth1/31
	MGMT0	GbE	GbE management switch	Any

For devices requiring GbE connectivity, use the GbE Copper SFP+s (GLC-T=).

Table 6 Cisco Nexus 9372-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 B	Eth1/1	10GbE	NetApp Controller 1	e0d
	Eth1/2	10GbE	NetApp Controller 2	e0d
	Eth1/17	10GbE	Nexus 1110-X 1	Eth8
	Eth1/18	10GbE	Nexus 1110-X 2	Eth8
	Eth1/45	10GbE	Cisco UCS 6248UP FI A	Eth1/32
	Eth1/46	10GbE	Cisco UCS 6248UP FI B	Eth1/32
	Eth1/49	40GbE	Cisco Nexus 9372 A	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 A	Eth1/50
	Eth1/51	40GbE	Cisco UCS 6332-16UP FI A	Eth1/32

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/52	40GbE	Cisco UCS 6332-16UP FI B	Eth1/32
	MGMT0	GbE	GbE management switch	Any

Table 7 NetApp Controller-1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 2	e0a
	e0b	10GbE	Cisco Nexus 9372 A	Eth1/1
	e0c	10GbE	NetApp Controller 2	e0c
	e0d	10GbE	Cisco Nexus 9372 B	Eth1/1
	0e	16Gb FC	Cisco UCS 6332-16UP FI A (only if using FC connectivity)	FC 1/1
	0f	16Gb FC	Cisco UCS 6332-16UP FI B (only if using FC connectivity)	FC 1/1
	0g	16Gb FC	Cisco UCS 6248UP FI A (only if using FC connectivity)	FC 2/15
	0h	16Gb FC	Cisco UCS 6248UP FI B (only if using FC connectivity)	FC 2/15

When the term e0M is used, the physical Ethernet port to which the table is referring is the port indicated by a wrench icon on the rear of the chassis.

Table 8 NetApp Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2	e0M	100MbE	100MbE management switch	Any
	e0i	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 1	e0a
	e0b	10GbE	Cisco Nexus 9372 A	Eth1/2
	e0c	10GbE	NetApp Controller 1	e0c
	e0d	10GbE	Cisco Nexus 9372 B	Eth1/2
	0e	16Gb FC	Cisco UCS 6332-16UP FI A (only if using FC connectivity)	FC 1/2
	0f	16Gb FC	Cisco UCS 6332-16UP FI B (only if using FC connectivity)	FC 1/2
	0g	16Gb FC	Cisco UCS 6248UP FI A (only if using FC connectivity)	FC 2/16
	0h	16Gb FC	Cisco UCS 6248UP FI B (only if using FC connectivity)	FC 2/16

Table 9 Cisco UCS 6332-16UP Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	FC 1/1	16Gb FC	NetApp controller 1 (only if using FC)	0e
	FC 1/2	16Gb FC	NetApp controller 2 (only if using FC)	0e
	Eth1/28	40GbE	Cisco UCS C-Series 1	Port 0
	Eth1/29	40GbE	Cisco UCS Chassis 1 2304 FEX A	IOM 1/1
	Eth1/30	40GbE	Cisco UCS Chassis 1 2304 FEX A	IOM 1/2
	Eth1/31	40GbE	Cisco Nexus 9372 A	Eth1/51
	Eth1/32	40GbE	Cisco Nexus 9372 B	Eth1/51
	MGMT0	GbE	GbE management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	L1	GbE	Cisco UCS 6332-16UP FI B	L1
	L2	GbE	Cisco UCS 6332-16UP FI B	L2

Table 10 Cisco UCS 6332-16UP Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	FC 1/1	16Gb FC	NetApp controller 1 (only if using FC)	0f
	FC 1/2	16Gb FC	NetApp controller 2 (only if using FC)	0f
	Eth1/28	40GbE	Cisco UCS C-Series 1	Port 1
	Eth1/29	40GbE	Cisco UCS Chassis 1 2304 FEX B	IOM 1/1
	Eth1/30	40GbE	Cisco UCS Chassis 1 2304 FEX B	IOM 1/2
	Eth1/31	40GbE	Cisco Nexus 9372 A	Eth1/52
	Eth1/32	40GbE	Cisco Nexus 9372 B	Eth1/52
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6332-16UP FI B	L1
	L2	GbE	Cisco UCS 6332-16UP FI B	L2

Table 11 Cisco UCS 6248UP Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	FC 2/15	8Gb FC	NetApp controller 1 (only if using FC)	0g
	FC 2/16	8Gb FC	NetApp controller 2 (only if using FC)	0g
	Eth1/28	10GbE	Cisco UCS C-Series 2	Port 0
	Eth1/29	10GbE	Cisco UCS Chassis 2 2204 FEX A	IOM 1/1
	Eth1/30	10GbE	Cisco UCS Chassis 2 2204 FEX A	IOM 1/2
	Eth1/31	10GbE	Cisco Nexus 9372 A	Eth1/45
	Eth1/32	10GbE	Cisco Nexus 9372 B	Eth1/45
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6248UP FI B	L1
	L2	GbE	Cisco UCS 6248UP FI B	L2

Table 12 Cisco UCS 6248UP Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	FC 2/15	8Gb FC	NetApp controller 1 (only if using FC)	0f
	FC 2/16	8Gb FC	NetApp controller 2 (only if using FC)	0f
	Eth1/28	10GbE	Cisco UCS C-Series 2	Port 1
	Eth1/29	10GbE	Cisco UCS Chassis 2204 FEX B	IOM 1/1
	Eth1/30	10GbE	Cisco UCS Chassis 2204 FEX B	IOM 1/2
	Eth1/31	10GbE	Cisco Nexus 9372 A	Eth1/46
	Eth1/32	10GbE	Cisco Nexus 9372 B	Eth1/46
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6248UP FI B	L1
	L2	GbE	Cisco UCS 6248UP FI B	L2

Table 13 Cisco UCS C-Series 1

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 1	Port 0	40GbE	Cisco UCS 6332-16UP FI A	Eth1/28
	Port 1	40GbE	Cisco UCS 6332-16UP FI B	Eth1/28

Table 14 Cisco UCS C-Series 2

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 2	Port 0	10GbE	Cisco UCS 6248UP FI A	Eth1/28
	Port 1	10GbE	Cisco UCS 6248UP FI B	Eth1/28

Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as covered in the section "FlexPod Cabling."

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Nexus 9000 7.0(3)I1(3).

The following procedure includes setup of NTP distribution on the In-Band Management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes the default VRF will be used to route the In-Band Management VLAN.

Set Up Initial Configuration

Cisco Nexus 9372PX A

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

1. Configure the switch.

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

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_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: <<var_nexus_A_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_A_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_A_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <<var_nexus_A_mgmt0_gw>>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

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

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus 9372PX B

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

1. Configure the switch.

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

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin": <<var_password>>
Confirm the password for "admin": <<var_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: <<var_nexus_B_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <<var_nexus_B_mgmt0_ip>>
Mgmt0 IPv4 netmask: <<var_nexus_B_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <<var_nexus_B_mgmt0_gw>>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

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

2. Review the configuration summary before enabling the configuration.

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

FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

1. Log in as admin.
2. Run the following commands:

```
config t
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

Set Global Configurations

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To set global configurations, complete the following step on both switches:

Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <<var_global_ntp_server_ip>> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <<var_ib-mgmt-vlan_gateway>>
copy run start
```

Create VLANs

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
exit
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
vlan <<var_vmotion_vlan_id>>
name vMotion-VLAN
exit
vlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
exit
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
exit
vlan <<var_iscsi-a_vlan_id>>
name iSCSI-A-VLAN
exit
vlan <<var_iscsi-b_vlan_id>>
name iSCSI-B-VLAN
exit
vlan <<var_packet-ctrl_vlan_id>>
name Packet-Ctrl-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus 9372PX A

From the global configuration mode, run the following commands:

```
ntp source <<var_switch_a_ntp_ip>>
interface Vlan<<var_ib-mgmt_vlan_id>>
ip address <<var_switch_a_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>
no shutdown
exit
```

Cisco Nexus 9372PX B

From the global configuration mode, run the following commands:

```
ntp source <<var_switch_b_ntp_ip>>
interface Vlan<<var_ib-mgmt_vlan_id>>
ip address <<var_switch_b_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>
no shutdown
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus 9372PX A

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

> ✍ In this step and in further sections, configure the <<var_ucs_6248_clustername>> and <<var_ucs_6332_clustername>> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_node01>>:e0b
exit
interface Eth1/2
description <<var_node02>>:e0b
exit
interface Eth1/17
description <<var_n1110-x>>-1:eth 7
exit
interface Eth1/18
description <<var_n1110-x>>-2:eth7
exit
interface Eth1/45
description <<var_ucs_6248_clustername>>-a:1/31
exit
interface Eth1/46
description <<var_ucs_6248_clustername>>-b:1/31
exit
```

```
interface Eth1/49
description <<var_nexus_B_hostname>>:1/49
exit
interface Eth1/50
description <<var_nexus_B_hostname>>:1/50
exit
interface Eth1/51
description <<var_ucs_6332_clustername>>-a:1/31
exit
interface Eth1/52
description <<var_ucs_6332_clustername>>-b:1/31
exit
```

Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
interface Eth1/1
description <<var_node01>>:e0d
exit
interface Eth1/2
description <<var_node02>>:e0d
exit
interface Eth1/17
description <<var_n1110-x>>-1:eth 8
exit
interface Eth1/18
description <<var_n1110-x>>-2:eth 8
exit
interface Eth1/45
description <<var_ucs_6248_clustername>>-a:1/32
exit
interface Eth1/46
description <<var_ucs_6248_clustername>>-b:1/32
exit
interface Eth1/49
description <<var_nexus_A_hostname>>:1/49
exit
interface Eth1/50
description <<var_nexus_A_hostname>>:1/50
exit
interface Eth1/51
description <<var_ucs_6332_clustername>>-a:1/32
exit
interface Eth1/52
description <<var_ucs_6332_clustername>>-b:1/32
exit
```

Create Port Channels

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

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

From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
exit
interface Eth1/49-50
channel-group 10 mode active
no shutdown
exit
interface Po11
description <<var_node01>>
exit
interface Eth1/1
channel-group 11 mode active
no shutdown
exit
interface Po12
description <<var_node02>>
exit
interface Eth1/2
channel-group 12 mode active
no shutdown
exit
interface Po117
description <<var_n1110-x>>-1
exit
interface Eth1/17
channel-group 117 mode active
no shutdown
exit
interface Po118
description <<var_n1110-x>>-2
exit
interface Eth1/18
channel-group 118 mode active
no shutdown
exit
copy run start
interface Po145
description <<var_ucs_6248_clustername>>-a
exit
interface Eth1/45
channel-group 145 mode active
no shutdown
exit
interface Po112
description <<var_ucs_6248_clustername>>-b
exit
interface Eth1/46
channel-group 146 mode active
no shutdown
exit
interface Po151
description <<var_ucs_6332_clustername>>-a
exit
interface Eth1/51
channel-group 151 mode active
no shutdown
exit
interface Po152
description <<var_ucs_6332_clustername>>-b
exit
interface Eth1/52
channel-group 152 mode active
no shutdown
exit
```

Configure Port Channel Parameters

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To configure port channel parameters, complete the following step on both switches:

From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>,
 <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>, <<var_packet-ctrl_vlan_id>>
spanning-tree port type network
exit
interface Po11
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po12
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po117
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_packet-ctrl_vlan_id>>
spanning-tree port type edge trunk
exit
interface Po118
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_packet-ctrl_vlan_id>>
spanning-tree port type edge trunk
exit
interface Po145
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>,
 <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po146
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>,
 <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po151
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>,
 <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit
interface Po152
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>,
 <<var_iscsi-a_vlan_id>>, <<var_iscsi-b_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
exit


copy run start
```

Configure Virtual Port Channels

Cisco Nexus 9372PX A

To configure virtual port channels (vPCs) for switch A, complete the following step:

From the global configuration mode, run the following commands:

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 10
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source <<var_nexus_A_mgmt0_ip>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po117
vpc 117
exit
interface Po118
vpc 118
exit
interface Po145
vpc 111
exit
interface Po146
vpc 112
exit
interface Po151
vpc 111
exit
interface Po152
vpc 112
exit
copy run start
```

Cisco Nexus 9372PX B

To configure vPCs for switch B, complete the following step:

From the global configuration mode, run the following commands.

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 20
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source <<var_nexus_B_mgmt0_ip>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
exit
interface Po10
vpc peer-link
exit
interface Po11
vpc 11
exit
interface Po12
vpc 12
exit
interface Po117
vpc 117
exit
interface Po118
vpc 118
exit
interface Po145
vpc 111
exit
interface Po146
vpc 112
exit
interface Po151
vpc 111
exit
```

```
interface Po152
vpc 112
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 9372PX switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Storage Configuration

AFF80XX Series Controllers

See the following sections in the Site Requirements Guide for planning the physical location of the storage systems:

- Site Preparation
- System Connectivity Requirements
- Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements
- 80xx Series Systems

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

> Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the HWU application at the NetApp Support site.

1. Access the HWU application to view the System Configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers found in the AFF8000 Series product documentation at the NetApp Support site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported by the AFF 80xx is available at the NetApp Support site.

When using SAS disk shelves with NetApp storage controllers, refer to the SAS Disk Shelves Universal SAS and ACP Cabling Guide for proper cabling guidelines.

Clustered Data ONTAP 8.**3.2**

Complete the Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the Clustered Data ONTAP 8.3 Software Setup Guide. You must have access to the NetApp Support site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the Clustered Data ONTAP 8.3 Software Setup Guide to learn about configuring ONTAP. Table 15 lists the information that you will need to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 15 ONTAP software installation prerequisites

Cluster Detail	Cluster Detail Value
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.3.2 URL	<<var_url_boot_software>>

Configure Node 01

To configure node 01, complete the following steps:

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 you see this message:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If Data ONTAP 8.3.2 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.2 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.

```
7
```

5. Enter y to perform an upgrade.

```
y
```

6. Select e0M for the network port you want to use for the download.

```
e0M
```

7. Enter y to reboot now.

```
y
```

8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_node02_mgmt_ip>> <<var_node02_mgmt_mask>> <<var_node02_mgmt_gateway>>
```

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

This web server must be pingable.

```
<<var_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.

```
y
```

12. Enter y to reboot the node.

```
y
```

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C when you see this message:

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

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

```
4
```

15. Enter y to zero disks, reset config, and install a new file system.

```
y
```

16. Enter y to erase all the data on the disks.

```
y
```

The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue with the node 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, complete the following steps:

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 you see this message:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

If Data ONTAP 8.3.2 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3.2 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.

```
7
```

5. Enter y to perform an upgrade.

```
y
```

6. Select e0M for the network port you want to use for the download.

```
e0M
```

7. Enter y to reboot now.

```
y
```

8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_node02_mgmt_ip>> <<var_node02_mgmt_mask>> <<var_node02_mgmt_gateway>>
```

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

✍ This web server must be pingable.

```
<<var_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.

```
y
```

12. Enter y to reboot the node.

```
y
```

✍ When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C when you see this message:

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

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

```
4
```

15. Enter y to zero disks, reset config, and install a new file system.

```
y
```

16. Enter y to erase all the data on the disks.

```
y
```

✍ The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when Data ONTAP 8.3.2 boots on the node for the first time.

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.

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: <<var_node01_mgmt_ip>>
Enter the node management interface netmask: <<var_node01_mgmt_mask>>
Enter the node management interface default gateway: <<var_node01_mgmt_gateway>>
A node management interface on port e0M with IP address <<var_node01_mgmt_ip>> has been created

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.
```

2. Press Enter and log in to the node with the admin user ID and no password.
3. At the node command prompt, enter the following commands to set HA mode for storage failover.

If the node responds that the HA mode was already set, then proceed with step 4.

```
::> storage failover modify -mode ha
Mode set to HA.  Reboot node to activate HA.

::> system node reboot

Warning: Are you sure you want to reboot node "localhost"? {y|n}: y
```

4. After reboot, set up the node with the preassigned values.

```
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.

To accept a default or omit a question, do not enter a value.



Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address [<<var_node01_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node01_mgmt_mask>>]: Enter
Enter the node management interface default gateway [<<var_node01_mgmt_gateway>>]: Enter

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.
```

5. Log in to the node as the admin user with no password.

Repeat this procedure for storage cluster node 02.

Create Cluster on Node 01

In ONTAP, the first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered node 01.

Table 16 Cluster create in ONTAP prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster node01 IP address	<<var_node01_mgmt_ip>>
Cluster node01 netmask	<<var_node01_mgmt_mask>>
Cluster node01 gateway	<<var_node01_mgmt_gateway>>

Run the cluster setup command to start the Cluster Setup wizard.

```
cluster setup
Welcome to the cluster setup wizard.
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 cluster 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.
Do you want to create a new cluster or join an existing cluster? {create, join}:
```

If a login prompt appears instead of the Cluster Setup wizard, start the wizard by logging in with the factory default settings and then enter the cluster setup command.

To create a new cluster, complete the following steps:

1. Run the following command to create a new cluster:

```
create
```

2. Enter no for the single-node cluster option.

```
Do you intend for this node to be used as a single node cluster? {yes, no} [no]: no
```

3. Enter no for a cluster network using network switches.

```
Will the cluster network be configured to use network switches? [yes]:no
```

4. The system defaults are displayed. Enter yes to use the system defaults. Use the following prompts to configure the cluster ports.

```
Existing cluster interface configuration found:


Port    MTU     IP              Netmask
e0a     9000    169.254.118.102 255.255.0.0
e0c     9000    169.254.191.92  255.255.0.0


Do you want to use this configuration? {yes, no} [yes]: no


System Defaults:
Private cluster network ports [e0a,e0c].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.


Do you want to use these defaults? {yes, no} [yes]: yes
```

If four ports are being used for the switchless cluster interconnect, enter e0a, e0b, e0c, and e0d for the private cluster network ports above.

5. The steps to create a cluster are displayed.

```
Enter the cluster administrators (username “admin”) password: <<var_password>>
Retype the password: <<var_password>>

It can take several minutes to create cluster interfaces...



Step 1 of 5: Create a Cluster
You can type "back", "exit", or "help" at any question.

Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter an additional license key []:<<var_iscsi_license>>
```

The cluster is created. This can take a few minutes.

For this validated architecture, NetApp recommends installing license keys for NetApp SnapRestore® data recovery software, NetApp FlexClone® data replication technology, and the NetApp SnapManager® suite. In addition, install all required storage protocol licenses and all licenses that came with the AFF bundle. After you finish entering the license keys, press Enter.

```
Enter the cluster management interface port [e0e]: e0i
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway: <<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

```
Enter the DNS domain names:<<var_dns_domain_name>>
Enter the name server IP addresses:<<var_nameserver_ip>>
```

If you have more than one name server IP address, separate the IP addresses with a comma.

7. Set up the node.

```
Where is the controller located []:<<var_node_location>>
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node01_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node01_mgmt_mask>>]: Enter
Enter the node management interface default gateway [<<var_node01_mgmt_gateway>>]: Enter

The node management interface has been modified to use port e0M with IP address <<var_node01_mgmt_ip>>.
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, please see: http://support.netapp.com/autosupport/
Press enter to continue: Enter
Cluster " <<var_clustername>>"  has been created.
To complete cluster setup, you must join each additional node to the cluster by running " cluster setup"  on each node.


Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for information about additional system configuration
tasks.  You can find the Software Setup Guide on the NetApp Support Site.


To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line interface.


To access OnCommand System Manager, point your web browser to the cluster management IP address (<<var_clustermgmt_ip>>).
To access the command-line interface, connect to the cluster management IP address (for example, ssh admin@<<var_clustermgmt_ip>>).




<<var_clustername>>::>
```

The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document it is assumed to be on the same subnet.

Join Node 02 to Cluster

The first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered node 01, and the node joining the cluster in this example is node 02.

Table 17 Cluster join in ONTAP prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster node02 IP address	<<var_node02_mgmt_ip>>
Cluster node02 netmask	<<var_node02_mgmt_mask>>
Cluster node02 gateway	<<var_node02_mgmt_gateway>>

To join node 02 to the existing cluster, complete the following steps:

1. If prompted, enter admin in the login prompt.

```
admin
```

2. Run the cluster setup command to start the Cluster Setup wizard.

```
cluster setup

This node's storage failover partner is already a member of a cluster.
Storage failover partners must be members of the same cluster.
The cluster setup wizard will default to the cluster join dialog.

Welcome to the cluster setup wizard.
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 cluster 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.
Do you want to create a new cluster or join an existing cluster?
{join}:
```

If a login prompt is displayed instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings, and then enter the cluster setup command.

3. Run the following command to join a cluster:

```
join
```

4. Data ONTAP detects the existing cluster and agrees to join the same cluster. Follow the prompts to join the cluster.

```
Existing cluster interface configuration found:


Port    MTU     IP              Netmask
e0a     9000    169.254.1.79    255.255.0.0
e0c     9000    169.254.100.157 255.255.0.0
Do you want to use this configuration? {yes, no} [yes]: no

System Defaults:
Private cluster network ports [e0a,e0c].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
```

If four ports are being used for the switchless cluster interconnect, enter e0a, e0b, e0c, and e0d for the private cluster network ports above.

```
Do you want to use these defaults? {yes, no} [yes]:Enter
It can take several minutes to create cluster interfaces...
```

5. The steps to join a cluster are displayed.

```
Step 1 of 3: Join an Existing Cluster
You can type "back", "exit", or "help" at any question.

Enter the name of the cluster you would like to join [<<var_clustername>>]:Enter
Joining cluster <<var_clustername>>
Starting cluster support services ..

This node has joined the cluster <<var_clustername>>.



Step 2 of 3: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.



SFO is enabled.



Step 3 of 3: Set Up the Node
You can type "back", "exit", or "help" at any question.

Notice: HA is configured in management.
```

The node should find the cluster name. Cluster joining can take a few minutes.

6. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node02_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node02_netmask>>]: Enter
Enter the node management interface default gateway [<<var_node02_gw>>]: Enter
The node management interface has been modified to use port e0M with IP address <<var_node02_mgmt_ip>>.
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, please see: http://support.netapp.com/autosupport/
Press enter to continue: Enter

This node has been joined to cluster "<<var_clustername>>".
To complete cluster setup, you must join each additional node to the cluster by running "cluster setup" on each node.

Once all nodes have been joined to the cluster, see the Clustered Data ONTAP Software Setup Guide for information about additional system configuration tasks.  You can find the Software Setup Guide on the NetApp Support Site.
```

To complete system configuration, you can use either OnCommand System Manager or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster management IP address (<<var_clustermgmt_ip>>).

To access the command-line interface, connect to the cluster management IP address (for example, ssh admin@<<var_clustermgmt_ip>>).

✍ The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document it is assumed to be on the same subnet.

Log In to the Cluster

To log in to the cluster, complete the following steps:

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

Zero All Spare Disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

✍ Disk autoassign should have assigned half of the connected disks to each node in the HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the disk option modify command. Spare disks can then be moved from one node to another by running the disk removeowner and disk assign commands.

Set Onboard Unified Target Adapter 2 Port Personality

To set the personality of the onboard Unified Target Adapter 2 (UTA2), complete the following steps:

1. Verify the Current Mode and Current Type properties of the ports by running the ucadmin show command.

```
ucadmin show
                    Current  Current    Pending  Pending    Admin
Node         Adapter  Mode     Type       Mode     Type       Status
------------ ------- ------- --------- ------- --------- -----------
<<var_node01>>
             0e       fc       target    -        -          online
<<var_node01>>
             0f       fc       target    -        -          online
<<var_node01>>
             0g       cna      target    -        -          online
<<var_node01>>
             0h       cna      target    -        -          online
<<var_node02>>
             0e       fc       target    -        -          online
<<var_node02>>
             0f       fc       target    -        -          online
<<var_node02>>
             0g       cna      target    -        -          online
<<var_node02>>
             0h       cna      target    -        -          online
8 entries were displayed.
```

2. Verify that the Current Mode and Current Type properties for all ports are set properly. Set ports used for Fibre Channel (FC) connectivity to mode fc; otherwise, set them to the mode cna. That includes FCoE ports, which should be set to the mode cna. The port type for all protocols should be set to target. Change the port personality with the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode {fc|cna} -type target
```

✍ The ports must be offline to run this command. To take an adapter offline, run the fcp adapter modify -node <home node of the port> -adapter <port name> -state down command. Ports must be converted in pairs (for example, 0e and 0f). After conversion, a reboot is required, and the ports must be brought back to the up state.

Set Auto-Revert on Cluster Management

To set the auto-revert parameter on the cluster management interface, complete the following step:

✍ A storage virtual machine (SVM) is referred to as a Vserver (or vserver) in the GUI and CLI.

Run the following command:

```
network interface modify -vserver <<var_clustername>> -lif cluster_mgmt -auto-revert true
```

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Ports used for data services (for example, e0b, e0d, e0e, e0f, e0g, e0h, e0j, e0k, and e0l) should be removed from the default broadcast domain, leaving just the management network ports (e0i and e0M). To perform this task, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_node01>>:e0c,<<var_node01>>:e0d,<<var_node01>>:e0e,<<var_node01>>:e0f,<<var_node01>>:e0g,<<var_node01>>:e0h,<<var_node01>>:e0j,<<var_node01>>
broadcast-domain show
```

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 <<var_node01>> –address-family IPv4 –enable true –dhcp none –ip-address <<var_node01_sp_ip>>
-netmask <<var_node01_sp_mask>> -gateway <<var_node01_sp_gateway>>


system service-processor network modify –node <<var_node02>> –address-family IPv4 –enable true –dhcp none –ip-address <<var_node02_sp_ip>>
-netmask <<var_node02_sp_mask>> -gateway <<var_node02_sp_gateway>>
```

The service processor IP addresses should be in the same subnet as the node management IP addresses.

Create Aggregates

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <<var_node01>> -diskcount <<var_num_disks>>
aggr create -aggregate aggr1_node02 -node <<var_node02>> -diskcount <<var_num_disks>>
```

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Start with five disks initially; you can add disks to an aggregate when additional storage is required. In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all but one remaining disk (spare) assigned to the controller.

The aggregate cannot be created until disk zeroing completes. Run the aggr show command to display aggregate creation status. Do not proceed until both aggr1_node1 and aggr1_node2 are online.

2. Disable NetApp Snapshot® copies for the two data aggregates recently created.

```
node run <<var_node01>> aggr options aggr1_node01 nosnap on
node run <<var_node02>> aggr options aggr1_node02 nosnap on
```

3. Delete any existing Snapshot copies for the two data aggregates.

```
node run <<var_node01>> snap delete –A –a –f aggr1_node01
node run <<var_node02>> snap delete –A –a –f aggr1_node02
```

4. Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02.

```
aggr show
aggr rename –aggregate aggr0 –newname <<var_node01_rootaggrname>>
```

Verify Storage Failover

To confirm that storage failover is enabled, run the following commands for a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Both the nodes <<var_node01>> and <<var_node02>> must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_node01>> -enabled true
```

Enabling failover on one node enables it for both nodes.

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

This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Continue with step 6 if high availability is configured.
5. 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
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
storage failover modify –hwassist-partner-ip <<var_node02_mgmt_ip>> -node <<var_node01>>
storage failover modify –hwassist-partner-ip <<var_node01_mgmt_ip>> -node <<var_node02>>
```

Disable Flow Control on UTA2 Ports

NetApp recommends disabling flow control on all of the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps:

1. Run the following commands to configure node 01:

```
network port modify -node <<var_node01>> -port e0b,e0d,e0e,e0f,e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <<var_node02>> -port e0b,e0d,e0e,e0f,e0g,e0h -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show –fields flowcontrol-admin
```

Disable Unused FC Capability on CNA Ports

If a UTA2 port is set to CNA mode, and is only expected to handle Ethernet data traffic (e.g., NFS), then the unused FC capability of the port should be disabled by setting the corresponding FCP adapter to state down, with the fcp adapter modify command. Here are some examples:

```
fcp adapter modify -node <<var_node01>> -adapter 0g –state down
fcp adapter modify -node <<var_node01>> -adapter 0h –state down
fcp adapter modify -node <<var_node02>> -adapter 0g –state down
fcp adapter modify -node <<var_node02>> -adapter 0h –state down
fcp adapter show –fields state
```

Configure NTP

To configure time synchronization on the cluster, complete the following steps:

1. Set the time zone for the cluster.

```
timezone <<var_timezone>>
```

For example, in the eastern United States, the time zone is America/New_York.

2. Set the date for the cluster.

```
date <ccyymmddhhmm.ss>
```

The format for the date is <[Century][Year][Month][Day][Hour][Minute].[Second]> (for example, 201309081735.17).

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

```
cluster time-service ntp server create -server <<var_switch_a_ntp_ip>>
cluster time-service ntp server create -server <<var_switch_b_ntp_ip>>
```

Configure SNMP

To configure the Simple Network Management Protocol (SNMP), complete the following steps:

4. Configure basic SNMP information, such as the location and contact. When polled, this information is visible as the sysLocation and sysContact variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location “<<var_snmp_location>>”
snmp init 1
options snmp.enable on
```

5. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <<var_oncommand_server_fqdn>>
```

Configure SNMPv1 Access

To configure SNMPv1 access, set the shared, secret plain-text password (called a community):

```
snmp community add ro <<var_snmp_community>>
```

Create SNMPv3 User

SNMPv3 requires that a user be defined and configured for authentication. To create and configure a user for SNMPv3, complete the following steps:

1. Create a user called snmpv3user.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

2. Enter the authoritative entity's engine ID and select md5 as the authentication protocol.
3. Run the security snmpusers command to view the engine ID.
4. When prompted, enter an eight-character minimum-length password for the authentication protocol.
5. Select des as the privacy protocol.
6. When prompted, enter an eight-character minimum-length password for the privacy protocol.

Configure AutoSupport

AutoSupport sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable –mail-hosts <<var_mailhost>> -transport https -support enable -noteto
<<var_storage_admin_email>>
```

Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command to enable CDP on ONTAP:

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

To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.

Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands to create a broadcast domain for NFS on ONTAP:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
```

If you are using iSCSI, run the following commands to create the broadcast domains for iSCSI on ONTAP:

```
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, run the following commands:

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0b
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0d

ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0b
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0d


ifgrp show
```

Create VLANs

To create VLANs, create NFS VLAN ports and add them to the NFS broadcast domain:

```
network port modify –node <<var_node01>> -port a0a –mtu 9000
network port modify –node <<var_node02>> -port a0a –mtu 9000
network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <<var_node01>>:a0a-<<var_nfs_vlan_id>>, <<var_node02>>:a0a-<<var_nfs_vlan_id>>
```

If you are using iSCSI, create iSCSI VLAN ports and add them to the iSCSI broadcast domains:

```
network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_node01>>:a0a-<<var_iscsi_vlan_A_id>>,<<var_node02>>:a0a-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_node01>>:a0a-<<var_iscsi_vlan_B_id>>,<<var_node02>>:a0a-
<<var_iscsi_vlan_B_id>>
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the vserver create command.

```
vserver create –vserver Infra-SVM –rootvolume rootvol –aggregate aggr1_node01 –rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping fcp, iscsi, and nfs.

```
vserver remove-protocols –vserver Infra-SVM -protocols cifs,ndmp
```

3. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp VSC.

```
vserver modify –vserver Infra-SVM –aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the NFS protocol in the Infra-SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the SVM vstorage parameter for the NetApp NFS VAAI plugin.

```
vserver nfs modify –vserver Infra-SVM –vstorage enabled
vserver nfs show
```

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, complete the following steps:

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

```
volume create -vserver Infra-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Infra-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

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

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path //Infra-SVM/rootvol -destination-path //Infra-SVM/rootvol_m01 -type LS -schedule 15min

snapmirror create -source-path //Infra-SVM/rootvol -destination-path //Infra-SVM/rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path //Infra-SVM/rootvol
snapmirror show
```

Create Block Protocol (iSCSI, FC) Service

If you are using iSCSI, run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

```
iscsi create -vserver Infra-SVM
iscsi show
```

If you are using FC, run the following command to create the FCP service on each SVM. This command also starts the FCP service and sets the WWN for the SVM.

```
fcp create -vserver Infra-SVM
fcp show
```

Configure HTTPS Access

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag

Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate, and obtain parameters (for example, <<serial_number>>) by running the following command:

```
security certificate show
```

For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the four default certificates and replace them with either self-signed certificates or certificates from a Certificate Authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial <<serial_number>>
```

Deleting expired certificates before creating new certificates is a best practice. Run the security certificate delete `command` to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

3. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Infra-SVM and the cluster SVM. Use TAB completion to aid in the completion of these commands.

```
security certificate create -common-name <<var_cert_common_name>> -type  server -size 2048 -country <<var_cert_country>> -state <<var_cert_state>> -locality <<var_cert_locality>> -organization <<var_cert_org>> -unit <<var_cert_unit>> -email-addr <<var_cert_email>> -expire-days <<var_cert_days>> -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

4. To obtain the values for the parameters required in step 6 (<<var_cert_ca>> and <<var_cert_serial>>), run the security certificate show command.

5. Enable each certificate that was just created by using the -server-enabled true and -client-enabled false parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <<var_clustername>> -server-enabled true -client-enabled false -ca <<var_cert_ca>> -serial <<var_cert_serial>> -common-name <<var_cert_common_name>>
```

6. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true

Warning: Modifying the cluster configuration will cause pending web service requests to be interrupted as the web servers are restarted.

Do you want to continue {y|n}: y

system services firewall policy delete -policy mgmt -service http -vserver <<var_clustername>>
```

It is normal for some of these commands to return an error message stating that the entry does not exist.

7. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin

vserver services web modify -name compat -vserver * -enabled true
```

Configure NFSv3

To configure NFSv3 on the SVM, complete the following steps:

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for each ESXi host created so that each host has its own rule index. For example, the first ESXi host has rule index 1, the second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create –vserver Infra-SVM –policyname default –ruleindex 2 –protocol nfs –clientmatch <<var_esxi_host2_nfs_ip>> -rorule sys
–rwrule sys –superuser sys –allow-suid false
vserver export-policy rule show
```

2. Assign the FlexPod export policy to the infrastructure SVM root volume.

```
volume modify –vserver Infra-SVM –volume rootvol –policy default
```

Create FlexVol Volumes

The following information is required to create a NetApp FlexVol® volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

To create a FlexVol volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node02 -size 500GB -state online -policy default -junction-
path /infra_datastore_1 -space-guarantee none -percent-snapshot-space 0


volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_node01 -size 100GB -state online -policy default -junction-path /infra_swap
-space-guarantee none -percent-snapshot-space 0 -snapshot-policy none


volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_node01 -size 100GB -state online -policy default -space-guarantee none
-percent-snapshot-space 0


snapmirror update-ls-set -source-path //Infra-SVM/rootvol
```

Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Prod-02 -size 15GB -ostype vmware -space-reserve disabled
```

Schedule Deduplication

On NetApp All Flash FAS systems, deduplication is enabled by default. To schedule deduplication, complete the following steps:

1. After the volumes are created, assign a once-a-day dedup schedule to esxi_boot:

```
efficiency modify –vserver Infra-SVM –volume esxi_boot –schedule sun-sat@0
```

2. Create the Always_On_Deduplication efficiency policy:

```
cron create -name 1min -minute
0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31,32,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47,48,48,50,51,52,53,54,55,56,5
efficiency policy create -vserver Infra-SVM -policy Always_On_Deduplication -type scheduled -schedule 1min -qos-policy background -enabled true
```

3. Optionally, assign the Always On Deduplication policy to infra_datastore_1:

```
efficiency modify -vserver Infra-SVM -volume infr_datastore_1 -policy Always-On-Deduplication
```

4. If you do not want to assign an Always On Deduplication policy to infra_datastore_1, assign the once-a-day deduplication schedule:

```
efficiency modify –vserver Infra-SVM –volume infra_datastore_1 -schedule sun-sat@0
```

Create iSCSI LIFs

If you are using iSCSI, run the following commands to create four iSCSI LIFs (two on each node):

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-node <<var_node01>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_node01_iscsi_lif01a_ip>> -netmask <<var_node01_iscsi_lif01a_mask>> –status-admin up –failover-policy
disabled –firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -home-node <<var_node01>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_node01_iscsi_lif01b_ip>> -netmask <<var_node01_iscsi_lif01b_mask>> –status-admin up –failover-policy
disabled –firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -home-node <<var_node02>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_node02_iscsi_lif01a_ip>> -netmask <<var_node02_iscsi_lif01a_mask>> –status-admin up –failover-policy
disabled –firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-node <<var_node02>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_node02_iscsi_lif01b_ip>> -netmask <<var_node02_iscsi_lif01b_mask>> –status-admin up –failover-policy
disabled –firewall-policy data -auto-revert false
network interface show
```

Create FCP LIFs

If you are using FCP, run the following commands to create four FCP LIFs (two on each node) per attached fabric interconnect:

```
network interface create -vserver Infra-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node <<var_node01>> -home-port
<<var_node01_fcp_port1>> –status-admin up
```

```
network interface create -vserver Infra-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node <<var_node01>> -home-port
<<var_node01_fcp_port2>> -status-admin up


network interface create -vserver Infra-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node <<var_node02>> -home-port
<<var_node02_fcp_port1>> -status-admin up


network interface create -vserver Infra-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node <<var_node02>> -home-port
<<var_node02_fcp_port2>> -status-admin up
network interface show
```

Create NFS LIF

To create an NFS LIF, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_infra_swap -role data -data-protocol nfs -home-node <<var_node01>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_node01_nfs_lif_infra_swap_ip>> -netmask <<var_node01_nfs_lif_infra_swap_mask>> -status-admin up -failover-
policy broadcast-domain-wide -firewall-policy data -auto-revert true


network interface create -vserver Infra-SVM -lif nfs_infra_datastore_1 -role data -data-protocol nfs -home-node <<var_node02>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_node02_nfs_lif_infra_datastore_1_ip>> -netmask <<var_node02_nfs_lif_infra_datastore_1_mask>> -status-admin
up –failover-policy broadcast-domain-wide –firewall-policy data -auto-revert true


network interface show
```

NetApp recommends creating a new LIF for each datastore.

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, complete the following steps:

1. Run the following commands:

```
network interface create –vserver Infra-SVM –lif vsmgmt –role data –data-protocol none –home-node <<var_node02>> -home-port  e0i –address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up –failover-policy broadcast-domain-wide –firewall-policy mgmt –auto-
revert true
```

The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create –vserver Infra-SVM -destination 0.0.0.0/0 –gateway <<var_svm_mgmt_gateway>>


network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password –username vsadmin –vserver Infra-SVM
Enter a new password:  <<var_password>>
Enter it again:  <<var_password>>


security login unlock –username vsadmin –vserver Infra-SVM
```

A cluster serves data through at least one and possibly multiple storage virtual machines (SVM). We have just gone through creating a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create additional SVMs.

Server Configuration

Cisco UCS Base **Configuration**

This FlexPod deployment will show configuration steps for both the Cisco UCS 6332-16UP and Cisco UCS 6248UP Fabric Interconnects (FI) in a design that will support iSCSI as well as Fibre Channel direct attached connectivity to the NetApp AFF. Implementation of both of these protocols simultaneously should not be considered mandatory, and the selection of one or the other should be acceptable depending upon your environment and preferences.

Configuration steps will be referenced for both fabric interconnects and will be called out by the specific model where steps have differed.

Perform Initial Setup of Cisco UCS 6332-16UP and 6248UP Fabric Interconnects for FlexPod Environments

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

Cisco UCS 6332-16UP A

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

1. Connect to the console port on the first Cisco UCS 6332-16UP fabric interconnect.

```
Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have chosen to setup a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y
```

```
Enter the password for "admin": <<var_password>>
Enter the same password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)? (yes/no) [n]: y
Which switch fabric (A|B): A
Enter the system name: <<var_ucs_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_cluster_ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? y
Default domain name: <<var_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 to make sure that the configuration has been saved.

Cisco UCS 6332-16UP B

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

1. Connect to the console port on the second Cisco UCS 6332-16UP fabric interconnect.

```
Enter the configuration method: 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 for the peer fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>
Apply and save the configuration (select 'no' if you want to re-enter)?
(yes/no): y
```

2. Wait for the login prompt to make sure that the configuration has been saved.
3. Repeat these steps for the 6248 Fabric Interconnects.

Cisco UCS **Setup**

Log in to Cisco UCS Manager

The steps are the same between the UCS 6332-16UP and the UCS 6248UP Fabric Interconnects unless otherwise noted

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

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

Upgrade Cisco UCS Manager Software to Version **3.1(1h)**

This document assumes the use of Cisco UCS 3.1(1h). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.1(1h), refer to Cisco UCS Manager Install and Upgrade Guides.

Anonymous Reporting

To create anonymous reporting, complete the following step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products:

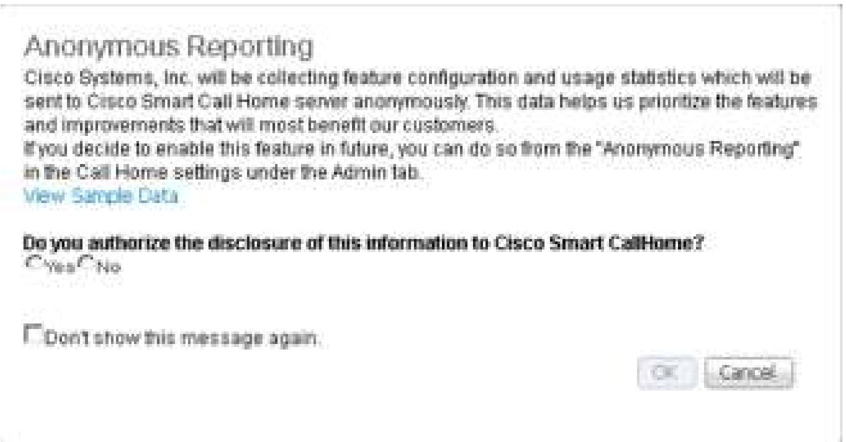

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.

If you decide to enable this feature in future, you can do so from the "Anonymous Reporting" in the Call Home settings under the Admin tab.

View Sample Data

Do you authorize the disclosure of this information to Cisco Smart CallHome?

○ Yes ○ No

□ Don't show this message again

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for in band server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.

4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

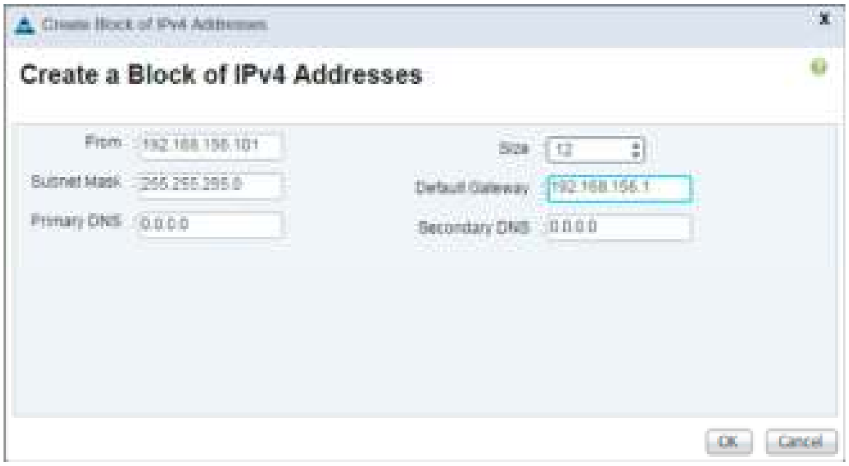

Create Block of IPv4 Addresses

Create a Block of IPv4 Addresses

From: 192.168.156.101

Size: 12

Subnet Mask: 255.255.255.0

Default Gateway: 192.168.156.1

Primary DNS: 0.0.0.0

Secondary DNS: 0.0.0.0

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

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, click the Admin tab in the navigation pane.
2. Select All > Timezone Management.

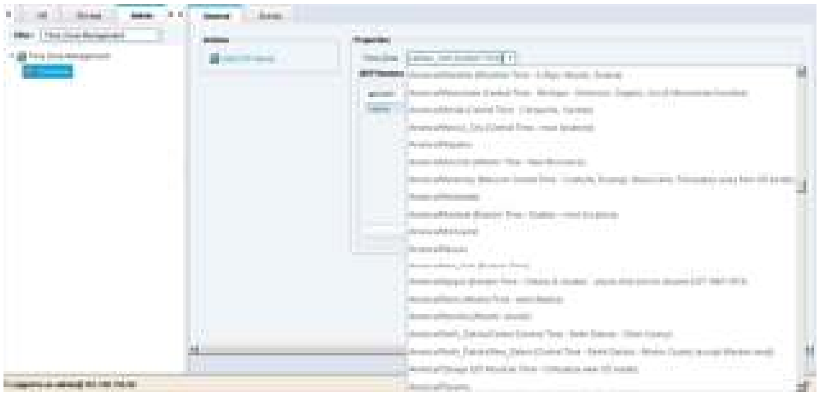

3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.
6. Enter <<var_switch_a_ntp_ip>> and click OK.

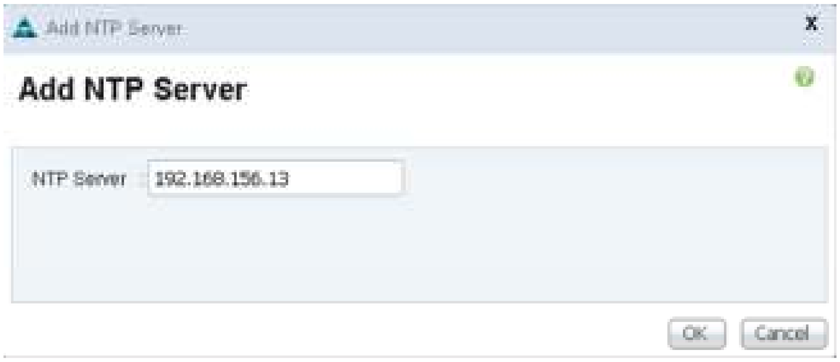

Add NTP Server

Add NTP Server

NTP Server: 192.168.156.13

7. Click Add NTP Server.
8. Enter <<var_switch_b_ntp_ip>> and click OK.

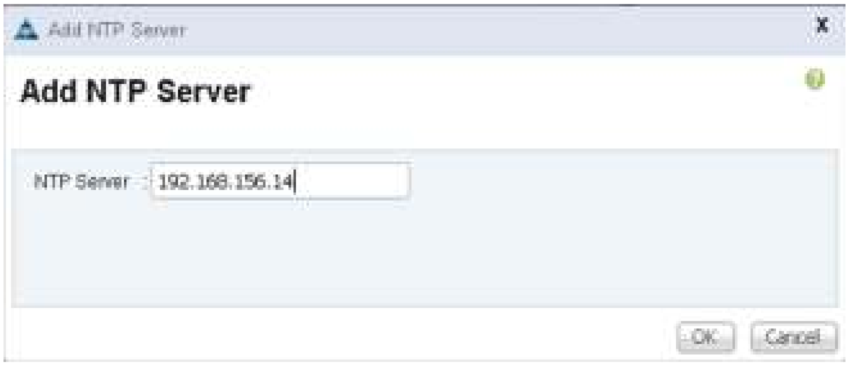

Add NTP Server

NTP Server : 192.168.156.14

9. Click OK.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and of additional fabric extenders for further C-Series connectivity. To modify the chassis discovery policy, complete the following steps:

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

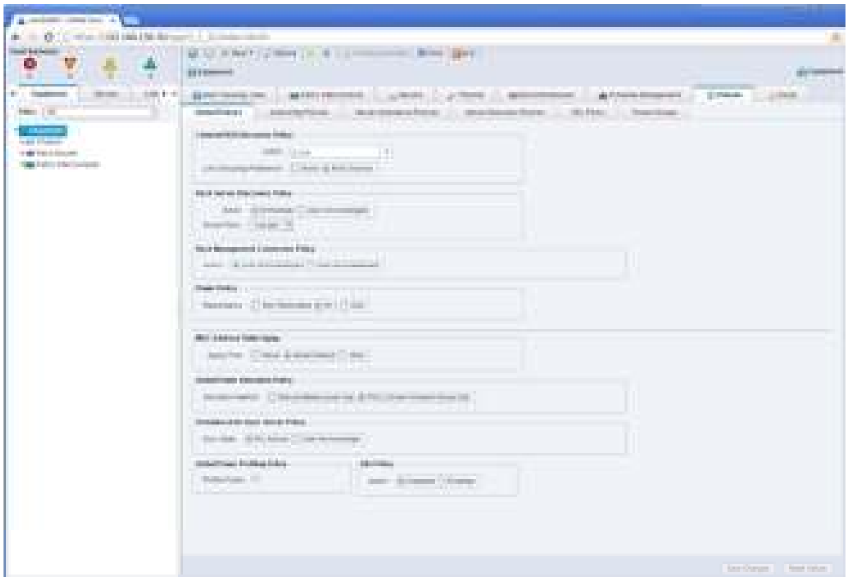

5. Click Save Changes.
6. Click OK.

Enable FC Switching

To use direct attached Fibre Channel connectivity, the fabric interconnects will need to be placed in Fibre Channel Switching mode, by completing the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Expand Fabric Interconnects and select either Fabric Interconnect.

This next step will reboot both UCS Fabric Interconnects. If any servers are running on this system, they should be shut down before this step is executed.

---

3. In the Actions pane, select Set FC Switching Mode. Click Yes. Click OK.

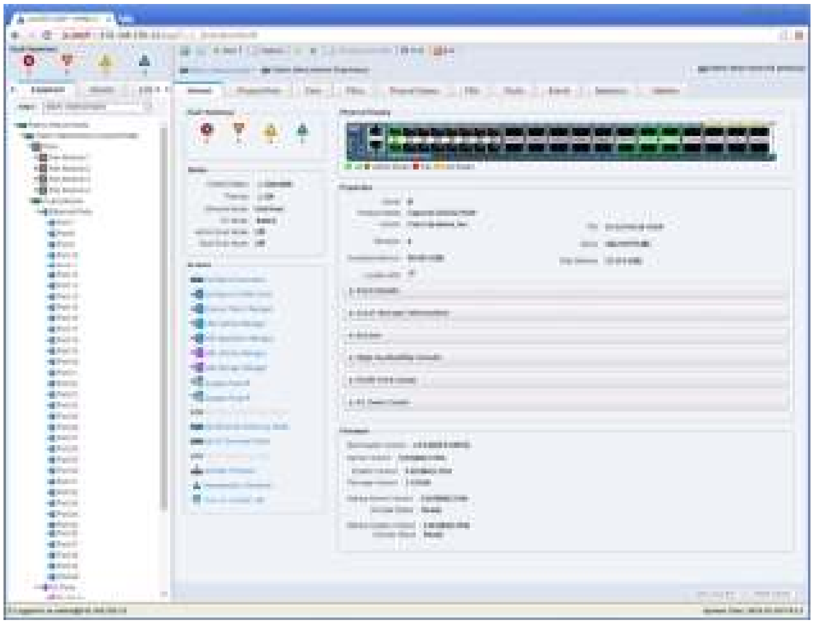

4. After the Fabric Interconnects have rebooted, log back into UCS Manager.
5. Expand Fabric Interconnects and select Fabric Interconnects.
6. For each Fabric Interconnect, verify under Status that the FC Mode is now Switch.

Configure Unified Ports

Fibre Channel port configurations will slightly differ between the 6332-16UP and the 6248UP Fabric Interconnects. Both Fabric Interconnects will have a slider mechanism within the UCSM GUI interface, but the fibre channel port selection options for the 6332-16UP will be from the first 16 ports starting from the first port, and configured in increments of the first 6, 12, or all 16 of the unified ports. With the 6248UP, the port selection options will start from the upper end of the 32 fixed ports, or the upper end of the 16 ports of the expansion module, going down in contiguous increments of 2.

To enable the fibre channel ports, complete the following steps for the 6332-16UP:

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

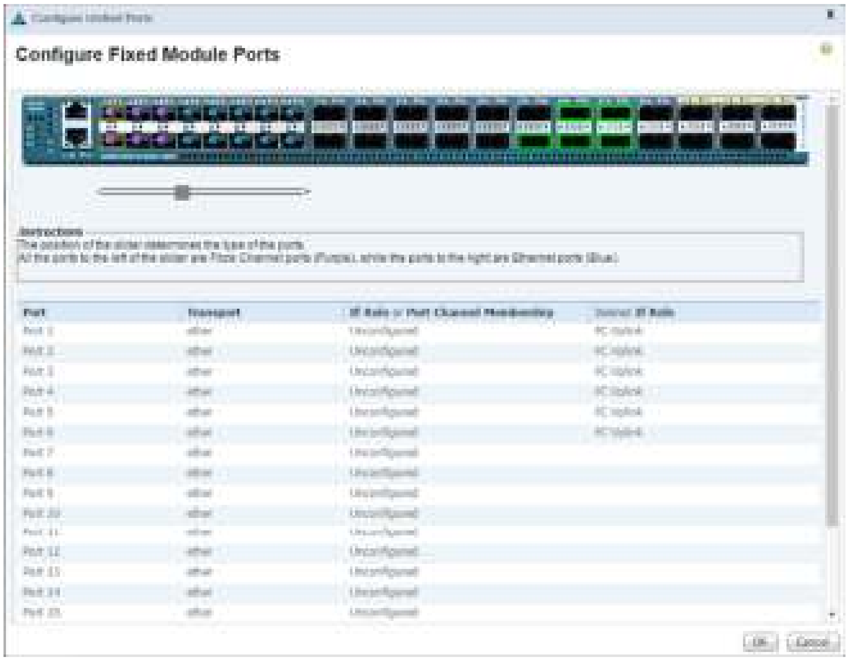

Configure Unified Ports

Configure Fixed Module Ports

Instructions
The position of the slider determines the type of the ports.
All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	FC Uplink
Port 6	ether	Unconfigured	FC Uplink
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	

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

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

To enable the fibre channel ports, complete the following steps for the 6248UP:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary)
3. Select Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
5. Configure the Fixed Module Ports from the subsequent Configure Fixed Module Ports pop-up window, or click on the Configure Expansion Module Ports button to select from expansion module ports.
6. Within either option (Expansion Module shown below) move the gray slider bar from the right to the left selecting ports in increments of two to set as FC Uplinks.

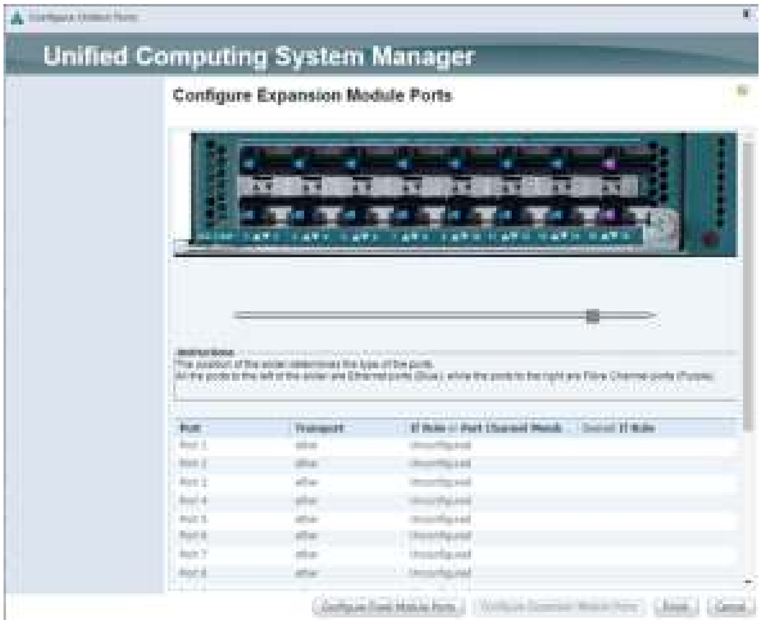
Configure Unified Ports
Unified Computing System Manager
Configure Expansion Module Ports
Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

Instructions
The position of the slider determines the type of the ports.
All the ports to the left of the slider are Ethernet ports (Blue), while the ports to the right are Fibre Channel ports (Purple).

Port	Transport	If Role or Port Channel Member	Desired If Role
Port 1	ether	Unconfigured	
Port 2	ether	Unconfigured	
Port 3	ether	Unconfigured	
Port 4	ether	Unconfigured	
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	

7. Click Finish.
8. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate)
9. Select Configure Unified Ports.
10. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
11. Configure the Fixed Module Ports from the subsequent Configure Fixed Module Ports pop-up window, or click on the Configure Expansion Module Ports button to select from expansion module ports.
12. Within either option move the gray slider bar from the right to the left selecting ports in increments of two to set as FC Uplinks.
13. Click Finish.

Enable FC Storage Ports

To enable FC ports, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module. (or Expansion Module if FC ports were selected from the Expansion Module in the 6248UP)
3. Expand FC Ports.
4. Select ports that are connected to the FC ports on the storage controllers, (this will be ports 1 and 2 in our 6332-16UP example and ports 15 and 16 of the Expansion Module in our 6248UP example), right-click them, and select Configure as FC Storage Port. Click Yes to confirm.
5. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module. (or Expansion Module if FC ports were selected from the Expansion Module in the 6248UP)
6. Expand FC Ports.
7. Select ports that are connected to the FC ports on the storage controllers, (this will be ports 1 and 2 in our 6332-16UP example and ports 15 and 16 of the Expansion Module in our 6248UP example), right-click them, and select Configure as FC Storage Port. Click Yes to confirm.

Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis, and Cisco FEX, and direct connect UCS C-Series servers, right-click them, and select “Configure as Server Port.”
5. Click Yes to confirm server ports and click OK.
6. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
7. Select ports 31 and 32 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

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

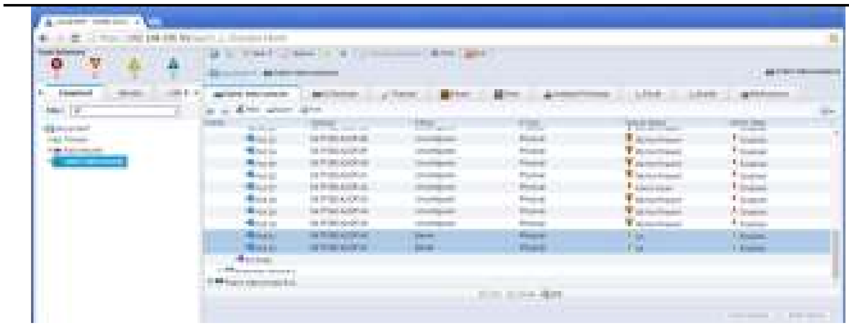

8. Click Yes to confirm uplink ports and click OK.
9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
10. Expand Ethernet Ports.
11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select ports 19 and 20 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external 2232 FEX modules, complete the following steps:

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

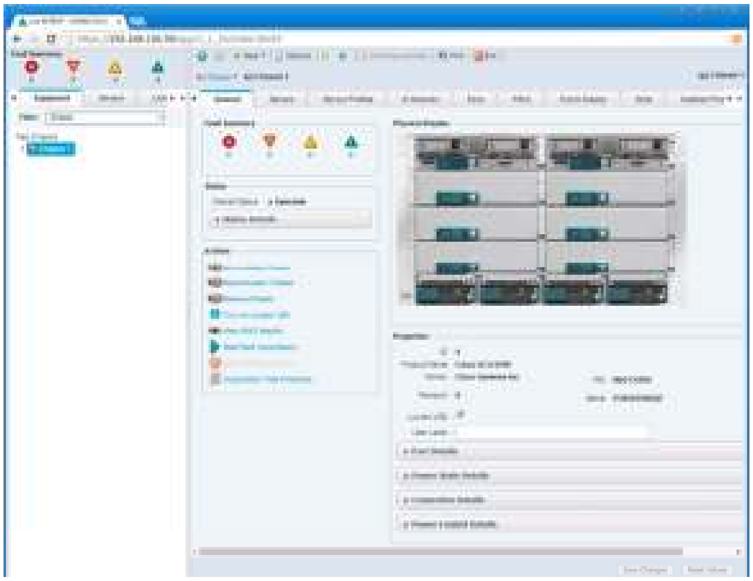

4. Click Yes and then click OK to complete acknowledging the chassis.
5. If the Nexus 2232 FEX is part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and select Acknowledge FEX.
7. Click Yes and then click OK to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus Switches

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

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

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

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 13 as the unique ID of the port channel.
6. Enter vPC-13-Nexus as the name of the port channel.
7. Click Next.

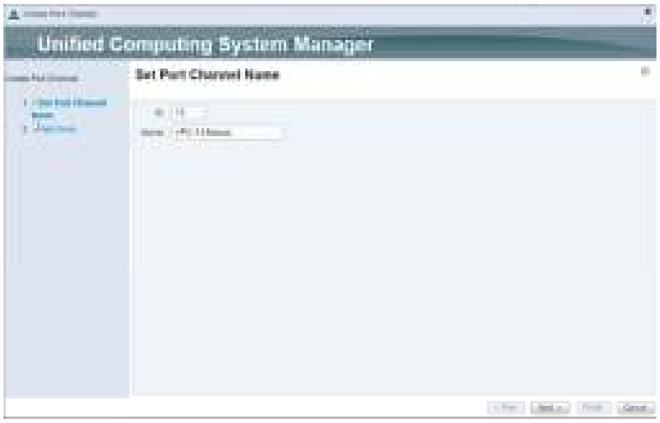

Set Port Channel Name

ID: 13

Name: vPC-13-Nexus

8. Select the following ports to be added to the port channel:
   - Slot ID 1 and port 31
   - Slot ID 1 and port 32
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter vPC-14-Nexus as the name of the port channel.
17. Click Next.
18. Select the following ports to be added to the port channel:
   - Slot ID 1 and port 31
   - Slot ID 1 and port 32
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create a WWNN Pool for FC Boot

To configure the necessary WWNN pool for the Cisco UCS environment, complete the following steps on Cisco UCS Manager.

1. Select the SAN tab on the left.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Select Create WWNN Pool to create the WWNN pool.
5. Enter `WWNN_Pool` for the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.
7. Select Sequential for Assignment Order.

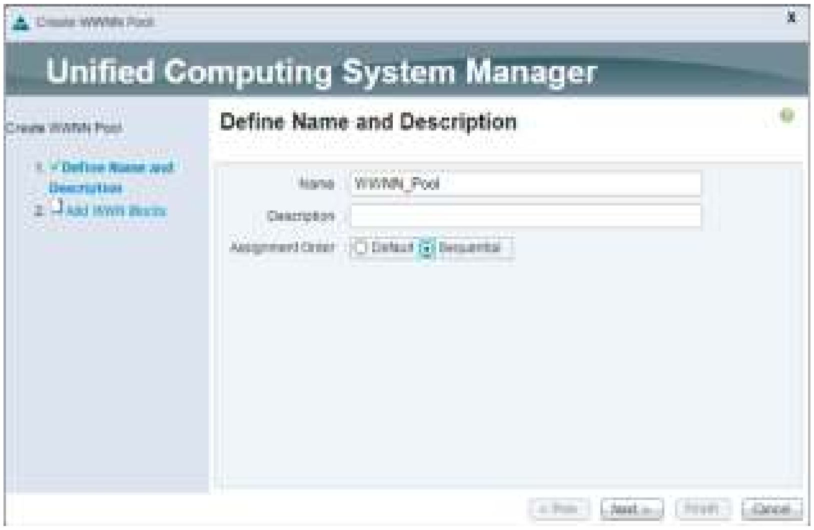

Define Name and Description

Name: WWNN_Pool

Description:

Assignment Order: ○ Default ● Sequential

8. Click Next.
9. Click Add.
10. Modify the From field as necessary for the UCS Environment.

Modifications of the WWN block, as well as the WWPN and MAC Addresses, can convey identifying information for the UCS domain. Within the From field in our example, the 6th octet was changed from 00 to 91 to represent as identifying information for this being in building 9 on the 1st floor, and the 7th octet was changed from 00 to 10 to represent our first UCS domain.

Also, when having multiple UCS domains sitting in adjacency, it is important that these blocks, the WWN, WWPN, and MAC hold differing values between each set.

11. Specify a size of the WWNN block sufficient to support the available server resources.

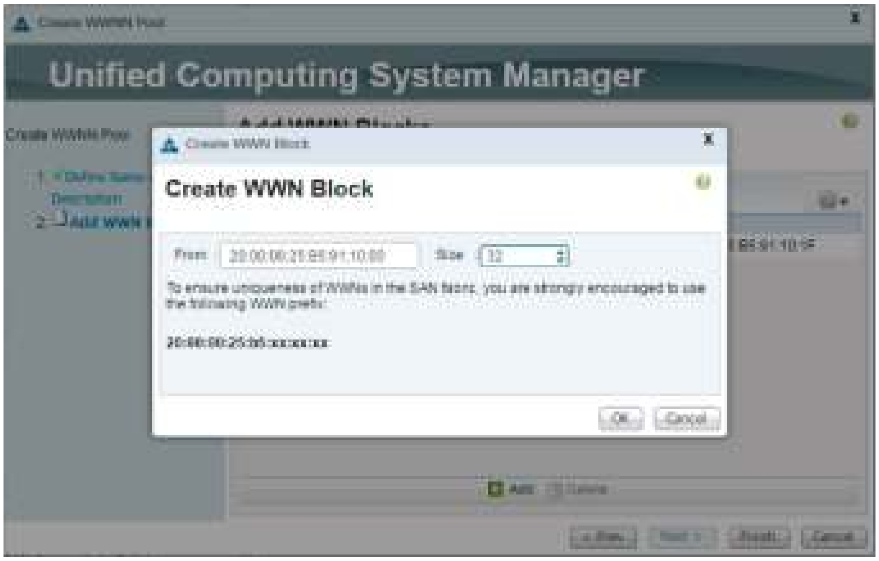

Create WWN Block

From: 20:00:00:25:B5:91:10:00 Size: 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

12. Click OK.

Create WWPN Pools

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. In this procedure, two WWPN pools are created, one for each switching fabric.

4. Right-click WWPN Pools under the root organization.
5. Select Create WWPN Pool to create the WWPN pool.
6. Enter `WWPN_Pool_A` as the name of the WWPN pool.
7. Optional: Enter a description for the WWPN pool.
8. Select Sequential for Assignment Order.

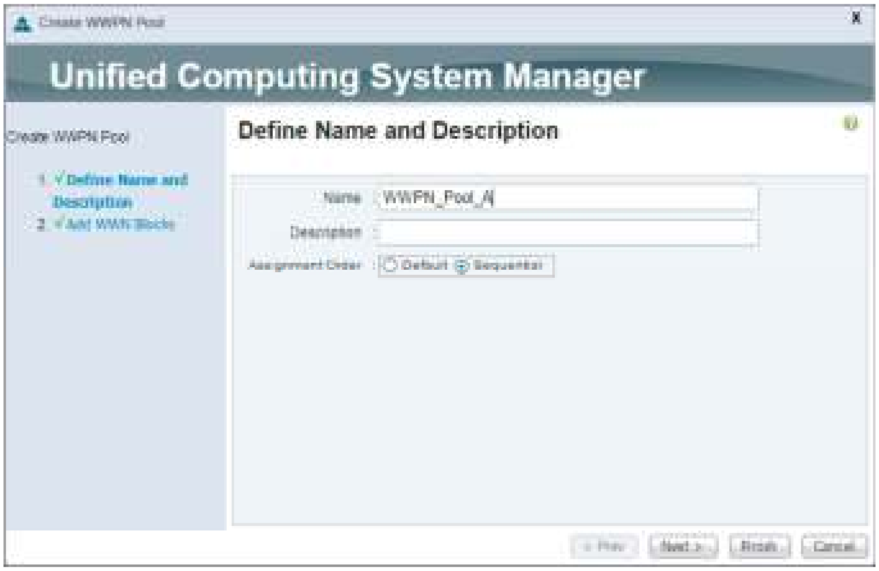

Define Name and Description

Name: WWPN_Pool_A

Description:

Assignment Order: ○ Default ● Sequential

9. Click Next.
10. Click Add.
11. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:91:1A:00`.

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

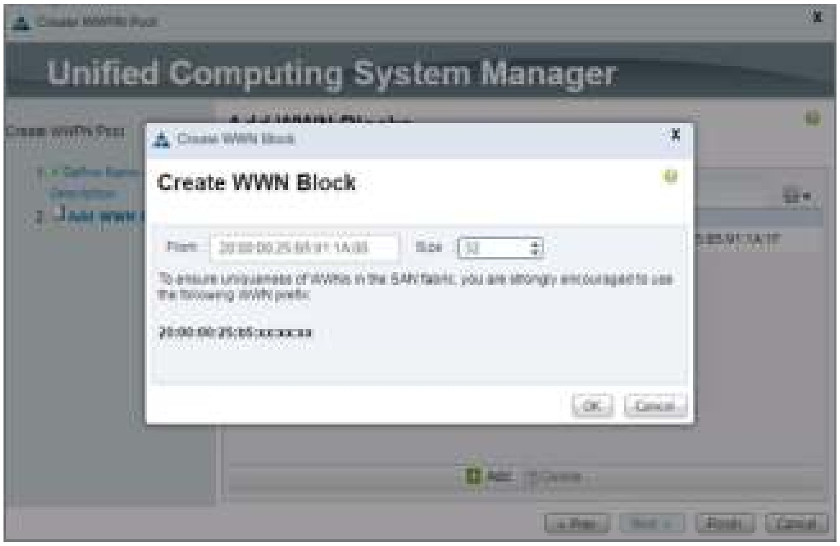

Create WWN Block

From: 20:00:00:25:B5:91:1A:00 Size: 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

13. Click OK.
14. Click Finish.
15. In the confirmation message, click OK.
16. Right-click WWPN Pools under the root organization.
17. Select Create WWPN Pool to create the WWPN pool.
18. Enter `WWPN_Pool_B` as the name of the WWPN pool.
19. Optional: Enter a description for the WWPN pool.
20. Select Sequential for Assignment Order.

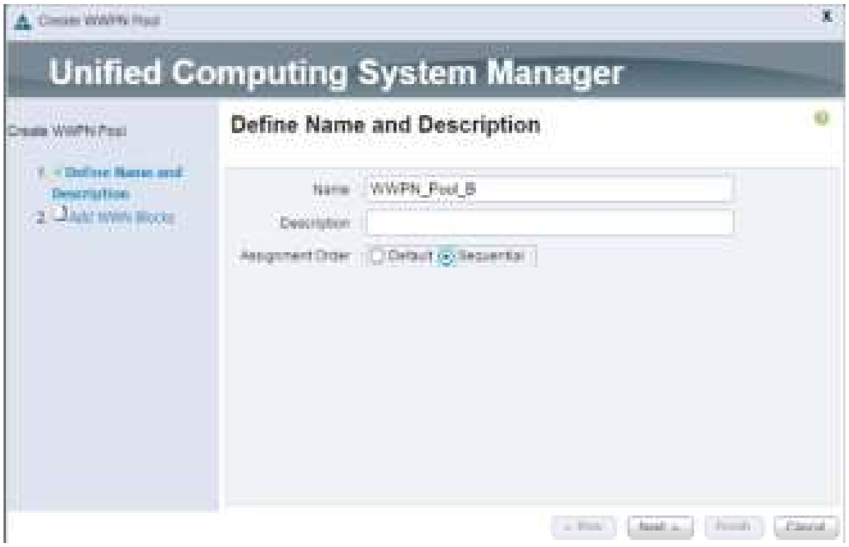

Define Name and Description

Name: WWPN_Pool_B

Description:

Assignment Order: ○ Default ● Sequential

21. Click Next.
22. Click Add.
23. Specify a starting WWPN.

For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:91:1AB:00`.

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

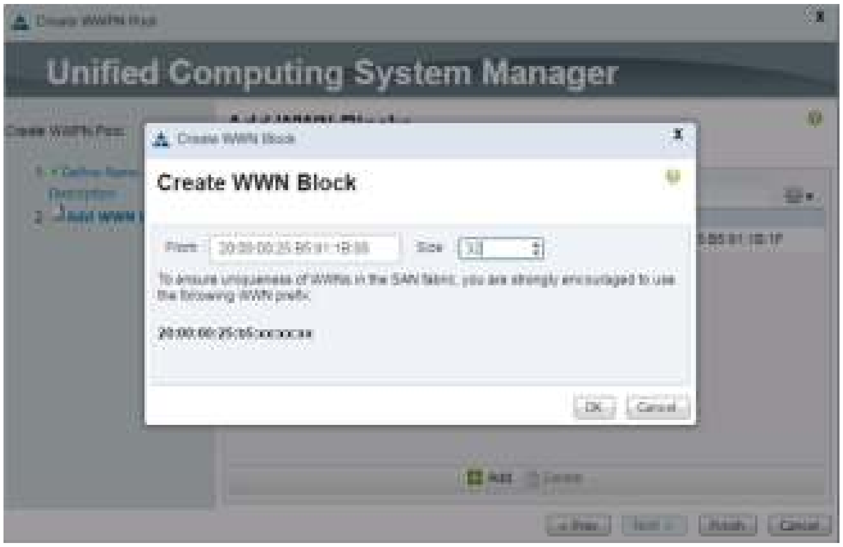

Create WWN Block

From: 20:00:00:25:B5:91:1B:00 Size: 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

25. Click OK.
26. Click Finish.
27. In the confirmation message, click OK.

Create VSANs

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

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

In this procedure, two VSANs are created.

2. Select SAN > SAN Cloud.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A
6. Select Enabled for FC Zoning.
7. Select Fabric A.
8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. It is recommended use the same ID for both parameters and to use something other than 1.

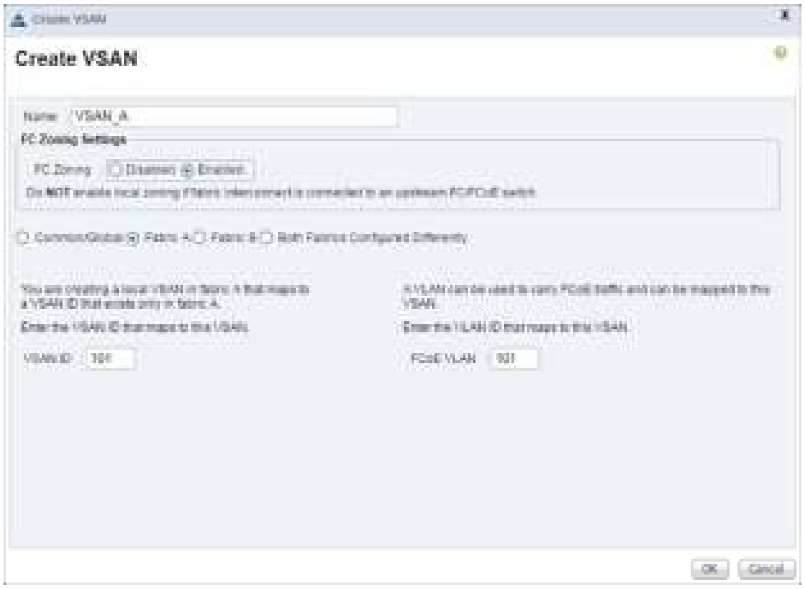

Create VSAN

Name: VSAN_A

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ● Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 101

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 101

9. Click OK, and then click OK again.
10. Under SAN Cloud, right-click VSANs.
11. Select Create VSAN.
12. Enter `VSAN_B` as the name of the VSAN to be used for Fabric B.
13. Select Enabled for FC Zoning.
14. Select Fabric B.
15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. It is recommended use the same ID for both parameters and to use something other than 1.

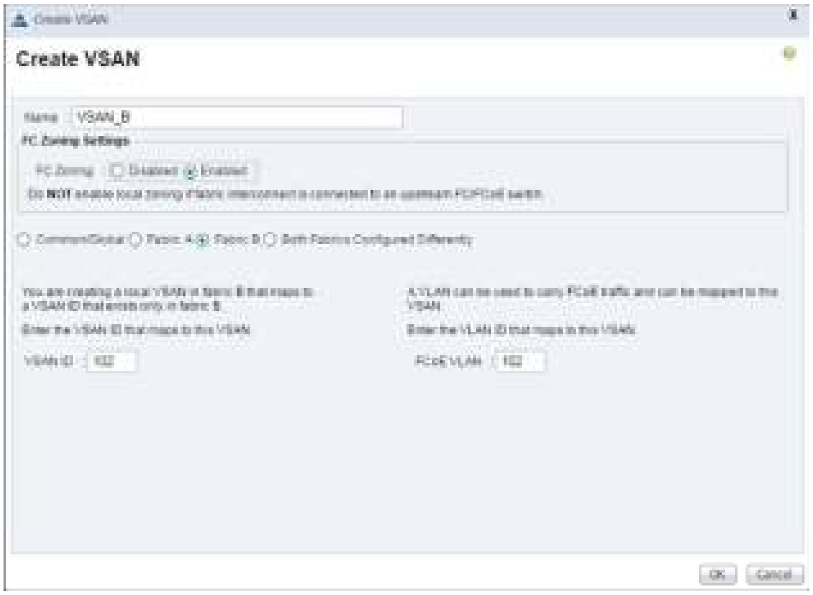

Create VSAN

Name: VSAN_B

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ○ Fabric A ● Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 102

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 102

16. Click OK, and then click OK again.
17. Under Storage Cloud, right-click VSANs.
18. Select Create Storage VSAN.
19. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A.
20. Select Enabled for FC Zoning.
21. Select Fabric A.
22. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for Fabric A above.

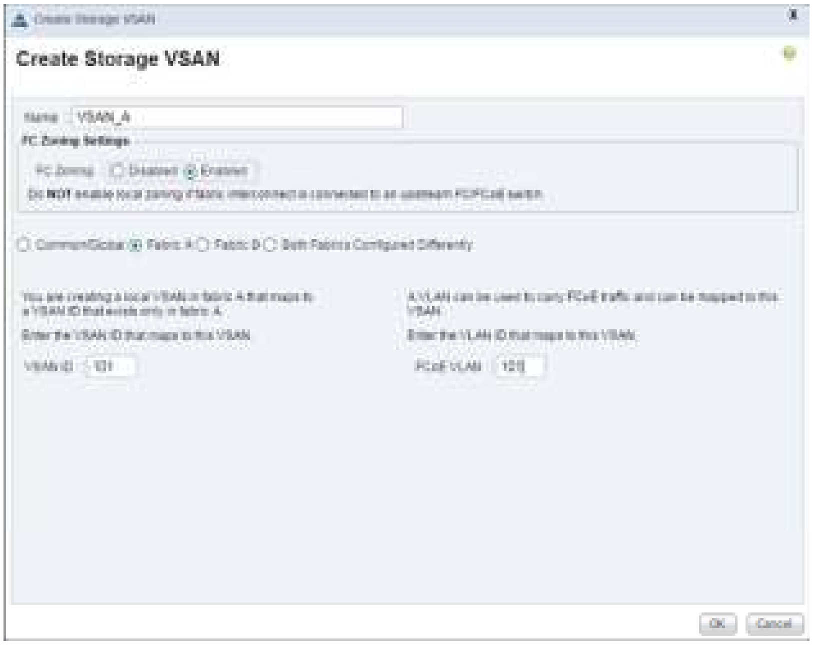

Create Storage VSAN

Create Storage VSAN

Name: VSAN_A

FC Zoning Settings

FC Zoning ○ Disabled ● Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ● Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 101

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 101

23. Click OK, and then click OK again.
24. Under Storage Cloud, right-click VSANs.
25. Select Create Storage VSAN.
26. Enter `VSAN_B` as the name of the VSAN to be used for Fabric B.
27. Select Enabled for FC Zoning.
28. Select Fabric B.
29. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for Fabric B above.

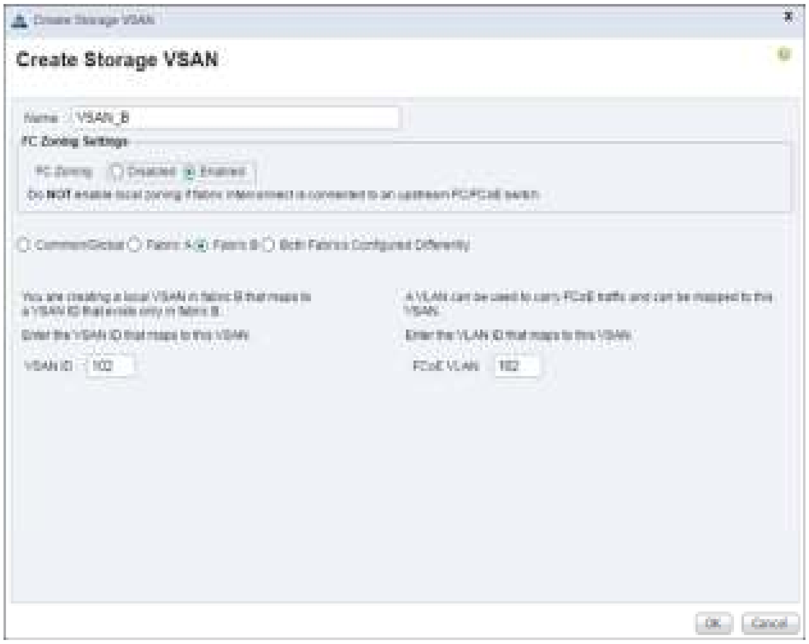

Create Storage VSAN

Create Storage VSAN

Name: VSAN_B

FC Zoning Settings

FC Zoning ○ Disabled ● Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ○ Fabric A ● Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 102

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 102

30. Click OK, and then click OK again.

Assign VSANs to FC Storage Ports

To assign the necessary virtual storage area networks (VSANs) to the FC Storage Ports 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 > Storage Cloud.
3. Expand Fabric A and Storage FC Interfaces.
4. Right-click FC Interface 1/1 for the 6332-16UP or 2/15 in the 6248UP and select Storage FC Interface.
5. Set the User Label to the storage controller name and port that this interface is connected to.
6. Select `VSAN_A(101)` as the VSAN.

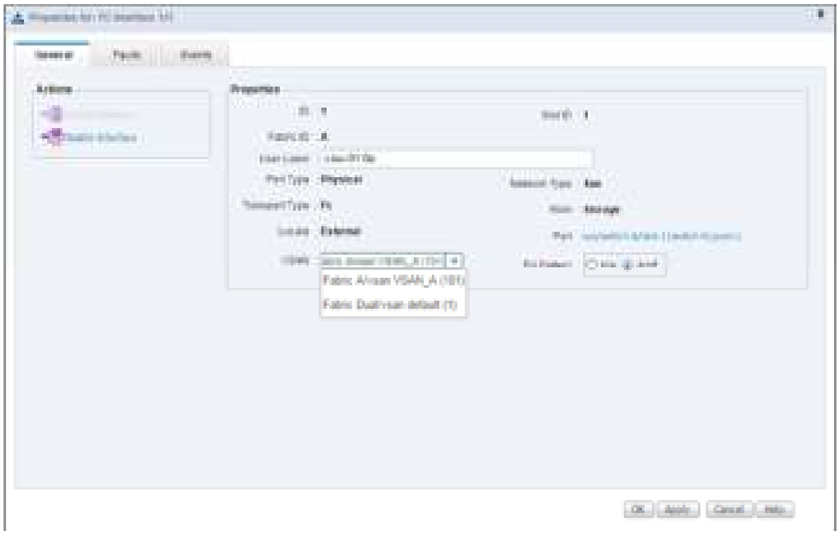

7. Click OK.
8. Expand Fabric A and Storage FC Interfaces.
9. Right-click FC Interface 1/2 for the 6332-16UP or 2/16 in the 6248UP and select Storage FC.
10. Set the User Label to the storage controller name and port that this interface is connected to.
11. Select `VSAN_A(101)` as the VSAN.

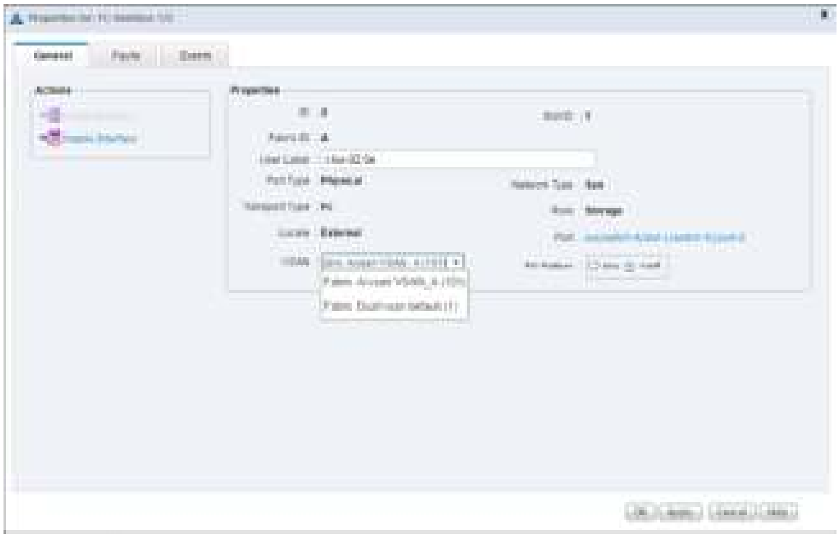

12. Click OK.
13. Expand Fabric B and Storage FC Interfaces.
14. Right-click FC Interface 1/1 for the 6332-16UP or 2/15 in the 6248UP and select Storage FC.
15. Set the User Label to the storage controller name and port that this interface is connected to.
16. Select `VSAN_B(102)` as the VSAN.

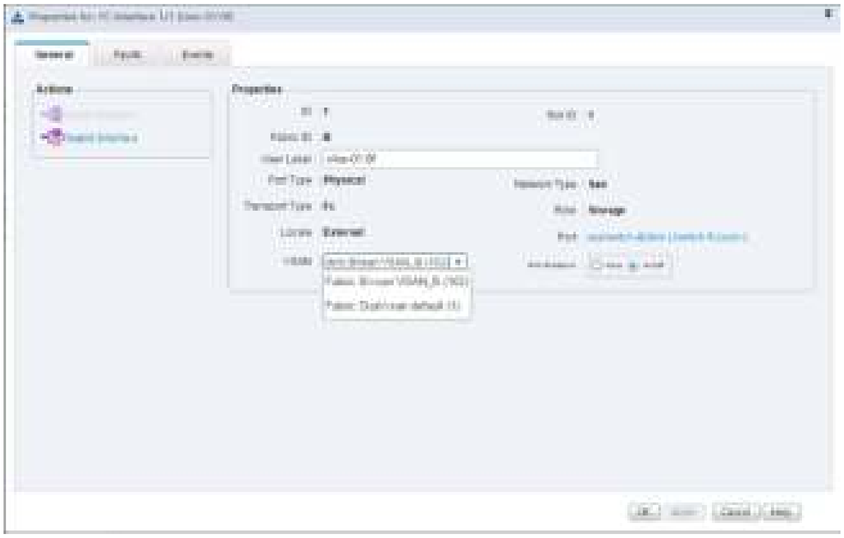

17. Click OK.
18. Expand Fabric B and Storage FC Interfaces.
19. Right-click FC Interface 1/2 for the 6332-16UP or 2/16 in the 6248UP and select Storage FC.
20. Set the User Label to the storage controller name and port that this interface is connected to.
21. Select `VSAN_B(102)` as the VSAN.

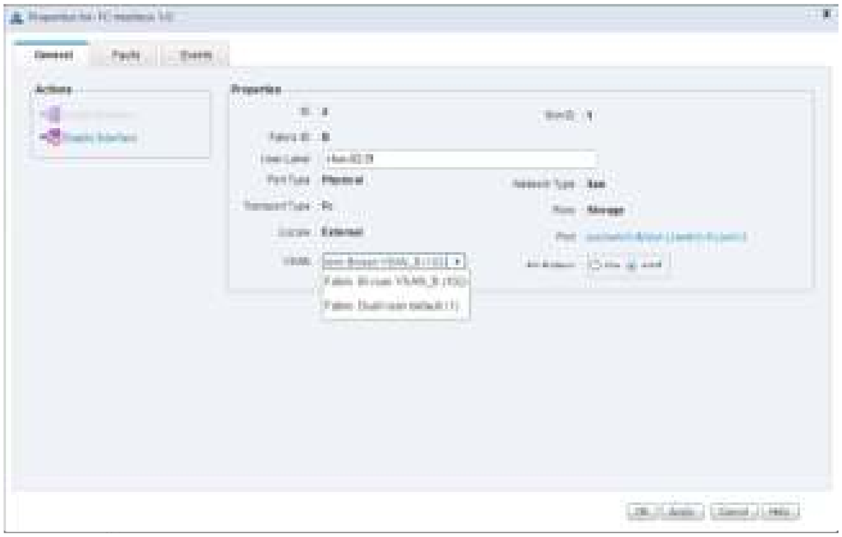

22. Click OK.

Create Storage Connection Policies for FC Zoning

To create Storage Connection Policies for the FC Zoning, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Storage Connection Policies.
4. Select Create Storage Connection Policy.
5. Enter `Infra-Fabric-A` as the name of the policy.
6. Select the Single Initiator Multiple Targets Zoning Type.
7. Click the Plus Sign on the right to add a zoning target.
8. Enter the WWPN for [fcp_lif01_63a or fcp_lif01_62a] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the network interface show -vserver Infra-SVM command.
9. Select Path A and VSAN_A.

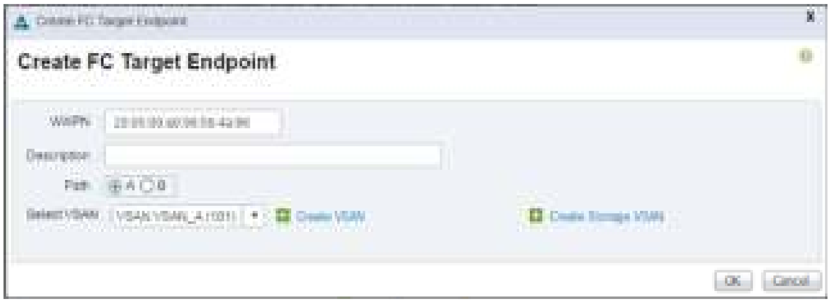

Create FC Target Endpoint

WWPN: 20:08:00:a0:98:5b:4a:86

Description:

Path: ● A ○ B

Select VSAN: VSAN VSAN_A (101)

10. Click OK.

11. Click the Plus Sign on the right to add a second zoning target.

12. Enter the WWPN for [fcp_lif02_63a or fcp_lif02_62a] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the network interface show -vserver Infra-SVM command.

13. Select Path A and VSAN_A.

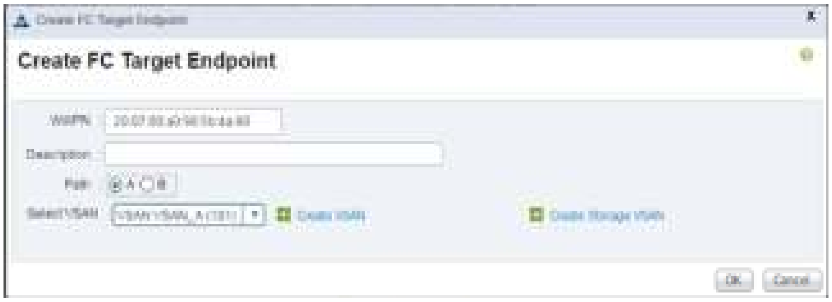

Create FC Target Endpoint

WWPN: 20:07:00:a0:98:5b:4a:86

Description:

Path: ● A ○ B

Select VSAN: VSAN VSAN_A (101)

14. Click OK.

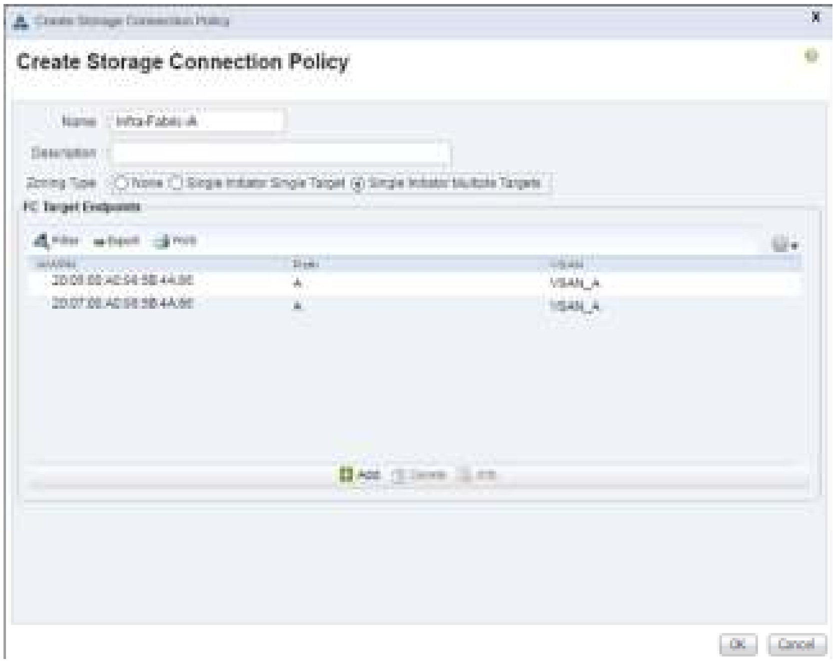

Create Storage Connection Policy

Name: Infra-Fabric-A

Description:

Zoning Type: ○ None ○ Single Initiator Single Target ● Single Initiator Multiple Targets

FC Target Endpoints

WWPN	Path	VSAN
20:08:00:A0:98:5B:4A:86	A	VSAN_A
20:07:00:A0:98:5B:4A:86	A	VSAN_A

15. Click OK, and then click OK again.

16. Right-click Storage Connection Policies.

17. Select Create Storage Connection Policy

18. Enter `Infra-Fabric-B` as the name of the policy.

19. Select the Single Initiator Multiple Targets Zoning Type.

20. Click the Plus Sign on the right to add a zoning target.

21. Enter the WWPN for [fcp_lif01_63b or fcp_lif01_62b] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the network interface show -vserver Infra-SVM command.

22. Select Path B and VSAN_B.

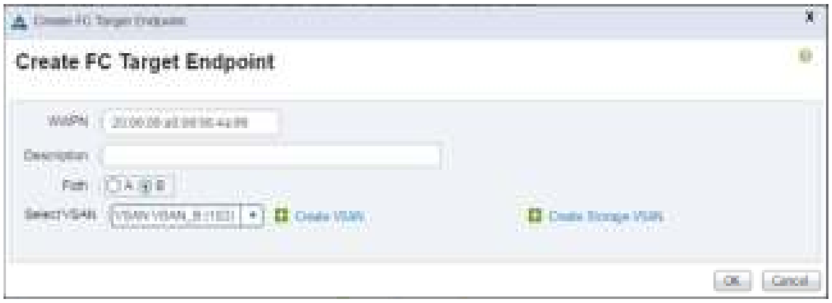

Create FC Target Endpoint

WWPN: [illegible]

Description:

Path: ○ A ● B

Select VSAN: VSAN VSAN_B (102)

23. Click OK.

24. Click the Plus Sign on the right to add a second zoning target.

25. Enter the WWPN for [fcp_lif02_63b or fcp_lif02_62b] from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the network interface show -vserver Infra-SVM command.

26. Select Path B and VSAN_B.

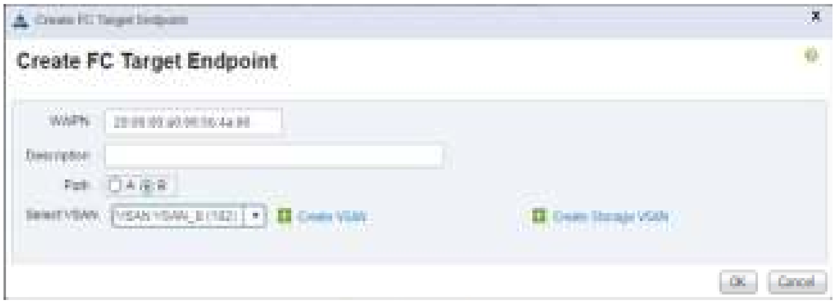

Create FC Target Endpoint

WWPN: [illegible]

Description:

Path: ○ A ● B

Select VSAN: VSAN VSAN_B (102)

27. Click OK.

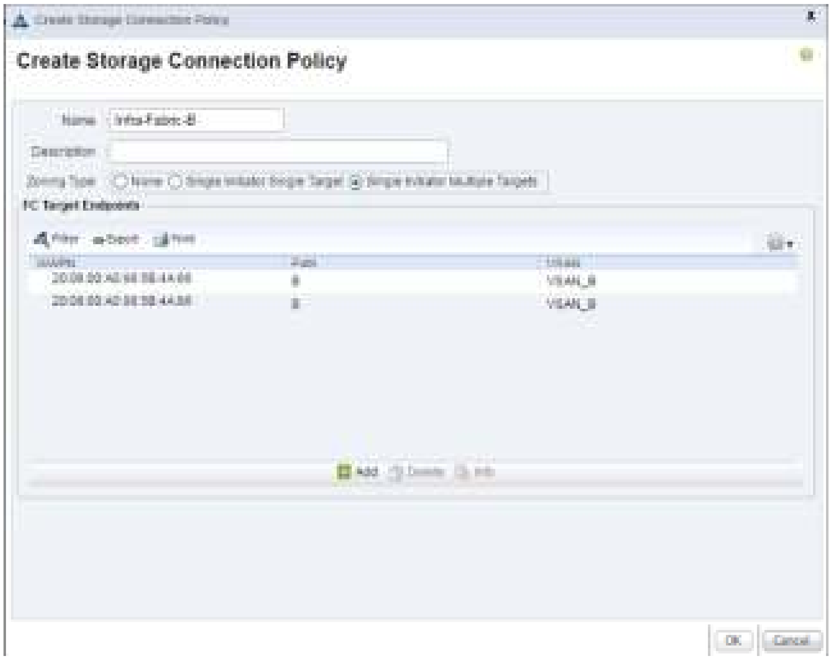

Create Storage Connection Policy

Name: Infra-Fabric-B

Description:

Zoning Type: ○ None ○ Single Initiator Single Target ● Single Initiator Multiple Targets

FC Target Endpoints

WWPN	Path	VSAN
[illegible]	B	VSAN_B
[illegible]	B	VSAN_B

28. Click OK, and then click OK again.

Create vHBA Templates

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter `vHBA_Template_A` as the vHBA template name.
6. Keep Fabric A selected.
7. Select VSAN_A.
8. Leave Initial Template as the Template Type.
9. Select `WWPN_Pool_A` as the WWPN Pool.

10. Click OK to create the vHBA template.
11. Click OK.

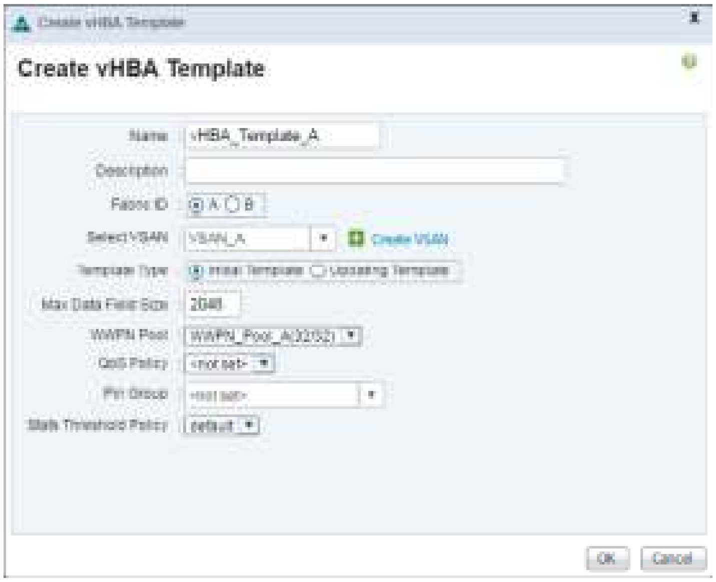

Create vHBA Template

Name: vHBA_Template_A
Description:
Fabric ID: ● A ○ B
Select VSAN: VSAN_A
Template Type: ● Initial Template ○ Updating Template
Max Data Field Size: 2048
WWPN Pool: WWPN_Pool_A(32/32)
QoS Policy: <not set>
Pin Group: <not set>
Stats Threshold Policy: default

12. Right-click vHBA Templates.
13. Select Create vHBA Template.
14. Enter `vHBA_Template_B` as the vHBA template name.
15. Select Fabric B as the Fabric ID.
16. Select VSAN_B.
17. Leave Initial Template as the Template Type.
18. Select `WWPN_Pool_B` as the WWPN Pool.
19. Click OK to create the vHBA template.
20. Click OK.

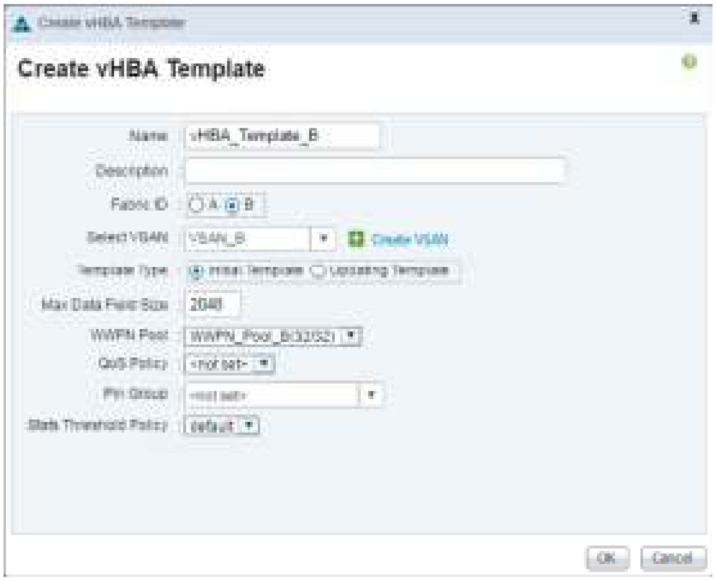

Create vHBA Template

Name: vHBA_Template_B
Description:
Fabric ID: ○ A ● B
Select VSAN: VSAN_B
Template Type: ● Initial Template ○ Updating Template
Max Data Field Size: 2048
WWPN Pool: WWPN_Pool_B(32/32)
QoS Policy: <not set>
Pin Group: <not set>
Stats Threshold Policy: default

Create SAN Connectivity Policy

During testing, iSCSI vNICs and FC vHBAs were deployed to all hosts regardless of boot policy used. The FC vHBA interfaces can be left out for environments configured to utilize iSCSI for boot and data.

To configure the necessary Infrastructure SAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > Policies > root.
3. Right-click SAN Connectivity Policies.
4. Select Create SAN Connectivity Policy.
5. Enter Infra-SAN-Policy as the name of the policy.
6. Select the previously created WWNN_Pool for the WWNN Assignment.

7. Click the Add button at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter Fabric-A as the name of the vHBA.
9. Select the Use vHBA Template checkbox.
10. In the vHBA Template list, select vHBA_Template_A.

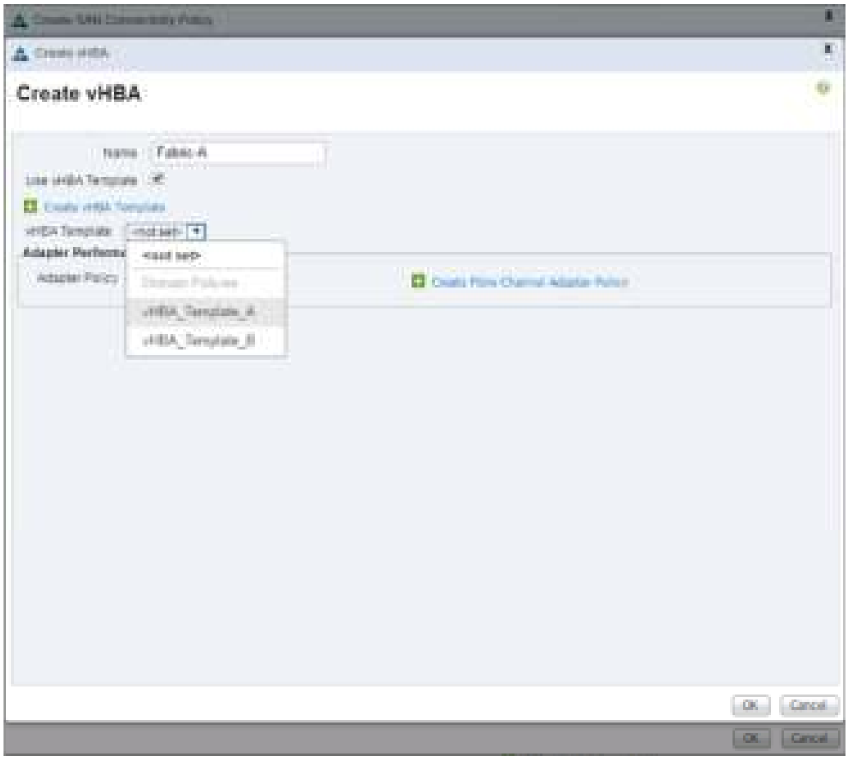

Create vHBA

Name: Fabric-A

Use vHBA Template: ☑

vHBA Template: <not set>

11. In the Adapter Policy list, select VMWare.
12. Click OK.
13. Click the Add button at the bottom to add a second vHBA.
14. In the Create vHBA dialog box, enter Fabric-B as the name of the vHBA.
15. Select the Use vHBA Template checkbox.
16. In the vHBA Template list, select vHBA_Template_B.

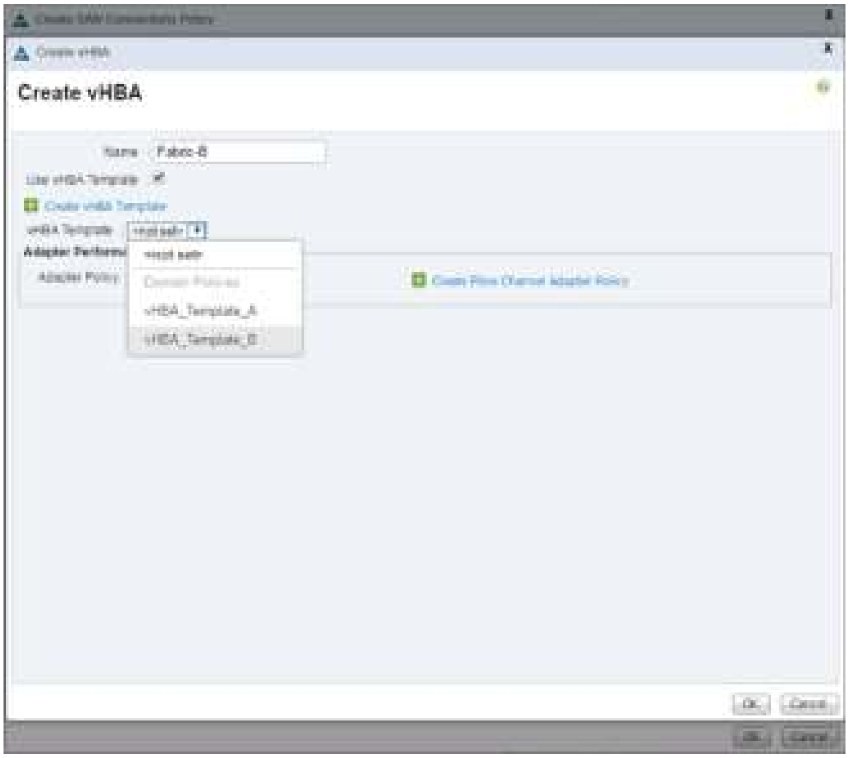

Create vHBA

Name: Fabric-B

Use vHBA Template: ☑

vHBA Template: <not set>

17. In the Adapter Policy list, select VMWare.
18. Click OK.

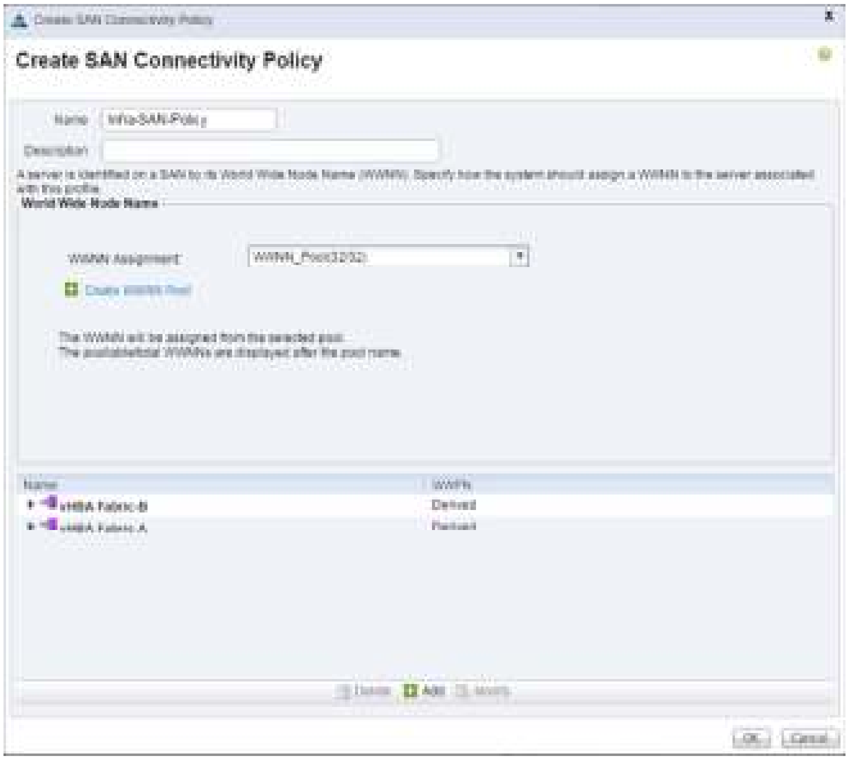

Create SAN Connectivity Policy

Name: Infra-SAN-Policy

Description:

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: WWNN_Pool(32/32)

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

Name	WWPN
vHBA Fabric-B	Derived
vHBA Fabric-A	Derived

19. Click OK to create the SAN Connectivity Policy.
20. Click OK to confirm creation.

Create MAC Address Pools

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

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

In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC_Pool_A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select **Sequential** as the option for Assignment Order.

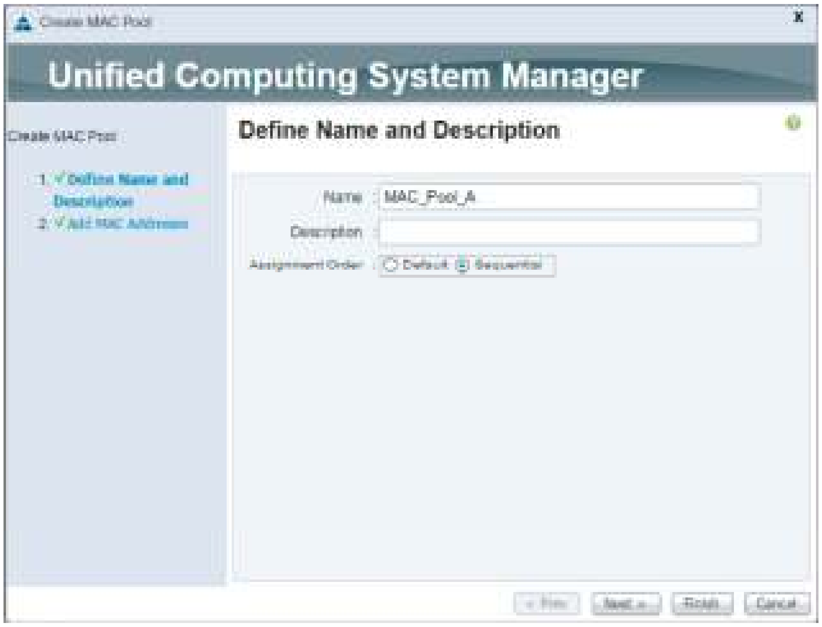

Define Name and Description

Name: MAC_Pool_A

Description:

Assignment Order: ○ Default ● Sequential

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

For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the of also embedding the extra building, floor and UCS domain number information giving us `00:25:B5:91:1A:00` as our first MAC address.

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

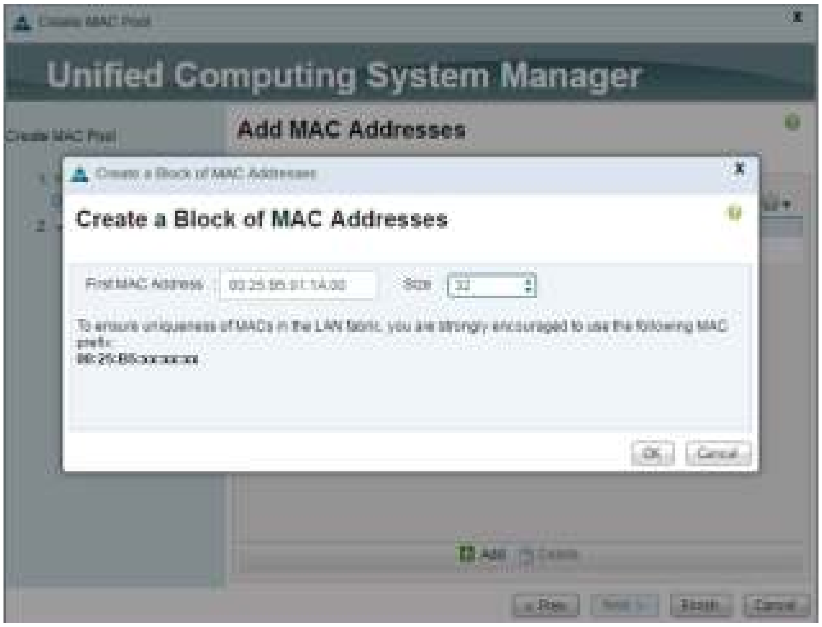

Create a Block of MAC Addresses

First MAC Address: 00:25:B5:91:1A:00 Size: 32

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

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

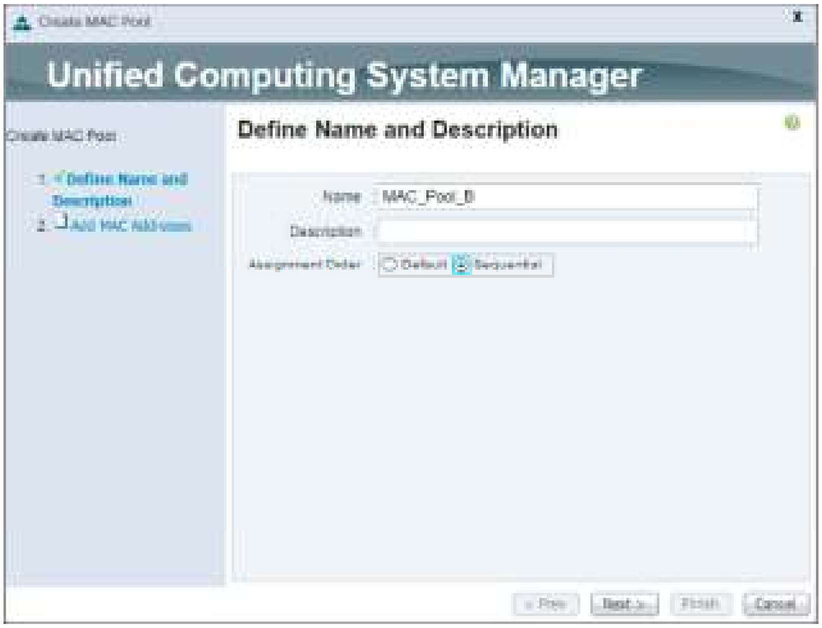

Define Name and Description

Name: MAC_Pool_B

Description:

Assignment Order: ○ Default ● Sequential

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

For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the extra building, floor and UCS domain number information giving us `00:25:B5:91:1B:00` as our first MAC address.

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

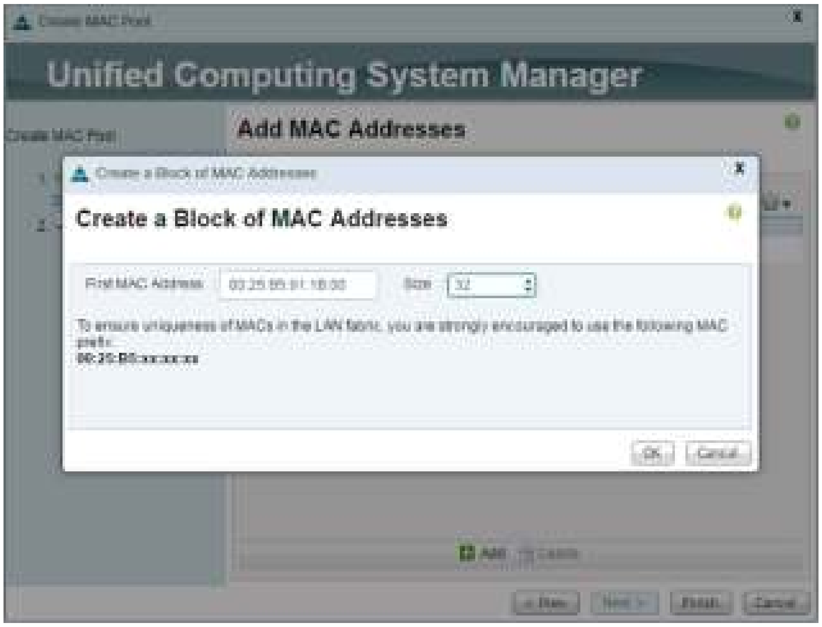

Create a Block of MAC Addresses

First MAC Address: 00:25:B5:91:1B:00 Size: 32

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

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

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps.

1. In the UCS Manager, select the SAN tab on the left.
2. Select Pools > root.
3. Right-click IQN Pools under the root organization.

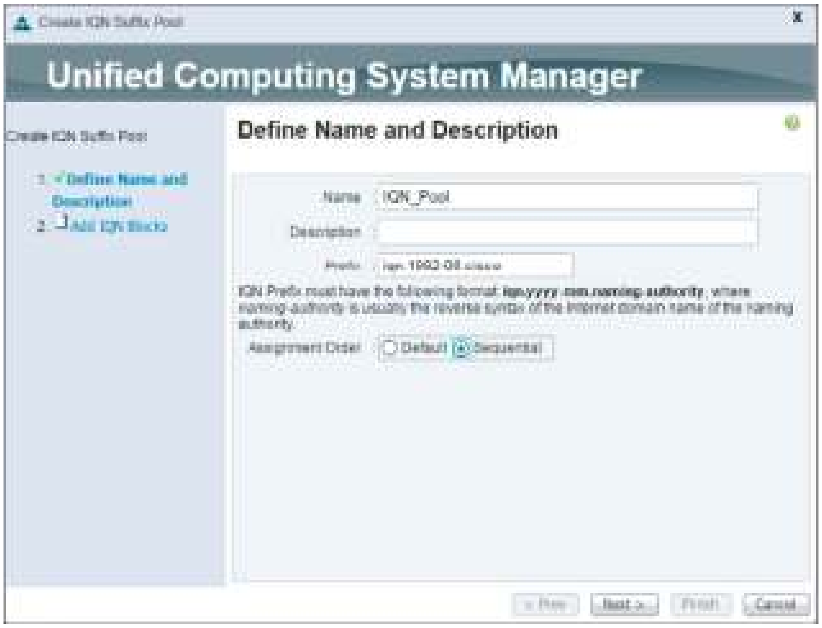

Define Name and Description

Name: IQN_Pool

Description:

Prefix: iqn.1992-08.cisco

IQN Prefix must have the following format: **iqn.yyyy-mm.naming-authority**, where naming-authority is usually the reverse syntax of the Internet domain name of the naming authority.

Assignment Order: ○ Default ● Sequential

4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN_Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter iqn.1992-08.com.cisco as the prefix
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter ucs-host as the suffix.
12. Enter 1 in the From field.
13. Specify a size of the IQN block sufficient to support the available server resources.
14. Click OK.

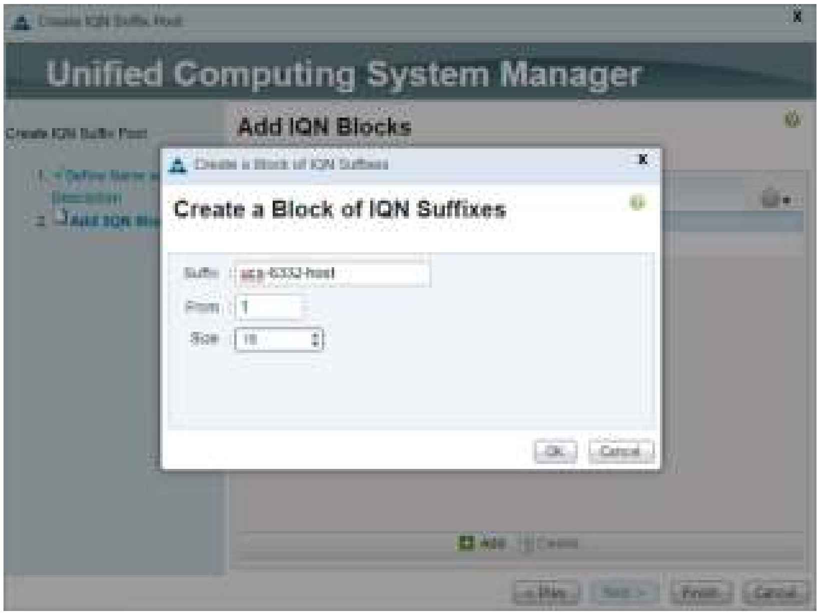

Create a Block of IQN Suffixes

Suffix: ucs-6332-host

From: 1

Size: 16

15. Click Finish.
16. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot

These steps provide details for configuring the necessary IP pools iSCSI boot for the Cisco UCS environment.

1. In Cisco UCS Manager, select the LAN tab on the left.
2. Select Pools > root.

> Two IP pools are created, one for each switching fabric.

3. Right-click IP Pools under the root organization.
4. Select Create IP Pool to create the IP pool.
5. Enter iSCSI_IP_Pool_A for the name of the IP pool.
6. Optional: Enter a description of the IP pool.
7. Select Sequential for Assignment Order.

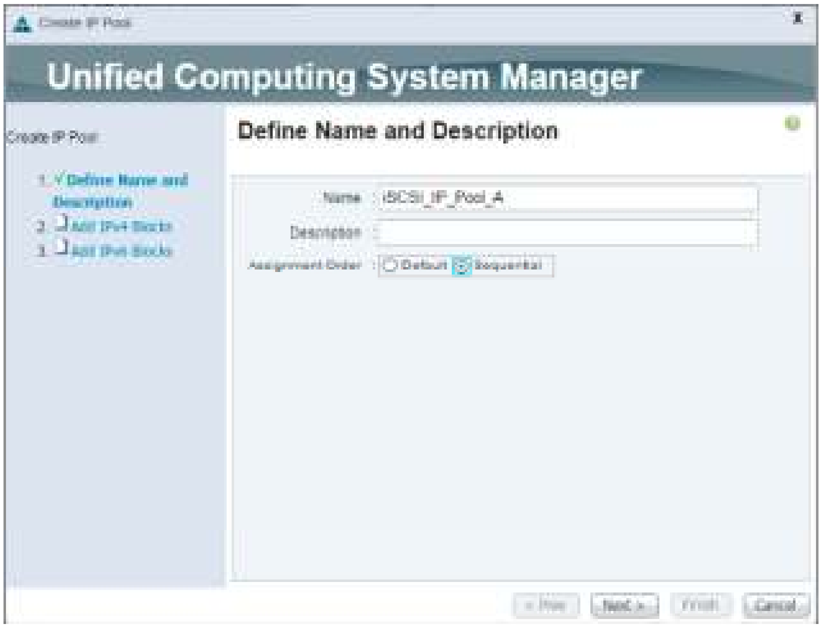

Define Name and Description

Name: iSCSI_IP_Pool_A

Description:

Assignment Order: ○ Default ● Sequential

8. Click Next.
9. Click Add.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Set the size to enough addresses to accommodate the servers.

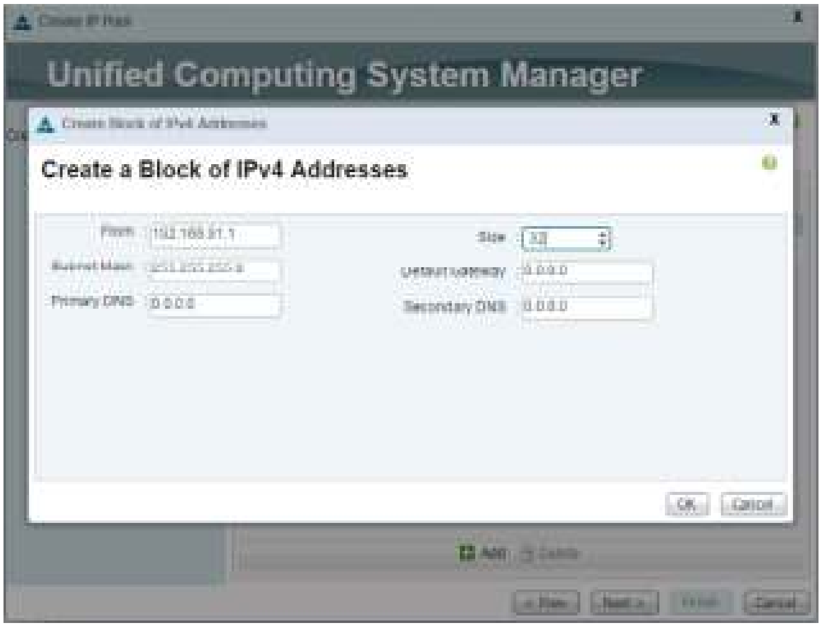
Create IP Pool
Unified Computing System Manager
Create Block of IPv4 Addresses
OK
Cancel
Add
Delete
< Prev
Next >
Finish
Cancel

Create a Block of IPv4 Addresses

From: 192.168.91.1

Size: 32

Subnet Mask: 255.255.255.0

Default Gateway: 0.0.0.0

Primary DNS: 0.0.0.0

Secondary DNS: 0.0.0.0

12. Click OK.
13. Click Finish.
14. Repeat these steps for an iSCSI_IP_Pool_B:
15. Right-click IP Pools under the root organization.
16. Select Create IP Pool to create the IP pool.
17. Enter iSCSI_IP_Pool_B for the name of the IP pool.
18. Optional: Enter a description of the IP pool.
19. Select Sequential for Assignment Order.
20. Click Next.

21. Click Add.
22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
23. Set the size to enough addresses to accommodate the servers.
24. Click OK.
25. Click Finish.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID_Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select **Sequential** for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

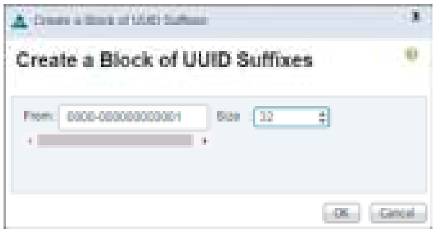

Create a Block of UUID Suffixes

From: 0000-000000000001 Size: 32

13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter Infra_Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra_Pool server pool.

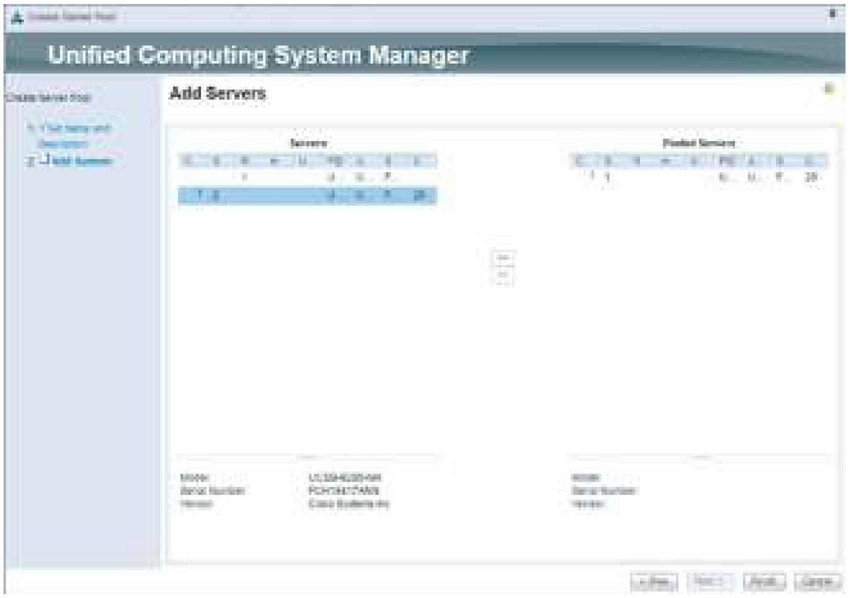

9. Click Finish.
10. Click OK.

Create VLANs

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

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

> In this procedure, four unique VLANs are created. See Table 2.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.

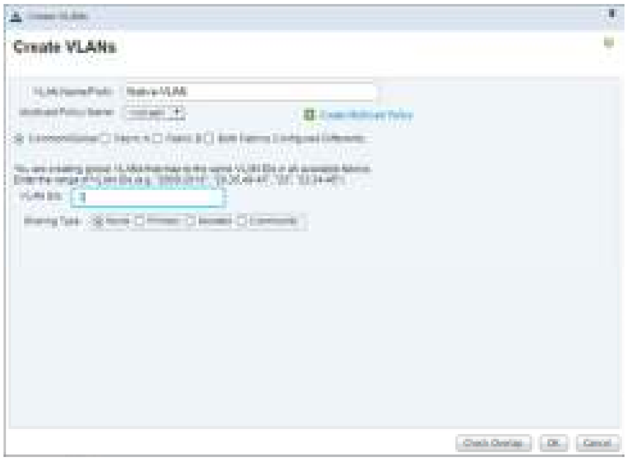

10. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs.
14. Enter iSCSI-A-VLAN as the name of the VLAN to be used for the first iSCSI VLAN.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the VLAN ID for the first iSCSI VLAN.
17. Click OK, then OK.

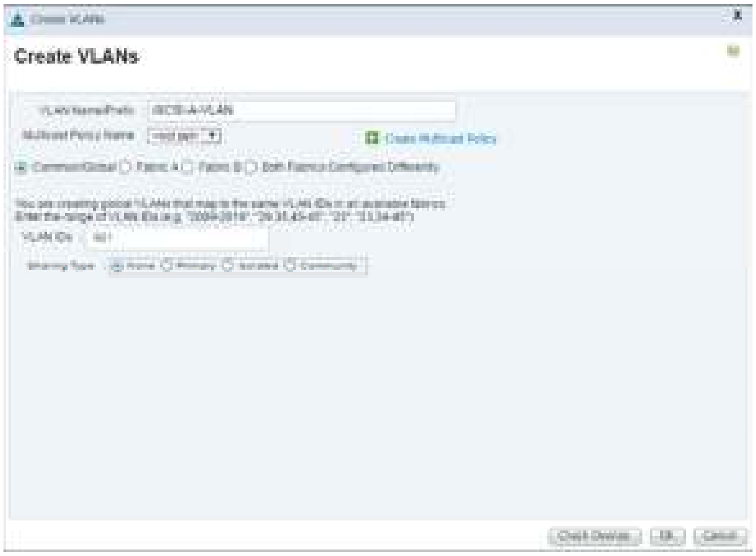

Create VLANs

VLAN Name/Prefix : iSCSI-A-VLAN

Multicast Policy Name : <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs : 901

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

18. Right-click VLANs.
19. Select Create VLANs.
20. Enter iSCSI-B-VLAN as the name of the VLAN to be used for the second iSCSI VLAN.
21. Keep the Common/Global option selected for the scope of the VLAN.
22. Enter the VLAN ID for the second iSCSI VLAN.
23. Click OK, then OK.

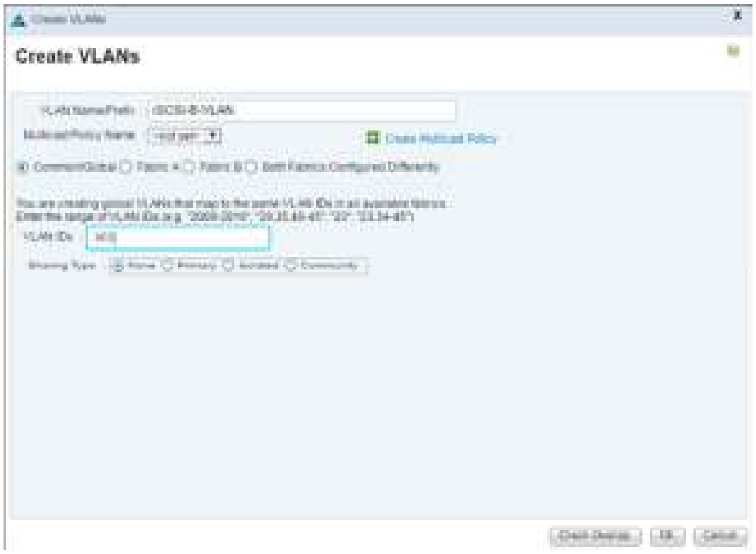

Create VLANs

VLAN Name/Prefix : iSCSI-B-VLAN

Multicast Policy Name : <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs : [illegible]

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

24. Right-click VLANs.
25. Select Create VLANs
26. Enter IB-Mgmt as the name of the VLAN to be used for management traffic.
27. Keep the Common/Global option selected for the scope of the VLAN.
28. Enter the In-Band management VLAN ID.
29. Keep the Sharing Type as None.
30. Click OK, and then click OK again.
31. Right-click VLANs.
32. Select Create VLANs.
33. Enter Infra-NFS as the name of the VLAN to be used for NFS.
34. Keep the Common/Global option selected for the scope of the VLAN.
35. Enter the NFS VLAN ID.
36. Keep the Sharing Type as None.
37. Click OK, and then click OK again.
38. Right-click VLANs.
39. Select Create VLANs.
40. Enter vMotion as the name of the VLAN to be used for vMotion.

41. Keep the Common/Global option selected for the scope of the VLAN.
42. Enter the vMotion VLAN ID.
43. Keep the Sharing Type as None.
44. Click OK, and then click OK again.
45. Right-click VLANs.
46. Select Create VLANs.
47. Enter VM-Traffic as the name of the VLAN to be used for VM Traffic.
48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VM-Traffic VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK, and then click OK again.

Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.1(1h) for both the Blade and Rack Packages.
7. Leave M-Series Package as <not set> and leave Excluded Components with only Local Disk selected.

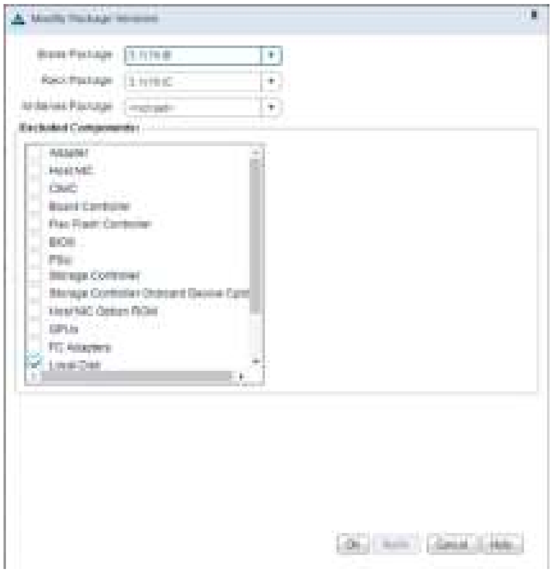

Modify Package Versions

Blade Package: 3.1(1h)B

Rack Package: 3.1(1h)C

M-Series Package: <not set>

Excluded Components:

☐ Adapter
☐ Host NIC
☐ CIMC
☐ Board Controller
☐ Flex Flash Controller
☐ BIOS
☐ PSU
☐ Storage Controller
☐ Storage Controller Onboard Device Cpld
☐ Host NIC Option ROM
☐ GPUs
☐ FC Adapters
☑ Local Disk

8. Click OK to modify the host firmware package.

Set Jumbo Frames in Cisco UCS Fabric

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes in the bottom of the window.
6. Click OK

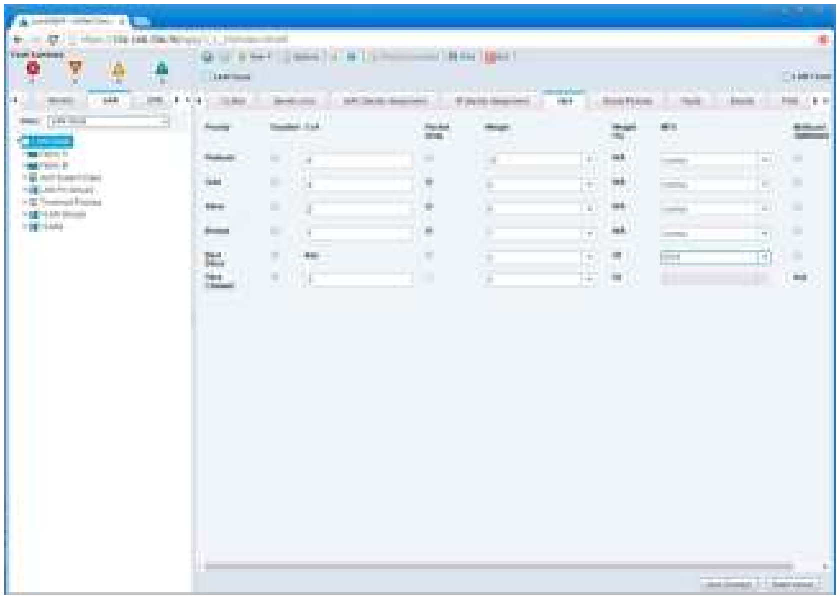

Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.

This policy should not be used on servers that contain local disks.

To create a local disk configuration policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter SAN-Boot as the local disk configuration policy name.
6. Change the mode to No Local Storage.
7. Click OK to create the local disk configuration policy.

Create Local Disk Configuration Policy

Name: SAN-Boot

Mode: No Local Storage

8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable_CDP as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.

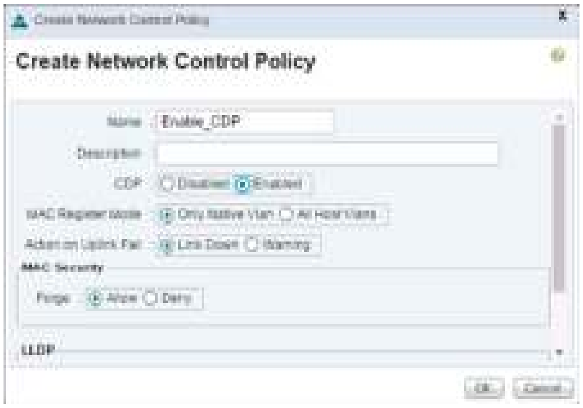

Create Network Control Policy

Name: Enable_CDP

Description:

CDP: ○ Disabled ● Enabled

MAC Register Mode: ● Only Native Vlan ○ All Host Vlans

Action on Uplink Fail: ● Link Down ○ Warning

MAC Security

Forge: ● Allow ○ Deny

LLDP

8. Click OK.

Create Power Control Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

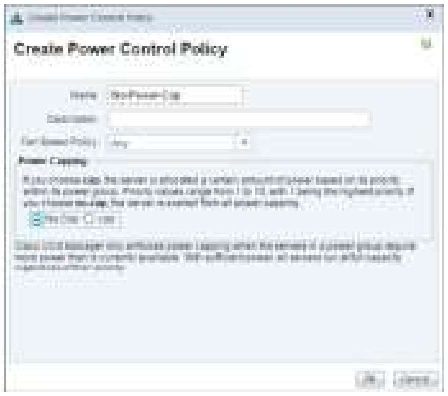

Create Power Control Policy

Name: No-Power-Cap

Description:

Fan Speed Policy: Any

Power Capping

If you choose cap, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose no-cap, the server is exempt from all power capping.

● No Cap ○ cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

Create Server Pool Qualification Policy (Optional)

To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:

This example creates a policy for Cisco UCS B-Series and Cisco UCS C-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCS-Broadwell.
6. Select Create CPU/Core Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Select UCS-CPU-E52660E as the PID.
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy then OK for the confirmation.

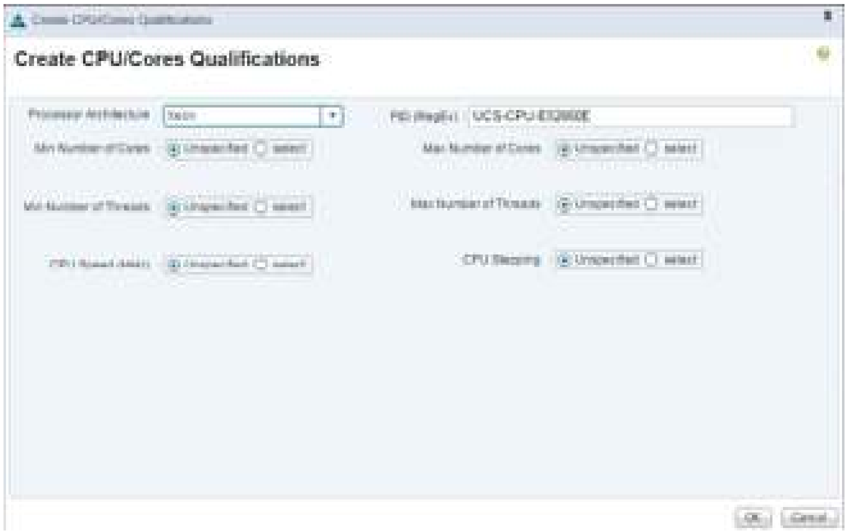

Create CPU/Cores Qualifications

Processor Architecture: Xeon

PID (RegEx): UCS-CPU-E52660E

Min Number of Cores: ● Unspecified ○ select

Max Number of Cores: ● Unspecified ○ select

Min Number of Threads: ● Unspecified ○ select

Max Number of Threads: ● Unspecified ○ select

CPU Speed (MHz): ● Unspecified ○ select

CPU Stepping: ● Unspecified ○ select

Create Server BIOS Policy

To create a server BIOS 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.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter VM-Host-Infra as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.
8. Click Finish to create the BIOS policy.

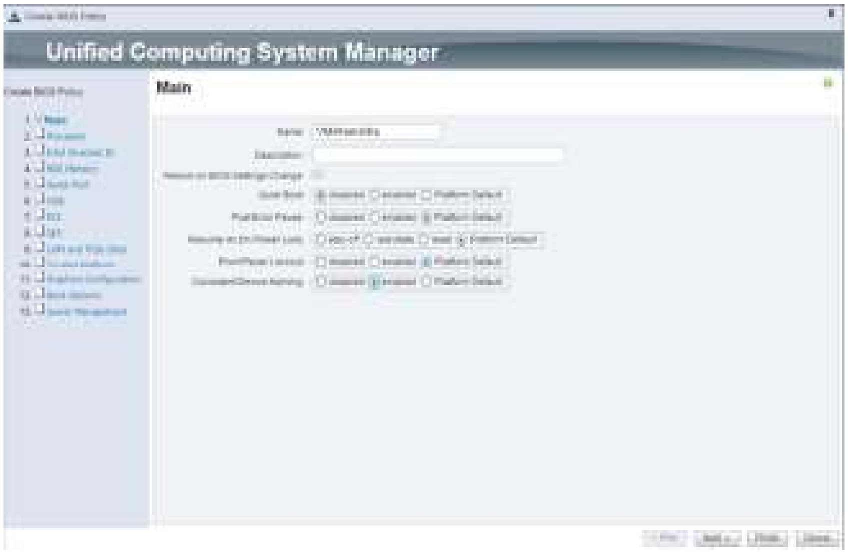

Main

Name: VM-Host-Infra

Description:

Reboot on BIOS Settings Change: □

Quiet Boot: ● disabled ○ enabled ○ Platform Default

Post Error Pause: ○ disabled ○ enabled ● Platform Default

Resume on AC Power Loss: ○ stay-off ○ last-state ○ reset ● Platform Default

Front Panel Lockout: ○ disabled ○ enabled ● Platform Default

Consistent Device Naming: ○ disabled ● enabled ○ Platform Default

9. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter VM-Host-Infra as the name of the placement policy.
6. Click 1 and select Assigned Only.
7. Click OK, and then click OK again.

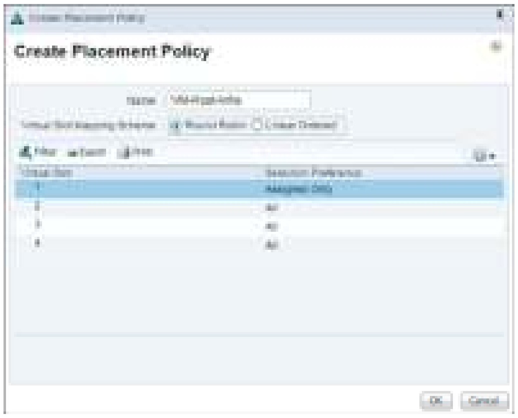

Create Placement Policy

Name: VM-Host-Infra

Virtual Slot Mapping Scheme: ● Round Robin ○ Linear Ordered

Virtual Slot	Selection Preference
1	Assigned Only
2	All
3	All
4	All

Update the 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.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. (Optional: Click “On Next Boot” to delegate maintenance windows to server owners)

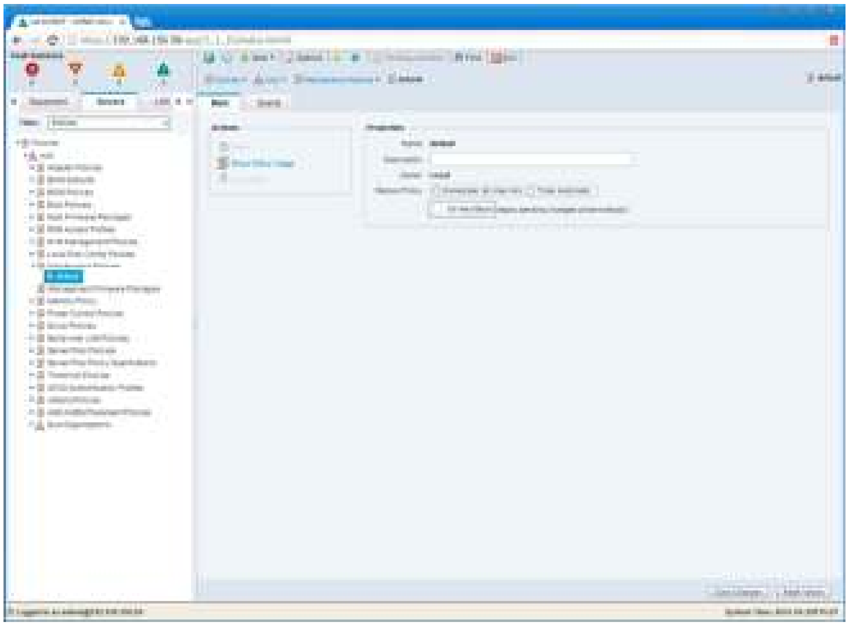

6. Click Save Changes.
7. Click OK to accept the change.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 4 vNIC Templates will be created.

The same Infra VLANs were used on both the infrastructure (Infra) and production (Prod) hosts deployed in this environment. If production networks were to differ in the VLANs used for infrastructure networks, differing vNIC Templates should be created for each.

Create Data vNICs

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.

10. Under VLANs, select the checkboxes for IB-MGMT, Infra-NFS, Native-VLAN,VM-Traffic, and vMotion VLANs.

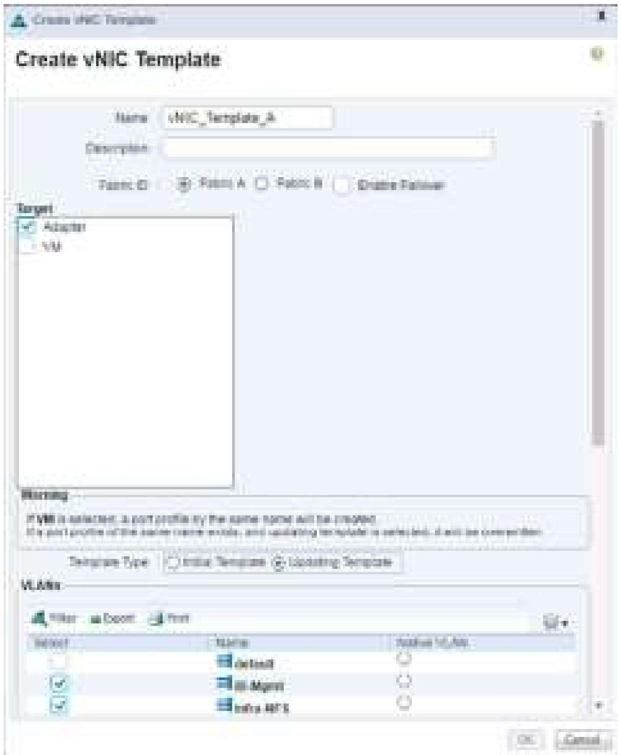

Create vNIC Template

Name: vNIC_Template_A

Description:

Fabric ID: ● Fabric A ○ Fabric B □ Enable Failover

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

Select	Name	Native VLAN
□	default	○
☑	IB-Mgmt	○
☑	Infra-NFS	○

11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_A.
14. In the Network Control Policy list, select Enable_CDP.

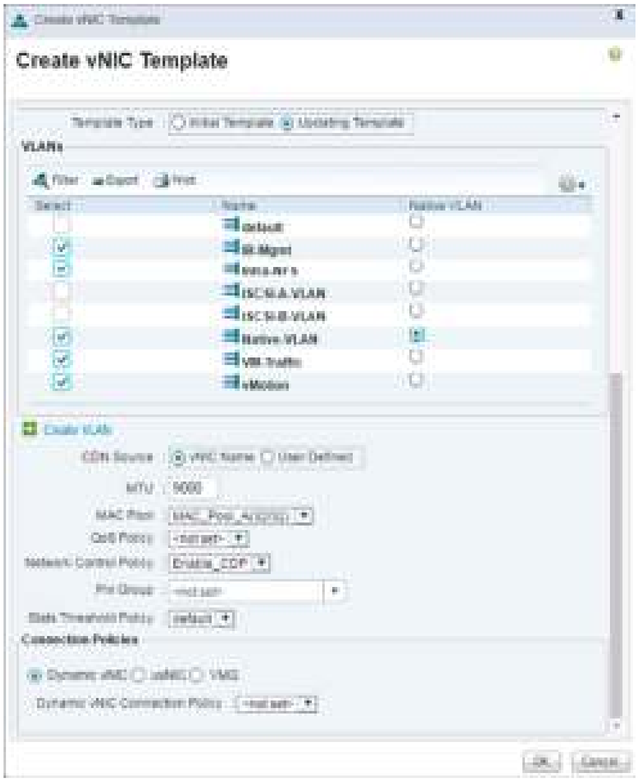

Create vNIC Template

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
□	default	○
☑	IB-Mgmt	○
☑	Infra-NFS	○
□	iSCSI-A-VLAN	○
□	iSCSI-B-VLAN	○
☑	Native-VLAN	●
☑	VM-Traffic	○
☑	vMotion	○

CDN Source: ● vNIC Name ○ User Defined

MTU: 9000

MAC Pool: MAC_Pool_A(32/32)

QoS Policy: <not set>

Network Control Policy: Enable_CDP

Pin Group: <not set>

Stats Threshold Policy: default

Connection Policies

● Dynamic vNIC ○ usNIC ○ VMQ

Dynamic vNIC Connection Policy: <not set>

15. Click OK to create the vNIC template.
16. Click OK.

Repeat these equivalent steps for vNIC_Template_B:

1. In the navigation pane, select the LAN tab.

2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter vNIC_Template_B as the vNIC template name.
6. Select Fabric B.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure the VM checkbox is not selected.
9. Select Updating Template as the template type.
10. Under VLANs, select the checkboxes for IB-MGMT, INFRA-NFS, Native-VLAN, and vMotion VLANs.

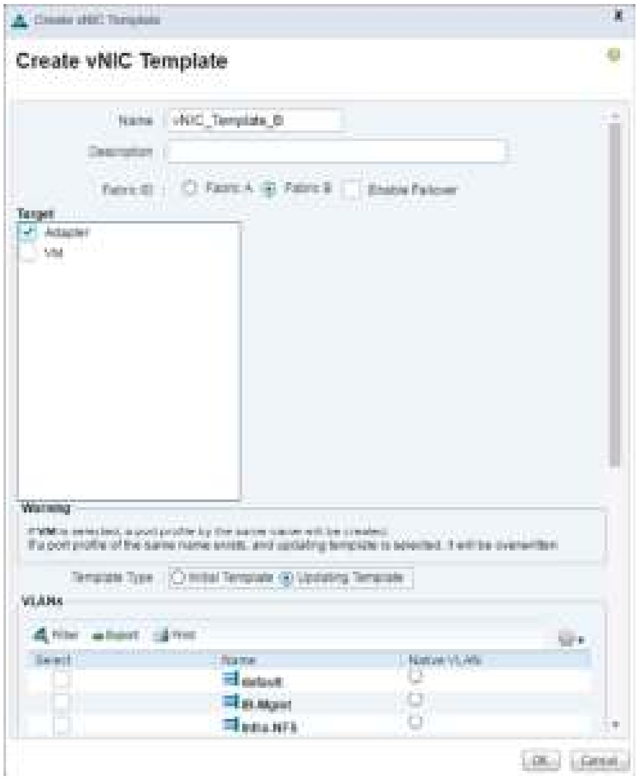

Create vNIC Template

Name: vNIC_Template_B

Description:

Fabric ID: ○ Fabric A ● Fabric B □ Enable Failover

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

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Infra-NFS	○

11. Set default as the native VLAN.
12. Select vNIC Name for the CDN Source.
13. For MTU, enter 9000.
14. In the MAC Pool list, select MAC_Pool_B.
15. In the Network Control Policy list, select Enable_CDP.

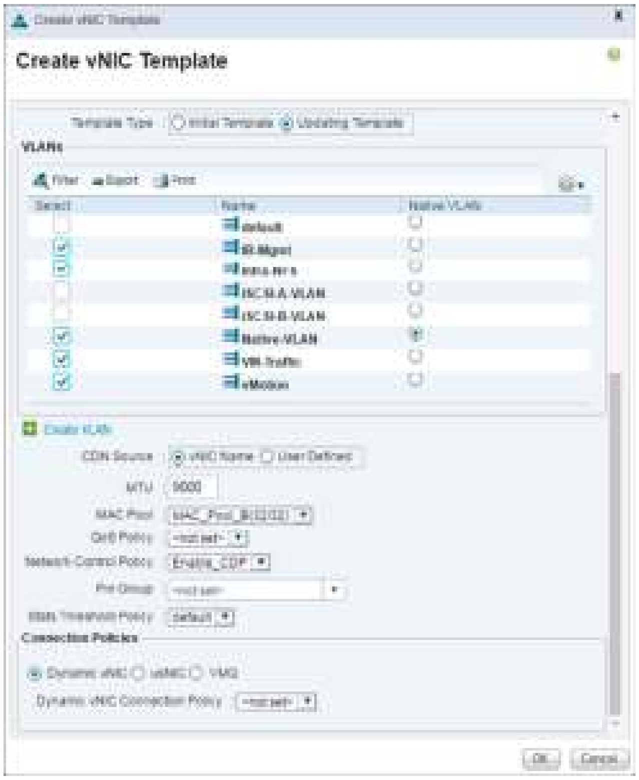

Create vNIC Template

Template Type ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
□	default	○
☑	IB-Mgmt	○
☑	Infra-NFS	○
□	iSCSI-A-VLAN	○
□	iSCSI-B-VLAN	○
☑	Native-VLAN	●
☑	VM-Traffic	○
☑	vMotion	○

CDN Source ● vNIC Name ○ User Defined

MTU 9000

MAC Pool MAC_Pool_B(32/32)

QoS Policy <not set>

Network Control Policy Enable_CDP

Pin Group <not set>

Stats Threshold Policy default

Connection Policies

● Dynamic vNIC ○ usNIC ○ VMQ

Dynamic vNIC Connection Policy <not set>

16. Click OK to create the vNIC template.
17. Click OK.

Create iSCSI vNICs

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter iSCSI_Template_A as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.

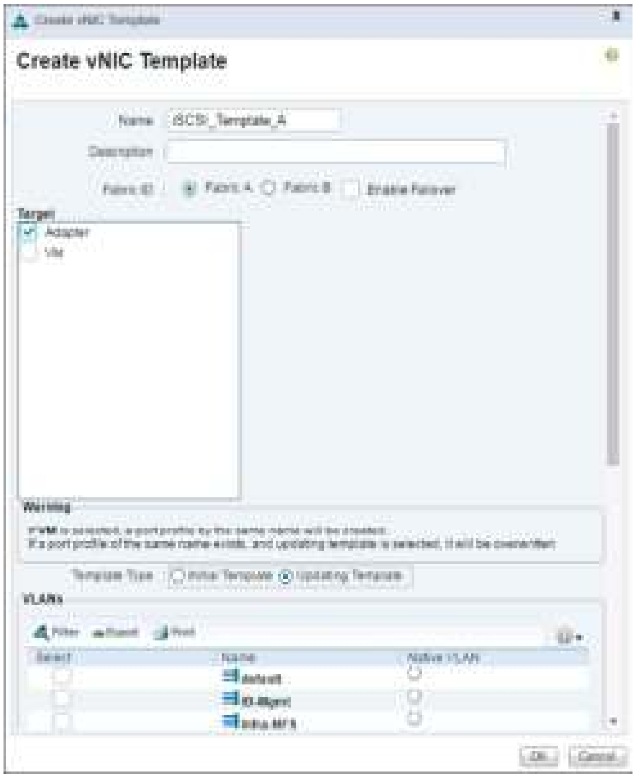

Create vNIC Template

Name: iSCSI_Template_A

Description:

Fabric ID: ● Fabric A ○ Fabric B □ Enable Failover

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

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Infra-NFS	○

9. Under VLANs, select iSCSI-A-VLAN.
10. Set iSCSI-A-VLAN as the native VLAN.
11. Under MTU, enter 9000.
12. From the MAC Pool list, select MAC_Pool_A.
13. From the Network Control Policy list, select Enable_CDP.

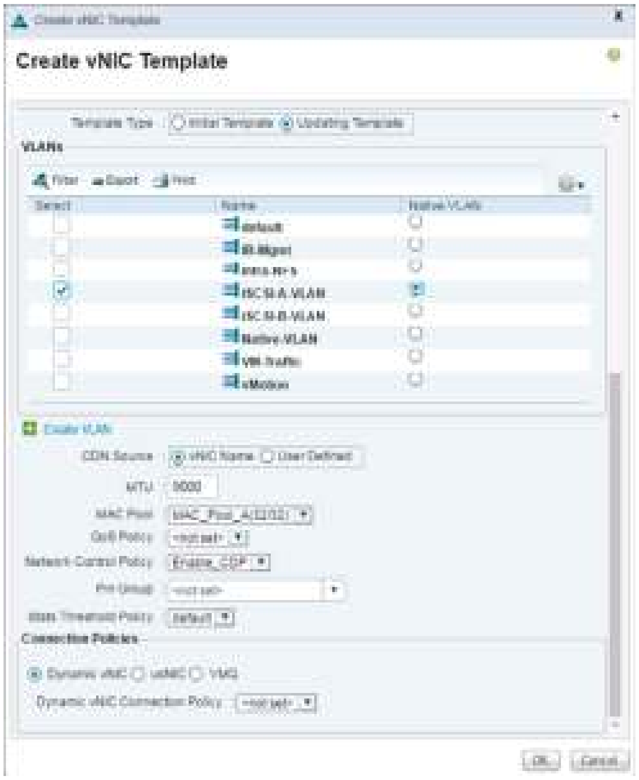

Create vNIC Template

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Infra-NFS	○
☑	iSCSI-A-VLAN	●
□	iSCSI-B-VLAN	○
□	Native-VLAN	○
□	VM-Traffic	○
□	vMotion	○

Create VLAN

CDN Source: ● vNIC Name ○ User Defined

MTU: 9000

MAC Pool: MAC_Pool_A(32/32)

QoS Policy: <not set>

Network Control Policy: Enable_CDP

Pin Group: <not set>

Stats Threshold Policy: default

Connection Policies

● Dynamic vNIC ○ usNIC ○ VMQ

Dynamic vNIC Connection Policy: <not set>

14. Click OK to complete creating the vNIC template.
15. Click OK.

Repeat these equivalent steps for iSCSI_Template_B:

1. Select the LAN tab on the left.

2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter iSCSI_Template_B as the vNIC template name.
6. Select Fabric B. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.

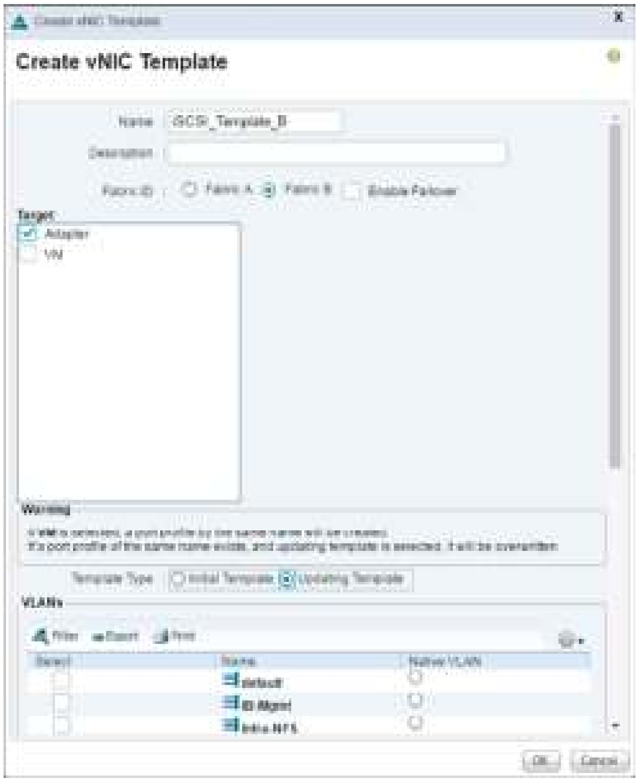

Create vNIC Template

Name: iSCSI_Template_B

Description:

Fabric ID: ○ Fabric A ● Fabric B □ Enable Failover

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

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Infra-NFS	○

9. Under VLANs, select iSCSI-B-VLAN.
10. Set iSCSI-B-VLAN as the native VLAN.
11. Under MTU, enter 9000.
12. From the MAC Pool list, select MAC_Pool_B.
13. From the Network Control Policy list, select Enable_CDP.

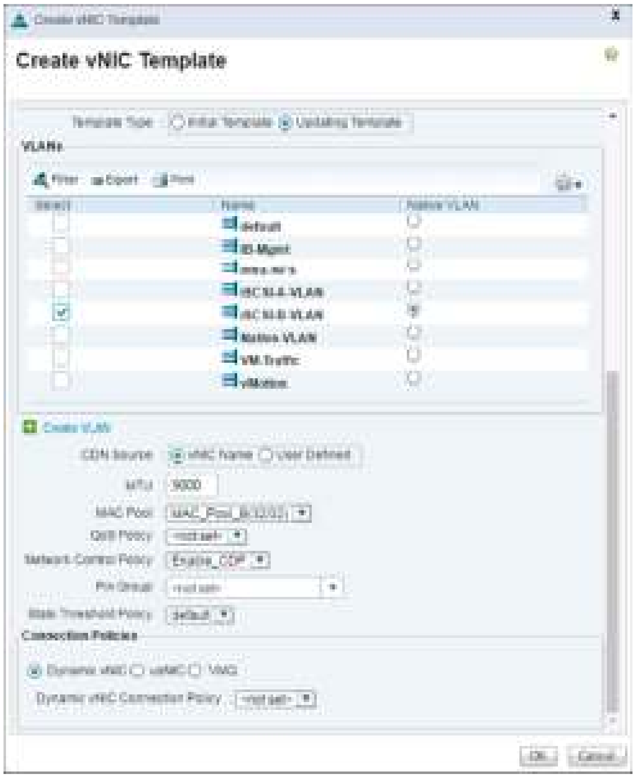

Template Type ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
□	default	○
□	IB-Mgmt	○
□	Infra-NFS	○
□	iSCSI-A-VLAN	○
☑	iSCSI-B-VLAN	●
□	Native-VLAN	○
□	VM-Traffic	○
□	vMotion	○

CDN Source ● vNIC Name ○ User Defined

MTU 9000

MAC Pool MAC_Pool_B(32/32)

QoS Policy <not set>

Network Control Policy Enable_CDP

Pin Group <not set>

Stats Threshold Policy default

Connection Policies

● Dynamic vNIC ○ usNIC ○ VMQ

Dynamic vNIC Connection Policy <not set>

14. Click OK to complete creating the vNIC template.
15. Click OK.

Create LAN Connectivity Policy

During testing, iSCSI vNICs and FC vHBAs were deployed to all hosts regardless of boot policy used. The iSCSI interfaces can be left out for environments configured to utilize fibre channel for boot and data.

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter Infra-LAN-Policy as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter 00-Infra-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select vNIC_Template_A.
10. In the Adapter Policy list, select VMWare.

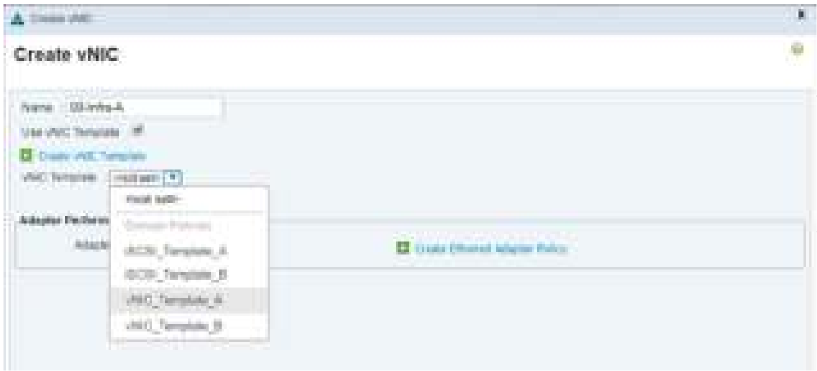

Name 00-Infra-A

Use vNIC Template ☑

vNIC Template <not set>

11. Click OK to add this vNIC to the policy.
12. Click the upper Add button to add another vNIC to the policy.
13. In the Create vNIC box, enter vNIC-01-Infra-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select vNIC_Template_B.
16. In the Adapter Policy list, select VMWare.
17. Click OK to add the vNIC to the policy.

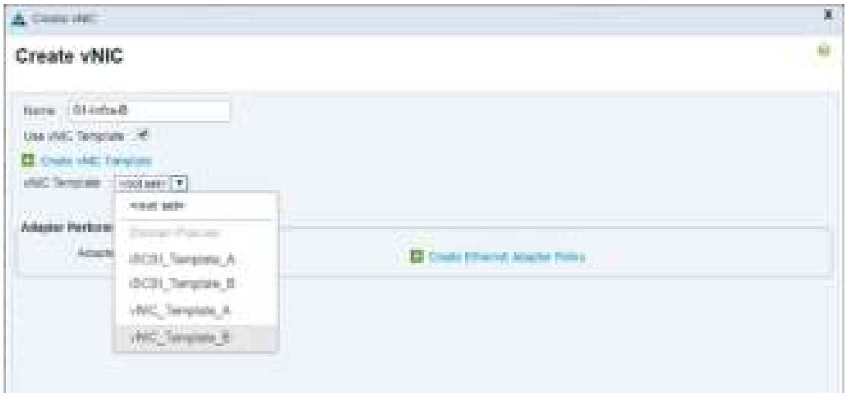
Create vNIC

Name 01-Infra-B

Use vNIC Template ☑

vNIC Template <not set>

18. Click the upper Add button to add a vNIC.
19. In the Create vNIC dialog box, enter vNIC-02-iSCSI-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select iSCSI_Template_A.
22. In the Adapter Policy list, select VMWare.
23. Click OK to add this vNIC to the policy.

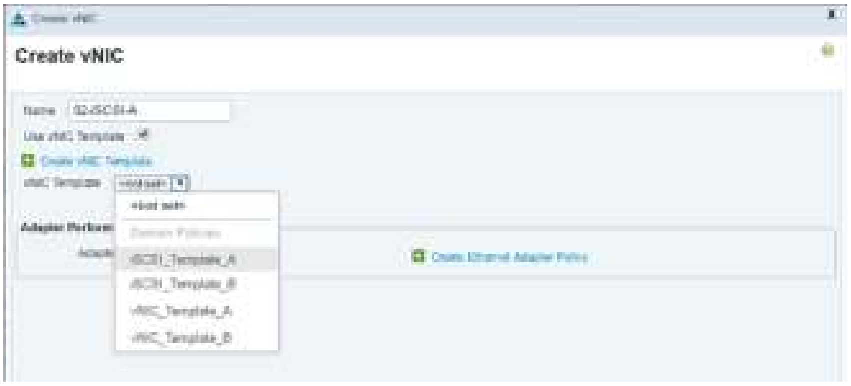
Create vNIC

Name 02-iSCSI-A

Use vNIC Template ☑

vNIC Template <not set>

24. Click the upper Add button to add a vNIC to the policy.
25. In the Create vNIC dialog box, enter vNIC-03-iSCSI-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select iSCSI_Template_B.
28. In the Adapter Policy list, select VMWare.

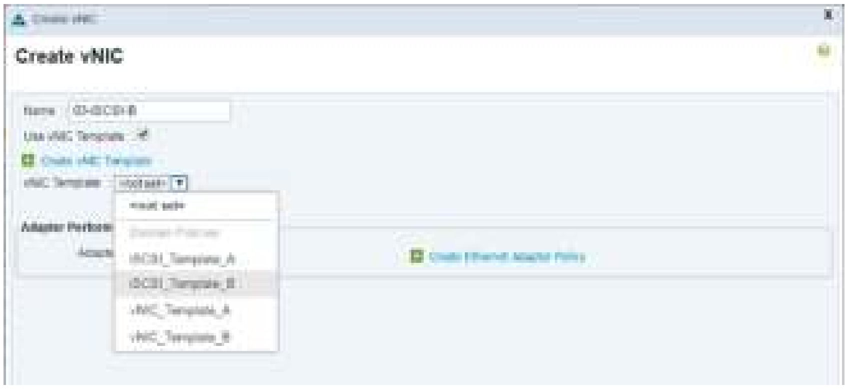
Create vNIC

Name 03-iSCSI-B

Use vNIC Template ☑

vNIC Template <not set>

29. Click OK to add this vNIC to the policy.

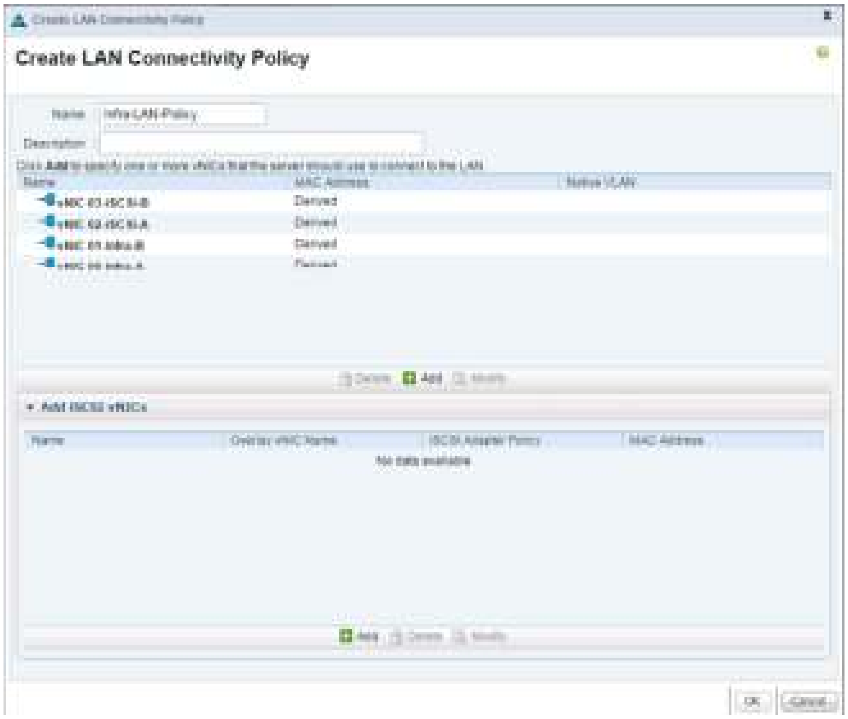

Create LAN Connectivity Policy

Name: Infra-LAN-Policy

Description:

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

Name	MAC Address	Native VLAN
vNIC-03-iSCSI-B	Derived	
vNIC-02-iSCSI-A	Derived	
vNIC-01-Infra-B	Derived	
vNIC-00-Infra-A	Derived	

Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address

No data available

30. Expand the Add iSCSI vNICs section.
31. Click the lower Add button in the iSCSI vNIC section to define an iSCSI boot vNIC.
32. Enter iSCSI-A-vNIC as the name of the vNIC.
33. Select vNIC-02-iSCSI-A for Overlay vNIC.
34. Set the iSCSI Adapter Policy to default.
35. Set the VLAN to Infra-iSCSI-A.
36. Leave the MAC Address set to None.

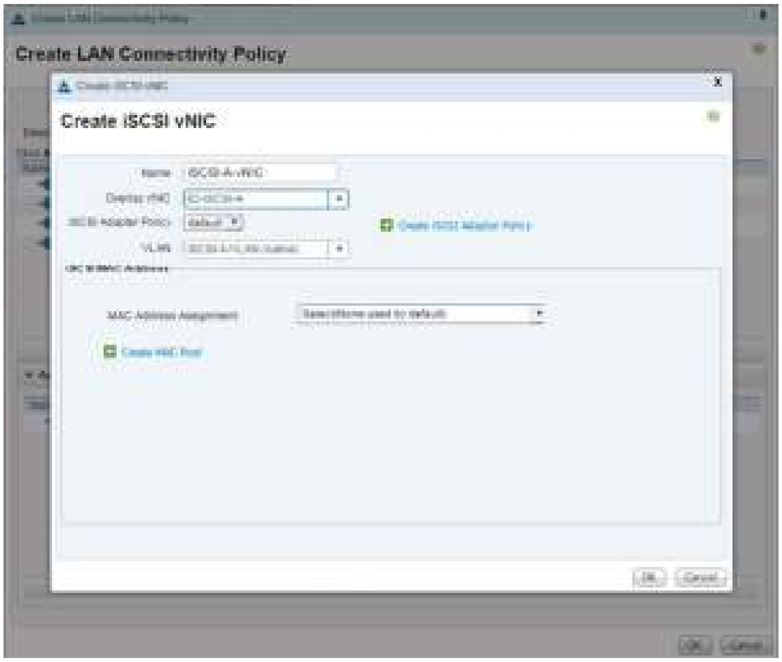

Create iSCSI vNIC

Name: iSCSI-A-vNIC

Overlay vNIC: 02-iSCSI-A

iSCSI Adapter Policy: default

VLAN: [illegible] (native)

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

37. Click OK.
38. Click the lower Add button in the iSCSI vNIC section to define an iSCSI boot vNIC.
39. Enter iSCSI-B-vNIC as the name of the vNIC.
40. Set the Overlay vNIC to vNIC-03-iSCSI-B.
41. Set the iSCSI Adapter Policy to default.
42. Set the VLAN to Infra-iSCSI-B.
43. Leave the MAC Address set to None.

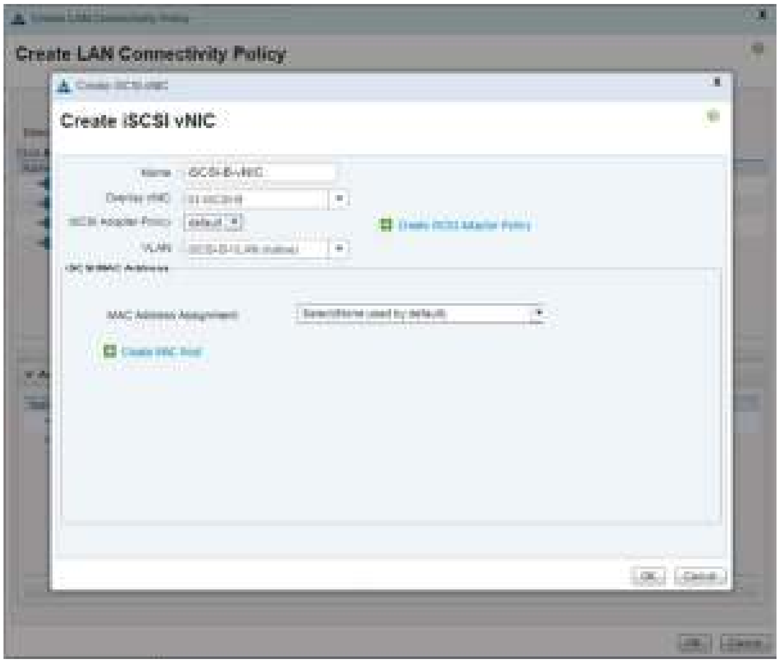

Create LAN Connectivity Policy

Create iSCSI vNIC

Name: iSCSI-B-vNIC

Overlay vNIC: 03-iSCSI-B

iSCSI Adapter Policy: default

VLAN: iSCSI-B-VLAN (native)

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

44. Click OK.

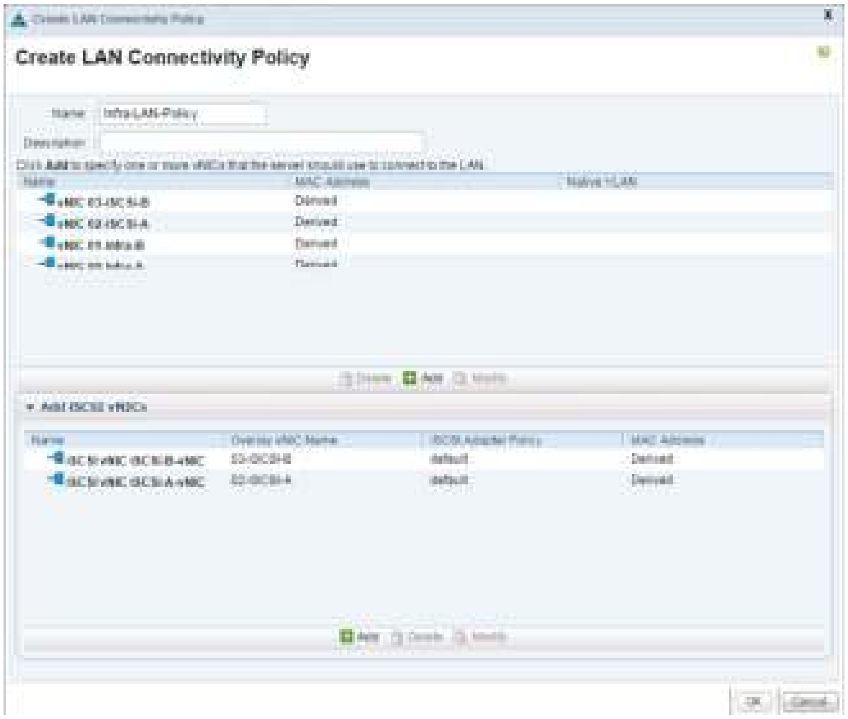

Create LAN Connectivity Policy

Name: Infra-LAN-Policy

Description:

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

Name	MAC Address	Native VLAN
vNIC 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-B-vNIC	03-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-A-vNIC	02-iSCSI-A	default	Derived

45. Click OK to create the LAN Connectivity Policy.

Create vMedia Policy for VMware ESXi 6.0 U1b Install Boot

In the NetApp Data ONTAP setup steps an HTTP web server is required, which will be used for hosting NetApp Data ONTAP as well as VMware software. The vMedia Policy created here will map the VMware ESXi 6.0u1b ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, complete the following steps:

1. In Cisco UCS Manager, select the Servers tab.
2. Select Policies > root.
3. Right-click vMedia Policies.
4. Select Create vMedia Policy.
5. Name the policy ESXi-6.0U1b-HTTP.
6. Enter “Mounts Cisco Custom ISO for ESXi 6.0U1b” in the Description field.
7. Click Add.
8. Name the mount ESXi-6.0U1b-HTTP.
9. Select the CDD Device Type.
10. Select the HTTP Protocol.
11. Enter the IP Address of the web server.

Since DNS server IPs were not entered into the KVM IP earlier, it is necessary to enter the IP of the web server instead of the host name.pool

12. Enter Vmware-ESXi-6.0.0-3380124-Custom-Cisco-6.0.1.2.iso as the Remote File name.
13. Enter the web server path to the ISO file in the Remote Path field.

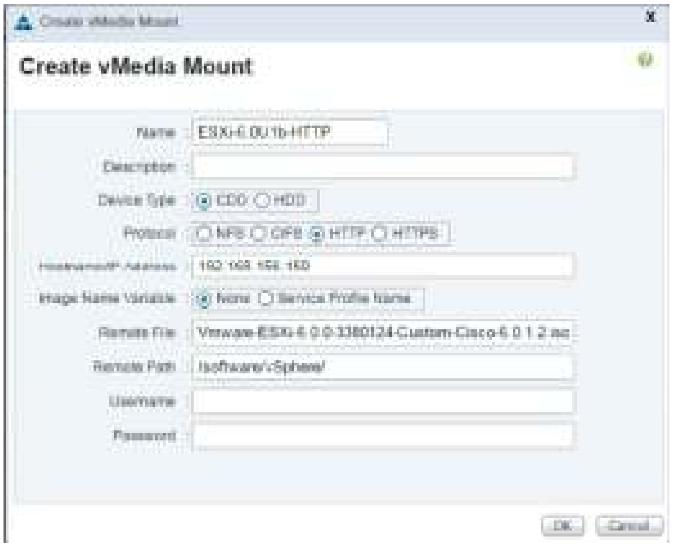

Create vMedia Mount

Create vMedia Mount

Name: ESXi-6.0U1b-HTTP

Description:

Device Type: ● CDD ○ HDD

Protocol: ○ NFS ○ CIFS ● HTTP ○ HTTPS

Hostname/IP Address: 192.168.156.150

Image Name Variable: ● None ○ Service Profile Name

Remote File: Vmware-ESXi-6.0.0-3380124-Custom-Cisco-6.0.1.2.iso

Remote Path: /software/vSphere/

Username:

Password:

14. Click OK to create the vMedia Mount.
15. Click OK then OK again to complete creating the vMedia Policy.

For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi host. On first boot the host will boot into the ESXi installer. After ESXi is installed, the vMedia will not be referenced as long as the boot disk is accessible.

Create Boot Policies **(iSCSI Boot)**

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi lif01a and iscsi lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi lif02a and iscsi lif02b). One boot policy is configured in this procedure. This policy configures the primary target to be iscsi_lif01a.

To create boot policies 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.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter Boot-iSCSI-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Add Remote CD/DVD`.
9. Expand the iSCSI vNICs section and select Add iSCSI Boot.
10. In the Add iSCSI Boot dialog box, enter iSCSI-A-vNIC.
11. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
14. Click OK.
15. Expand CIMC Mounted vMedia.

16. Select Add CIMC Mounted CD/DVD.

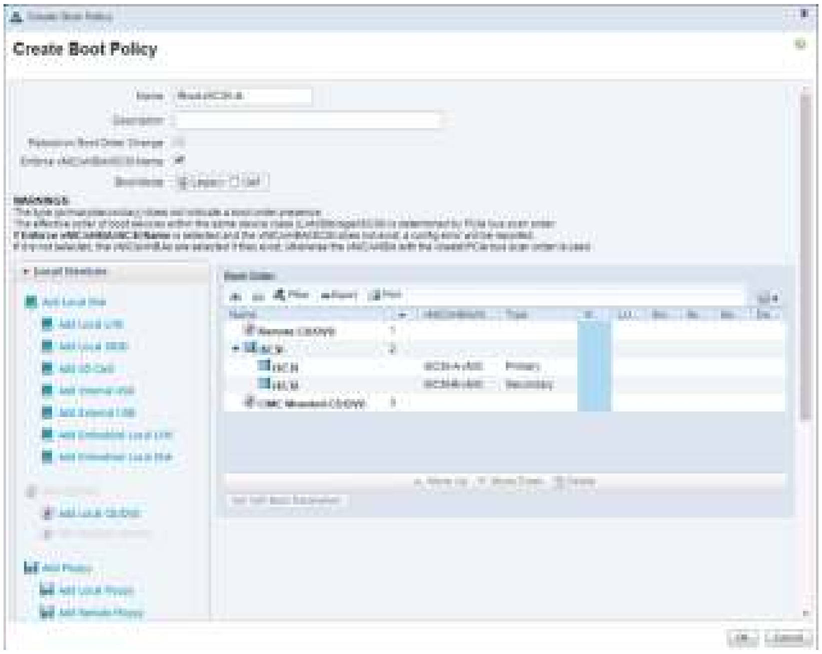

Create Boot Policy

17. Click OK to save the boot policy. Click OK to close the Boot Policy window.

Create Boot Policies (FC Boot)

This procedure applies to a Cisco UCS environment in which two FC logical interfaces (LIFs) are on cluster node 1 and two FC LIFs are on cluster node 2 for each Cisco UCS Fabric Interconnect:

	6332-16UP Fabric A	6332-16UP Fabric B	6248UP Fabric A	6248UP Fabric B
AFF Cluster Node 1 LIF	fcp_lif01_63a (fcp_lif01a)	fcp_lif01_63b (fcp_lif01b)	fcp_lif01_62a (fcp_lif01a)	fcp_lif01_63b (fcp_lif01b)
AFF Cluster Node 2 LIF	fcp_lif02_63a (fcp_lif02a)	fcp_lif02_63b (fcp_lif02b)	fcp_lif02_62a (fcp_lif02a)	fcp_lif02_62a (fcp_lif02b)

Deployments utilizing iSCSI and not FC should ignore these steps.

To create boot policies 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.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter Boot-FC-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.

Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Add Remote CD/DVD`.
9. Expand the vHBAs drop-down menu and select Add SAN Boot.
10. In the Add SAN Boot dialog box, enter Fabric-A in the vHBA field.
11. Confirm that Primary is selected for the Type option.

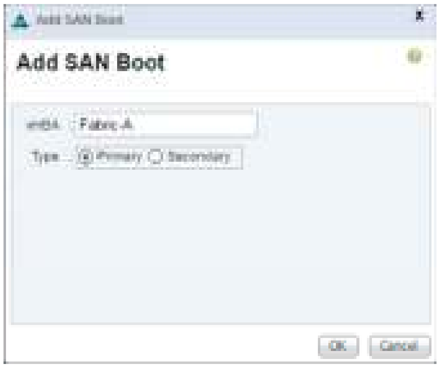

Add SAN Boot

vHBA: Fabric-A

Type: ● Primary ○ Secondary

12. Click OK to add the SAN boot initiator.
13. From the vHBA drop-down menu, select Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for fcp_lif01a.

To obtain this information, log in to the storage cluster and run the network interface show command.

16. Select Primary for the SAN boot target type.

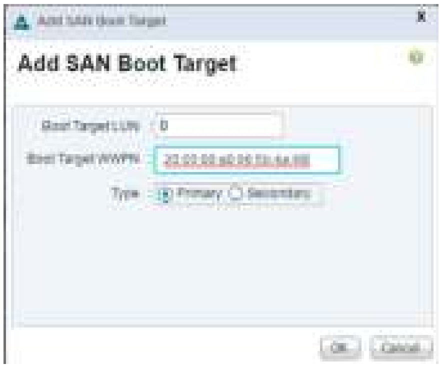

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:03:00:a0:98:fb:4a:8b

Type: ● Primary ○ Secondary

17. Click OK to add the SAN boot target.
18. From the vHBA drop-down menu, select Add SAN Boot Target.
19. Enter 0 as the value for Boot Target LUN.
20. Enter the WWPN for fcp_lif02a.

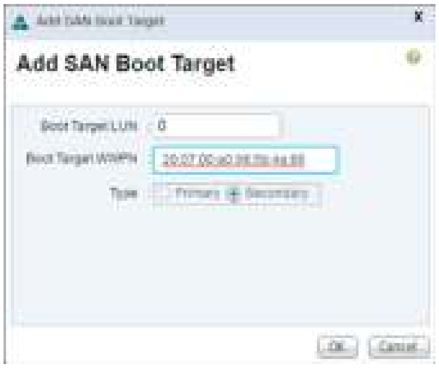

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:07:00:a0:98:fb:4a:8b

Type: ○ Primary ● Secondary

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

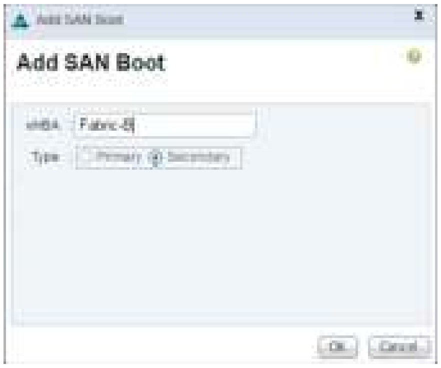

Add SAN Boot

vHBA: Fabric-B

Type: ○ Primary ● Secondary

25. Click OK to add the SAN boot initiator.
26. From the vHBA drop-down menu, select Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN for fcp_lif01b.
29. Select Primary for the SAN boot target type.

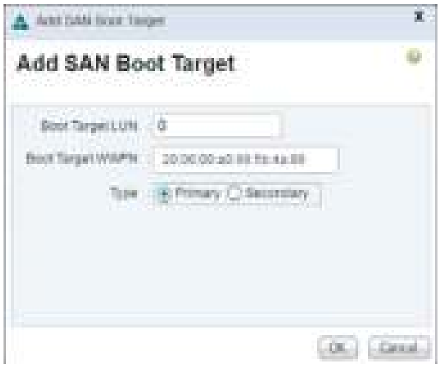

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:00:00:a0:98:fb:4a:88

Type: ● Primary ○ Secondary

30. Click OK to add the SAN boot target.
31. From the vHBA drop-down menu, select Add SAN Boot Target.
32. Keep 0 as the value for Boot Target LUN.
33. Enter the WWPN for fcp_lif02b.

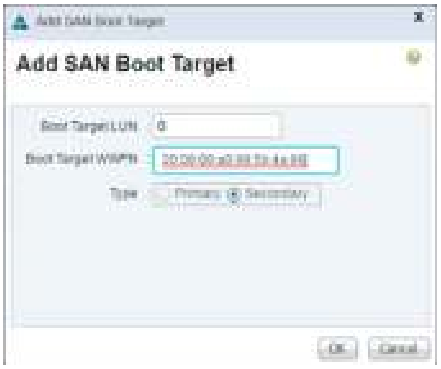

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:00:00:a0:98:fb:4a:88

Type: ○ Primary ● Secondary

34. Click OK to add the SAN boot target.

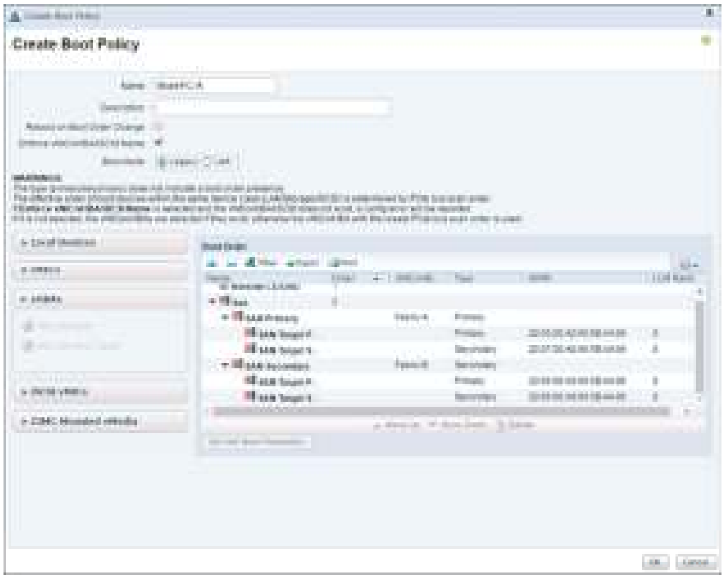

Create Boot Policy

18. Click OK, then click OK again to create the boot policy.

Create Service Profile Template (iSCSI Boot)

Service Profile Templates can be created for Fibre Channel (Fabric) boot or iSCSi boot, with VLANs appropriate infrastructure (Infra) or production (Prod) workloads differentiate by vNIC Templates allowing for VLAN presentation to the interfaces appropriate to the workload. In our example environment, the Service Profile Templates created were:

- VM-Host-Infra-Fabric-A
- VM-Host-Infra-iSCSI-A
- VM-Host-Prod-Fabric-A
- VM-Host-Prod-iSCSI-A

Examples shown in this section are primarily detail one Infra host that was provisioned from a B-Series server on the UCS 6248UP Fabric Interconnect, and one Prod host that was provisioned from a C-Series server on the UCS 6332-16UP Fabric Interconnect.

In this procedure, one service profile template for Infrastructure ESXi hosts is created for fabric A boot.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.

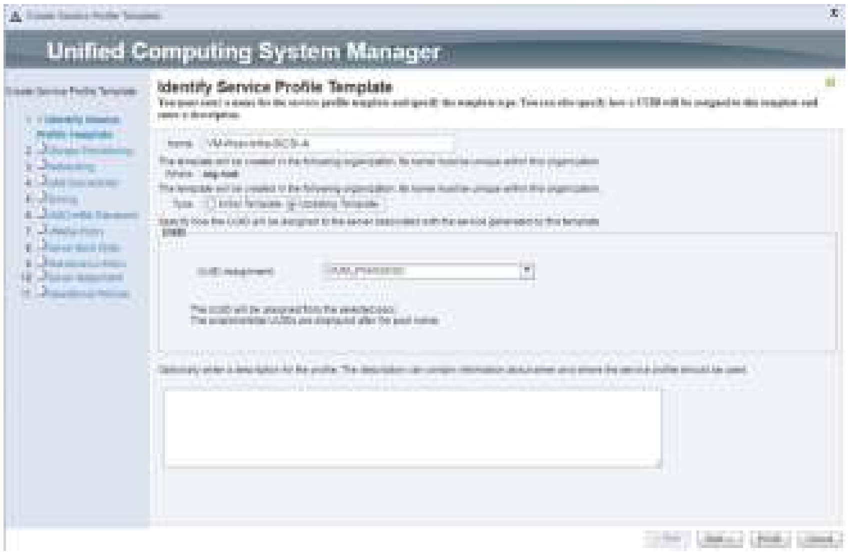
Unified Computing System Manager
Identify Service Profile Template

8. Click Next.

Configure Storage Provisioning

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy and select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select Infra-LAN-Policy from the LAN Connectivity Policy pull-down.
4. Select IQN_Pool within the Initiator Name Assignment pull-down.

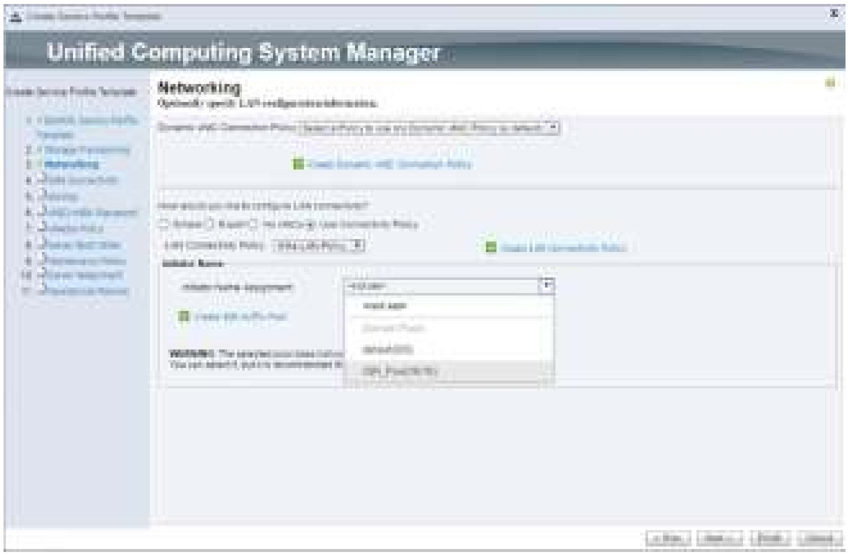

5. Click Next.

Configure Storage Options

1. Select the `Use Connectivity Policy` option for the “How would you like to configure SAN connectivity?” field.
2. Pick the Infra-SAN-Policy option from the SAN Connectivity Policy pull-down.

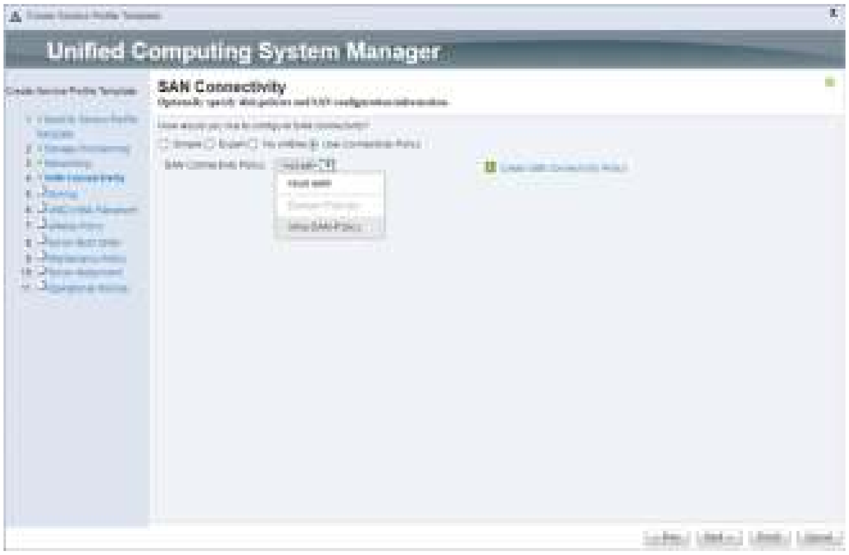

6. Click Next.

Configure Zoning Options

Set no Zoning options and click Next.

Configure vNIC/HBA Placement

1. In the “Select Placement” list, leave the placement policy as “Let System Perform Placement”.
2. Click Next.

Configure vMedia Policy

1. From the vMedia Policy pulldown select “ESXi-6.0U1b-HTTP”
2. Click Next.

Configure Server Boot Order

1. Select Boot-iSCSI-A for Boot Policy.
2. In the Boot Order pane, select iSCSI-A-vNIC.
3. Click the “Set iSCSI Boot Parameters” button.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.
5. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps
6. Set iSCSI_IP_Pool_A as the “Initiator IP address Policy”.
7. Keep the “iSCSI Static Target Interface” button selected and click the + button at the bottom right.
8. Log in to the storage cluster management interface and run the following command:

```
iscsi show

           Target                             Target                         Status

Vserver    Name                               Alias                          Admin

---------- -------------------------------- ---------------------------- ------

Infra-SVM  iqn.1992-08.com.netapp:sn.cbc5f0dff5b911e5aaa600a0985b4a74:vs.3

                                              Infra-SVM                      up
```

9. Note or copy the iSCSI target name for Infra-SVM shown in highlight above.
10. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra-SVM.
11. Enter the IP address of `iscsi_lif02a` for the IPv4 Address field.

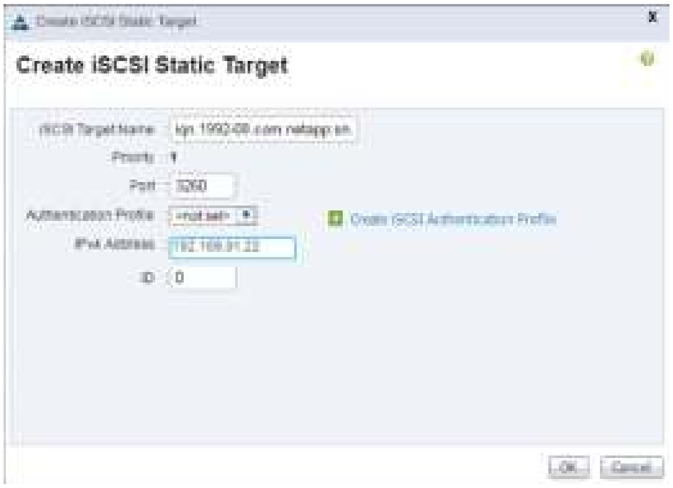

Create iSCSI Static Target

iSCSI Target Name: iqn.1992-08.com.netapp:sn.

Priority: 1

Port: 3260

Authentication Profile: <not set>

IPv4 Address: 192.168.91.22

ID: 0

12. Click OK to add the iSCSI static target.
13. Keep the iSCSI Static Target Interface option selected and click the ➕ button.
14. In the Create iSCSI Static Target window, paste the iSCSI target node name from `Infra-SVM` into the iSCSI Target Name field.
15. Enter the IP address of `iscsi_lif01a` in the IPv4 Address field.

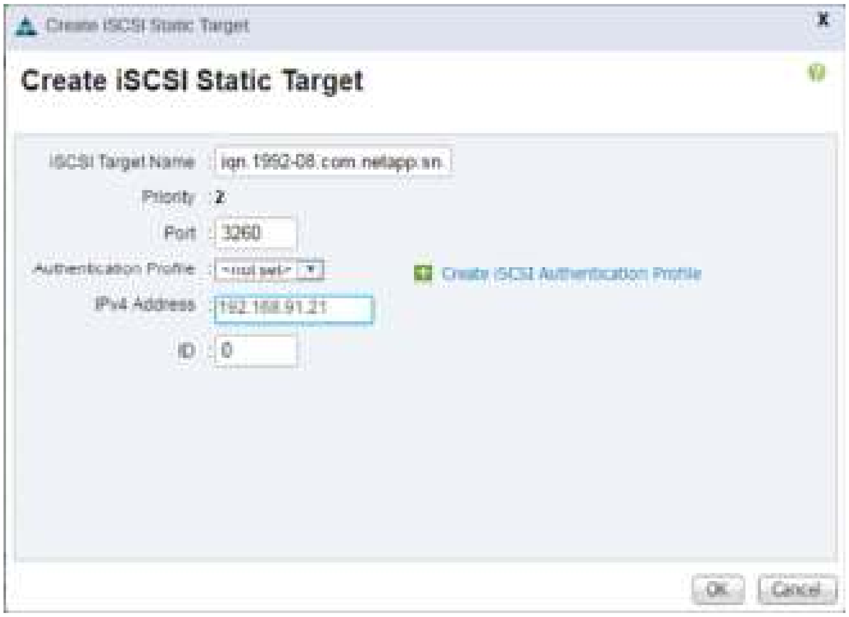

Create iSCSI Static Target

iSCSI Target Name: iqn.1992-08.com.netapp:sn.

Priority: 2

Port: 3260

Authentication Profile: <not set>

IPv4 Address: 192.168.91.21

ID: 0

16. Click OK.

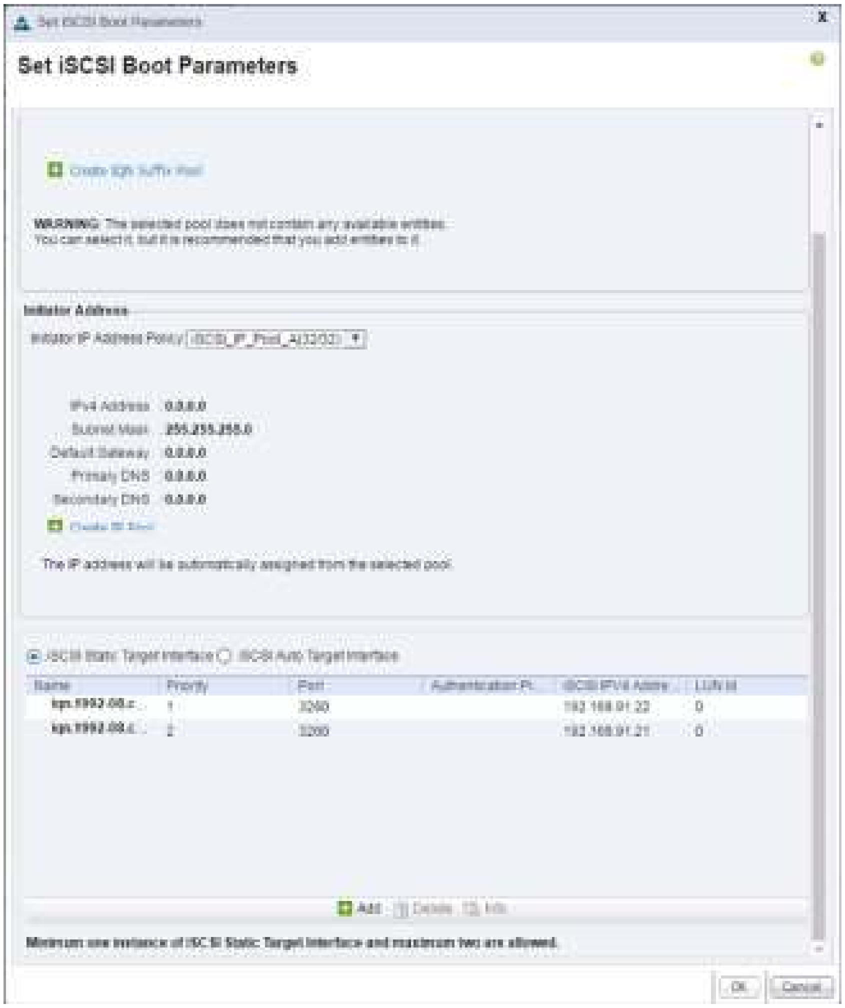

WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_A(32/32)

IPv4 Address: 0.0.0.0
Subnet Mask: 255.255.255.0
Default Gateway: 0.0.0.0
Primary DNS: 0.0.0.0
Secondary DNS: 0.0.0.0

The IP address will be automatically assigned from the selected pool.

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...	LUN Id
iqn.1992-08.c...	1	3260		192.168.91.22	0
iqn.1992-08.c...	2	3260		192.168.91.21	0

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

17. Click OK.
18. In the Boot Order pane, select iSCSI-vNIC-B.
19. Click the Set iSCSI Boot Parameters button.
20. In the Set iSCSI Boot Parameters dialog box, set the leave the “Initiator Name Assignment” to <not set>.
21. In the Set iSCSI Boot Parameters dialog box, set the initiator IP address policy to iSCSI_IP_Pool_B.
22. Keep the iSCSI Static Target Interface option selected and click the ⊞ button at the bottom right.
23. In the Create iSCSI Static Target window, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field (same target name as above).
24. Enter the IP address of iscsi_lif02b in the IPv4 address field.

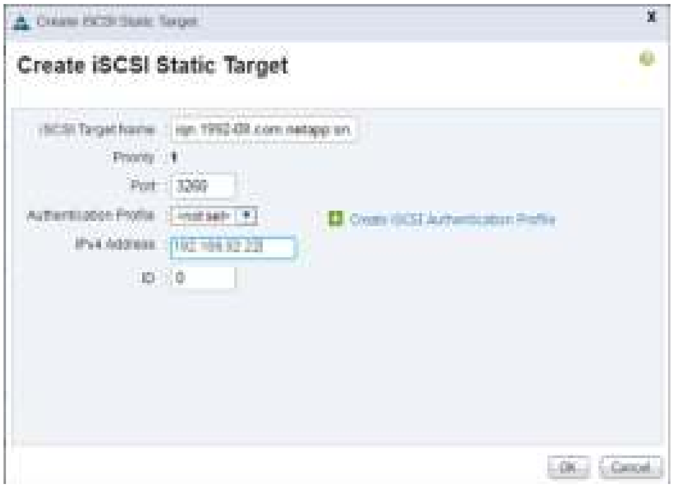

Create iSCSI Static Target
Create iSCSI Static Target

iSCSI Target Name: iqn.1992-08.com.netapp:sn.
Priority: 1
Port: 3260
Authentication Profile: <not set>
IPv4 Address: 192.168.92.23
ID: 0

Create iSCSI Authentication Profile
OK
Cancel

25. Click OK to add the iSCSI static target.

26. Keep the iSCSI Static Target Interface option selected and click the ➕ button.
27. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field.
28. Enter the IP address of iscsi_lif01b in the IPv4 Address field.

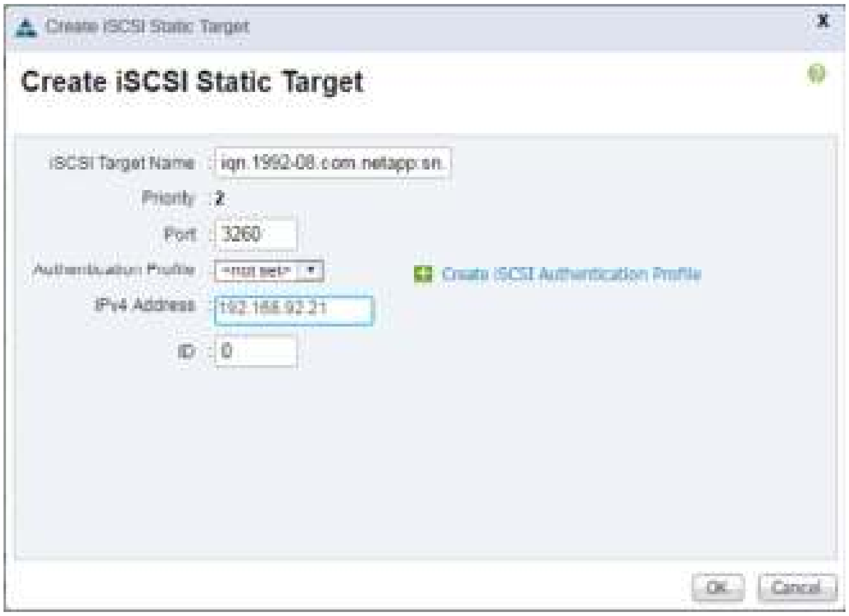

Create iSCSI Static Target

Create iSCSI Static Target

iSCSI Target Name: iqn.1992-08.com.netapp:sn.
Priority: 2
Port: 3260
Authentication Profile: <not set>
IPv4 Address: 192.168.92.21
ID: 0

29. Click OK.

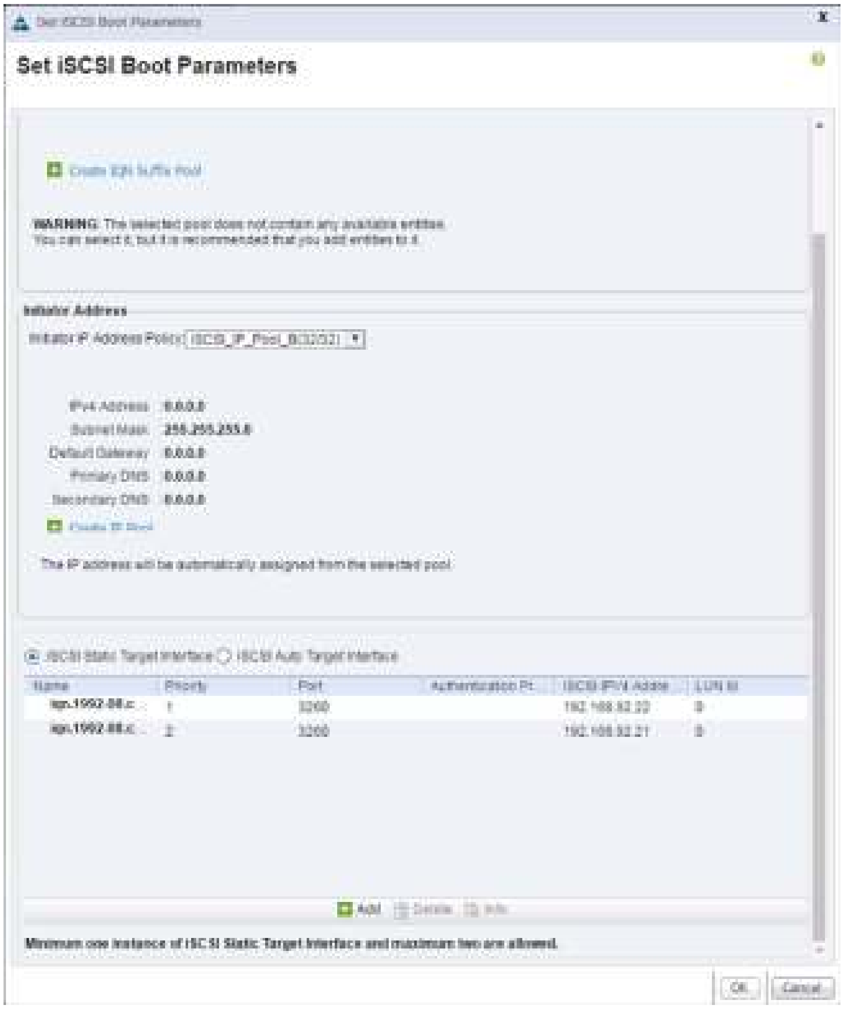

Set iSCSI Boot Parameters

Set iSCSI Boot Parameters

WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_B(32/32)

IPv4 Address: 0.0.0.0
Subnet Mask: 255.255.255.0
Default Gateway: 0.0.0.0
Primary DNS: 0.0.0.0
Secondary DNS: 0.0.0.0

The IP address will be automatically assigned from the selected pool.

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...	LUN Id
iqn.1992-08.c...	1	3260		192.168.92.20	0
iqn.1992-08.c...	2	3260		192.168.92.21	0

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

30. Click OK.
31. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

32. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

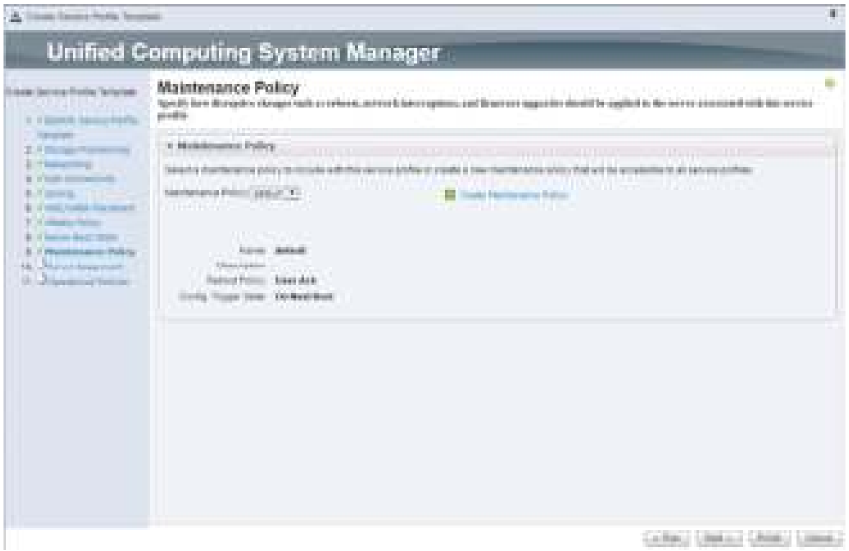

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select Infra_Pool.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Select “UCS-Broadwell” for the Server Pool Qualification.
5. Firmware Management at the bottom of the page can be left alone as it will use default from the Host Firmware list.

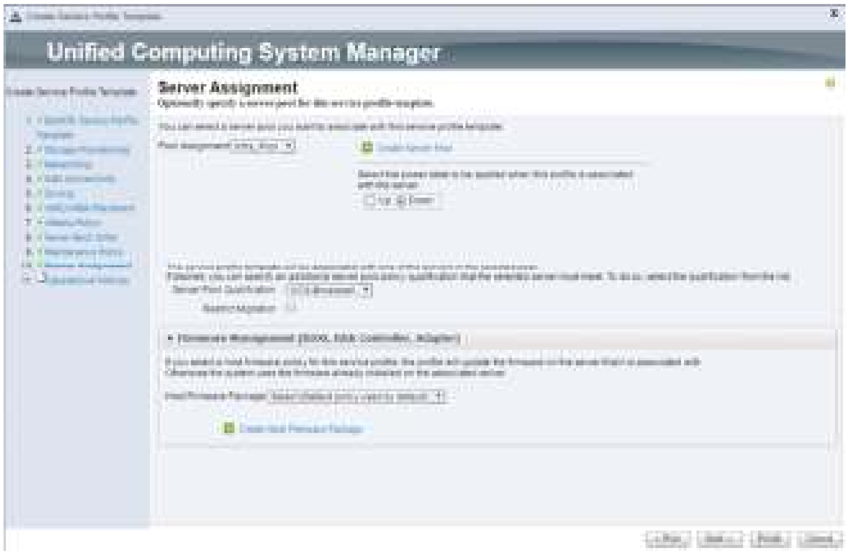

6. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. In the BIOS Policy list, select VM-Host-Infra.
2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

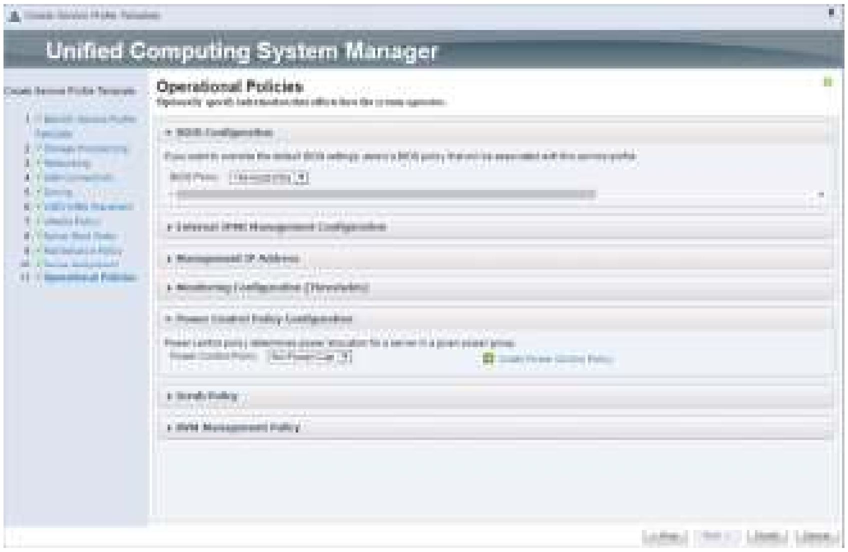

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Service Profile Template (FC Boot)

In this procedure, one service profile template for Infrastructure ESXi hosts is created for fabric A boot.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Prod-FC-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.

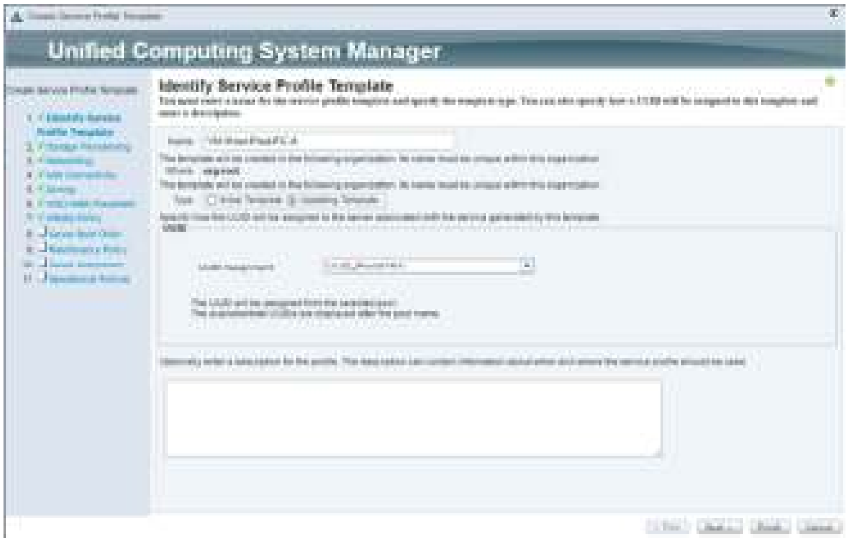
Unified Computing System Manager
Identify Service Profile Template

8. Click Next.

Configure Storage Provisioning

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy and select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the “Use Connectivity Policy” option to configure the LAN connectivity.
3. Select Infra-LAN-Policy from the LAN Connectivity Policy pull-down.

4. Select IQN_Pool within the Initiator Name Assignment pull-down.

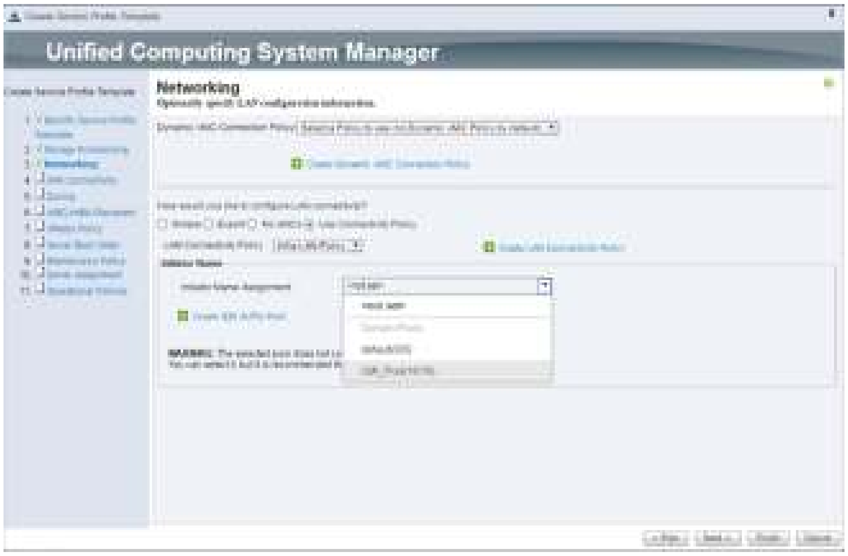

5. Click Next.

Configure Storage Options

1. Select the Use Connectivity Policy option for the "How would you like to configure SAN connectivity?" field.
2. Pick the Infra-SAN-Policy option from the SAN Connectivity Policy pull-down.

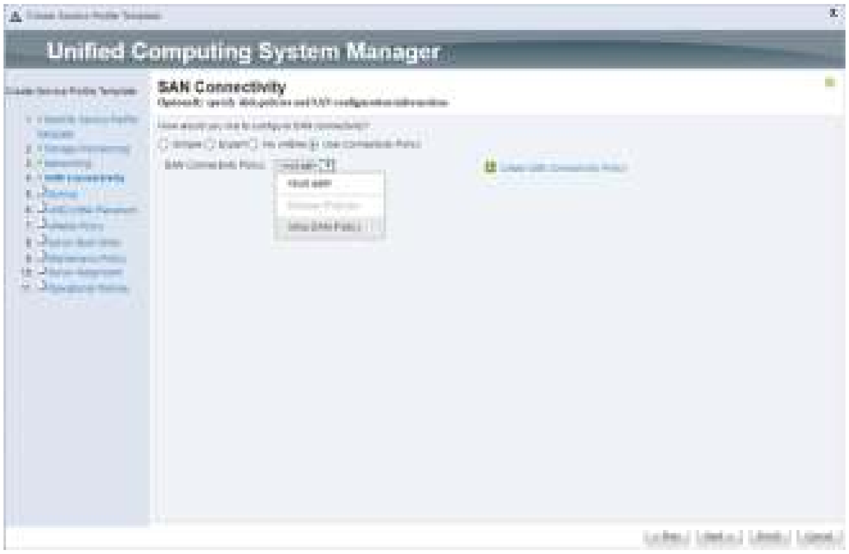

The SAN Connectivity policy created earlier will work for FC or iSCSI environments and should be changed if one of these protocols is not being used.

3. Click Next.

Configure Zoning Options

Set no Zoning options and click Next.

Configure vNIC/HBA Placement

1. In the "Select Placement" list, leave the placement policy as "Let System Perform Placement".
2. Click Next.

Configure vMedia Policy

1. From the vMedia Policy pulldown select "ESXi-6.0U1b-HTTP"
2. Click Next.

Configure Server Boot Order

1. Select `Boot-FC-A` for Boot Policy.

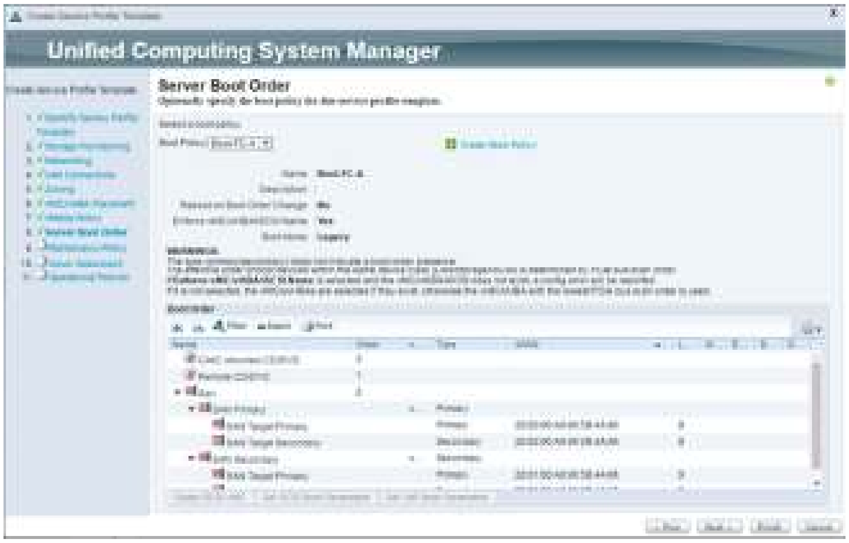

2. Click Next to continue to the next section.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

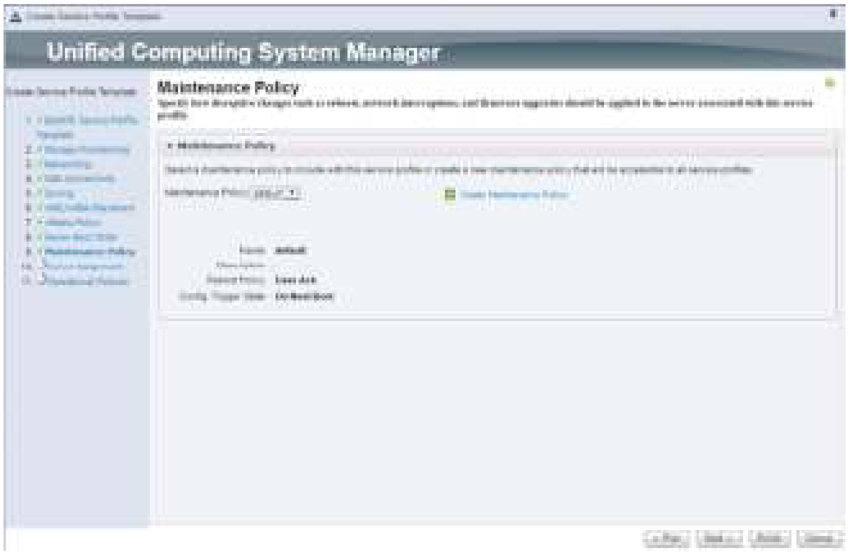

Name: default
Reboot Policy: User Ack
Config. Trigger State: On Next Boot

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select `Infra_Pool`.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Select “UCS-Broadwell” for the Server Pool Qualification.
5. Firmware Management at the bottom of the page can be left alone as it will use default from the Host Firmware list.

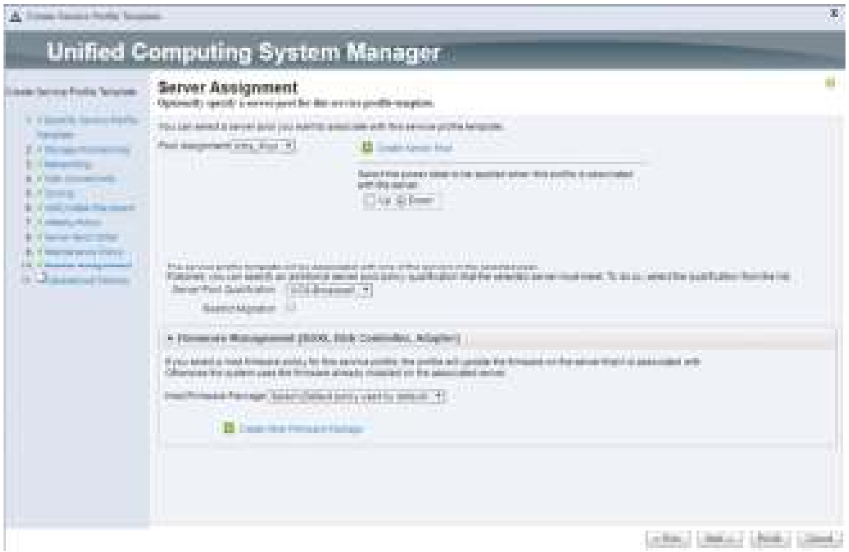

6. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. In the BIOS Policy list, select VM-Host-Infra.
2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

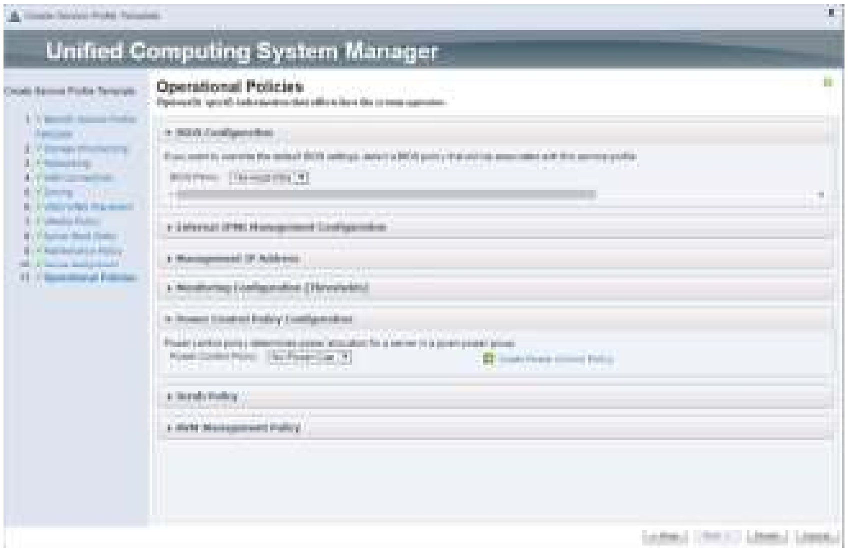

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Service Profiles

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

1. Connect to the UCS 6248UP Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and select Create Service Profiles from Template.
4. Enter VM-Host-Infra-0 as the service profile prefix.
5. Enter 1 as “Name Suffix Starting Number.”
6. Enter 1 as the “Number of Instances.”
7. Click OK to create the service profile.

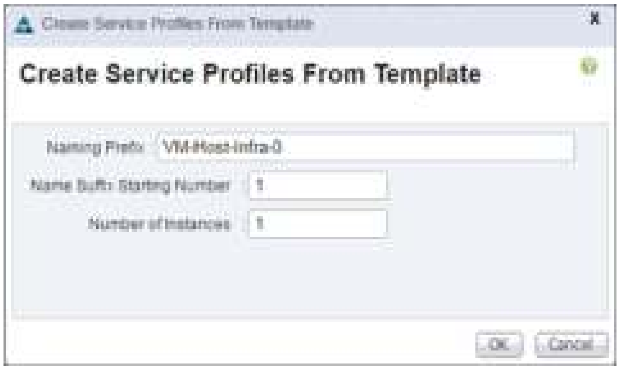

Create Service Profiles From Template

Naming Prefix: VM-Host-Infra-0

Name Suffix Starting Number: 1

Number of Instances: 1

8. Click OK in the confirmation message.
9. Connect to the UCS 6332-16UP Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.
10. Select Service Profile Templates > root > Service Template VM-Host-Prod-Fabric-A.
11. Right-click `VM-Host-Prod-Fabric-A` and select Create Service Profiles from Template.
12. Enter `VM-Host-Prod-0` as the service profile prefix.
13. Enter `1` as "Name Suffix Starting Number."
14. Enter `1` as the "Number of Instances."
15. Click OK to create the service profile.

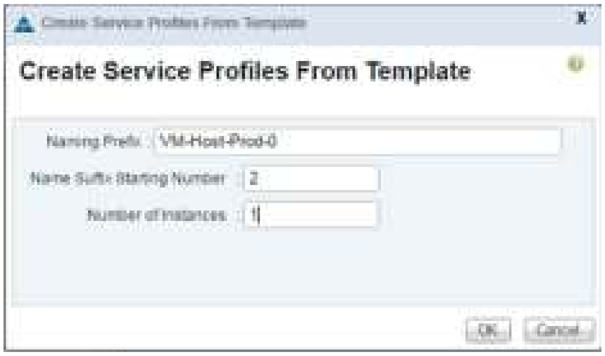

Create Service Profiles From Template

Naming Prefix: VM-Host-Prod-0

Name Suffix Starting Number: 2

Number of Instances: 1

16. Click OK in the confirmation message.

Add More Servers to FlexPod Unit

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

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into Table 18 and Table 19.

Table 18 iSCSI LIFs for iSCSI IQN

Vserver	Target: WWPN (FC), or IQN (iSCSI)
Infra-SVM	

To obtain the FC WWPN, run the `fcp show` command on the storage cluster management interface.

To obtain the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface.

Table 19 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	Initiator: WWPNs (FC) or IQN (iSCSI)	Variables
VM-Host-Infra-01		<<var_vm_host_infra_01_wwpn1>> and <<var_vm_host_infra_01_wwpn2>>; or <<var_vm_host_infra_01_iqn>>
VM-Host-Prod-02		<<var_vm_host_prod_02_wwpn1>> and <<var_vm_host_prod_02_wwpn2>>; or <<var_vm_host_prod_02_iqn>>

To obtain the FC vHBA WWPN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the "Storage" tab, then "vHBAs" tab on the right. The WWPNs are displayed in the table at the bottom of the page.

To obtain the iSCSI vNIC IQN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the "iSCSI vNICs" tab on the right. The "Initiator Name" is displayed at the top of the page under the "Service Profile Initiator Name."

Storage Configuration – Boot LUNs and Igroups

Clustered Data ONTAP Boot Storage Setup

Create igroups

If you are using FC connectivity, create igroups by entering the following commands from the cluster management node SSH connection:

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol fc –ostype vmware –initiator <<var_vm_host_infra_01_wwpn1>>,
<<var_vm_host_infra_01_wwpn2>>

igroup create –vserver Infra-SVM –igroup VM-Host-Prod-02 –protocol fc –ostype vmware –initiator <<var_vm_host_prod_02_wwpn1>>,
<<var_vm_host_prod_02_wwpn2>>

igroup create –vserver Infra-SVM –igroup MGMT-Hosts –protocol fc –ostype vmware –initiator <<var_vm_host_infra_01_wwpn1>>,
<<var_vm_host_infra_01_wwpn2>>, <<var_vm_host_prod_02_wwpn1>>, <<var_vm_host_prod_02_wwpn2>>
```

If you are using iSCSI connectivity, create igroups by entering the following commands from the cluster management node SSH connection:

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol iscsi –ostype vmware –initiator <<var_vm_host_infra_01_iqn>>

igroup create –vserver Infra-SVM –igroup VM-Host-Prod-02 –protocol iscsi –ostype vmware –initiator <<var_vm_host_prod_02_iqn>>

igroup create –vserver Infra-SVM –igroup MGMT-Hosts –protocol iscsi –ostype vmware –initiator <<var_vm_host_infra_01_iqn>>,
<<var_vm_host_prod_02_iqn>>
```

Use the values listed in Table 18 and Table 19 for the WWPN and IQN information.

To view the three igroups just created, type `igroup` show.

Map Boot LUNs to igroups

From the storage cluster management SSH connection, enter the following commands:

```
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-01 –igroup VM-Host-Infra-01 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Prod-02 –igroup VM-Host-Prod-02 –lun-id 0
```

VMware vSphere 6.0 U1 Setup

VMware ESXi **6.0 U1**

This section provides detailed instructions for installing VMware ESXi 6.0 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi **6.0 U1**

1. Click the following link vmware login page.
2. Type your email or customer number and the password and then click Log in.
3. Click on the following link CiscoCustomImage6.0 U1b.
4. Click Download Now.
5. Save it to your destination folder.

This ESXi 6.0 U1b Cisco custom image includes updates for the fnic and eNIC drivers. The versions that are part of this image are: eNIC: 2.3.0.6; fNIC: 1.6.0.24

Log in to Cisco UCS 6300/6200 Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the UCS environment to run the IP KVM.

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

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. To download the Cisco UCS Manager software, click the Launch UCS Manager link.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click the Servers tab.
7. Select Servers > Service Profiles > root > VM-Host-Infra-01.
8. Right-click VM-Host-Infra-01 and select KVM Console.
9. If prompted to accept an Unencrypted KVM session, accept as necessary.
10. Select Servers > Service Profiles > root > VM-Host-Prod-02.
11. Right-click VM-Host-Prod-02. and select KVM Console.
12. If prompted to accept an Unencrypted KVM session, accept as necessary.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

Skip this step if using vMedia policies. ISO file will already be connected to KVM.

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.
8. Boot the server by selecting Boot Server and clicking OK. Then click OK again.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To install VMware ESXi to the iSCSI-bootable or FC-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
4. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, click on the Virtual Media tab and clear the ✓ mark next to the ESXi installation media. Click Yes.

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

9. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host:

ESXi Host VM-Host-Infra-01

To configure the VM-Host-Infra-01 ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root, enter the corresponding password, and press Enter to log in.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the <<var_ib_mgmt_vlan_id>> and press Enter.
6. Select Network Adapters option and select vmnic04 and press Enter.
7. From the Configure Management Network menu, select IP Configuration and press Enter.
8. Select the Set Static IP Address and Network Configuration option by using the space bar.
9. Enter the IP address for managing the first ESXi host: <<var_vm_host_infra_01_ip>>.
10. Enter the subnet mask for the first ESXi host.
11. Enter the default gateway for the first ESXi host.
12. Press Enter to accept the changes to the IP configuration.
13. Select the IPv6 Configuration option and press Enter.
14. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
15. Select the DNS Configuration option and press Enter.

Because the IP address is assigned manually, the DNS information must also be entered manually.

16. Enter the IP address of the primary DNS server.
17. Optional: Enter the IP address of the secondary DNS server.
18. Enter the fully qualified domain name (FQDN) for the first ESXi host.
19. Press Enter to accept the changes to the DNS configuration.
20. Press Esc to exit the Configure Management Network submenu.
21. Press Y to confirm the changes and return to the main menu.
22. The ESXi host reboots. After reboot, press F2 and log back in as root.
23. Select Test Management Network to verify that the management network is set up correctly and press Enter.
24. Press Enter to run the test.
25. Press Enter to exit the window.
26. Press Esc to log out of the VMware console.

ESXi Host VM-Host-Prod-02

To configure the VM-Host-Prod-02 ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the <<var_ib-mgmt_vlan_id>> and press Enter.
6. Select Network Adapters option and select vmnic4 (defined earlier as OOB vNIC) and press Enter.
7. From the Configure Management Network menu, select IP Configuration and press Enter.
8. Select the Set Static IP Address and Network Configuration option by using the space bar.
9. Enter the IP address for managing the second ESXi host: <<var_vm_host_prod_02_ip>>.
10. Enter the subnet mask for the second ESXi host.
11. Enter the default gateway for the second ESXi host.
12. Press Enter to accept the changes to the IP configuration.
13. Select the IPv6 Configuration option and press Enter.
14. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
15. Select the DNS Configuration option and press Enter.

Because the IP address is assigned manually, the DNS information must also be entered manually.

16. Enter the IP address of the primary DNS server.
17. Optional: Enter the IP address of the secondary DNS server.
18. Enter the FQDN for the second ESXi host.
19. Press Enter to accept the changes to the DNS configuration.
20. Press Esc to exit the Configure Management Network submenu.
21. Press Y to confirm the changes and return to the main menu.
22. The ESXi host reboots. After reboot, press F2 and log back in as root.
23. Select Test Management Network to verify that the management network is set up correctly and press Enter.
24. Press Enter to run the test.
25. Press Enter to exit the window.
26. Press Esc to log out of the VMware console.

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-01 management IP address.
2. Download and install the vSphere Client.

This application is downloaded from the VMware website and Internet access is required on the management workstation.

Log in to VMware ESXi Hosts by Using VMware vSphere Client

ESXi Host VM-Host-Infra-01

To log in to the VM-Host-Infra-01 ESXi host (which was provisioned from an iSCSI boot Service Profile Template) by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-Infra-01 as the host you are trying to connect to: <<var_vm_host_infra_01_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

ESXi Host VM-Host-Prod-02

To log in to the VM-Host-Prod-02 ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-Prod-02 as the host you are trying to connect to: <<var_vm_host_prod_02_ip>>.
2. Enter root for the user name.
3. Enter the root password.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

To set up the VMkernel ports and the virtual switches on the `VM-Host-Infra-01`. ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of `vSwitch0`, click Properties.
5. Select the vSwitch configuration and click Edit.

6. From the General tab, change the MTU to 9000.
7. Click OK.
8. Select the Management Network configuration and click Edit.
9. Change the network label to VMkernel-MGMT and select the Management Traffic checkbox.
10. Click OK to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to MGMT Network and enter <<var_ib-mgmt_vlan_id>> in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.
14. Click Close.
15. On the right side of iScsiBootvSwitch, click Properties.
16. Select the vSwitch configuration and click Edit.
17. Change the MTU to 9000.
18. Click OK.
19. Select iScsiBootPG and click Edit.
20. Change the Network Label to VMkernel-iSCSI-A.
21. Change the MTU to 9000.
22. Click OK.
23. Click Close.
24. In the vSphere Standard Switch view, click Add Networking.
25. Select VMkernel and click Next.
26. Select Create a vSphere standard switch to create a new vSphere standard switch.
27. Select the check boxes for the network adapter vmnic3.
28. Click Next.
29. Change the network label to VMkernel-iSCSI-B.
30. Click Next.
31. Enter the IP address and the subnet mask for the iSCSI VLAN B interface for VM-Host-Infra-01.

✍ To obtain the iSCSI IP address information; login to the Cisco UCS Manager, in the servers tab select the corresponding service profiles. In the right pane, click the boot order and select the iSCSI-B-vNIC; click set iSCSI boot parameters; the IP address should appear as the initiator IP address.

32. Click Next.
33. Click Finish.
34. On the right side of vSwitch1, click Properties.
35. Select the vSwitch configuration and click Edit.
36. Change the MTU to 9000.
37. Click OK.
38. Select VMkernel-iSCSI-B and click Edit.
39. Change the MTU to 9000.
40. Click OK.
41. Click Close.
42. On the right side of vSwitch0, click Properties.
43. Click Add.
44. Change the network label to VMkernel-NFS and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.
45. Click Next.
46. Enter the IP address <<var_nfs_vlan_ip_host_01>> and the subnet mask <<var_nfs_vlan_ip_mask_host_01>> for the NFS VLAN interface for VM-Host-Infra-01.
47. To continue with the NFS VMkernel creation, click Next.
48. To finalize the creation of the NFS VMkernel interface, click Finish.
49. Select the VMkernel-NFS configuration and click Edit.
50. Change the MTU to 9000.
51. Click OK to finalize the edits for the VMkernel-NFS network.
52. Click Add.
53. Change the network label to VMkernel-vMotion and enter <<var_vmotion_vlan_id>> in the VLAN ID (Optional) field.
54. Click Next.
55. Enter the IP address <<var_vmotion_vlan_ip_host_01>> and the subnet mask <<var_vmotion_vlan_ip_mask_host_01>> for the vMotion VLAN interface for VM-Host-Infra-01.
56. To continue with the vMotion VMkernel creation, click Next.
57. To finalize the creation of the vMotion VMkernel interface, click Finish.
58. Select the VMkernel-vMotion configuration and click Edit.
59. Change the MTU to 9000.
60. Click OK to finalize the edits for the VMkernel-vMotion network.
61. The properties for vSwitch0 should be similar to the following example:

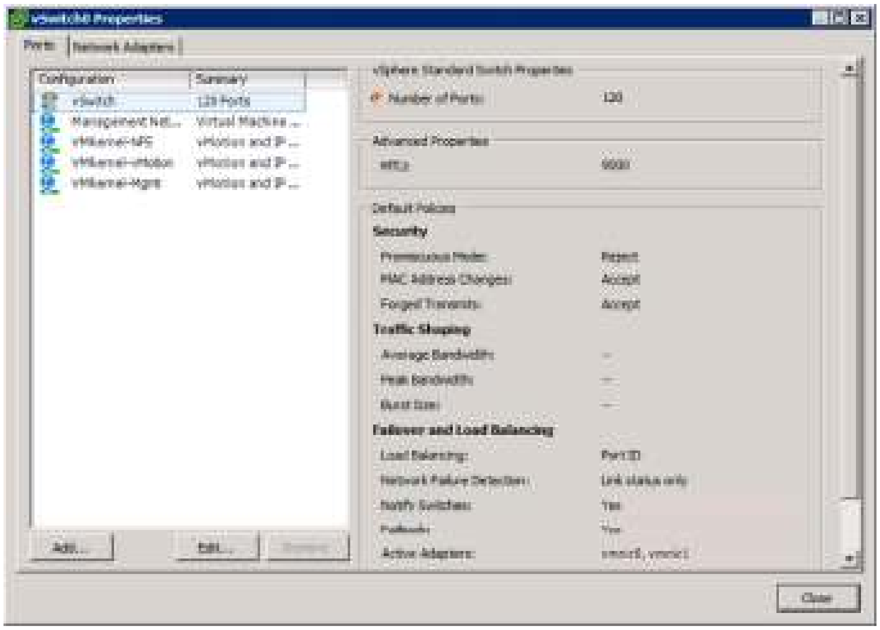

vSwitch0 Properties

Ports | Network Adapters

Configuration	Summary
vSwitch	120 Ports
Management Net...	Virtual Machine ...
VMkernel-NFS	vMotion and IP ...
VMkernel-vMotion	vMotion and IP ...
VMkernel-Mgmt	vMotion and IP ...

vSphere Standard Switch Properties

Number of Ports: 120

Advanced Properties

MTU: 9000

Default Policies

Security

Promiscuous Mode: Reject

MAC Address Changes: Accept

Forged Transmits: Accept

Traffic Shaping

Average Bandwidth: --

Peak Bandwidth: --

Burst Size: --

Failover and Load Balancing

Load Balancing: Port ID

Network Failure Detection: Link status only

Notify Switches: Yes

Failback: Yes

Active Adapters: vmnic0, vmnic1

62. To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:

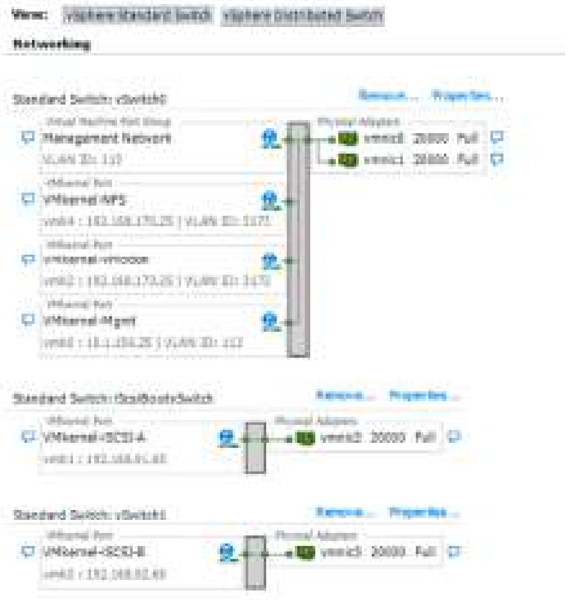

View: vSphere Standard Switch | vSphere Distributed Switch

Networking

Standard Switch: vSwitch0

Virtual Machine Port Group: Management Network, VLAN ID: 113

VMkernel Port: VMkernel-NFS, vmk4 : 192.168.170.25 | VLAN ID: 3171

VMkernel Port: VMkernel-vMotion, vmk2 : 192.168.173.25 | VLAN ID: 3173

VMkernel Port: VMkernel-Mgmt, vmk0 : 10.1.156.25 | VLAN ID: 113

Physical Adapters: vmnic0 20000 Full; vmnic1 20000 Full

Standard Switch: iScsiBootvSwitch

VMkernel Port: VMkernel-iSCSI-A, vmk1 : 192.168.91.65

Physical Adapters: vmnic2 20000 Full

Standard Switch: vSwitch1

VMkernel Port: VMkernel-iSCSI-B, vmk3 : 192.168.92.65

Physical Adapters: vmnic3 20000 Full

ESXi Host VM-Host-Prod-02

To set up the VMkernel ports and the virtual switches on the `VM-Host-Prod-02, which was provisioned as a Fibre Channel booted` ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of `vSwitch0`, click Properties.
5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to `9000`.
7. Click OK.
8. Select the Management Network configuration and click Edit.
9. Change the network label to <`VMkernel-MGMT`> and select the Management Traffic checkbox.
10. Click OK to finalize the edits for Management Network.
11. Select the VM Network configuration and click Edit.
12. Change the network label to <`MGMT Network`> and enter `<<var_ib-mgmt_vlan_id>>` in the VLAN ID (Optional) field.
13. Click OK to finalize the edits for VM Network.

14. Click Close.
15. In the vSphere Standard Switch view, click Add Networking.
16. Select VMkernel and click Next.
17. Select Create a vSphere standard switch to create a new vSphere standard switch.
18. Select the check boxes for the network adapter vmnic2.
19. Click Next.
20. Change the network label to <`VMkernel-iSCSI-A>.`
21. Enter the IP address and the subnet mask for the iSCSI VLAN interface for `VM-Host-Prod-02.`
22. Click Next.
23. Click Finish.
24. On the right side of `vSwitch1`, click Properties.
25. Change the MTU to 9000.
26. Click OK.
27. In the vSphere Standard Switch view, click Add Networking.
28. Select VMkernel and click Next.
29. Select Create a vSphere standard switch to create a new vSphere standard switch.
30. Select the check boxes for the network adapter vmnic3.
31. Click Next.
32. Change the network label to <`VMkernel-iSCSI-B>.`
33. Click Next.
34. Enter the IP address and the subnet mask for the iSCSI VLAN interface for `VM-Host-Prod-02.`
35. Click Next.
36. Click Finish.
37. On the right side of `vSwitch2`, click Properties.
38. Select the vSwitch configuration and click Edit.
39. Change the MTU to 9000.
40. Click OK.
41. Select VMkernel-iSCSI-B and click Edit.
42. Change the MTU to 9000.
43. Click OK.
44. Click Close.
45. On the right side of `vSwitch0`, click Properties.
46. Click Add.
47. Select VMkernel and click Next.
48. Change the network label to `VMkernel-NFS` and enter `<<var_nfs_vlan_id>>` in the VLAN ID (Optional) field.
49. Click Next.
50. Enter the IP address <<var_nfs_vlan_ip_host_02>> and the subnet mask <<var_nfs_vlan_ip_mask_host_02>> for the NFS VLAN interface for VM-Host-Prod-02.
51. To continue with the NFS VMkernel creation, click Next.
52. To finalize the creation of the NFS VMkernel interface, click Finish.
53. Select the `VMkernel-NFS` configuration and click Edit.
54. Change the MTU to `9000`.
55. Click OK to finalize the edits for the VMkernel-NFS network.
56. Click Add.
57. Change the network label to `VMkernel-vMotion` and enter `<<var_vmotion_vlan_id>>` in the VLAN ID (Optional) field.
58. Click Next.
59. Enter the IP address <<var_vmotion_vlan_ip_host_02>> and the subnet mask <<var_vmotion_vlan_ip_mask_host_02>> for the vMotion VLAN interface for VM-Host-Prod-02.
60. To continue with the vMotion VMkernel creation, click Next.
61. To finalize the creation of the vMotion VMkernel interface, click Finish.
62. Select the `VMkernel-vMotion` configuration and click Edit.
63. Change the MTU to `9000`.
64. Click OK to finalize the edits for the VMkernel-vMotion network.

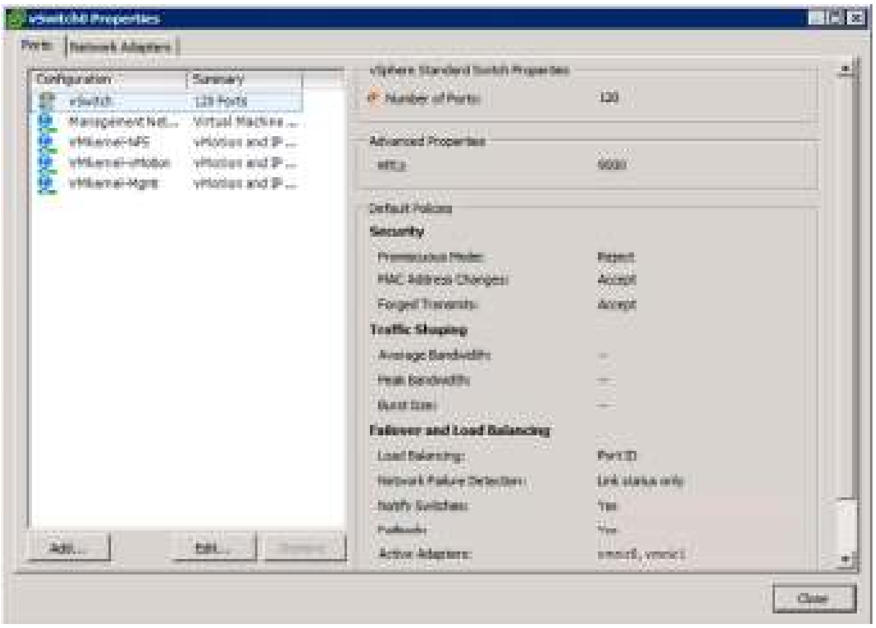

Configuration	Summary
vSwitch	120 Ports
Management Net...	Virtual Machine ...
VMkernel-NFS	vMotion and IP ...
VMkernel-vMotion	vMotion and IP ...
VMkernel-Mgmt	vMotion and IP ...

vSphere Standard Switch Properties

Number of Ports: 120

Advanced Properties

MTU: 9000

Default Policies

Security

Promiscuous Mode: Reject

MAC Address Changes: Accept

Forged Transmits: Accept

Traffic Shaping

Average Bandwidth: --

Peak Bandwidth: --

Burst Size: --

Failover and Load Balancing

Load Balancing: Port ID

Network Failure Detection: Link status only

Notify Switches: Yes

Failback: Yes

Active Adapters: vmnic0, vmnic1

65. To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:

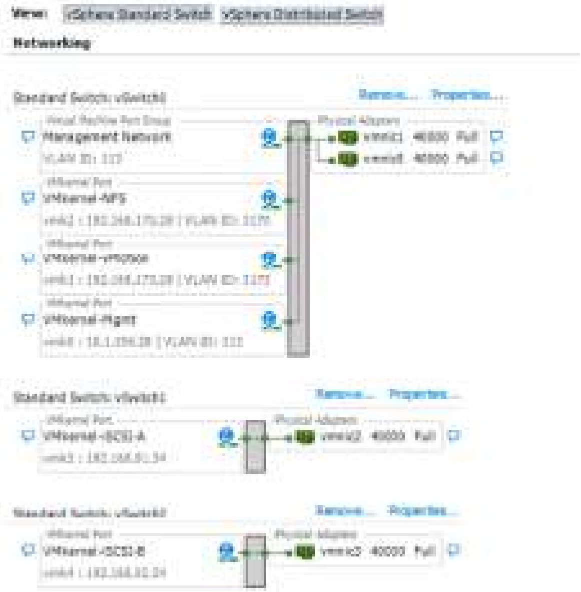

Standard Switch: vSwitch0

Management Network
VLAN ID: 113

VMkernel-NFS
vmk2 : 192.168.170.28 | VLAN ID: 3170

VMkernel-vMotion
vmk1 : 192.168.173.28 | VLAN ID: 3173

VMkernel-Mgmt
vmk0 : 10.1.156.28 | VLAN ID: 113

vmnic1 40000 Full
vmnic0 40000 Full

Standard Switch: vSwitch1

VMkernel-iSCSI-A
vmk3 : 192.168.81.34

vmnic2 40000 Full

Standard Switch: vSwitch2

VMkernel-iSCSI-B
vmk4 : 192.168.82.34

vmnic3 40000 Full

Setup iSCSI Multipathing

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To setup 4 iSCSI paths between storage and the ESXi host, complete the following steps on each ESXi host:

1. For the FC booted host, connect to the vSphere Client, select the configuration tab of the host.
2. Click Storage Adapters in the Hardware pane.
3. Select Add to create an iSCSI Software Adapter.
4. Specify the iqn assigned to the host by the Service Profile and create the iSCSI Software Adapter.
5. For both hosts, continue within the Storage Adapters section in the Hardware pane of the host Configuration tab.
6. Select the iSCSI Software Adapter and click Properties.
7. Select the Dynamic Discovery tab and click Add.
8. Enter the IP address of iscsi_lif01a.
9. Click OK.
10. Repeat putting in the IP addresses of iscsi_lif01b, iscsi_lif02a and iscsi_lif02b.

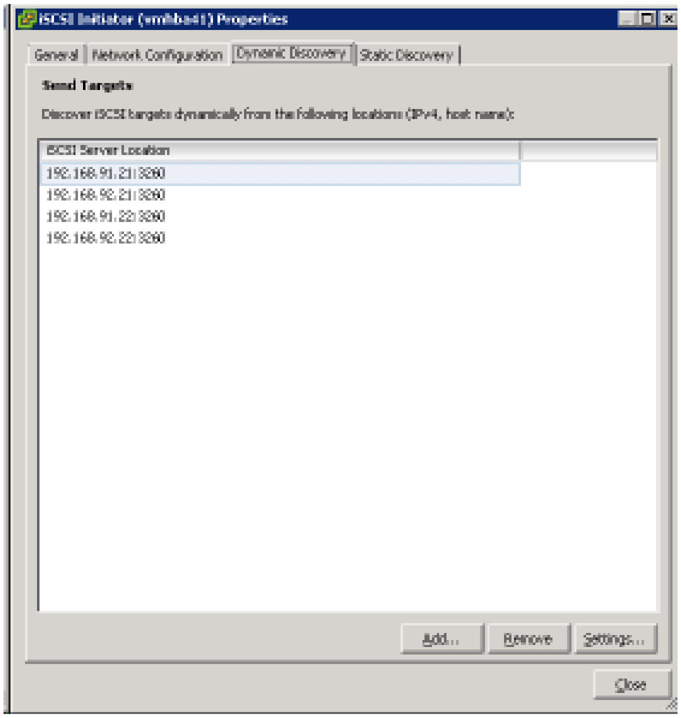

Send Targets

Discover iSCSI targets dynamically from the following locations (IPv4, host name):

iSCSI Server Location
192.168.91.21:3260
192.168.92.21:3260
192.168.91.22:3260
192.168.92.22:3260

11. Click Close and then click yes to rescan the host bus adapter.
12. You should now see 4 connected paths in the Details pane.

Install VMware Drivers for the Cisco Virtual Interface Card (VIC)

Download and extract the following VMware VIC Drivers to the Management workstation:

fnic Driver version 1.6.0.25
enic Driver version 2.3.0.7

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Prod-02, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. Click Enter Maintenance Mode within the Commands section of the Summary tab.
4. Click Yes for any dialogue box presented.
5. From Resources > Storage, right-click datastore1 and select Browse Datastore.
6. Click the fourth button and select Upload File.
7. Navigate to the saved location for the downloaded VIC drivers and select fnic_driver_1.6.0.25-3741467.zip.
8. Click Open and Yes to upload the file to datastore1.
9. Click the fourth button and select Upload File.
10. Navigate to the saved location for the downloaded VIC drivers and select ESXi60-enic-2.3.0.7-3642661.zip.
11. Click Open and Yes to upload the file to datastore1.
12. Make sure the files have been uploaded to both ESXi hosts.
13. Within the vSphere Client select the Configuration tab and click Security Profile within the Software section.

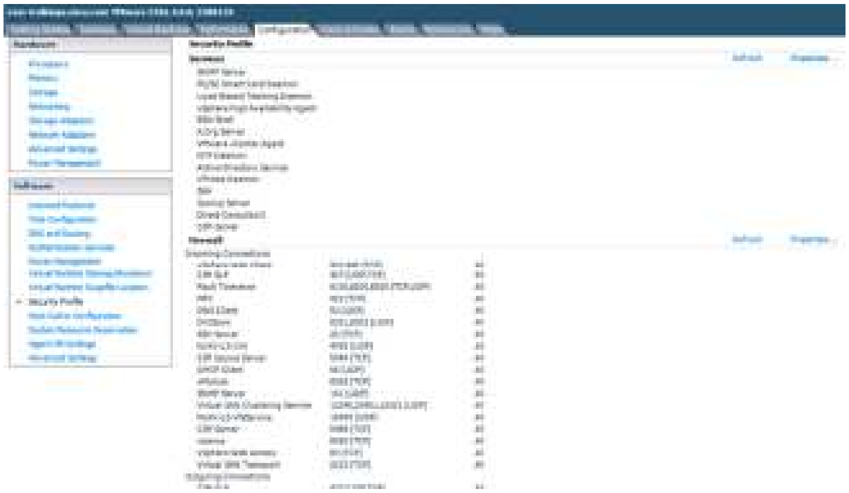

14. Click Properties within the Services section at the top.

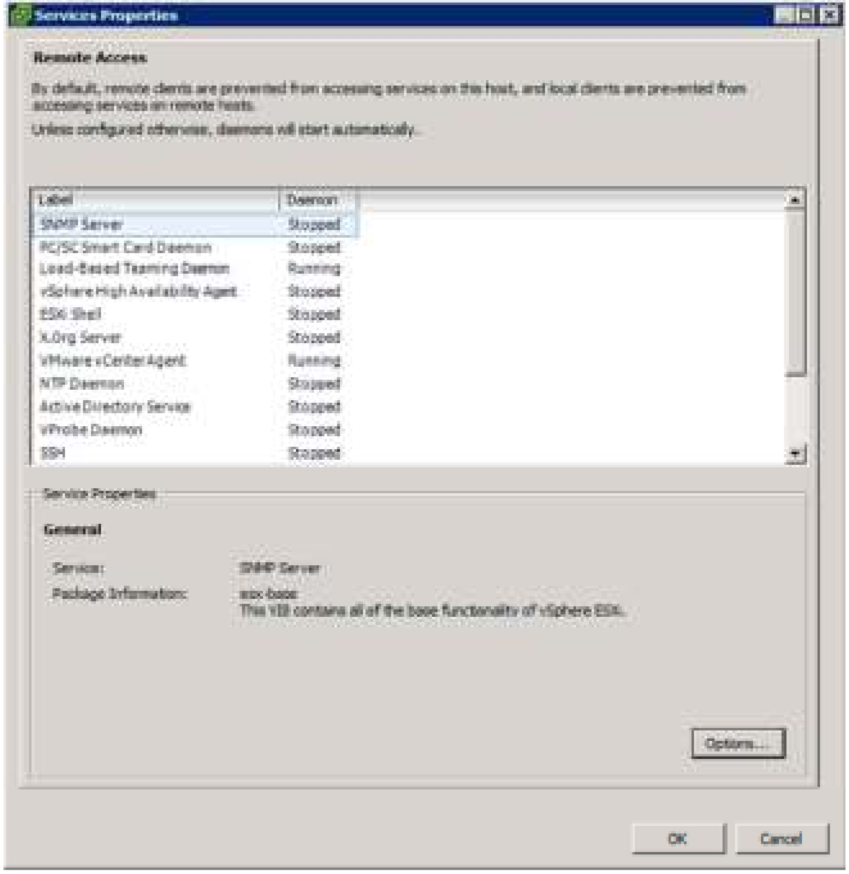

Remote Access

By default, remote clients are prevented from accessing services on this host, and local clients are prevented from accessing services on remote hosts.

Unless configured otherwise, daemons will start automatically.

Label	Daemon
SNMP Server	Stopped
PC/SC Smart Card Daemon	Stopped
Load-Based Teaming Daemon	Running
vSphere High Availability Agent	Stopped
ESXi Shell	Stopped
X.Org Server	Stopped
VMware vCenter Agent	Running
NTP Daemon	Stopped
Active Directory Service	Stopped
VProbe Daemon	Stopped
SSH	Stopped

Service Properties

General

Service: SNMP Server

Package Information: esx-base
This VIB contains all of the base functionality of vSphere ESXi.

15. Select ESXi Shell and click the Options... button.

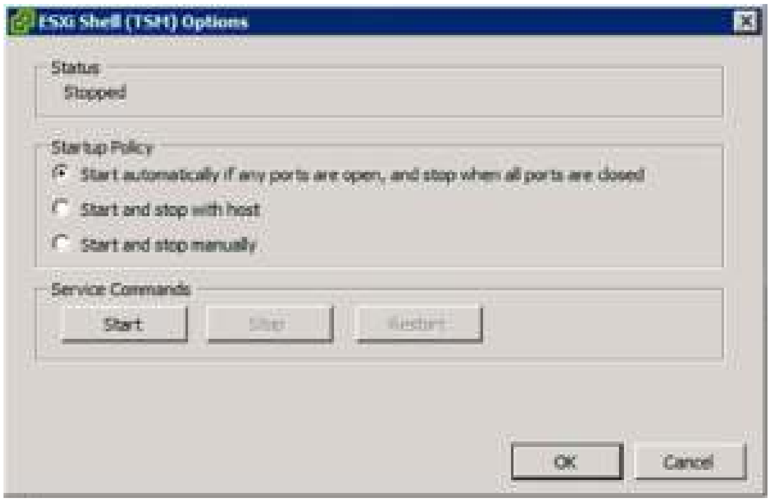

ESXi Shell (TSM) Options

Status
Stopped

Startup Policy
● Start automatically if any ports are open, and stop when all ports are closed
○ Start and stop with host
○ Start and stop manually

16. Select Start automatically if any ports are open, and stop when all ports are closed within the Startup Policy section.
17. Click Start, and OK.
18. Select SSH and click the Options... button.
19. Select Start automatically if any ports are open, and stop when all ports are closed within the Startup Policy section.
20. Click Start, and OK.
21. Click OK to exit from the Service Properties configuration window.

Enabling ssh can be considered optional if the VMware vSphere Remote CLI is installed and used.

22. Connect to each ESXi hosts through ssh from a shell connection or putty terminal.
23. Login as root with the password specified for <<var_password>>.
24. Run the following commands on each host:
25. esxcli software vib update -d /vmfs/volumes/datastore/fnic_driver_1.6.0.25-offline_bundle-3642682.zip
26. esxcli software vib install -d /vmfs/volumes/datastore/ESXi60-enic-2.3.0.7-offline_bundle-3642661.zip
27. Reboot each host after both commands have been run.

Mount Required Datastores

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To mount the required datastores, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the Configuration tab.
3. Click Storage in the Hardware pane.
4. From the Datastores area, click Add Storage to open the Add Storage wizard.

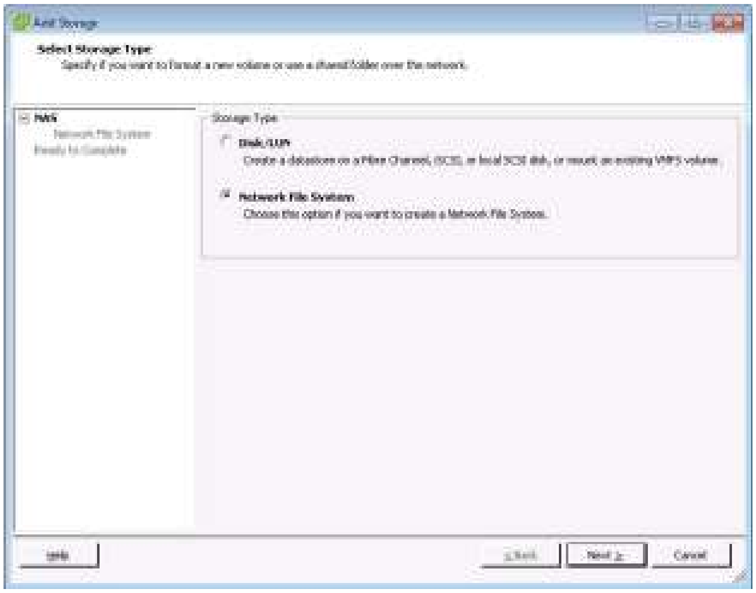

Add Storage

Select Storage Type

Specify if you want to format a new volume or use a shared folder over the network.

NAS
Network File System
Ready to Complete

Storage Type

○ Disk/LUN

Create a datastore on a Fibre Channel, iSCSI, or local SCSI disk, or mount an existing VMFS volume.

● Network File System

Choose this option if you want to create a Network File System.

Help | < Back | Next > | Cancel

5. Select Network File System and click Next.
6. The wizard prompts for the location of the NFS export. Enter <<var_node02_nfs_lif_infra_datastore_1_ip>> as the IP address for nfs_lif_infra_datastore_1.
7. Enter /infra_datastore_1 as the path for the NFS export.
8. Confirm that the Mount NFS read only checkbox is not selected.
9. Enter infra_datastore_1 as the datastore name.

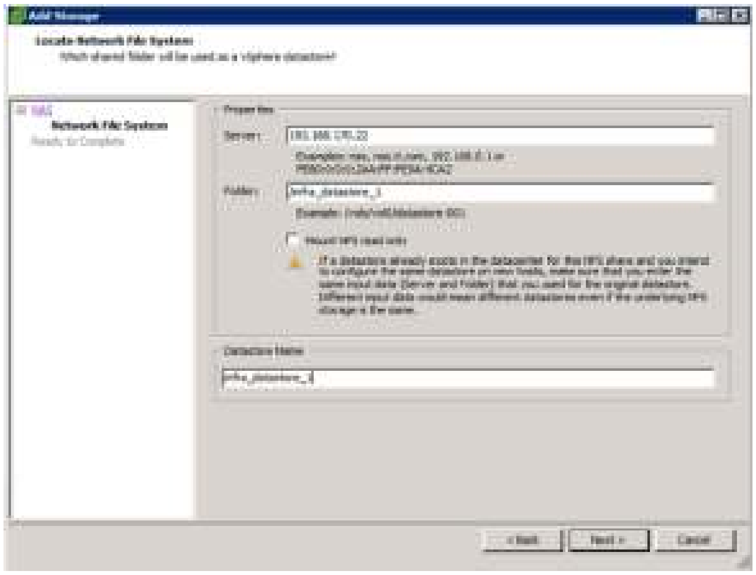

Add Storage

Locate Network File System

Which shared folder will be used as a vSphere datastore?

NAS
Network File System
Ready to Complete

Properties

Server: 192.168.170.22

Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2

Folder: infra_datastore_1

Example: /vols/vol0/datastore-001

□ Mount NFS read only

If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.

Datastore Name

infra_datastore_1

< Back | Next > | Cancel

10. To continue with the NFS datastore creation, click Next.
11. To finalize the creation of the NFS datastore, click Finish.
12. From the Datastores area, click Add Storage to open the Add Storage wizard.

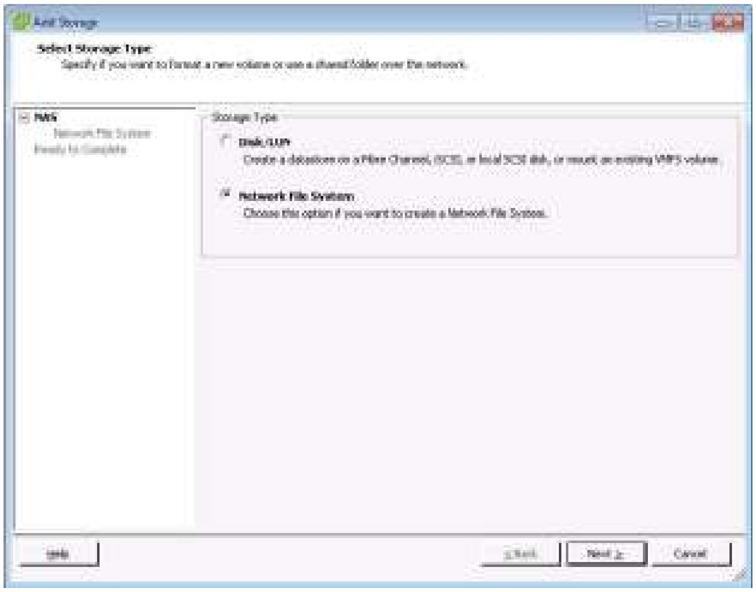

Select Storage Type
Specify if you want to format a new volume or use a shared folder over the network.

Storage Type

○ Disk/LUN
Create a datastore on a Fibre Channel, iSCSI, or local SCSI disk, or mount an existing VMFS volume.

● Network File System
Choose this option if you want to create a Network File System.

13. Select Network File System and click Next.
14. The wizard prompts for the location of the NFS export. Enter <<var_node01_nfs_lif_infra_swap_ip>> as the IP address for nfs_lif_infra_swap.
15. Enter /infra_swap as the path for the NFS export.
16. Confirm that the Mount NFS read only checkbox is not selected.
17. Enter infra_swap as the datastore name.

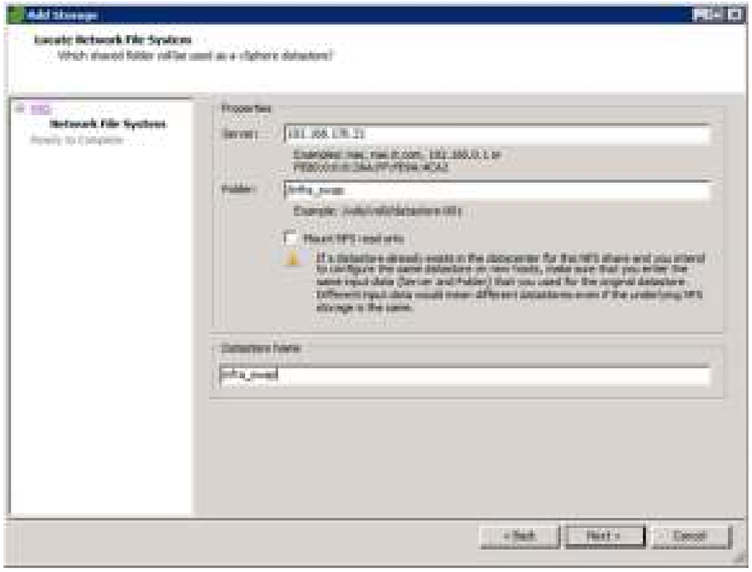

Locate Network File System
Which shared folder will be used as a vSphere datastore?

Properties

Server: 192.168.170.21

Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2

Folder: /infra_swap

Example: /vols/vol0/datastore-001

□ Mount NFS read only

If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.

Datastore Name

infra_swap

18. To continue with the NFS datastore creation, click Next.
19. To finalize the creation of the NFS datastore, click Finish.

Configure NTP on ESXi Hosts

ESXi Hosts VM-Host-Infra-01 and VM-Host-Prod-02

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. Click **Time Configuration** in the Software pane.
4. Click **Properties** at the upper-right side of the window.
5. At the bottom of the Time Configuration dialog box, click **Options**.
6. In the NTP Daemon (ntpd) Options dialog box, complete the following steps:

a. Click **General** in the left pane and select Start and stop with host.

b. Click **NTP Settings** in the left pane and click **Add**.

7. In the Add NTP Server dialog box, enter <<var_switch_a_ntp_ip>> as the IP address of the NTP server and click **OK**.
8. Click **Add**.
9. In the Add NTP Server dialog box, enter <<var_switch_b_ntp_ip>> as the IP address of the NTP server and click **OK**.
10. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes checkbox and click **OK**.
11. In the Time Configuration dialog box, complete the following steps:

a. Select the NTP Client Enabled checkbox and click **OK**.

b. Verify that the clock is now set to approximately the correct time.

The NTP server time may vary slightly from the host time.

Move VM Swap File Location

ESXi VM-Host-Infra-01 and VM-Host-Prod-02

To move the VM swap file location, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.
2. To enable configurations, click the **Configuration** tab.
3. Click **Virtual Machine Swapfil**e Location in the Software pane.
4. Click **Edit** at the upper-right side of the window.
5. Select “Store the swapfile in a swapfile datastore selected below.”
6. Select the <infra_swap> datastore in which to house the swap files.

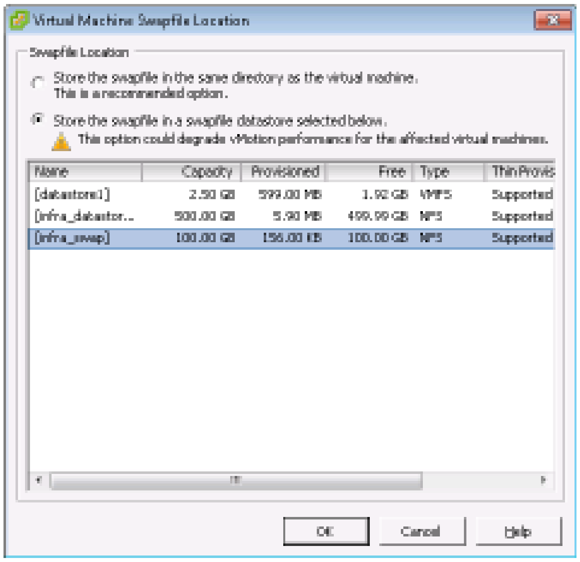

Virtual Machine Swapfile Location

Swapfile Location

○ Store the swapfile in the same directory as the virtual machine.
This is a recommended option.

● Store the swapfile in a swapfile datastore selected below.
This option could degrade vMotion performance for the affected virtual machines.

Name	Capacity	Provisioned	Free	Type	Thin Provis
[datastore1]	2.50 GB	599.00 MB	1.92 GB	VMFS	Supported
[infra_datastor...	500.00 GB	5.90 MB	499.99 GB	NFS	Supported
[infra_swap]	100.00 GB	156.00 KB	100.00 GB	NFS	Supported

7. Click OK to finalize moving the swap file location.

VMware vCenter **6.0 U1b**

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.0 U1b Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Install the Client Integration Plug-in

To install the client integration plug-in, complete the following steps:

1. Download the .iso installer for the vCenter Server Appliance and Client Integration Plug-in.
2. Mount the ISO image to the Windows virtual machine or physical server on which you want to install the Client Integration Plug-In to deploy the vCenter Server Appliance.
3. In the software installer directory, navigate to the vcsa directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe. The Client Integration Plug-in installation wizard appears.

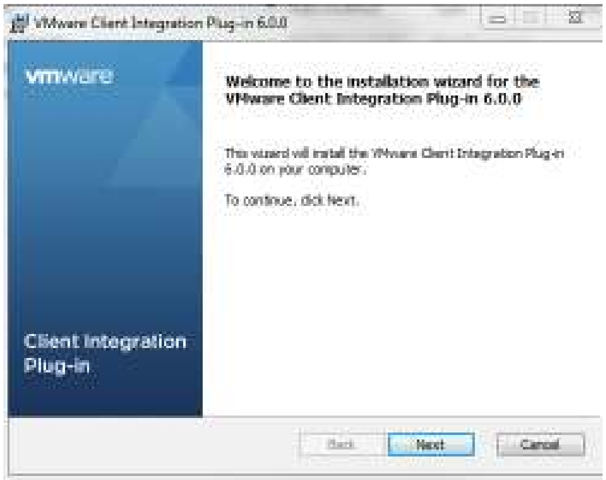
VMware Client Integration Plug-in 6.0.0
vmware
Client Integration Plug-in
Back
Next
Cancel

Welcome to the installation wizard for the VMware Client Integration Plug-in 6.0.0

This wizard will install the VMware Client Integration Plug-in 6.0.0 on your computer.

To continue, click Next.

4. On the Welcome page, click Next.
5. Read and accept the terms in the End-User License Agreement and click Next.
6. Click Next.
7. Click Install.

Building the VMware vCenter Server Appliance

To build the VMware vCenter virtual machine, complete the following steps:

1. In the software installer directory, double-click vcsa-setup.html.
2. Allow the plug-in to run on the browser when prompted.

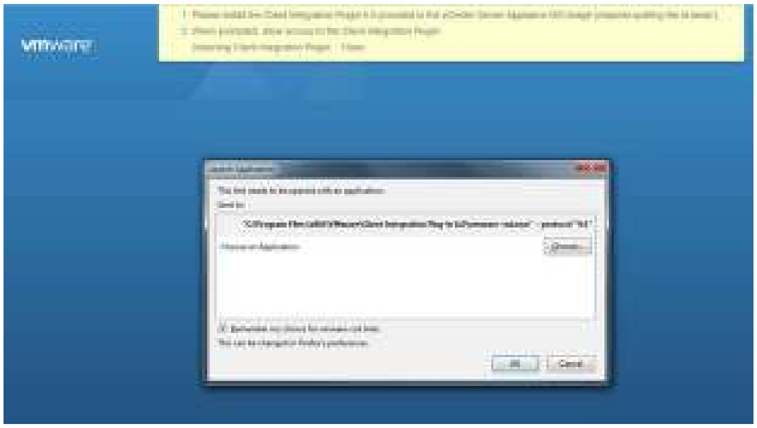

3. On the Home page, click Install to start the vCenter Server Appliance deployment wizard.

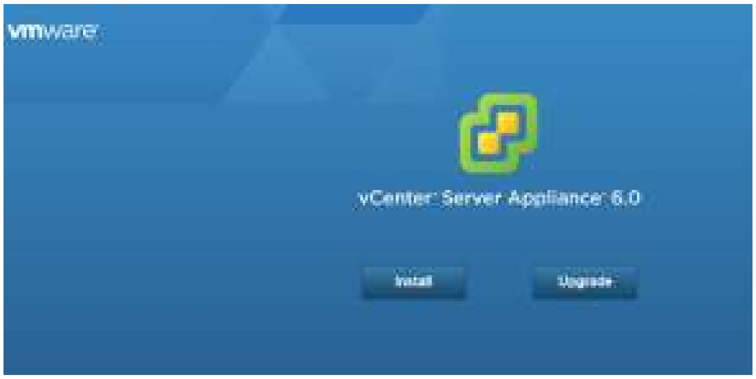

4. Read and accept the license agreement, and click Next.

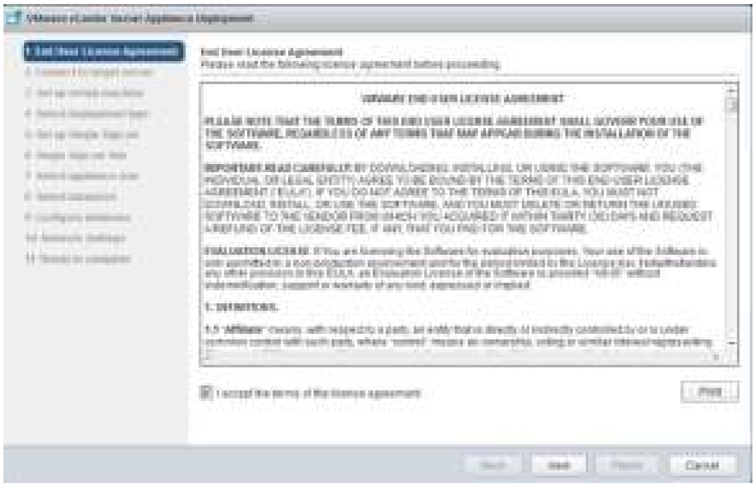

VMware vCenter Server Appliance Deployment

1 End User License Agreement
2 Connect to target server
3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Single Sign-on Site
7 Select appliance size
8 Select datastore
9 Configure database
10 Network Settings
11 Ready to complete

End User License Agreement
Please read the following license agreement before proceeding.

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE VENDOR FROM WHICH YOU ACQUIRED IT WITHIN THIRTY (30) DAYS AND REQUEST A REFUND OF THE LICENSE FEE, IF ANY, THAT YOU PAID FOR THE SOFTWARE.

EVALUATION LICENSE. If You are licensing the Software for evaluation purposes, Your use of the Software is only permitted in a non-production environment and for the period limited by the License Key. Notwithstanding any other provision in this EULA, an Evaluation License of the Software is provided "AS-IS" without indemnification, support or warranty of any kind, expressed or implied.

1. DEFINITIONS.

1.1 "Affiliate" means, with respect to a party, an entity that is directly or indirectly controlled by or is under common control with such party, where "control" means an ownership, voting or similar interest representing

☑ I accept the terms of the license agreement.

5. In the “Connect to target server” page, enter the ESXi host name, User name and Password.

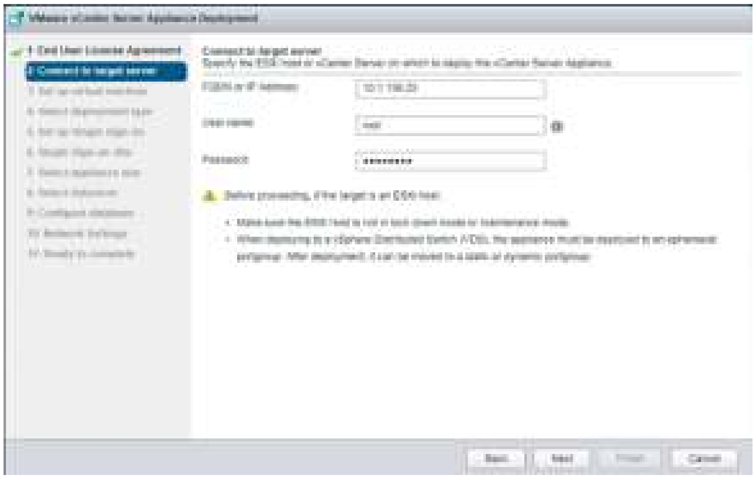

VMware vCenter Server Appliance Deployment

1 End User License Agreement
2 Connect to target server
3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Single Sign-on Site
7 Select appliance size
8 Select datastore
9 Configure database
10 Network Settings
11 Ready to complete

Connect to target server
Specify the ESXi host or vCenter Server on which to deploy the vCenter Server Appliance.

FQDN or IP Address: 10.1.156.20
User name: root
Password: *********

Before proceeding, if the target is an ESXi host:

- Make sure the ESXi host is not in lock down mode or maintenance mode.
- When deploying to a vSphere Distributed Switch (VDS), the appliance must be deployed to an ephemeral portgroup. After deployment, it can be moved to a static or dynamic portgroup.

6. Click Yes to accept the certificate.

7. Enter the Appliance name and password details in the “Set up virtual machine” page.

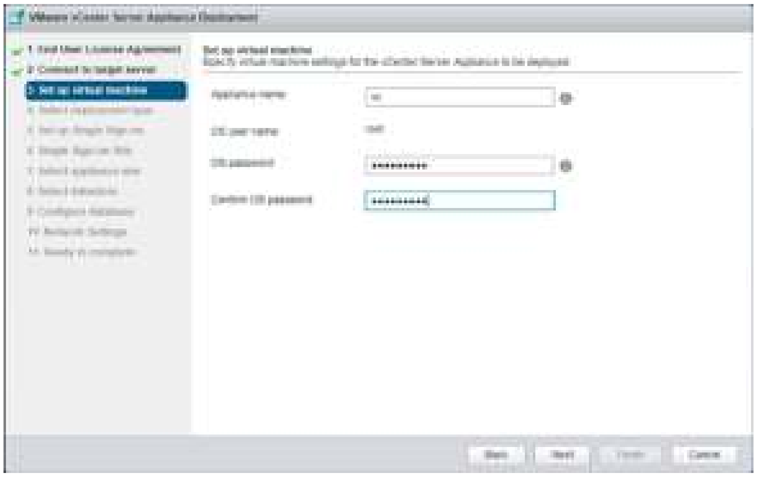

VMware vCenter Server Appliance Deployment

1 End User License Agreement
2 Connect to target server
3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Single Sign-on Site
7 Select appliance size
8 Select datastore
9 Configure database
10 Network Settings
11 Ready to complete

Set up virtual machine
Specify virtual machine settings for the vCenter Server Appliance to be deployed.

Appliance name: vc
OS user name: root
OS password: **********
Confirm OS password: **********

8. In the “Select deployment type” page, choose “Install vCenter Server with an embedded Platform Services Controller.”

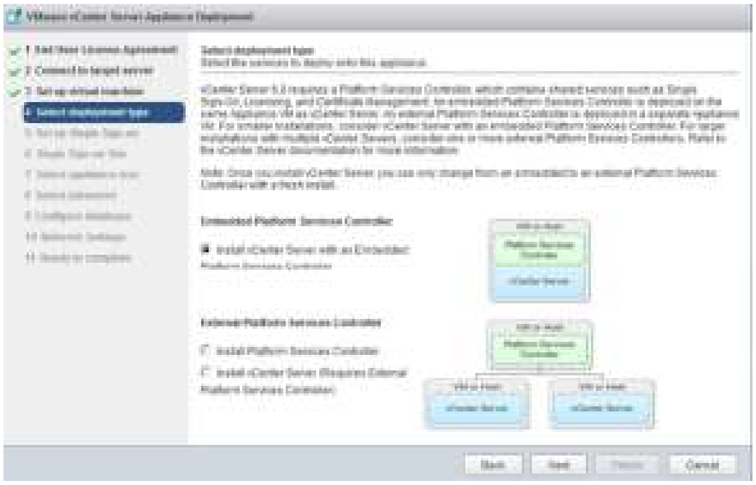

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- 4 Select deployment type
- 5 Set up Single Sign-on
- 6 Single Sign-on Site
- 7 Select appliance size
- 8 Select datastore
- 9 Configure database
- 10 Network Settings
- 11 Ready to complete

Select deployment type
Select the services to deploy onto this appliance.

vCenter Server 6.0 requires a Platform Services Controller, which contains shared services such as Single Sign-On, Licensing, and Certificate Management. An embedded Platform Services Controller is deployed on the same Appliance VM as vCenter Server. An external Platform Services Controller is deployed in a separate appliance VM. For smaller installations, consider vCenter Server with an embedded Platform Services Controller. For larger installations with multiple vCenter Servers, consider one or more external Platform Services Controllers. Refer to the vCenter Server documentation for more information.

Note: Once you install vCenter Server, you can only change from an embedded to an external Platform Services Controller with a fresh install.

Embedded Platform Services Controller

● Install vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

○ Install Platform Services Controller

○ Install vCenter Server (Requires External Platform Services Controller)

9. Click Next.
10. In the “Set up Single Sign-On” page, select “Create a new SSO domain.”
11. Enter the SSO password, Domain name and Site name.

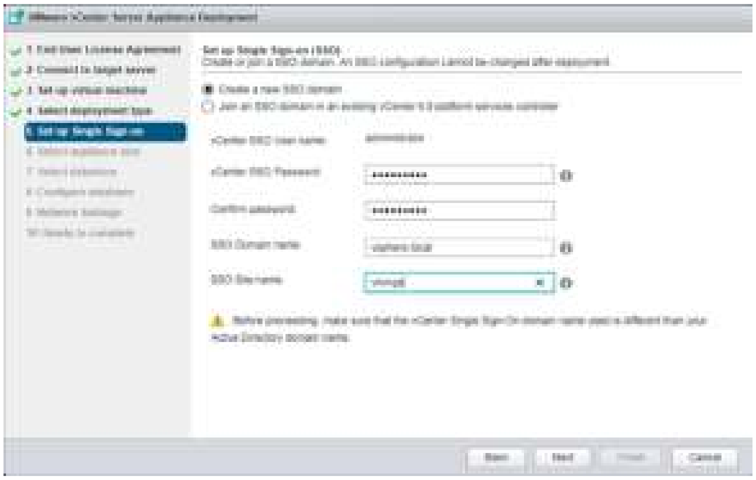

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- ✓ 4 Select deployment type
- 5 Set up Single Sign-on
- 6 Select appliance size
- 7 Select datastore
- 8 Configure database
- 9 Network Settings
- 10 Ready to complete

Set up Single Sign-on (SSO)
Create or join a SSO domain. An SSO configuration cannot be changed after deployment.

● Create a new SSO domain

○ Join an SSO domain in an existing vCenter 6.0 platform services controller

vCenter SSO User name: administrator

vCenter SSO Password: ••••••••••

Confirm password: ••••••••••

SSO Domain name: vsphere.local

SSO Site name: [illegible]

Before proceeding, make sure that the vCenter Single Sign-On domain name used is different than your Active Directory domain name.

12. Click Next.
13. Select the appliance size. For example, “Tiny (up to 10 hosts, 100 VMs).”

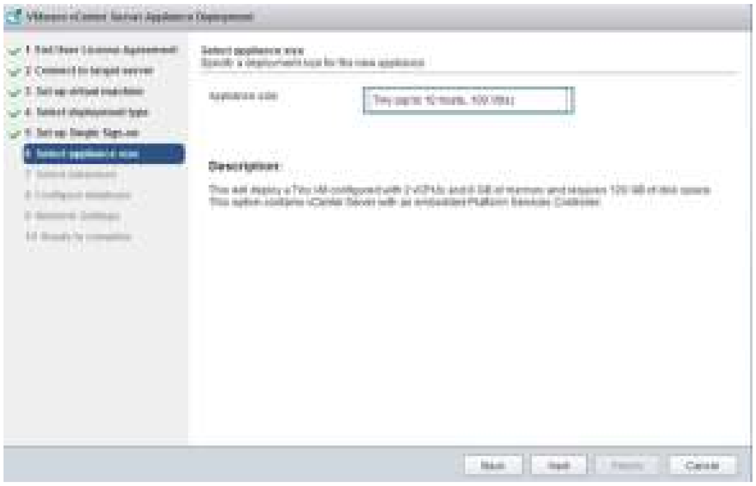

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- ✓ 4 Select deployment type
- ✓ 5 Set up Single Sign-on
- 6 Select appliance size
- 7 Select datastore
- 8 Configure database
- 9 Network Settings
- 10 Ready to complete

Select appliance size
Specify a deployment size for the new appliance

Appliance size: Tiny (up to 10 hosts, 100 VMs)

Description:

This will deploy a Tiny VM configured with 2 vCPUs and 8 GB of memory and requires 120 GB of disk space. This option contains vCenter Server with an embedded Platform Services Controller.

14. Click Next.
15. In the “Select datastore” page, choose infra_datastore_1.

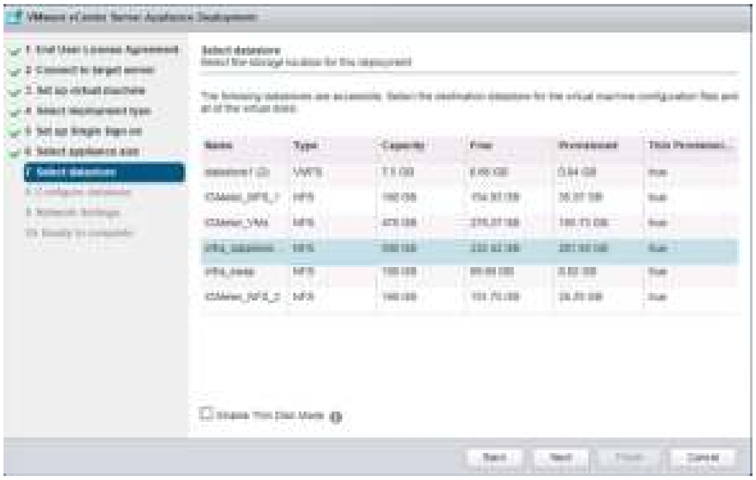

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- ✓ 4 Select deployment type
- ✓ 5 Set up Single Sign-on
- ✓ 6 Select appliance size
- 7 Select datastore
- 8 Configure database
- 9 Network Settings
- 10 Ready to complete

Select datastore
Select the storage location for this deployment

The following datastores are accessible. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Type	Capacity	Free	Provisioned	Thin Provisioni...
datastore1 (2)	VMFS	7.5 GB	6.66 GB	0.84 GB	true
ICMetal_NFS_1	NFS	160 GB	154.93 GB	35.37 GB	true
ICMetal_VMs	NFS	475 GB	275.27 GB	180.73 GB	true
infra_datastore...	NFS	500 GB	232.42 GB	287.93 GB	true
infra_swap	NFS	100 GB	99.89 GB	0.02 GB	true
ICMetal_NFS_2	NFS	160 GB	151.75 GB	24.25 GB	true

□ Enable Thin Disk Mode

16. Click Next.
17. Select embedded database in the “Configure database” page. Click Next.

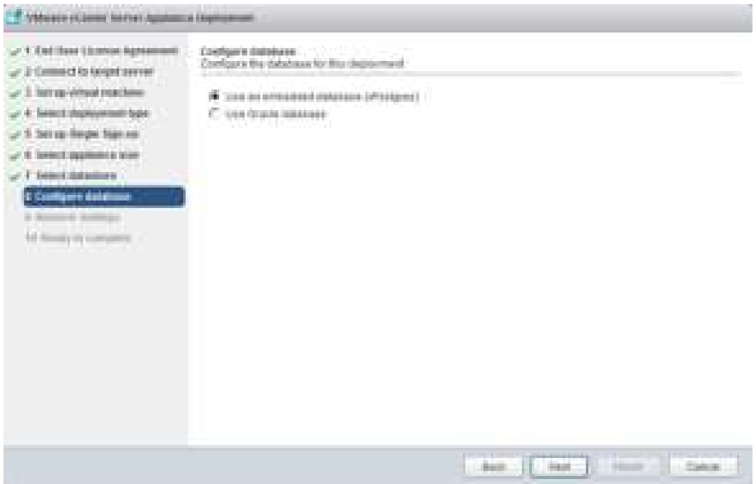

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- ✓ 4 Select deployment type
- ✓ 5 Set up Single Sign-on
- ✓ 6 Select appliance size
- ✓ 7 Select datastore
- 8 Configure database
- 9 Network Settings
- 10 Ready to complete

Configure database
Configure the database for this deployment

● Use an embedded database (vPostgres)
○ Use Oracle database

18. In the “Network Settings” page, configure the below settings:

a. Choose a Network: MGMT-Network

b. IP address family: IPV4

c. Network type: static

d. Network address: <<var_vcenter_ip>>

e. System name: <<var_vcenter_fqdn>>

f. Subnet mask: <<var_vcenter_subnet_mask>>

g. Network gateway: <<var_vcenter_gateway>>

h. Network DNS Servers: <<var_dns_server>>

i. Configure time sync: Use NTP servers

j. (Optional). Enable SSH

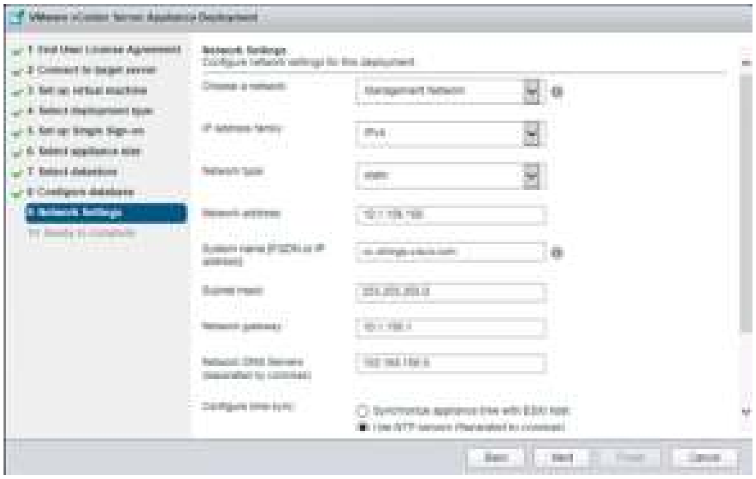

VMware vCenter Server Appliance Deployment

1 End User License Agreement
2 Connect to target server
3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Select appliance size
7 Select datastore
8 Configure database
9 Network Settings
10 Ready to complete

Network Settings
Configure network settings for this deployment

Choose a network: Management Network
IP address family: IPv4
Network type: static
Network address: 10.1.156.100
System name (FQDN or IP address): vc.vikings.cisco.com
Subnet mask: 255.255.255.0
Network gateway: 10.1.156.1
Network DNS Servers (separated by commas): 192.168.156.6
Configure time sync: ○ Synchronize appliance time with ESXi host
● Use NTP servers (Separated by commas)

19. Review the configuration and click Finish.

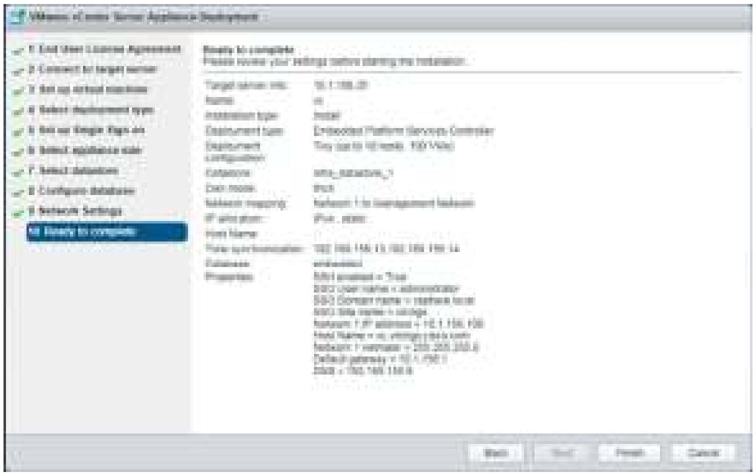

VMware vCenter Server Appliance Deployment

1 End User License Agreement
2 Connect to target server
3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Select appliance size
7 Select datastore
8 Configure database
9 Network Settings
10 Ready to complete

Ready to complete
Please review your settings before starting the installation.

Target server info: 10.1.156.20
Name: vc
Installation type: Install
Deployment type: Embedded Platform Services Controller
Deployment configuration: Tiny (up to 10 hosts, 100 VMs)
Datastore: infra_datastore_1
Disk mode: thick
Network mapping: Network 1 to Management Network
IP allocation: IPv4 , static
Host Name:
Time synchronization: 192.168.156.13,192.168.156.14
Database: embedded
Properties: SSH enabled = True
SSO User name = administrator
SSO Domain name = vsphere.local
SSO Site name = vikings
Network 1 IP address = 10.1.156.100
Host Name = vc.vikings.cisco.com
Network 1 netmask = 255.255.255.0
Default gateway = 10.1.156.1
DNS = 192.168.156.6

20. The vCenter appliance installation will take few minutes to complete.

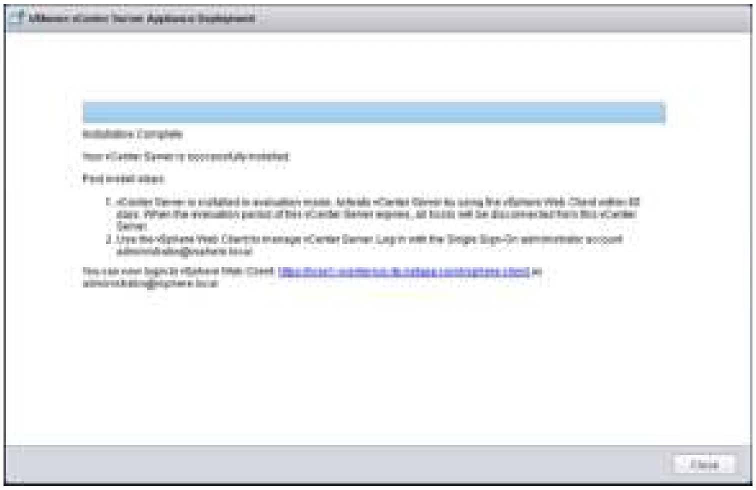

VMware vCenter Server Appliance Deployment

Installation Complete

Your vCenter Server is successfully installed.

Post install steps:

1. vCenter Server is installed in evaluation mode. Activate vCenter Server by using the vSphere Web Client within 60 days. When the evaluation period of this vCenter Server expires, all hosts will be disconnected from this vCenter Server.
2. Use the vSphere Web Client to manage vCenter Server. Log in with the Single Sign-On administrator account administrator@vsphere.local.

You can now login to vSphere Web Client: [illegible] as administrator@vsphere.local

Setting Up VMware vCenter Server

1. Using a web browser, navigate to https://<<var_vcenter_ip>.

2. Click Log in to vSphere Web Client.
3. Click OK if “Launch Application” window appears.

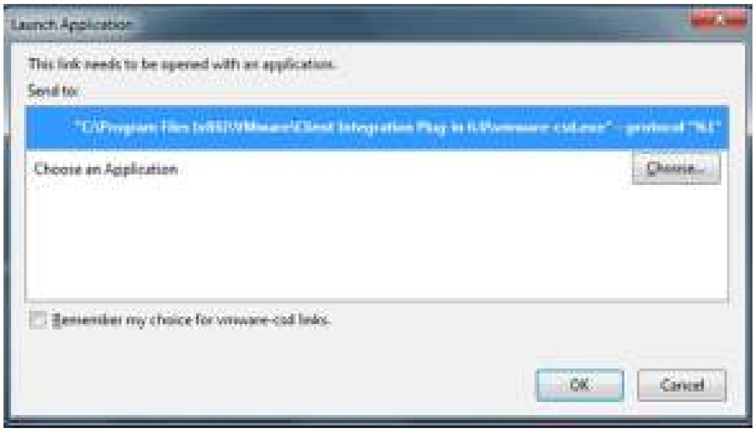

Launch Application

This link needs to be opened with an application.

Send to:

"C:\Program Files (x86)\VMware\Client Integration Plug-in 6.0\vmware-csd.exe" --protocol "%1"

Choose an Application

□ Remember my choice for vmware-csd links.

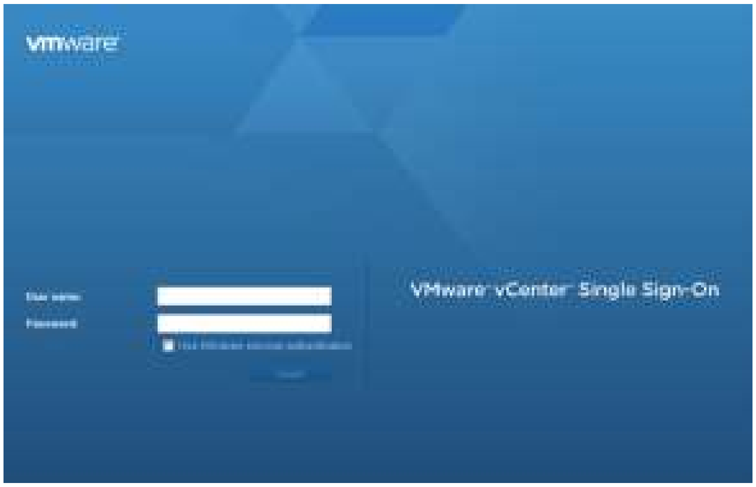

vmware
User name:
Password:
Use Windows session authentication
Login
VMware vCenter Single Sign-On

4. Log in using Single Sign-On username and password created during the vCenter installation.

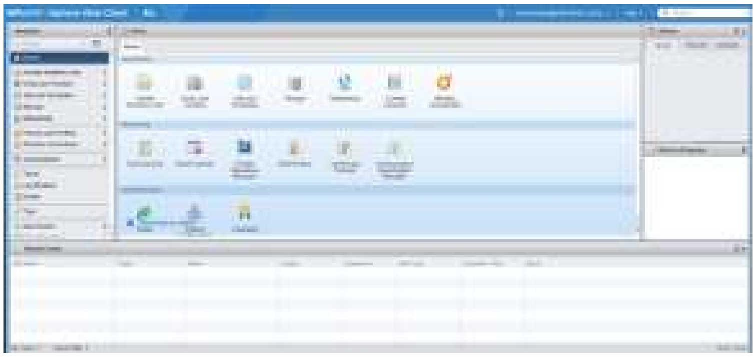

5. Navigate to vCenter Inventory Lists on the left pane.

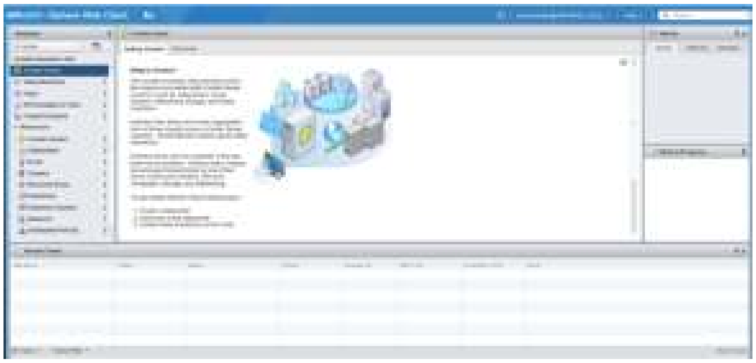

6. Under Resources, click Datacenters in the left plane.

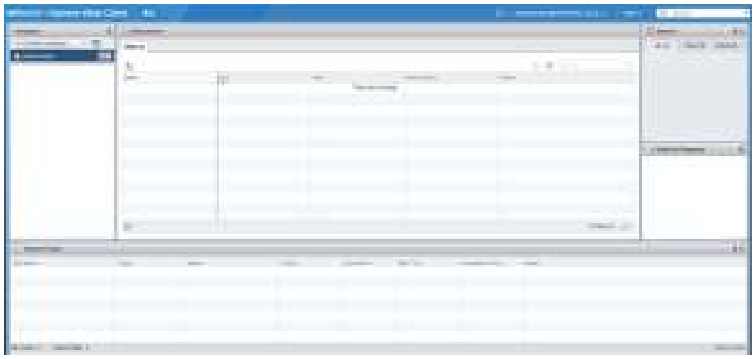

7. To create a Data center, click the leftmost icon in the center pane that has a green plus symbol above it.
8. Type “FlexPod_DC” in the Datacenter name field.
9. Select the vCenter Name/IP option.
10. Click OK.

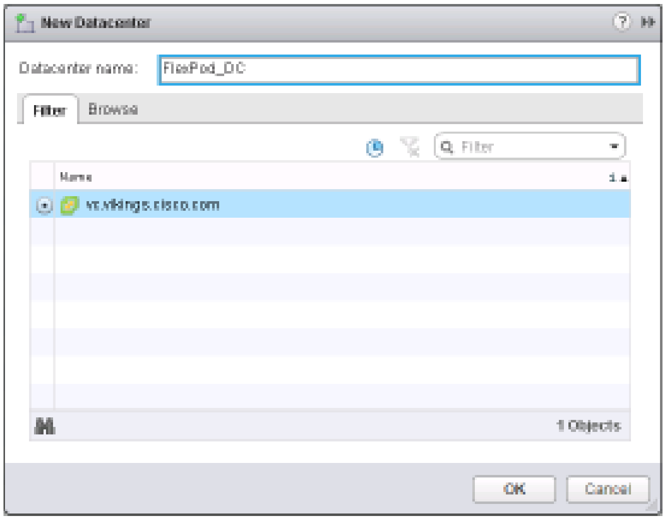

Datacenter name: FlexPod_DC

vc.vikings.cisco.com

11. Right-click the data center FlexPod_DC in the list in the center pane. Click New Cluster.

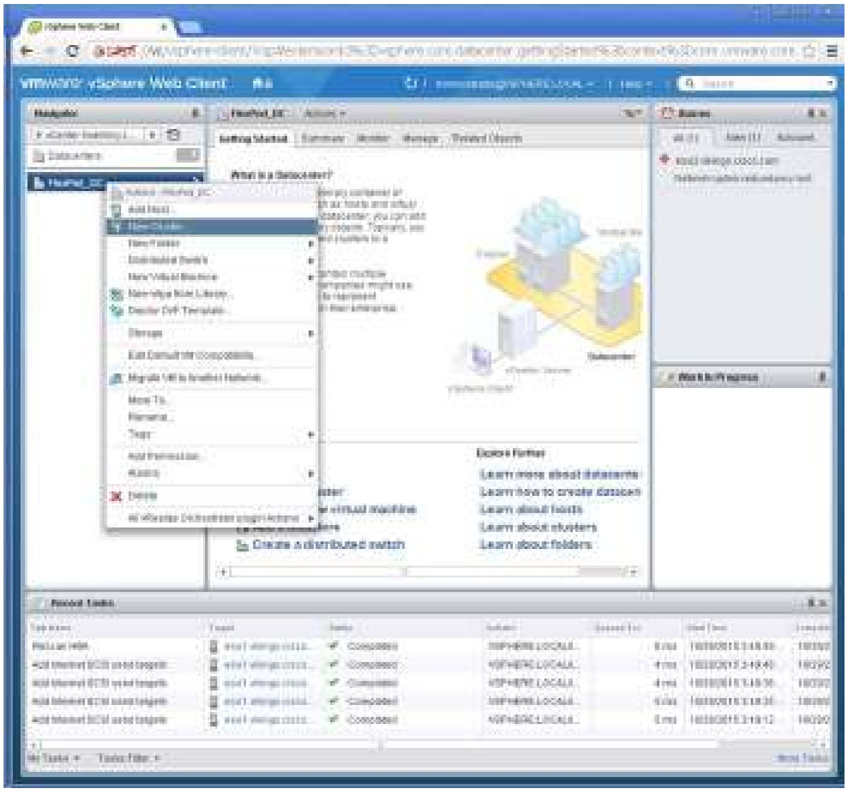

12. Name the cluster FlexPod_Management.
13. Check the box beside DRS. Leave the default values.
14. Check the box beside vSphere HA. Leave the default values.

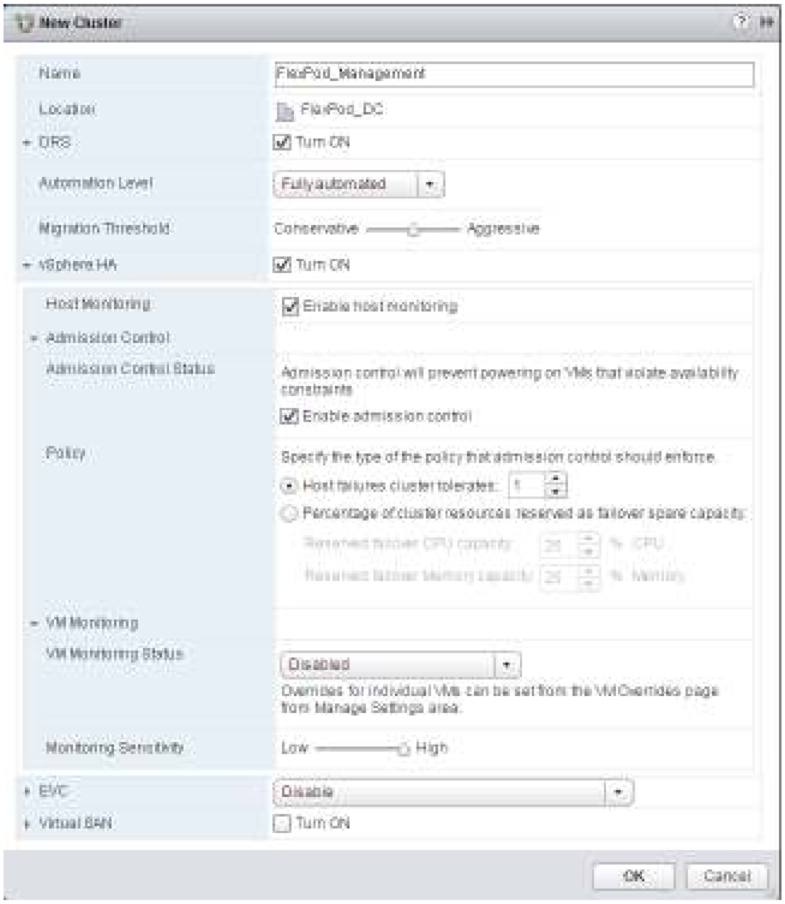
New Cluster

Name	FlexPod_Management
Location	FlexPod_DC
▾ DRS	☑ Turn ON
Automation Level	Fully automated ▾
Migration Threshold	Conservative ——○—— Aggressive
▾ vSphere HA	☑ Turn ON
Host Monitoring	☑ Enable host monitoring
▾ Admission Control	
Admission Control Status	Admission control will prevent powering on VMs that violate availability constraints ☑ Enable admission control
Policy	Specify the type of the policy that admission control should enforce. ● Host failures cluster tolerates: 1 ○ Percentage of cluster resources reserved as failover spare capacity: Reserved failover CPU capacity: 25 % CPU Reserved failover Memory capacity: 25 % Memory
▾ VM Monitoring	
VM Monitoring Status	Disabled ▾ Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low ————○ High
▸ EVC	Disable ▾
▸ Virtual SAN	□ Turn ON

OK Cancel

15. Click OK to create the new cluster.

16. On the left pane, double click the “FlexPod_DC”.
17. Click Clusters.

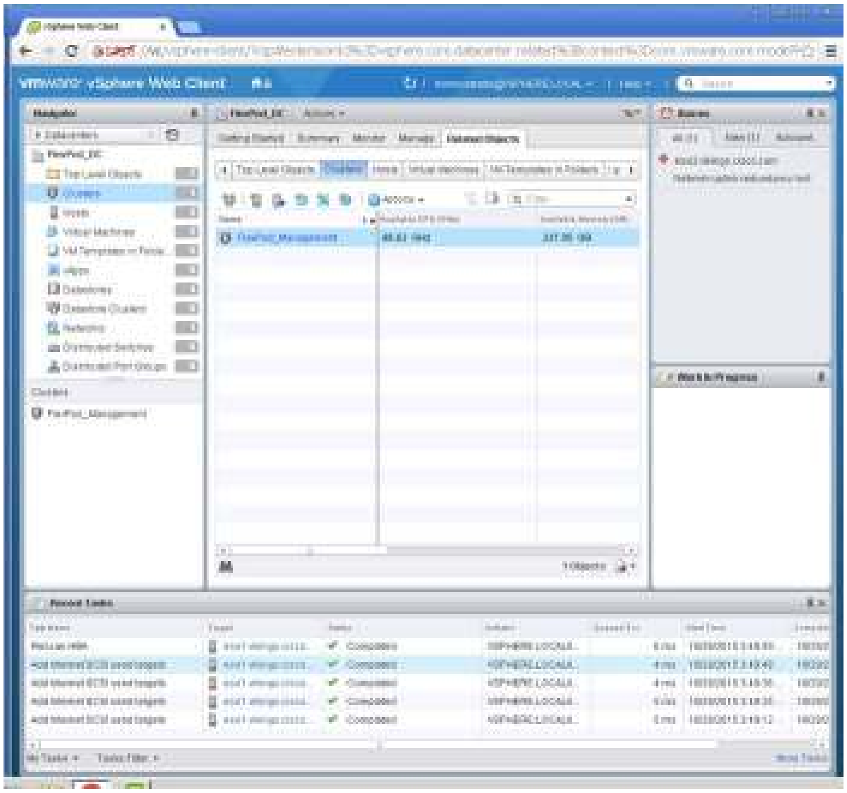

18. Under the Clusters pane, right click the “FlexPod_Management” and select Settings.
19. Select Configuration > General in the list on the left and select Edit to the right of General.
20. Select Datastore specified by host and click OK.

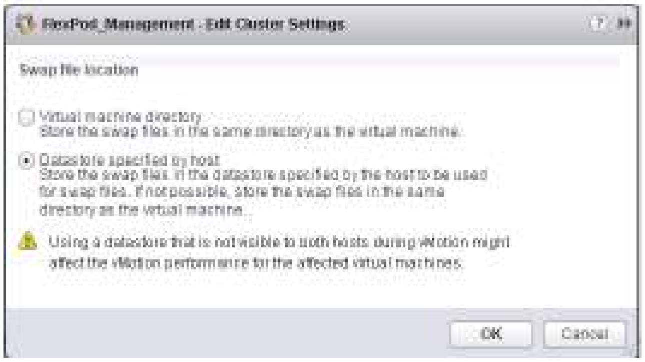

FlexPod_Management - Edit Cluster Settings

Swap file location

○ Virtual machine directory
Store the swap files in the same directory as the virtual machine.

● Datastore specified by host
Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine.

Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

OK Cancel

21. Under the Clusters pane, right click the “FlexPod_Management” and click Add Host.

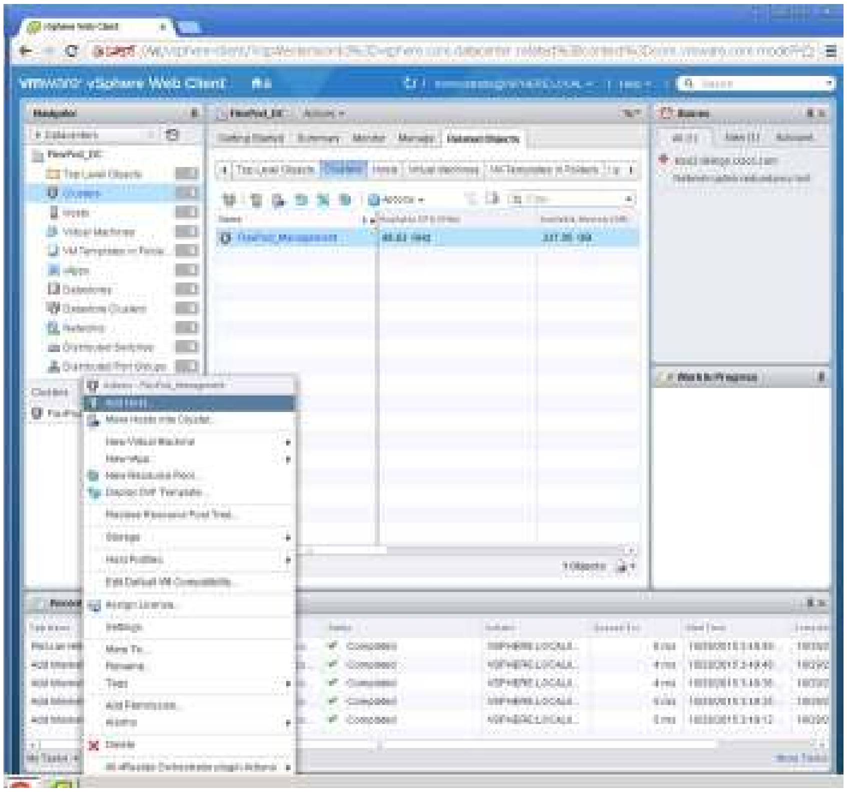

22. In the Host field, enter either the IP address or the host name of one of the VMware ESXi hosts. Click Next.
23. Type root as the user name and the root password. Click Next to continue.
24. Click Yes to accept the certificate.
25. Review the host details and click Next to continue.
26. Assign a license and click Next to continue.
27. Click Next to continue.
28. Click Next to continue.
29. Review the configuration parameters. Then click Finish to add the host.

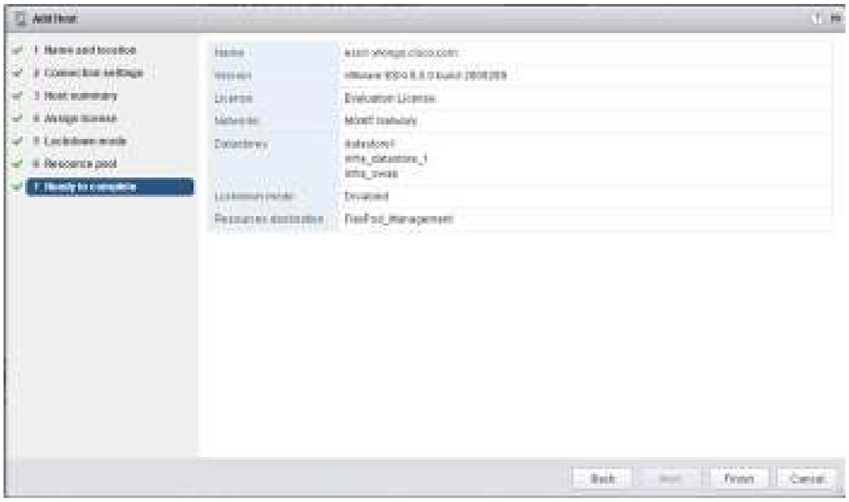

Name	[illegible]
Version	VMware ESXi 6.0.0 build-2809209
License	Evaluation License
Networks	MGMT Network
Datastores	datastore1 infra_datastore_1 infra_swap
Lockdown mode	Disabled
Resources destination	FlexPod_Management

30. Repeat the steps 21 to 29 to add the remaining VMware ESXi hosts to the cluster.

Two VMware ESXi hosts will be added to the cluster.

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. To setup the ESXi Dump Collector, complete the following steps:

1. In the vSphere web client, select Home.
2. In the center pane, click System Configuration.
3. In the left hand pane, click VMware vSphere ESXi Dump Collector.
4. In the Actions menu, choose Start.
5. In the Actions menu, click Edit Startup Type.
6. Select Automatic.
7. Click OK.
8. Connect to each ESXi host via ssh as root
9. Run the following commands:

```
esxcli system coredump network set --interface-name=vmk0 - -server-ipv4=10.1.156.100 --server-port=6500
esxcli  system coredump network set --enable=true
esxcli system coredump network check
```

FlexPod Cisco Nexus 1110-X and 1000V vSphere

This section provides detailed procedures for installing a pair of high-availability (HA) Cisco Nexus 1110-X Virtual Services Appliances (VSAs) in a FlexPod configuration. This validation effort used a preexisting management infrastructure to support the VSA devices and therefore does not document the cabling configuration.

Primary and standby Cisco Nexus 1000V Virtual Supervisor Modules (VSMs) are installed on the 1110-Xs and Cisco Nexus 1000V distributed virtual switch (DVS) will be provisioned. This procedure assumes that the Cisco Nexus 1000V software version 5.2(1)SV3(1.5b) has been downloaded from Cisco Nexus 1000V Download Link and expanded. It is recommended to install software version 5.2(1)SP1(7.3) on the Nexus 1110-Xs using Cisco Nexus Cloud Services Platform Software Installation and Upgrade Guide. Additionally, this procedure assumes that Cisco Virtual Switch Update Manager (VSUM) version 1.5.3 has been downloaded from Cisco VSUM Download Link and expanded. This procedure also assumes that VMware vSphere 6.0 Enterprise Plus licensing is installed.

Configure CIMC Interface on Both Cisco Nexus 1110-Xs

Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure the Cisco Integrated Management Controller (CIMC) interface on the Cisco Nexus 1110-X VSAs, complete the following steps:

1. Using the supplied dongle, connect a monitor and USB keyboard to the KVM console port on the front of the Cisco Nexus 1110-X virtual appliance.
2. Reboot the virtual appliance.
3. Press F8 when prompted to configure the CIMC interface.
4. Using the spacebar, set the NIC mode to Dedicated.
5. Clear the checkbox for DHCP enabled.
6. Set the CIMC IP address (<<var_cimc_ip>>) in the out-of-band management VLAN.
7. Set the CIMC subnet mask (<<var_cimc_mask>>).
8. Set the CIMC gateway (<<var_cimc_gateway>>).
9. Set the NIC redundancy to None.
10. Set and reenter the CIMC default password (<<var_password>>).
11. Press F10 to save the configuration.
12. Continue pressing F5 until Network settings configured is shown.
13. Press Esc to reboot the virtual appliance.

Configure Serial over LAN for Both Cisco Nexus 1110-Xs

Cisco Nexus 1110-X A and Cisco Nexus 1110-X B

To configure serial over LAN on the Cisco Nexus 1110-X VSAs, complete the following steps:

1. Use a Web browser to open the URL at http://<<var_cimc_ip>>.
2. Log in to the CIMC with the admin user id and the CIMC default password (<<var_password>>).
3. In the left column, click Remote Presence.
4. Click the Serial over LAN tab.
5. Select the Enabled checkbox for Serial over LAN Properties.
6. From the Baud Rate drop-down menu, select 9600 bps.
7. Click Save Changes.

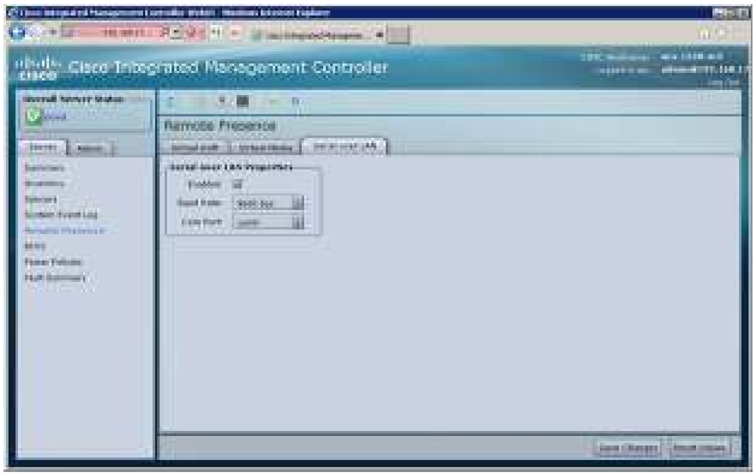

8. Log out of the CIMC Web interface.
9. Use an SSH client to connect to <<var_cimc_ip>> with the default CIMC user name and password.
10. Enter “connect host.”

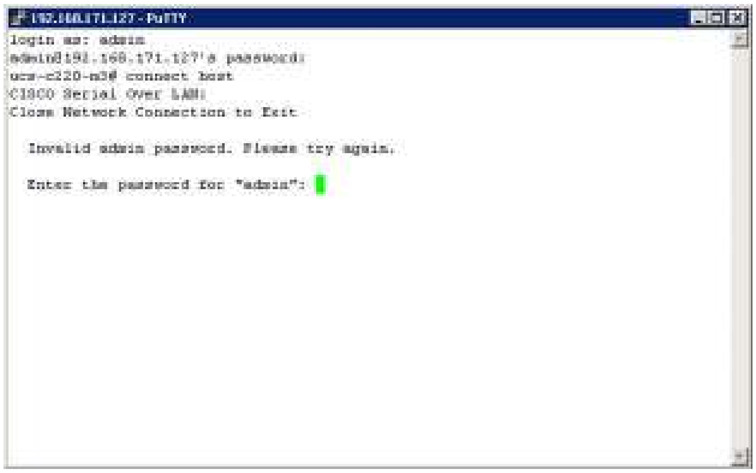

```
login as: admin
admin@192.168.171.127's password:
ucs-c220-m3# connect host
CISCO Serial Over LAN:
Close Network Connection to Exit

  Invalid admin password. Please try again.

  Enter the password for "admin":
```

Configure Cisco Nexus 1110-X Virtual Appliances

Cisco Nexus 1110-X A

To configure Cisco Nexus 1110-X A, complete the following steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.

```
Enter the password for “admin”: <<var_password>>
Confirm the password for “admin”: <<var_password>>
Enter HA role[primary/secondary]: primary
Enter the domain id<1-4095>: <<var_vsa_domain_id>>
Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
 Control Channel Setup.
Choose Uplink: < Gig:1,2 10Gig:7,8 NewPortChannel:0 >[0]: Enter
 Choose type of portchannel <ha/lacp>[ha]: lacp
PortChannel1 - Choose uplinks < Gig:1,2 10Gig:7,8 >[1,2]: 7,8
Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>
  Management Channel setup
  Choose Uplink: < Gig:1,2 Po1:9 NewPortChannel:0 >[9]: Enter
Would you like to enter the basic system configuration dialogue (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 VSA name : <<var_1110x_vsa>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IP address type V4/V6? (V4): V4
Mgmt0 IPv4 address : <<var_1110x_vsa_ip>>
Mgmt0 IPv4 netmask : <<var_1110x_vsa_mask>>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway : <<var_1110x_vsa_gateway>>
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 (das/rsa) [rsa]: Enter
Number of rsa key bits <768-2048> [1024]: Enter
Enable the http server? (yes/no) [y]: Enter
Configure the ntp server? (yes/no) [n]: y
  NTP server IPv4 address: <<var_switch_a_ntp_ip>>
```

2. Review the configuration summary. If everything is correct, enter no to skip editing the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter
```

3. The Cisco Nexus 1110-X saves the configuration and reboots. After reboot, log back in as admin.

Cisco Nexus 1110-X B

To configure the Cisco Nexus 1110-X B, complete the following steps:

1. Reboot the virtual appliance. The appliance should boot into a setup mode.3

```
Enter the password for “admin”: <<var_password>>
```

This is the same password that you entered on the primary Cisco Nexus 1110-X.

2. Enter the admin password again to confirm: <<var_password>>.

```
Enter HA role[primary/secondary]: secondary
Enter the domain id<1-4095>: <<var_vsa_domain_id>>
```

This is the same domain id that you entered on the primary Cisco Nexus 1110-X.

```
Enter control vlan <1-3967, 4048-4093>: <<var_pkt-ctrl_vlan_id>>
  Control Channel Setup.
 Choose Uplink: < Gig:1,2 10Gig:7,8 NewPortChannel:0 >[0]: Enter
  Choose type of portchannel <ha/lacp>[ha]: lacp
 PortChannel1 – Choose uplinks < Gig:1,2 10Gig:7,8 >[1,2]: 7,8
Enter management vlan <1-3967, 4048-4093>: <<var_ib-mgmt_vlan_id>>
  Management Channel setup
  Choose Uplink: < Gig:1,2 Po1:9 NewPortChannel:0 >[9]: Enter
```

3. The Cisco Nexus 1110-X saves the configuration and reboots.

Set Up the Primary Cisco Nexus 1000V VSM

Cisco Nexus 1110-X A

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, complete the following steps:

These steps are completed from the primary Nexus 1110-X A

1. Continue periodically running the following command until module 2 (Cisco Nexus 1110-X B) has a status of ha-standby.

```
show module
```

2. Enter the global configuration mode and create a virtual service blade.

```
config t
virtual-service-blade VSM-1
dir /repository
```

3. If the desired Cisco Nexus 1000V ISO file (n1000v-dk9.5.2.1.SV3.1.5b.iso) is not present on the Cisco Nexus 1110-X, run the copy command to copy it to the Cisco Nexus 1110-X disk. You must place the file either on an FTP server or on a UNIX or Linux® machine (using scp) that is accessible from the Cisco Nexus 1110-X management interface. An example copy command from an FTP server is copy ftp://<<var_ftp_server>>/n1000v-dk9.5.2.1.SV3.1.5b.iso /repository/.

```
virtual-service-blade-type new n1000v-dk9.5.2.1.SV3.1.5b.iso
interface control vlan <<var_pkt-ctrl_vlan_id>>
interface packet vlan <<var_pkt-ctrl_vlan_id>>
enable primary
Enter vsb image:[n1000v-dk9.5.2.1.SV3.1.5b.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
```

This domain ID should be different than the VSA domain ID.

```
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_mgmt_ip>>
Enter Management subnet mask: <<var_vsm_mgmt_mask>>
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>
Enter HostName: <<var_vsm_hostname>>
Enter the password for 'admin': <<var_password>>
```

This password must be entered with only uppercase and lowercase letters. No special characters can be used in this password.

```
Do you want to continue with installation with entered details (Y/N)? [Y} Enter
copy run start
```

4. Run show virtual-service-blade summary. Continue periodically entering this command until the primary VSM-1 has a state of VSB POWERED ON.
5. Modify the management, control and packet interface and set PortChannel 1 as the uplink interface (if needed):

```
virtual-service-blade VSM-1
  interface control uplink PortChannel1
  interface management uplink PortChannel1
  interface packet uplink PortChannel1
```

Set Up the Secondary Cisco Nexus 1000V VSM

To set up the secondary Cisco Nexus 1000V VSM on Cisco Nexus 1110-X B, complete the steps in the following two subsections:

Cisco Nexus 1110-X A

```
enable secondary
Enter vsb image: [n1000v-dk9.5.2.1.SV3.1.5b.iso] Enter
Enter domain id[1-4095]: <<var_vsm_domain_id>>
Enter SVS Control mode (L2 / L3): [L3] Enter
Management IP version [V4/V6]: [V4] Enter
Enter Management IP address: <<var_vsm_ mgmt_ip>>
Enter Management subnet mask: <<var_vsm_ mgmt_mask>>
IPv4 address of the default gateway: <<var_vsm_mgmt_gateway>>
Enter HostName: <<var_vsm_hostname>>
Enter the password for ‘admin’: : <<var_password>>
This password must be entered with only uppercase and lowercase letters. No special characters can be used in this password. Do you want to continue installation with entered details (Y/N)? [Y}
```

Type show virtual-service-blade summary. Continue periodically entering this command until both the primary and secondary VSM-1s have a state of VSB POWERED ON and Roles are correctly identified.

```
copy run start
```

Install Cisco Virtual Switch Update Manager

VMware vSphere Web Client

To install the Cisco Virtual Switch Upgrade Manager from OVA in the VMware virtual environment, complete the following steps:

1. Log into the VMware vSphere Web Client.
2. In the pane on the right, click VMs and Templates.
3. In the center pane, select Actions > Deploy OVF Template.
4. Select Browse and browse to and select the Nexus1000v-vsum.2.0.ova file.
5. Click Open.
6. Click Next.

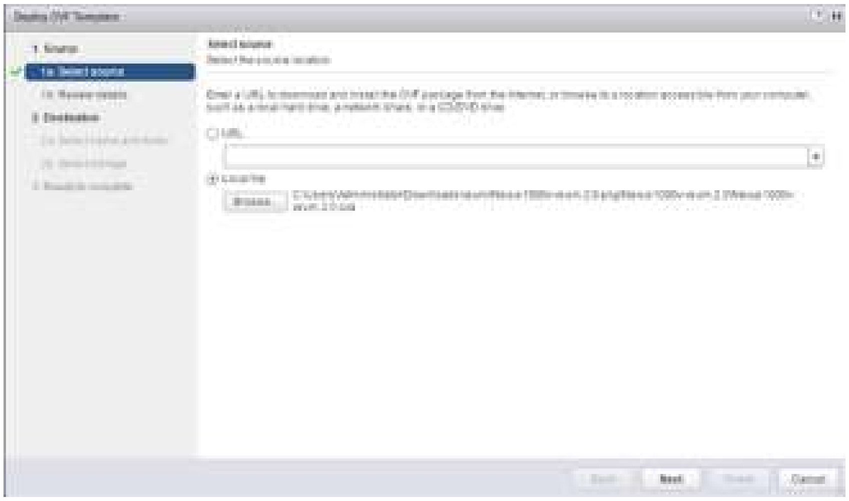

7. Review the details and click Next.

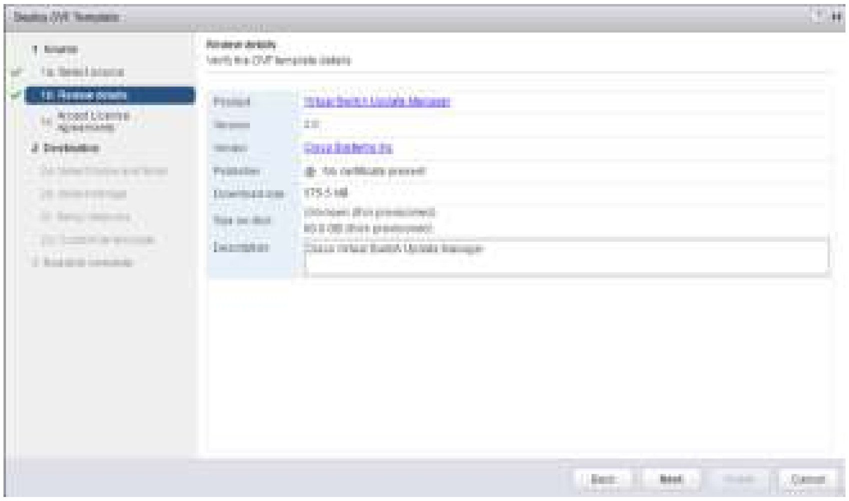

Product: Virtual Switch Update Manager
Version: 2.0
Vendor: Cisco Systems Inc
Publisher: No certificate present
Download size: 175.5 MB
Description: Cisco Virtual Switch Update Manager

8. Click Accept to accept the License Agreement and click Next.
9. Name the Virtual Machine, select the FlexPod_DC datacenter and click Next.
10. Select the FlexPod_Management cluster and click Next.
11. Select infra_datastore_1 and the Thin Provision virtual disk format and click Next.

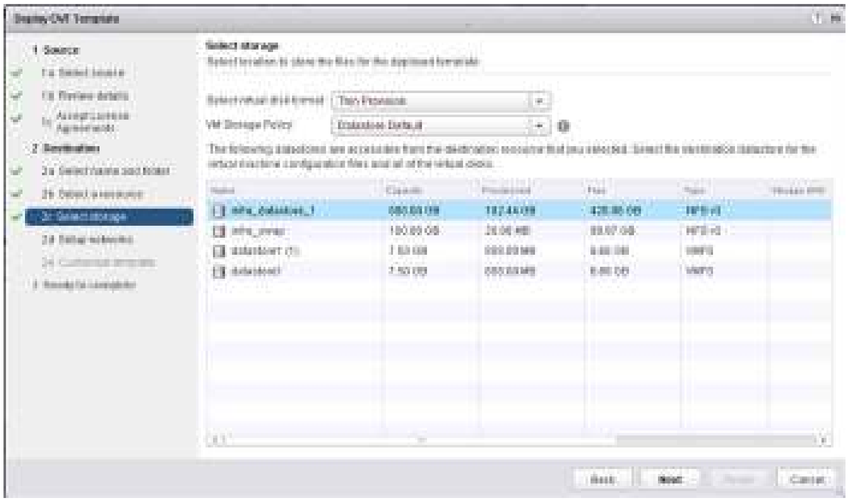

Deploy OVF Template

Select storage

Select virtual disk format: Thin Provision

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type
infra_datastore_1	[illegible]	182.44 GB	420.80 GB	NFS v3
infra_swap	100.00 GB	20.06 MB	99.97 GB	NFS v3
datastore1 (1)	7.50 GB	[illegible]	[illegible]	VMFS
datastore1	7.50 GB	[illegible]	[illegible]	VMFS

12. Select the IB-MGMT-VLAN Network and click Next.

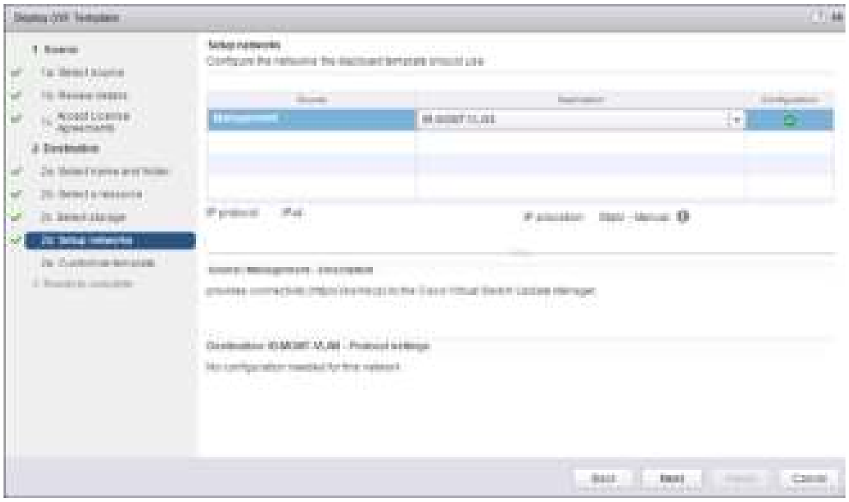

Deploy OVF Template

Setup networks

Configure the networks the deployed template should use

Source	Destination
Management	IB-MGMT-VLAN

Destination: IB-MGMT-VLAN - Protocol settings

No configuration needed for this network

13. Fill in the Networking Properties within the Customize template dialog.

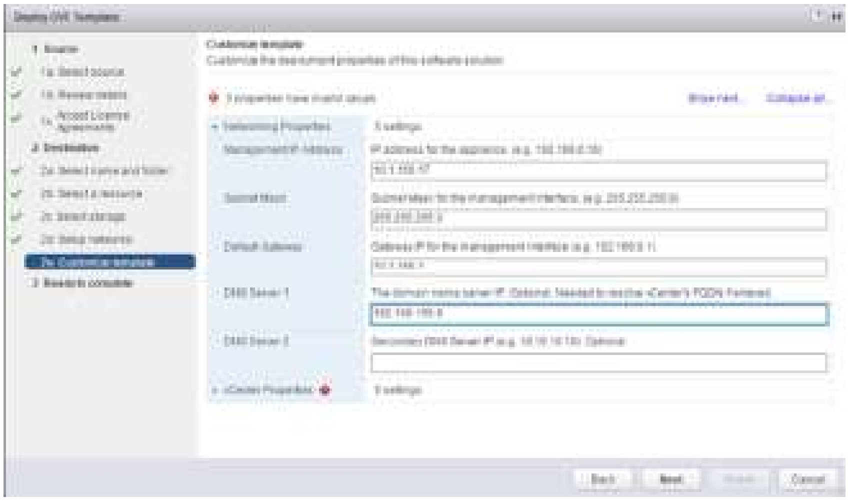

Deploy OVF Template

Customize template

Networking Properties

Management IP Address

Subnet Mask

Default Gateway

DNS Server 1

DNS Server 2

vCenter Properties

14. Expand the vCenter Properties and fill in vCenter address, Username, and Password fields.

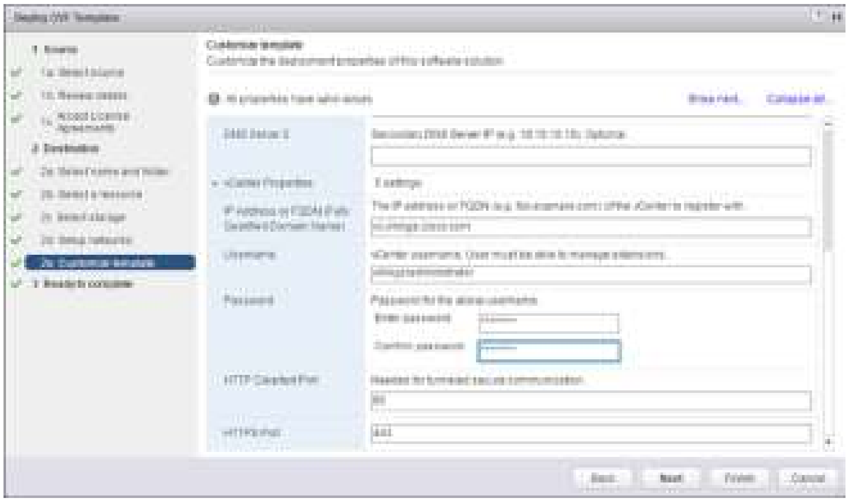

15. Click Next.
16. Review all settings and click Finish.
17. Wait for the Deploy OVF template task to complete.
18. Select the Home button in VMware vSphere Web Client and select Hosts and Clusters.
19. Expand the FlexPod_Management cluster and select the Virtual Switch Update Manager VM from the Summary tab.
20. Click the Console graphic at the top left of the Summary tab. If a security warning pops up, click Allow.
21. If a security certificate warning pops up, click Connect Anyway.
22. Power on the Virtual Switch Update Manager VM.
23. Once the VM has completely booted up, log out and log back into the VMware vSphere Web Client.

Register the Cisco Nexus 1000V in VMware vCenter

VMware vSphere Web Client

To register the Cisco Nexus 1000V, complete the following steps:

1. After logging back into the VMware vSphere Web Client, Cisco Virtual Switch Update Manager should now appear under the Home tab. Select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Click Install.
4. In the pane on the right, select FlexPod_DC.
5. Under Nexus1000v Switch Deployment Process, select I already have a control place (VSM) deployed.
6. Enter the IP Address of the VSM and the admin password.
7. Click Finish.
8. Click the Home button.
9. Select Cisco Virtual Switch Update manager.
10. Under Basic tasks, select Nexus 1000v.
11. Click Configure.
12. In the pane on the right, select FlexPod_DC.
13. The Nexus 1000v Switch should appear under the *Choose an associated Distributed Virtual Switch* section.

Perform Base Configuration of the Primary VSM

SSH Connection to Primary VSM

To perform the base configuration of the primary VSM, complete the following steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands.

```
config t
ntp server <<var_switch_a_ntp_ip>> use-vrf management
ntp server <<var_switch_b_ntp_ip>> use-vrf management
vlan <<var_ib-mgmt_vlan_id>>
name IB-MGMT-VLAN
vlan <<var_nfs_vlan_id>>
name NFS-VLAN
vlan <<var_vmotion_vlan_id>>
name vMotion-VLANvlan <<var_vm-traffic_vlan_id>>
name VM-Traffic-VLAN
vlan <<var_native_vlan_id>>
name Native-VLAN
vlan <<var_iscsi_a_vlan_id>>
name iSCSI-A-VLAN
vlan <<var_iscsi_b_vlan_id>>
name iSCSI-B-VLAN
vlan <<var_pkt-ctrl_vlan_id>>
name Pkt-Ctrl-VLAN
exit
port-profile type ethernet system-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-traffic_vlan_id>>
```

Any VLAN that has a VMKernal port should be assigned as a system vlan on both the uplink and the vEthernet ports of the virtual switch.

```
system mtu 9000
state enabled
port-profile type ethernet iscsi-a-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_a_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_a_vlan_id>>
no shutdown
system vlan <<var_iscsi_a_vlan_id>>
system mtu 9000
state enabled
port-profile type ethernet iscsi-b-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_iscsi_b_vlan_id>>
switchport trunk allowed vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
system mtu 9000
state enabled
port-profile type vethernet IB-MGMT-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet NFS-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_nfs_vlan_id>>
no shutdown
system vlan <<var_nfs_vlan_id>>
state enabled
port-profile type vethernet vMotion-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vmotion_vlan_id>>
```

```
no shutdown
system vlan <<var_vmotion_vlan_id>>
state enabled
port-profile type vethernet VM-Traffic-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vm-traffic_vlan_id>>
no shutdown
system vlan <<var_vm-traffic_vlan_id>>
state enabled
port-profile type vethernet n1kv-L3
capability l3control
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>>
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>
state enabled
port-profile type vethernet iSCSI-A-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_a_vlan_id>>
no shutdown
system vlan <<var_iscsi_a_vlan_id>>
state enabled
port-profile type vethernet iSCSI-B-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_iscsi_b_vlan_id>>
no shutdown
system vlan <<var_iscsi_b_vlan_id>>
state enabled
exit
copy run start
```

Add VMware ESXi Hosts to Cisco Nexus 1000V

VMware vSphere Web Client

To and VMware ESXi hosts, complete the following steps:

1. Back in the Vmware vSphere Web Client, from the Home tab, select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Select Configure.
4. Select the FlexPod_DC datacenter on the right.
5. Select the VSM on the lower right.
6. Click Manage.
7. In the center pane, select the Add Host tab.
8. Expand the FlexPod_Management ESXi Cluster and select both FlexPod Management Hosts.
9. Click Suggest.
10. Scroll down to Physical NIC Migration and expand each ESXi host.
11. On both hosts, unselect vmnic0, and select vmnic1. For vmnic1, select the system-uplink Profile. Select vmnic2 and select the iscsi-a-uplink Profile. Select vmnic3 and select the iscsi-b-uplink Profile.

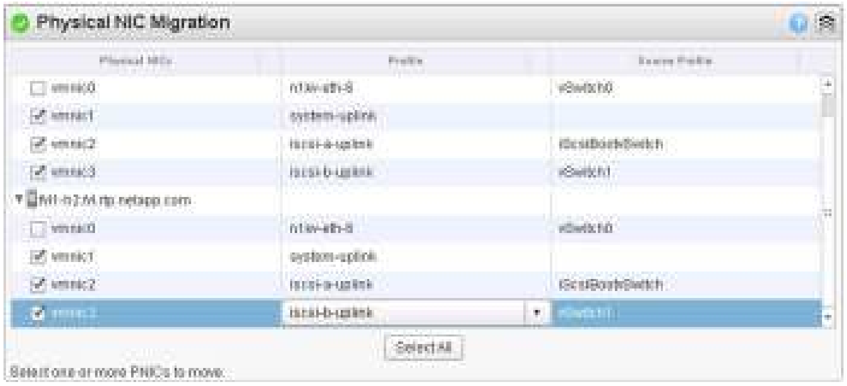

Physical NIC Migration

Physical NICs	Profile	Source Profile
☐ vmnic0	n1kv-eth-8	vSwitch0
☑ vmnic1	system-uplink	
☑ vmnic2	iscsi-a-uplink	iScsiBootvSwitch
☑ vmnic3	iscsi-b-uplink	vSwitch1
▾ M1-n2-M.rtp.netapp.com		
☐ vmnic0	n1kv-eth-8	vSwitch0
☑ vmnic1	system-uplink	
☑ vmnic2	iscsi-a-uplink	iScsiBootvSwitch
☑ vmnic3	iscsi-b-uplink ▾	vSwitch1

Select one or more PNICs to move.

12. Scroll down to VM Kernel NIC Setup and expand both ESXi hosts.
13. All VMkernel ports should already have the appropriate checkboxes selected.

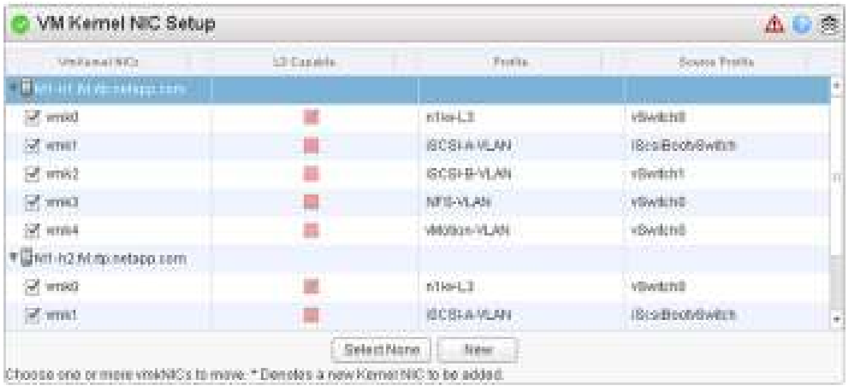

VM Kernel NIC Setup

VMkernel NICs	L3 Capable	Profile	Source Profile
M1-h1.M.rtp.netapp.com			
☑ vmk0	■	n1kv-L3	vSwitch0
☑ vmk1	■	iSCSI-A-VLAN	iScsiBootvSwitch
☑ vmk2	■	iSCSI-B-VLAN	vSwitch1
☑ vmk3	■	NFS-VLAN	vSwitch0
☑ vmk4	■	vMotion-VLAN	vSwitch0
M1-h2.M.rtp.netapp.com			
☑ vmk0	■	n1kv-L3	vSwitch0
☑ vmk1	■	iSCSI-A-VLAN	iScsiBootvSwitch

Choose one or more vmkNIC's to move. * Denotes a new Kernel NIC to be added.

14. Scroll down to VM Migration and expand both ESXi hosts.

15. Select the IB-MGMT-VLAN profile for the VSUM and vCenter Virtual Machines.

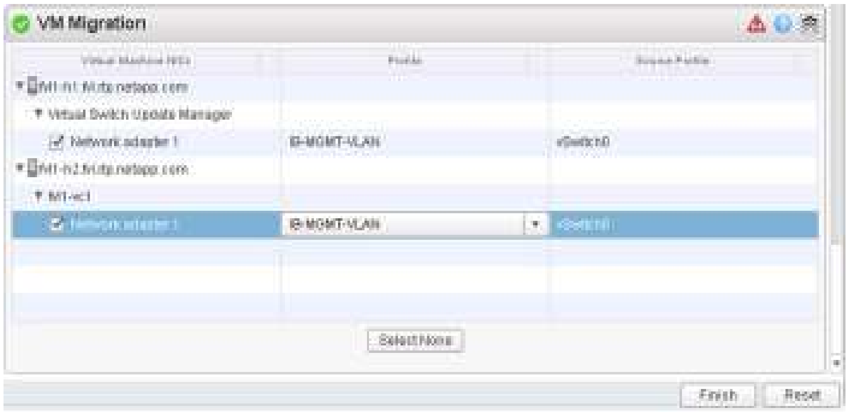

VM Migration

Virtual Machine NICs	Profile	Source Profile
M1-h1.M.rtp.netapp.com		
Virtual Switch Update Manager		
☑ Network adapter 1	IB-MGMT-VLAN	vSwitch0
M1-h2.M.rtp.netapp.com		
M1-vc1		
☑ Network adapter 1	IB-MGMT-VLAN	vSwitch0

16. Click Finish.

The progress of the virtual switch installation can be monitored from the c# interface.

Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V

To migrate the ESXi host redundant network ports, complete the following steps:

1. In the VMware vSphere Web Client window, select Home > Hosts and Clusters.
2. On the left expand the Datacenter and cluster, and select the first VMware ESXi host.
3. In the center pane, select the Manage tab, then select Networking.
4. Select vSwitch0. All of the port groups on vSwitch0 should be empty. Click the red X under Virtual switches to delete vSwitch0.
5. Click Yes to remove vSwitch0. It may be necessary to refresh the Web Client to see the deletion.
6. Delete iScsiBootvSwitch and vSwitch1.
7. The Nexus 1000V VSM should now be the only virtual switch. Select it and select the third icon above it under Virtual switches (Manage the physical network adapters connected to the selected switch).
8. Click the green plus sign to add an adapter.
9. For UpLink03, select the system-uplink port group and make sure vmnic0 is the Network adapter. Click OK.
10. Click OK to complete adding the Uplink. It may be necessary to refresh the Web Client to see the addition.
11. Repeat this procedure for the second ESXi host.
12. From the SSH client that is connected to the Cisco Nexus 1000V, run show interface status to verify that all interfaces and port channels have been correctly configured.

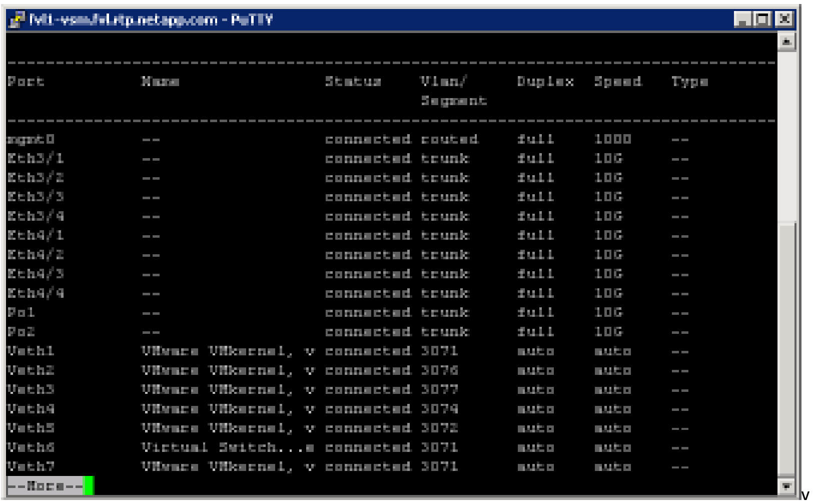

```
--------------------------------------------------------------------------------
Port          Name                 Status     Vlan/      Duplex  Speed   Type
                                              Segment
--------------------------------------------------------------------------------
mgmt0         --                   connected routed     full    1000    --
Eth3/1        --                   connected trunk      full    10G     --
Eth3/2        --                   connected trunk      full    10G     --
Eth3/3        --                   connected trunk      full    10G     --
Eth3/4        --                   connected trunk      full    10G     --
Eth4/1        --                   connected trunk      full    10G     --
Eth4/2        --                   connected trunk      full    10G     --
Eth4/3        --                   connected trunk      full    10G     --
Eth4/4        --                   connected trunk      full    10G     --
Po1           --                   connected trunk      full    10G     --
Po2           --                   connected trunk      full    10G     --
Veth1         VMware VMkernel, v connected 3071       auto    auto    --
Veth2         VMware VMkernel, v connected 3076       auto    auto    --
Veth3         VMware VMkernel, v connected 3077       auto    auto    --
Veth4         VMware VMkernel, v connected 3074       auto    auto    --
Veth5         VMware VMkernel, v connected 3072       auto    auto    --
Veth6         Virtual Switch...e connected 3071       auto    auto    --
Veth7         VMware VMkernel, v connected 3071       auto    auto    --
--More--
```

v

13. Run show module and verify that the two ESXi hosts are present as modules.

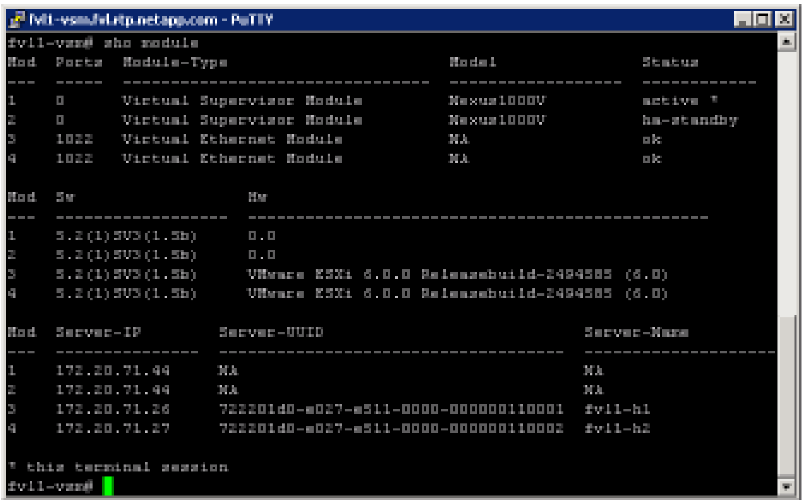

```
fvl1-vsm# sho module
Mod  Ports  Module-Type                       Model               Status
---  -----  --------------------------------  ------------------  ------------
1    0      Virtual Supervisor Module         Nexus1000V          active *
2    0      Virtual Supervisor Module         Nexus1000V          ha-standby
3    1022   Virtual Ethernet Module           NA                  ok
4    1022   Virtual Ethernet Module           NA                  ok

Mod  Sw                  Hw
---  ------------------  ------------------------------------------------
1    5.2(1)SV3(1.5b)     0.0
2    5.2(1)SV3(1.5b)     0.0
3    5.2(1)SV3(1.5b)     VMware ESXi 6.0.0 Releasebuild-2494585 (6.0)
4    5.2(1)SV3(1.5b)     VMware ESXi 6.0.0 Releasebuild-2494585 (6.0)

Mod  Server-IP        Server-UUID                           Server-Name
---  ---------------  ------------------------------------  --------------------
1    172.20.71.44     NA                                    NA
2    172.20.71.44     NA                                    NA
3    172.20.71.26     722201d0-e027-e511-0000-000000110001  fvl1-h1
4    172.20.71.27     722201d0-e027-e511-0000-000000110002  fvl1-h2

* this terminal session
fvl1-vsm#
```

14. Run copy run start.

Cisco Nexus 1000V vTracker

SSH Connection to Primary VSM

The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network. environment. To connect SSH to the primary VSM, complete the following steps:

1. From an ssh interface connected to the Cisco Nexus 1000V VSM, enter the following:

```
config t
feature vtracker
copy run start
show vtracker upstream-view
show vtracker vm-view vnic
show vtracker vm-view info
show vtracker module-view pnic
show vtracker vlan-view
```

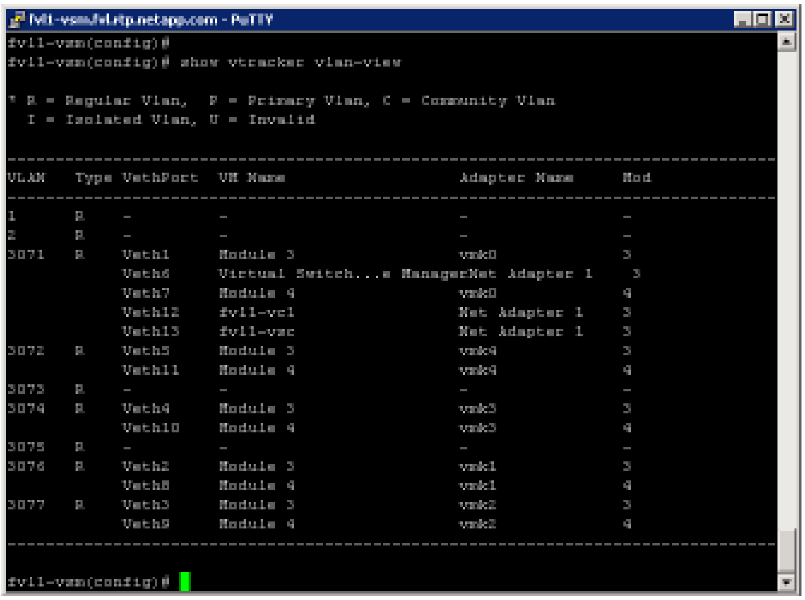

```
fvl1-vsm(config)#
fvl1-vsm(config)# show vtracker vlan-view

* R = Regular Vlan,  P = Primary Vlan, C = Community Vlan
  I = Isolated Vlan, U = Invalid

--------------------------------------------------------------------------------
VLAN   Type VethPort  VM Name                  Adapter Name     Mod
--------------------------------------------------------------------------------
1      R    -         -                        -                -
2      R    -         -                        -                -
3071   R    Veth1     Module 3                 vmk0             3
            Veth6     Virtual Switch...e ManagerNet Adapter 1    3
            Veth7     Module 4                 vmk0             4
            Veth12    fvl1-vc1                 Net Adapter 1    3
            Veth13    fvl1-vsc                 Net Adapter 1    3
3072   R    Veth5     Module 3                 vmk4             3
            Veth11    Module 4                 vmk4             4
3073   R    -         -                        -                -
3074   R    Veth4     Module 3                 vmk3             3
            Veth10    Module 4                 vmk3             4
3075   R    -         -                        -                -
3076   R    Veth2     Module 3                 vmk1             3
            Veth8     Module 4                 vmk1             4
3077   R    Veth3     Module 3                 vmk2             3
            Veth9     Module 4                 vmk2             4
--------------------------------------------------------------------------------

fvl1-vsm(config)#
```

FlexPod Management Tools Setup

Cisco UCS Performance Manager

This section describes the deployment and initial configuration of Cisco UCS Performance Manager within a FlexPod.

For full requirements and installation options, download and review the Cisco UCS Performance Manager Installation Guide, Release 2.0.0.

Cisco UCS Performance Manager OVA Deployment

1. Download the Cisco UCS Performance Manager OVA file from the Cisco UCS Performance Manager site to your workstation.
2. Use the VMware vSphere Web Client to log in to vCenter as root, or as a user with superuser privileges, and then select Hosts and Clusters from the Home view.

   The standard vSphere Client can be used for installation, but instructions will vary slightly.

3. Right-click the Cluster to deploy to, and select Deploy OVF Template from the pulldown options.

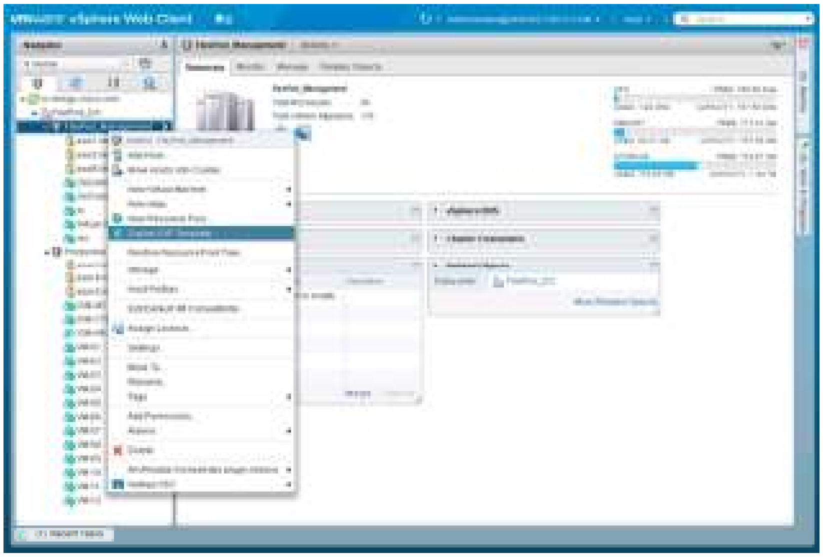

4. In the Selectsource panel, specify the path of the Cisco UCS Performance Manager package as a Local file and browse to find it’s location, and then click Next.

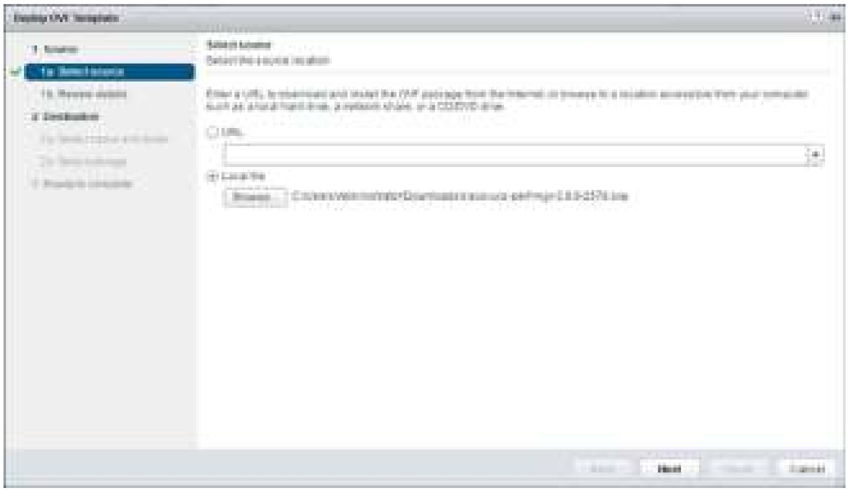

5. Click Next to continue past the Review details pane.

6. In the Select name and folder pane, accept the default name, or specify one appropriate to your environment, pick the data center to deploy to and alternately a folder within that datacenter.

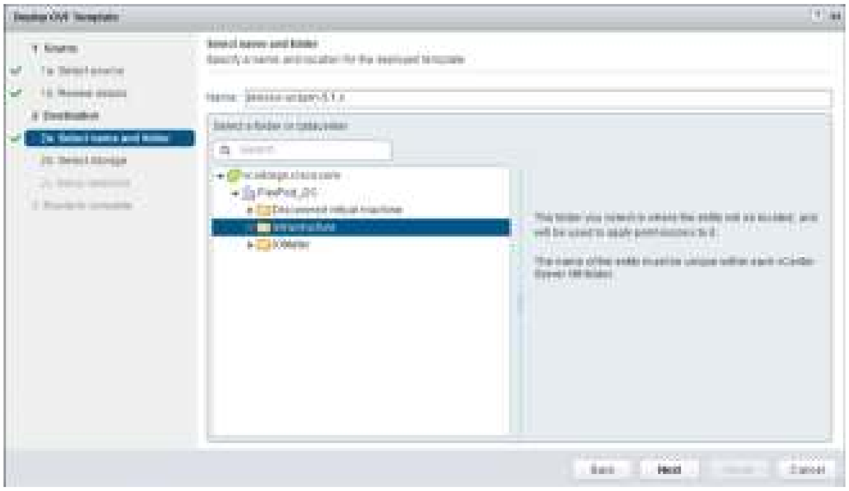

7. Click Next.

8. Specify the datastore to deploy to within the Select storage pane

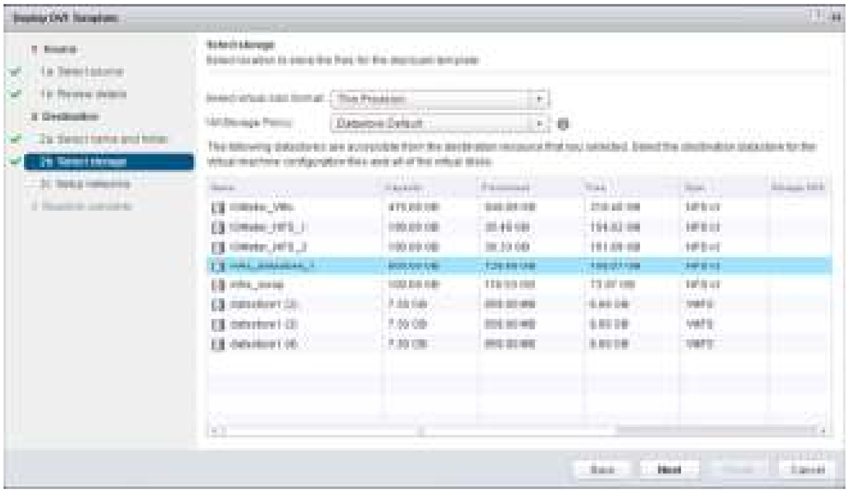

9. Click Next.

10. Specify the Destination port group within the Setup Networks pane.

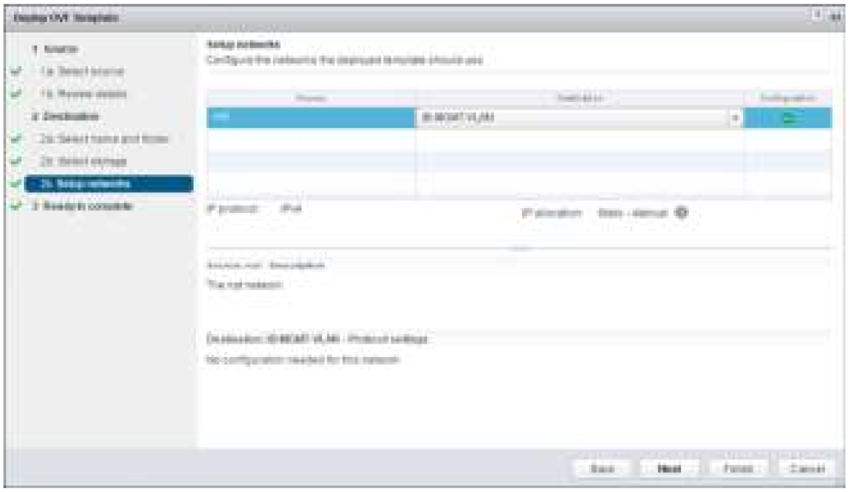

Deploy OVF Template

1 Source
1a Select source
1b Review details
2 Destination
2a Select name and folder
2b Select storage
2c Setup networks
3 Ready to complete

Setup networks
Configure the networks the deployed template should use

Source	Destination	Configuration
[illegible]	IB-MGMT VLAN	

IP protocol: IPv4
IP allocation: Static - Manual

Destination: IB-MGMT VLAN - Protocol settings
No configuration needed for this network

11. Click Next
12. Review he options in the Ready to complete pane, and make sure that “Power on after deployment” is not selected.

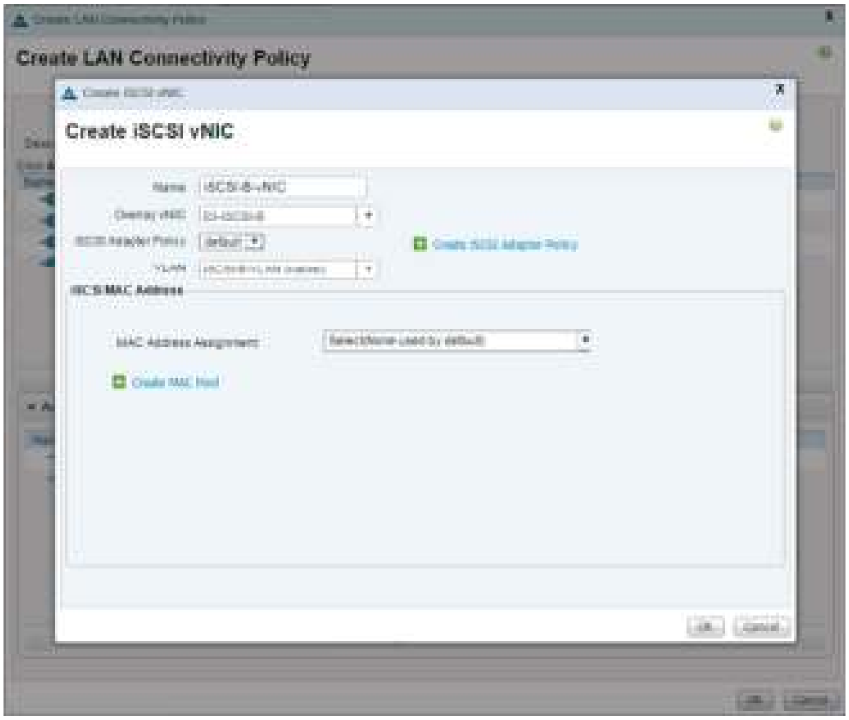

Create LAN Connectivity Policy

Create iSCSI vNIC

Name: iSCSI-B-vNIC
Overlay vNIC: 03-iSCSI-B
iSCSI Adapter Policy: default
VLAN: iSCSI-B-VLAN (native)

iSCSI MAC Address
MAC Address Assignment: Select(None used by default)

13. Click Finish to deploy the appliance from the Deploy OVF Template wizard.
14. Once the appliance has finished deploying, right click the VM from within the Hosts and Clusters or VMs and Templates section of the vSphere Web Client and select Edit Settings from the pulldown options.
15. Adjust the allocated Memory from 40G to 64G.

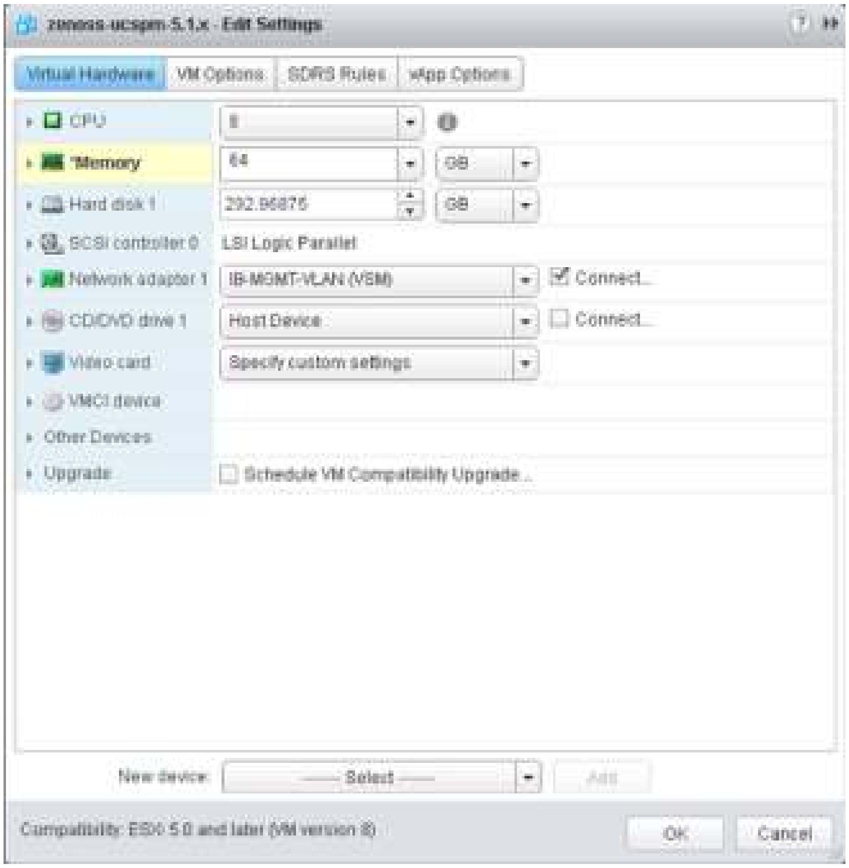

zenoss-ucspm-5.1.x - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

CPU	8	
*Memory	64	GB
Hard disk 1	292.96875	GB
SCSI controller 0	LSI Logic Parallel	
Network adapter 1	IB-MGMT-VLAN (VSM)	☑ Connect...
CD/DVD drive 1	Host Device	□ Connect...
Video card	Specify custom settings	
VMCI device		
Other Devices		
Upgrade	□ Schedule VM Compatibility Upgrade...	

New device: ------- Select ------- Add

Compatibility: ESXi 5.0 and later (VM version 8)

16. Click OK and power on the VM.

Cisco UCS Performance Manager Initial Configuration

1. Open up the Console of the newly provisioned UCS Performance Manager VM from the vSphere web client.
2. Login to the root account with the password "ucsm"

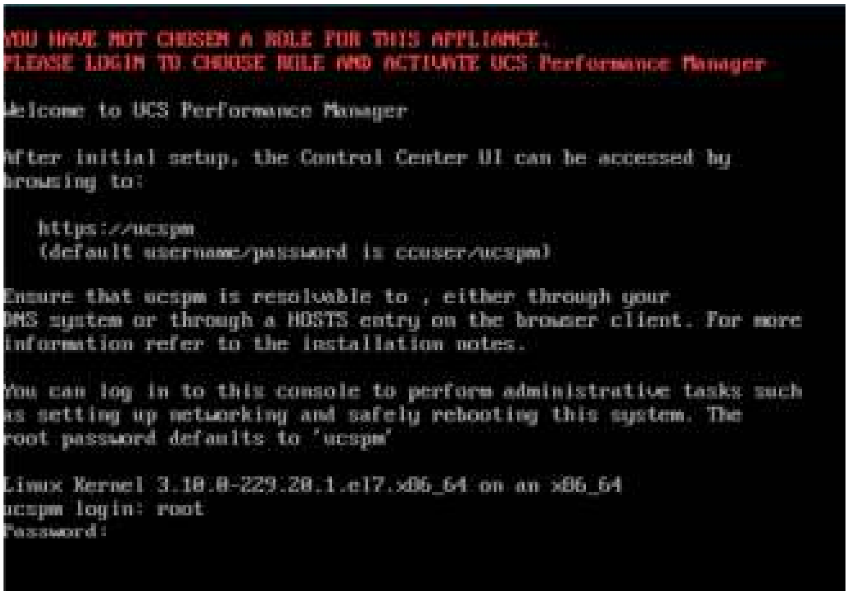

```
YOU HAVE NOT CHOSEN A ROLE FOR THIS APPLIANCE.
PLEASE LOGIN TO CHOOSE ROLE AND ACTIVATE UCS Performance Manager

Welcome to UCS Performance Manager

After initial setup, the Control Center UI can be accessed by
browsing to:

   https://ucspm
   (default username/password is ccuser/ucspm)

Ensure that ucspm is resolvable to , either through your
DNS system or through a HOSTS entry on the browser client. For more
information refer to the installation notes.

You can log in to this console to perform administrative tasks such
as setting up networking and safely rebooting this system. The
root password defaults to 'ucspm'

Linux Kernel 3.10.0-229.20.1.el7.x86_64 on an x86_64
ucspm login: root
Password:
```

3. Set a new root password and a password for the account "ccuser" when prompted.
4. Within the following screen, leave Master selected to configure the appliance as as the Control Center master host.

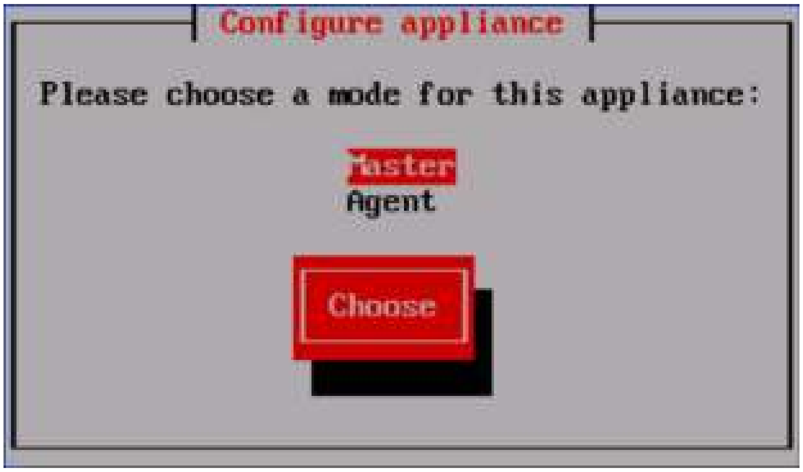

Configure appliance

Please choose a mode for this appliance:

Master
Agent

5. Press Enter to initiate a reboot and log back in with the root account after the appliance is back up.
6. Press Return from the Appliance Administration screen with Configure Network and DNS selected.

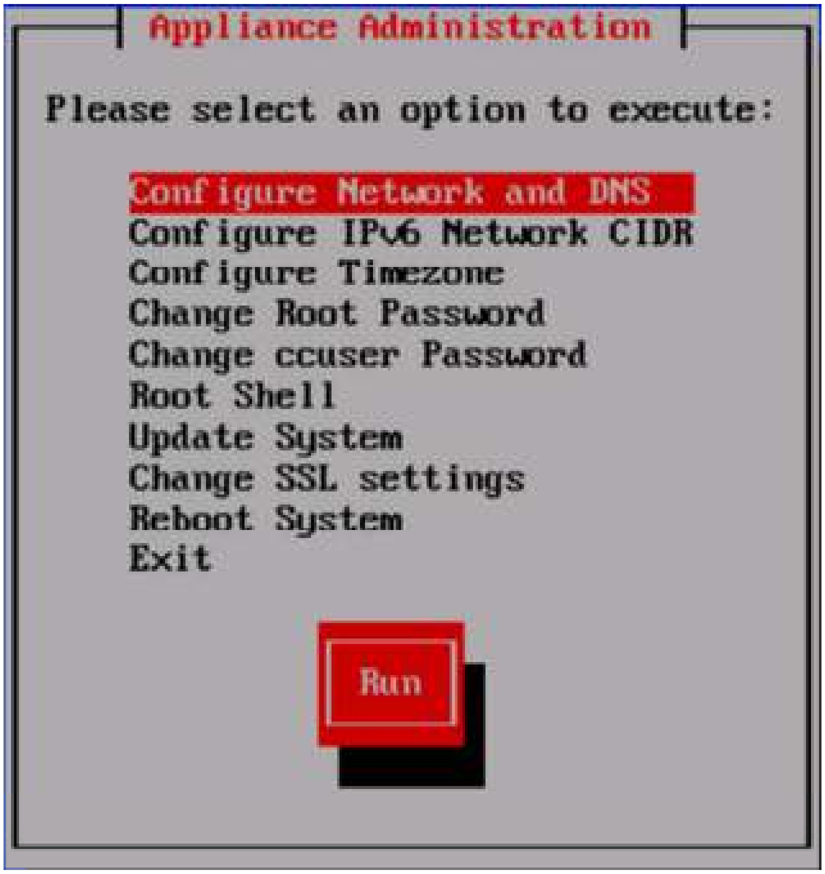

Appliance Administration

Please select an option to execute:

Configure Network and DNS
Configure IPv6 Network CIDR
Configure Timezone
Change Root Password
Change ccuser Password
Root Shell
Update System
Change SSL settings
Reboot System
Exit

7. Leave Edit a connection selected with the NetworkManager TUI screen and hit return.

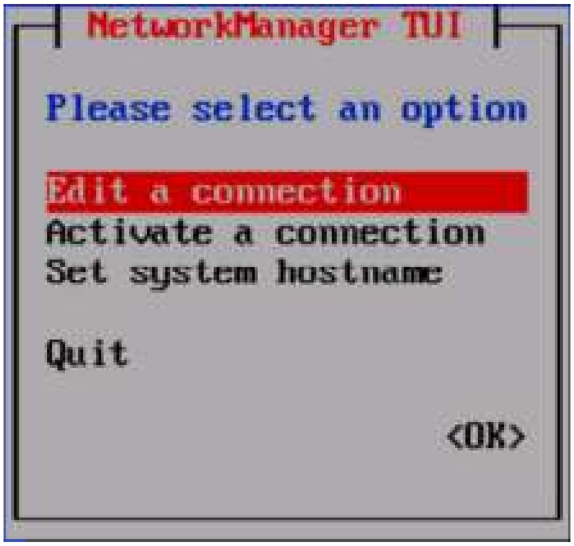

8. Select Wired connection 1, and press return to edit.

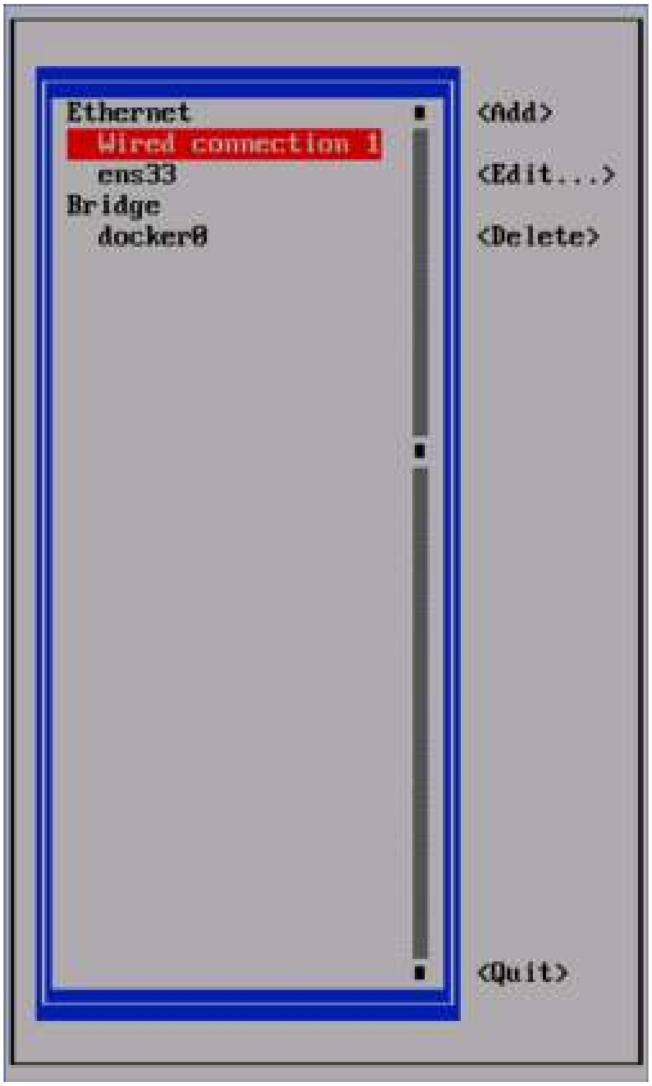

9. Arrow down to the IPV4 CONFIGURATION option and hit enter to change the option from Automatic to Manual.

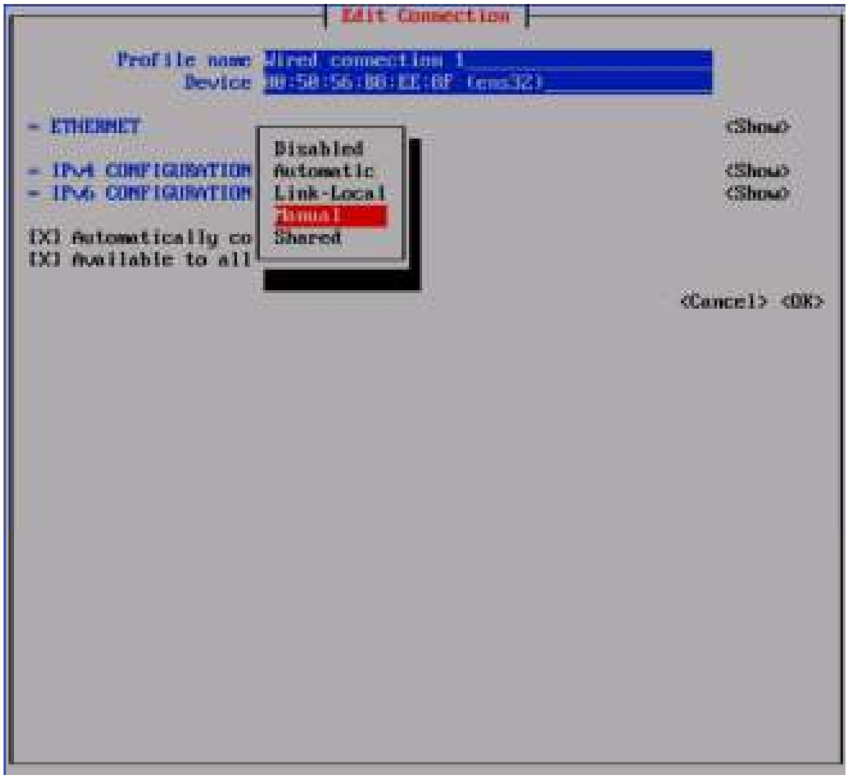

```
                         ┤ Edit Connection ├

      Profile name Wired connection 1
            Device 00:50:56:B0:EE:8F (ens32)

  = ETHERNET                                                  <Show>
                     Disabled
  = IPv4 CONFIGURATION Automatic                              <Show>
  = IPv6 CONFIGURATION Link-Local                             <Show>
                     Manual
  [X] Automatically co Shared
  [X] Available to all

                                                        <Cancel> <OK>
```

10. Right arrow over to <Show> next to the IPV4 CONFIGURATION and hit return to bring up the options.

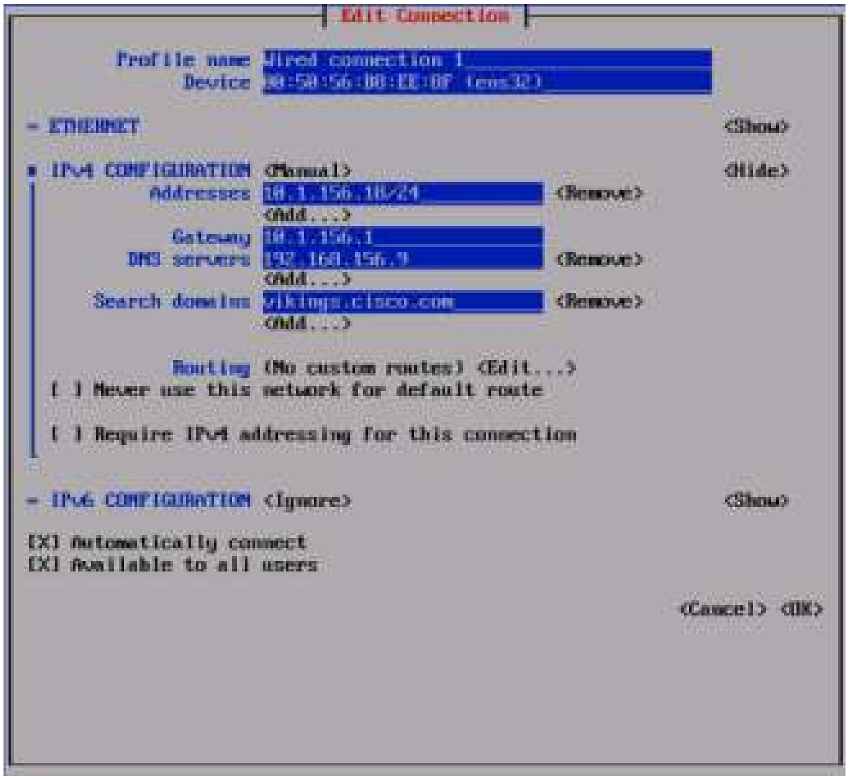

```
                         ┤ Edit Connection ├

      Profile name Wired connection 1
            Device 00:50:56:B0:EE:8F (ens32)

  = ETHERNET                                                  <Show>

  ╤ IPv4 CONFIGURATION <Manual>                               <Hide>
  │          Addresses 10.1.156.18/24         <Remove>
  │                    <Add...>
  │            Gateway 10.1.156.1
  │        DNS servers 192.168.156.9          <Remove>
  │                    <Add...>
  │     Search domains vikings.cisco.com      <Remove>
  │                    <Add...>
  │
  │            Routing (No custom routes) <Edit...>
  │  [ ] Never use this network for default route
  │
  │  [ ] Require IPv4 addressing for this connection
  └
  = IPv6 CONFIGURATION <Ignore>                               <Show>

  [X] Automatically connect
  [X] Available to all users

                                                        <Cancel> <OK>
```

11. Enter the assigned IP address/netmask, Gateway, DNS server(s), and any Search Domains.
12. Optionally change the IPV6 CONFIGURATION option from Automatic to Ignore.
13. Arrow down till <OK> is highlighted and hit return.
14. Hit the right arrow within the interface selection menu you are returned to, arrow further down until <Quit> is selected and hit return.
15. Arrow down to have Reboot highlighted and hit return to initiate a reboot.

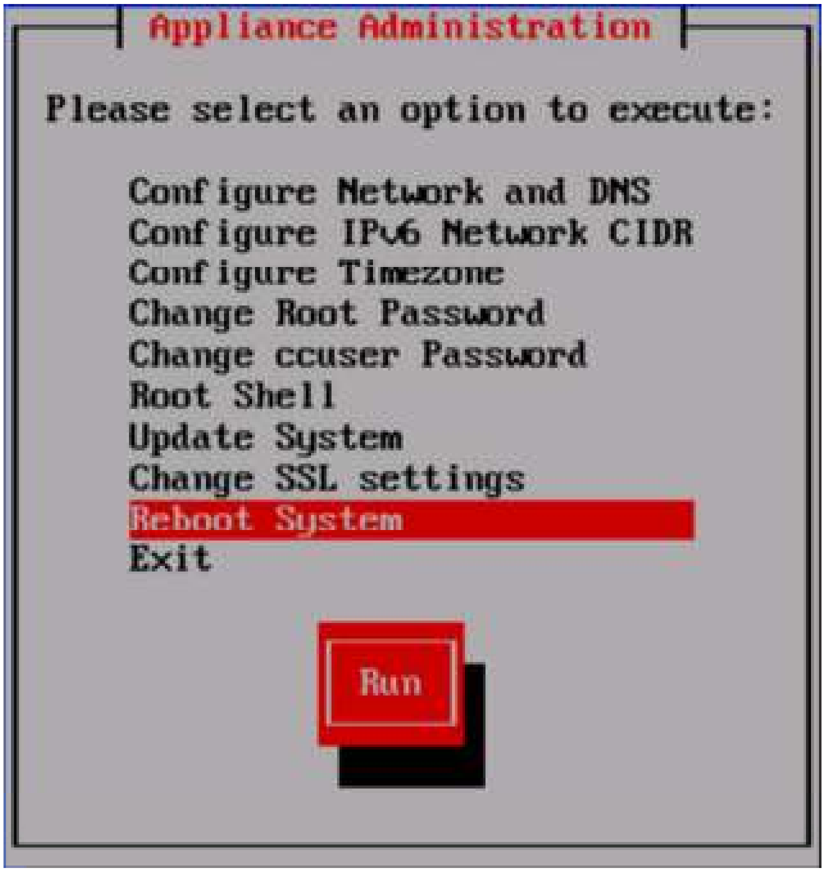

```
Appliance Administration

Please select an option to execute:

    Configure Network and DNS
    Configure IPv6 Network CIDR
    Configure Timezone
    Change Root Password
    Change ccuser Password
    Root Shell
    Update System
    Change SSL settings
    Reboot System
    Exit

            Run
```

At this point further console operations can continue from with the remote console of the vSphere Web Client, or can be initiated through an ssh connection via shell or putty.

16. Re-login as root once the appliance as finished rebooting and select Configure Network and DNS from the Appliance Administration menu

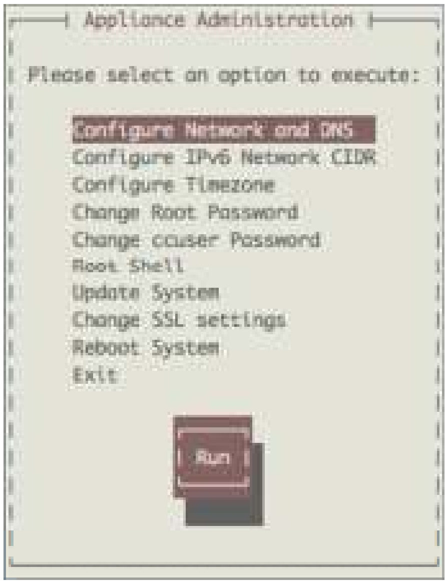

```
Appliance Administration

Please select an option to execute:

    Configure Network and DNS
    Configure IPv6 Network CIDR
    Configure Timezone
    Change Root Password
    Change ccuser Password
    Root Shell
    Update System
    Change SSL settings
    Reboot System
    Exit

            Run
```

17. Within the NetworkManager TUI menu select Set system hostname and hit return.

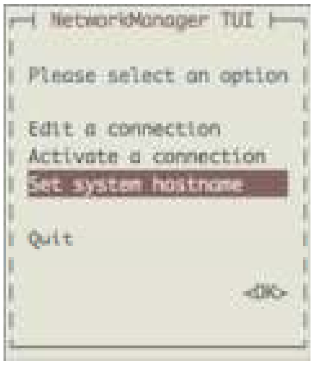

```
NetworkManager TUI

Please select an option

Edit a connection
Activate a connection
Set system hostname

Quit

<OK>
```

18. Enter the name desired for the UCS Performance Manager master host and press return.

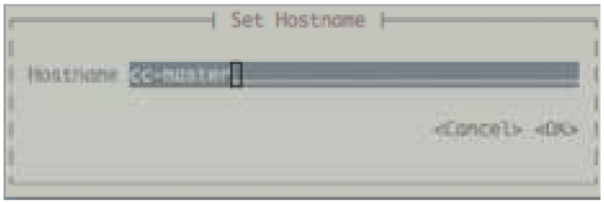

```
Set Hostname

Hostname cc-master

<Cancel> <OK>
```

19. Hit return to move past the confirmation screen.

20. Arrow down within the Appliance Administration menu to the Reboot System option and press return to change the system hostname.

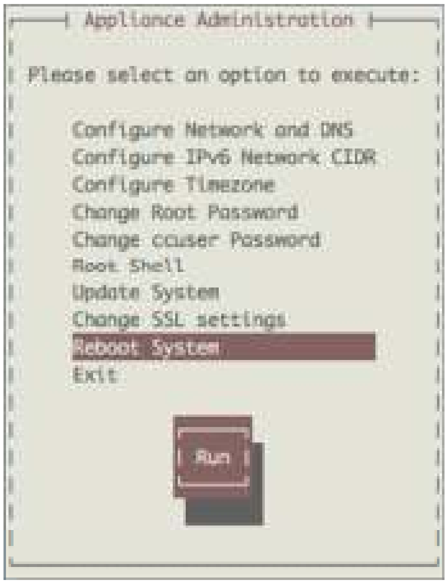

```
Appliance Administration

Please select an option to execute:

    Configure Network and DNS
    Configure IPv6 Network CIDR
    Configure Timezone
    Change Root Password
    Change ccuser Password
    Root Shell
    Update System
    Change SSL settings
    Reboot System
    Exit

        Run
```

Cisco UCS Performance Manager Deployment

The base IP or hostname of the UCS Performance Manager master host will allow a connection the Control Center that the UCS Performance Manager virtual host is provisioned from.

In our example *ControlCenter* was registered in our DNS as cc-master.vikings.cisco.com, and was accessible as https://cc-master.vikings.cisco.com or via the IP assigned to it. The UCS Performance Manger will need to to resolve in DNS as a CNAME alias or local /etc/hosts entry from the browsing system as ucspm.*ControlCenter* which was “ucspm.cc-master.vikings.cisco.com” in our example.

1. Login to the Control Center page with the ccuser account, confirm and Security Exceptions identified when connecting.

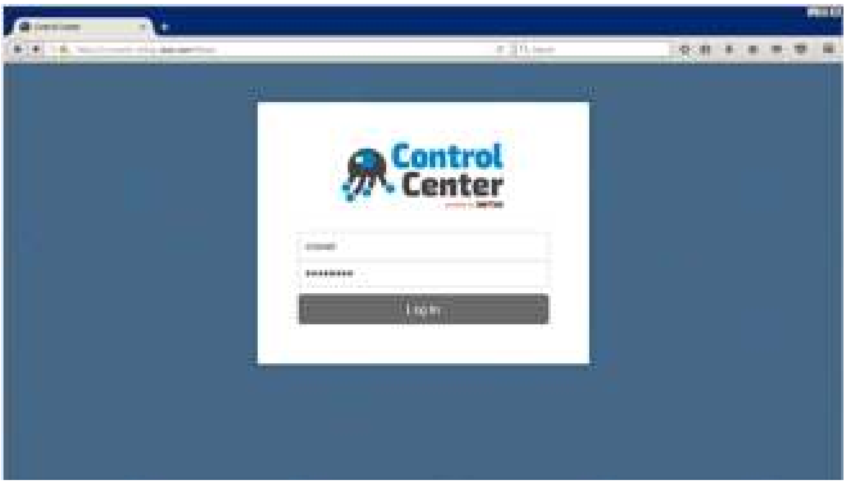

2. The initial login will bring up a Deployment Wizard to provision the UCS Performance Manager virtual host. If this does not come up, it can be initiated by selecting the +Application button in the mid to upper right of the screen.

3. Enter the Host and port values (ucspm.cc-master.vikings.cisco.com:4979 in our example), select default for the Resource Pool ID, and enter 75% for the RAM Commitment.

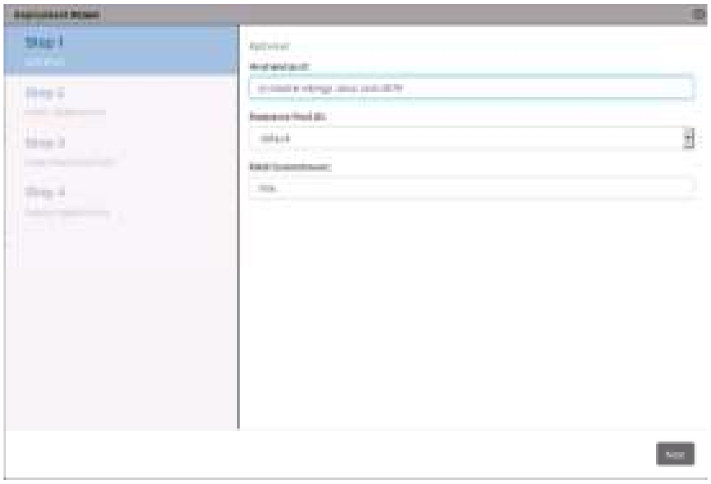

4. Click Next.

5. Select ucspm (v2.0.0) as the application to install.

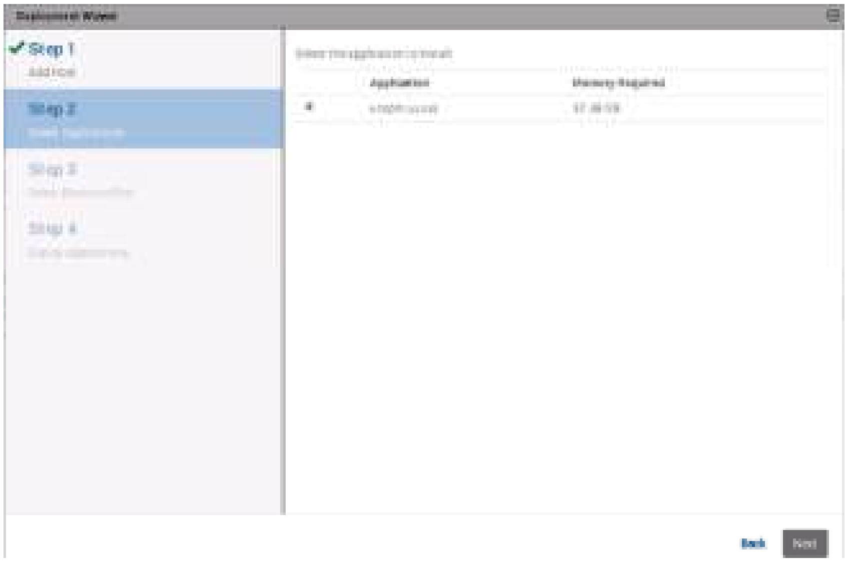

6. Click Next.

7. Select default as the resource pool.

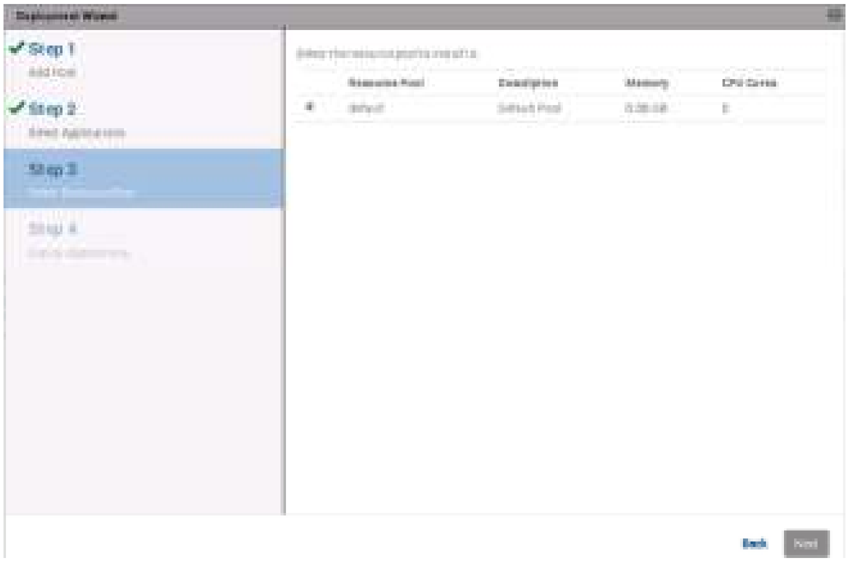

8. Click Next.
9. Specify a Deployment ID name to provision the application.

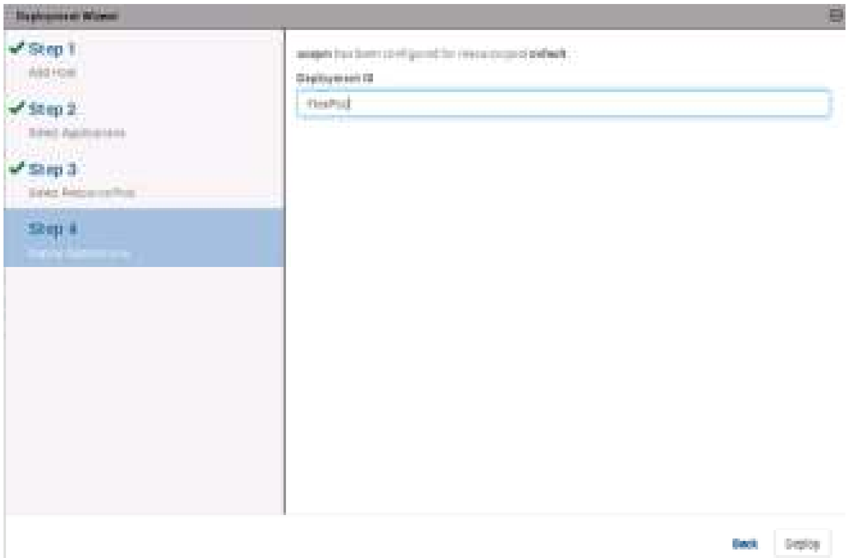

10. Click Deploy to provision.
11. In the Actions column of the Applications table, click the Start option of the ucspm (v2.0.0) row.

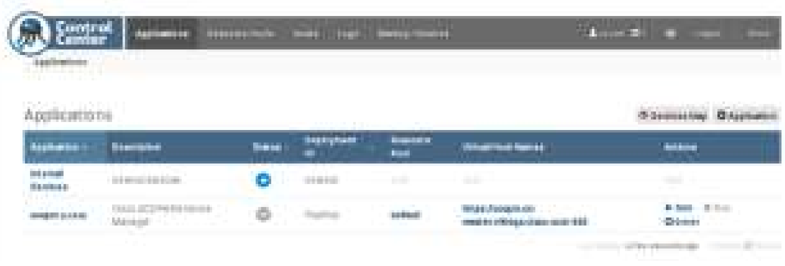

12. A Start Service dialog window will pop up.

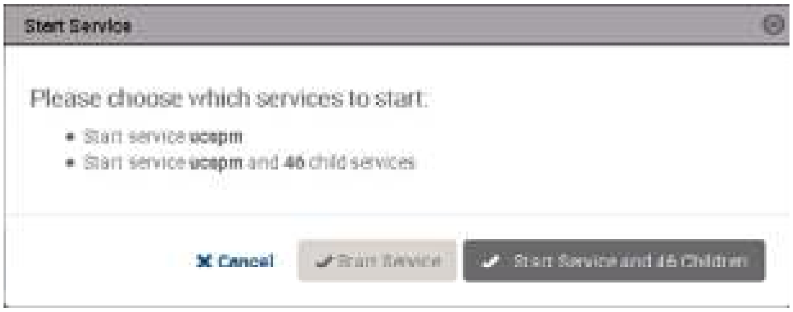

Start Service

Please choose which services to start:

- Start service **ucspm**
- Start service **ucspm** and **46** child services

13. Select Start Service and 46 Children.

14. In the Application column of the Applications table, click ucspm in the ucspm row

15. Scroll down to watch child services starting. Typically, child services take 4-5 minutes to start. When no child service shows a red exclamation point icon, Cisco UCS Performance Manager is running.

Cisco UCS Performance Manager Configuration of FlexPod Infrastructure

1. Using another web browser session, connect to the Cisco UCS Performance Manager virtual host using the DNS CNAME or local entry configured within the host the browser is running from (https://ucspm.vikings.cisco.com/).

2. Scroll down to the licensing screen that first appears, select the “Click this box to verify you agree to the License.” in the bottom left, and click Accept License in the bottom right of the page.

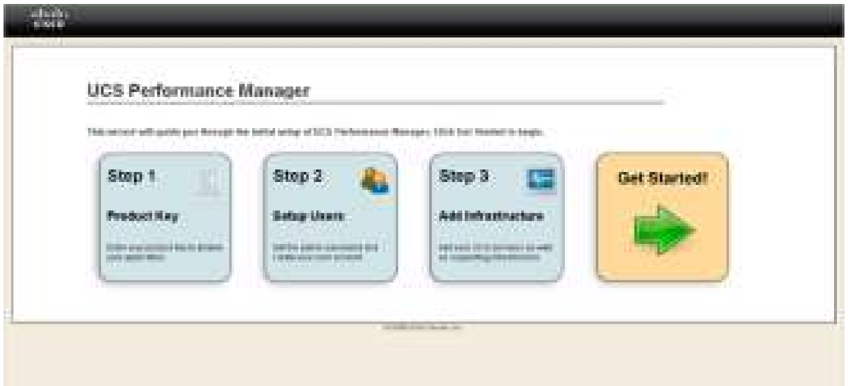

UCS Performance Manager

3. Click the “Get Started!” option in the initial screen.

4. Click Add License File in the following screen if you are adding one, or go directly to Click Next on the Licensing screen if you are going to run with the 30 day trial.

5. Specify an admin password and create a local account on the following screen.

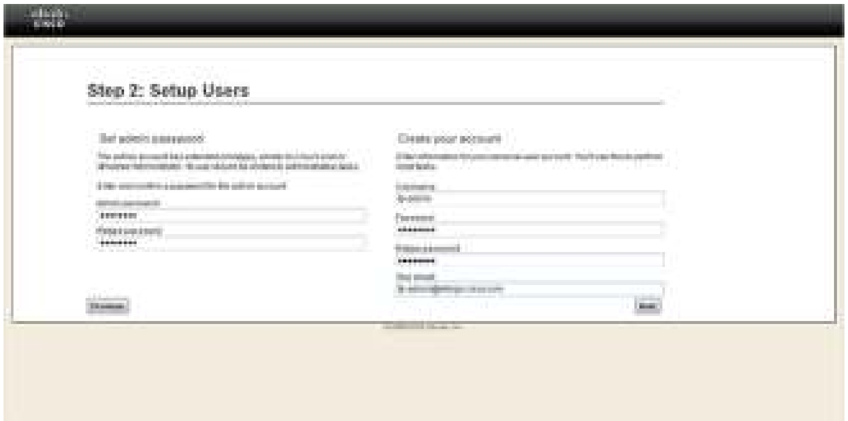

Step 2: Setup Users

6. Click Next.

Administrator and admin accounts were used as the access accounts in this section. In a production environment it may be more appropriate to use dedicated read only accounts.

7. Add a UCS Central IP, administrative account and password if you have one independently deployed in your environment (UCS Central was not covered in this document.)

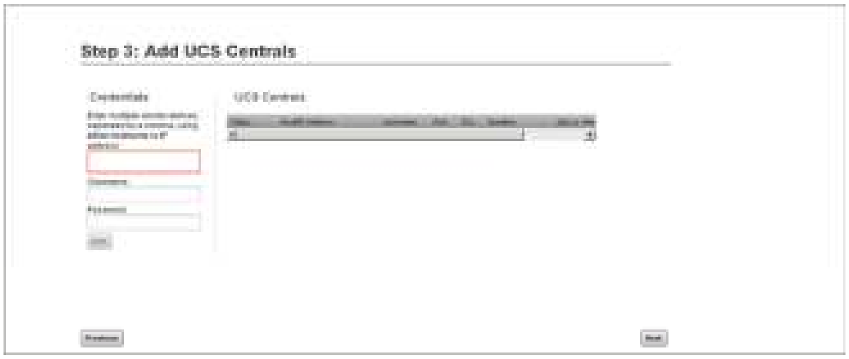

Step 3: Add UCS Centrals

8. Click Next.
9. Add in the UCS Fabric Interconnect(s) virtual IPs for UCS Domains to be monitored in your environment.

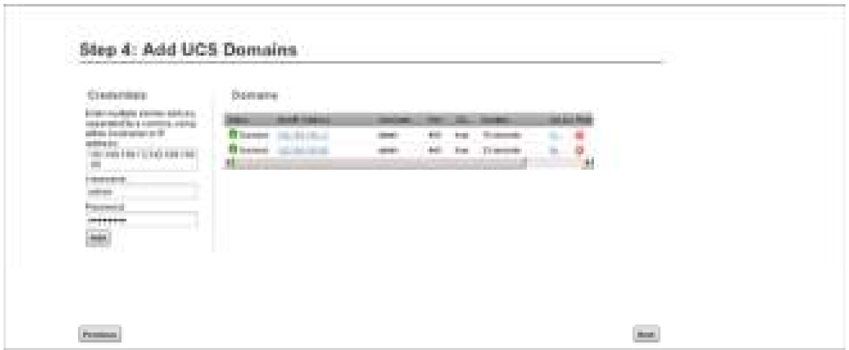

Step 4: Add UCS Domains

10. Click Add, then click Next.
11. Select Network from the Category column to add in the Nexus Switches to be monitored.
12. Select Cisco Nexus 9000 and enter the IP and credentials for access.

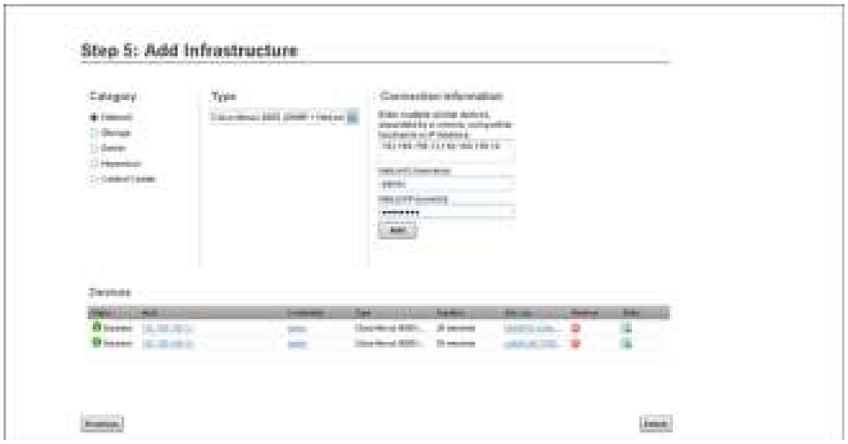

Step 5: Add Infrastructure

13. Click Add.
14. Click Storage from the Category column to add in the NetApp AFF.
15. Select NetApp C-Mode Filer (ZAPI) and enter the IP and credentials for access.
16. Click the “Use SSL?” checkmark box.

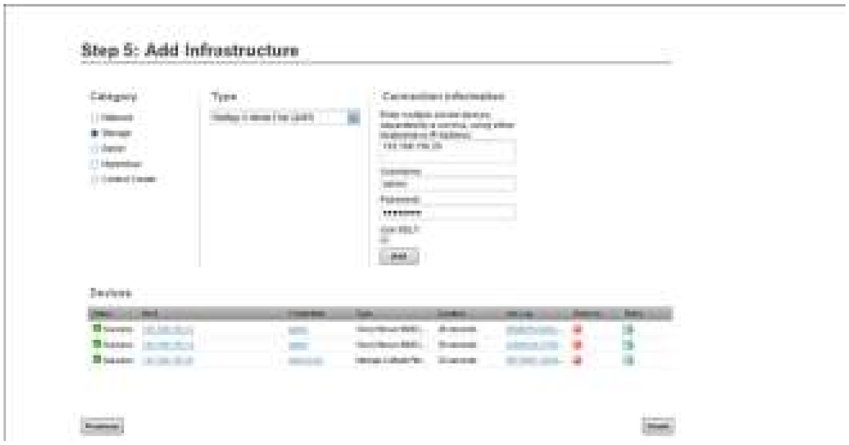

17. Click Add.
18. Click Hypervisor from the Category column to add in the vCenter Server.
19. Select vSphere EndPoint (SOAP) and enter the IP and credentials for access vCenter.
20. Click the “Use SSL?” checkmark box.

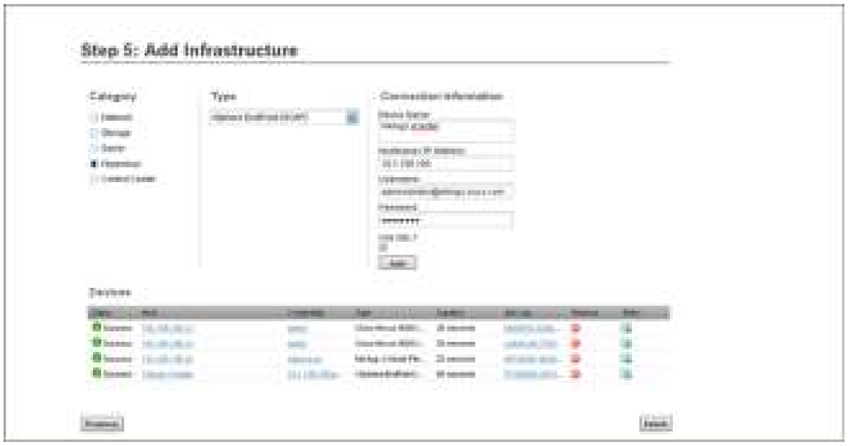

21. Click Finish.

The initial configuration of UCS Performance Manger to access the FlexPod environment is now complete.

Reference the Cisco UCS Performance Manager Administration Guide, Release 2.0.0 for further use and configuration of Cisco UCS Performance Manager.

NetApp Virtual Storage Console 6.2P1 Deployment Procedure

This section describes the deployment procedures for the NetApp VSC.

Virtual Storage Console 6.2P1 Pre-Installation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.3.2:

- Protocol licenses (NFS and FCP)
- FlexClone (for provisioning and cloning only)
- SnapRestore (for backup and recovery)
- SnapManager suite

Install Virtual Storage Console 6.2P1

To install the VSC 6.2P1 software, complete the following steps:

1. Build a VSC VM with Windows Server 2012 R2, 4GB RAM, two CPUs, and one virtual network interface in the <<var_ib_mgmt_vlan_id>> VLAN. The virtual network interface should be a VMXNET 3 adapter.
2. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain.
3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the User Interfaces and Infrastructure Feature on the VM.
4. Install all Windows updates on the VM.
5. Log in to the VSC VM as FlexPod admin user.
6. Download the x64 version of the Virtual Storage Console 6.2P1 from the NetApp Support site.
7. From the VMware Console, right-click the VSC .exe file downloaded in step 6 and select Run as administrator.
8. Select the appropriate language and click OK.

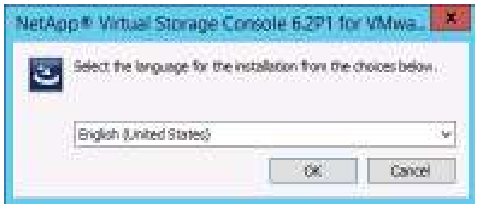

Select the language for the installation from the choices below.

English (United States)

9. On the Installation wizard Welcome page, click Next.

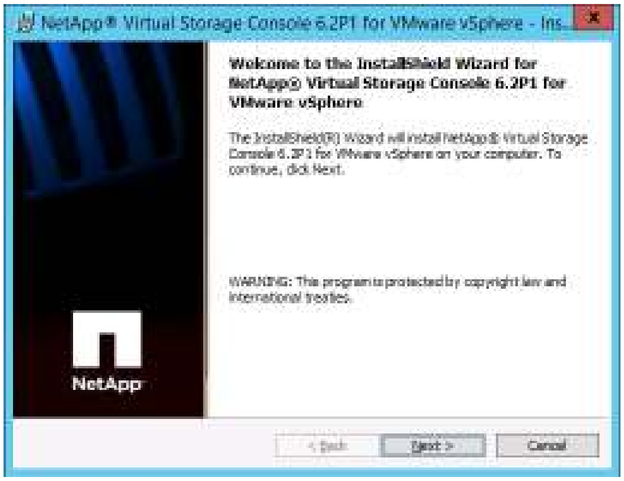
NetApp® Virtual Storage Console 6.2P1 for VMware vSphere - Ins...
X
NetApp
< Back
Next >
Cancel

Welcome to the InstallShield Wizard for NetApp® Virtual Storage Console 6.2P1 for VMware vSphere

The InstallShield(R) Wizard will install NetApp® Virtual Storage Console 6.2P1 for VMware vSphere on your computer. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

10. Select the checkbox to accept the message, click Next.

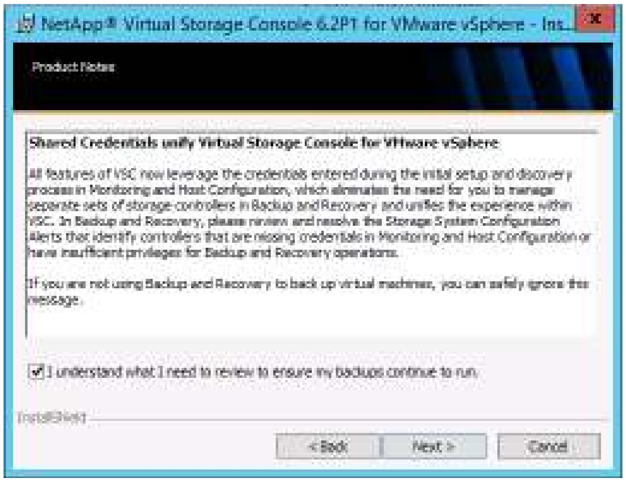
NetApp® Virtual Storage Console 6.2P1 for VMware vSphere - Ins...
X
Product Notes
InstallShield
< Back
Next >
Cancel

Shared Credentials unify Virtual Storage Console for VMware vSphere

All features of VSC now leverage the credentials entered during the initial setup and discovery process in Monitoring and Host Configuration, which eliminates the need for you to manage separate sets of storage controllers in Backup and Recovery and unifies the experience within VSC. In Backup and Recovery, please review and resolve the Storage System Configuration Alerts that identify controllers that are missing credentials in Monitoring and Host Configuration or have insufficient privileges for Backup and Recovery operations.

If you are not using Backup and Recovery to back up virtual machines, you can safely ignore this message.

☑ I understand what I need to review to ensure my backups continue to run.

The Backup and Recovery capability requires an additional license.

11. Click Next to accept the default installation location.

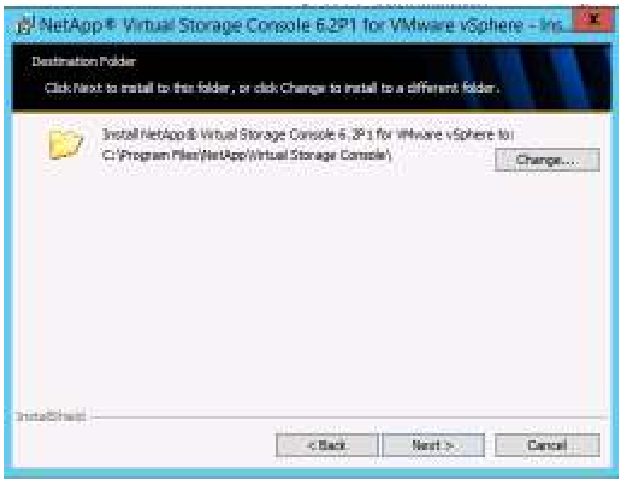

NetApp® Virtual Storage Console 6.2P1 for VMware vSphere - Ins...

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install NetApp® Virtual Storage Console 6.2P1 for VMware vSphere to:
C:\Program Files\NetApp\Virtual Storage Console\

12. Click Install.

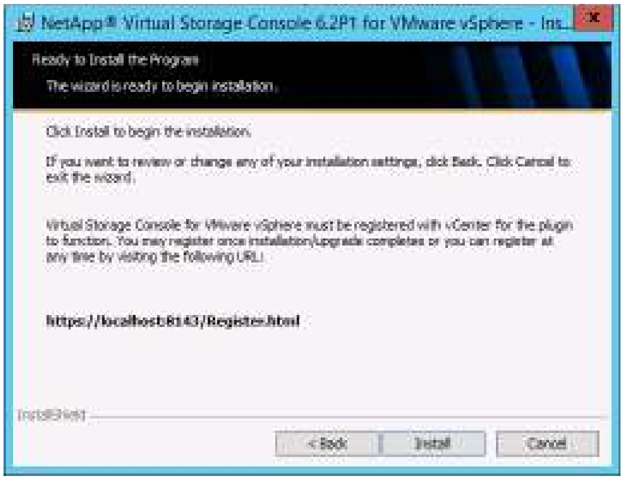

NetApp® Virtual Storage Console 6.2P1 for VMware vSphere - Ins...

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

Virtual Storage Console for VMware vSphere must be registered with vCenter for the plugin to function. You may register once installation/upgrade completes or you can register at any time by visiting the following URL:

https://localhost:8143/Register.html

13. Click Finish.

Register VSC with vCenter Server

To register the VSC with the vCenter server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open https://localhost:8143/Register.html in Internet Explorer.
2. Click Continue to This Website (Not Recommended).
3. In the Plug-in Service Information section, select the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.
4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user or root), and user password for the vCenter server. Click Register to complete the registration.

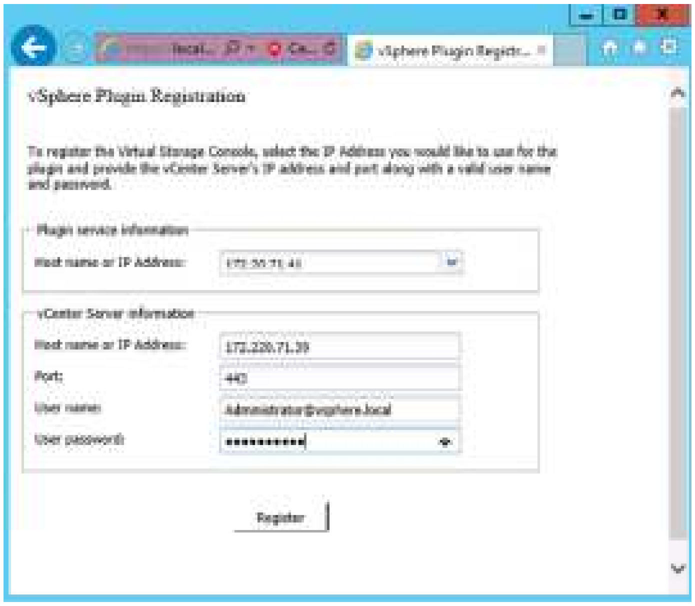

vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.

Plugin service information

Host name or IP Address: 172.20.71.41

vCenter Server information

Host name or IP Address: 172.220.71.39

Port: 443

User name: Administrator@vsphere.local

User password: ●●●●●●●●●●●

5. After successful registration, the storage controllers are discovered automatically.

The storage discovery process may take some time.

Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, complete the following steps:

1. Using the vSphere web client, log in to the vCenter Server as FlexPod admin user or root. If the vSphere web client was previously opened, close it and then reopen it.

2. In the Home screen, click the Home tab and click Virtual Storage Console.

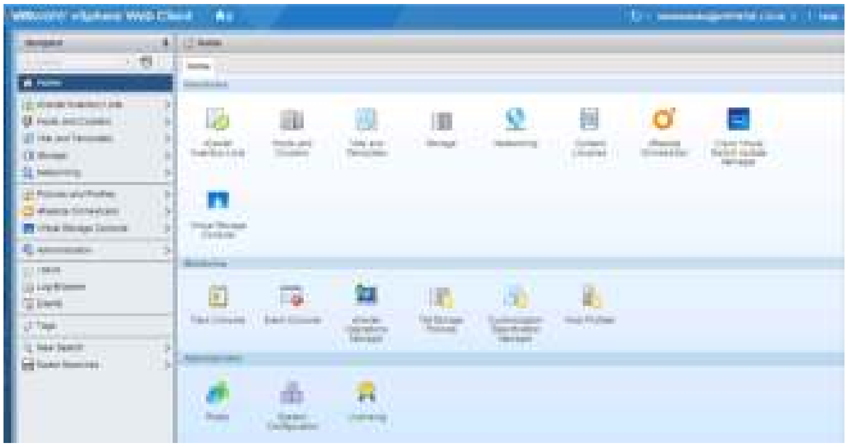

3. Select Storage Systems. Under the Objects tab, click Actions > Modify.

4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name and the admin password for password. Confirm that Use SSL to connect to this storage system is selected. Click OK.

5. Click OK to accept the controller privileges.

Optimal Storage Settings for ESXi Hosts

VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the Home screen, click vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for These Hosts.

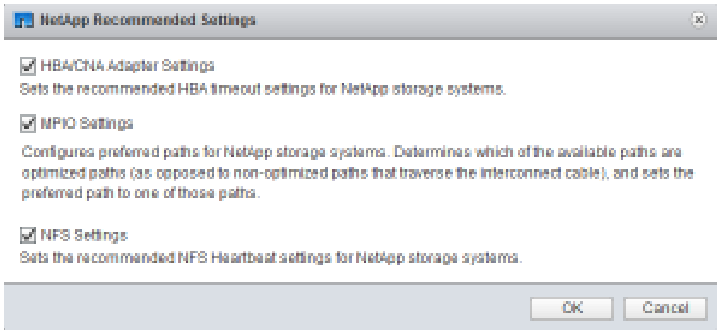

NetApp Recommended Settings

☑ HBA/CNA Adapter Settings
Sets the recommended HBA timeout settings for NetApp storage systems.

☑ MPIO Settings
Configures preferred paths for NetApp storage systems. Determines which of the available paths are optimized paths (as opposed to non-optimized paths that traverse the interconnect cable), and sets the preferred path to one of those paths.

☑ NFS Settings
Sets the recommended NFS Heartbeat settings for NetApp storage systems.

2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings. This functionality sets values for HBAs and CNAs, sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).
3. Click OK.

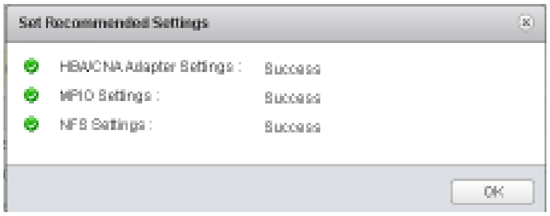

Set Recommended Settings

- HBA/CNA Adapter Settings : Success
- MPIO Settings : Success
- NFS Settings : Success

4. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.

VSC 6.2P1 Backup and Recovery

Prerequisites to Use Backup and Recovery Capability

Before you begin using the backup and recovery capability to schedule backups and restore your datastores, VMs, or virtual disk files, you must confirm that the storage systems that contain the datastores and VMs for which you are creating backups have valid storage credentials.

If you plan to use the SnapMirror update option, add all of the destination storage systems with valid storage credentials.

Backup and Recovery Configuration

1. From the Home screen, select the Home tab and click Storage.
2. On the left, expand the Datacenter and select Datastores.
3. Right-click the datastore that you want to backup. Select NetApp VSC > Backup > Schedule Backup Job.

If you prefer a one-time backup, then select Backup Now instead of Schedule Backup.

4. Type a backup job name and description.

If you want to create a VMware snapshot for each backup, select Perform VMware Consistency Snapshot in the options pane.

5. Click Next.
6. Select any options to include in the backup.

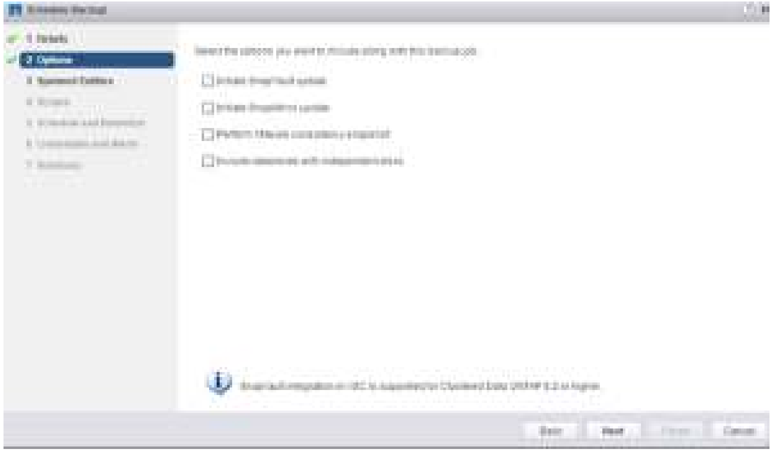

Schedule Backup

Select the options you want to include along with this backup job

□ Initiate SnapVault update
□ Initiate SnapMirror update
□ Perform VMware consistency snapshot
□ Include datastores with independent disks

SnapVault integration in VSC is supported for Clustered Data ONTAP 8.2 or higher

7. Click Next on the Options screen.
8. Click Next on the Spanned Entities screen.
9. Select one or more backup scripts if available and click Next in the Scripts screen.

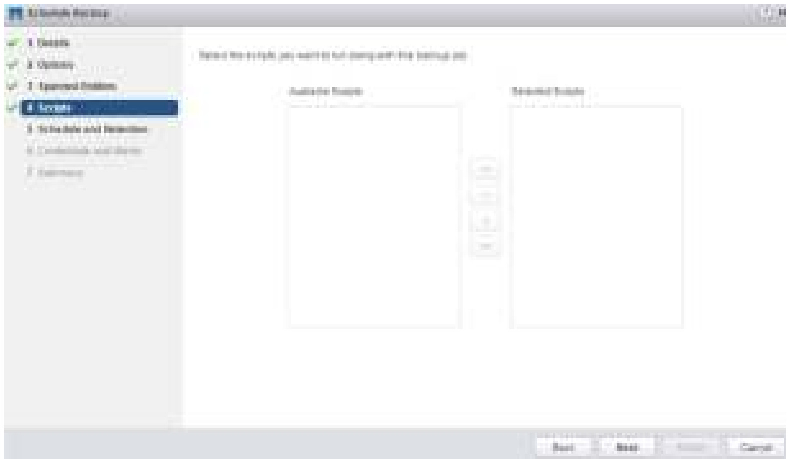

Schedule Backup

Select the scripts you want to run along with the backup job

Available Scripts

Selected Scripts

10. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click Next.

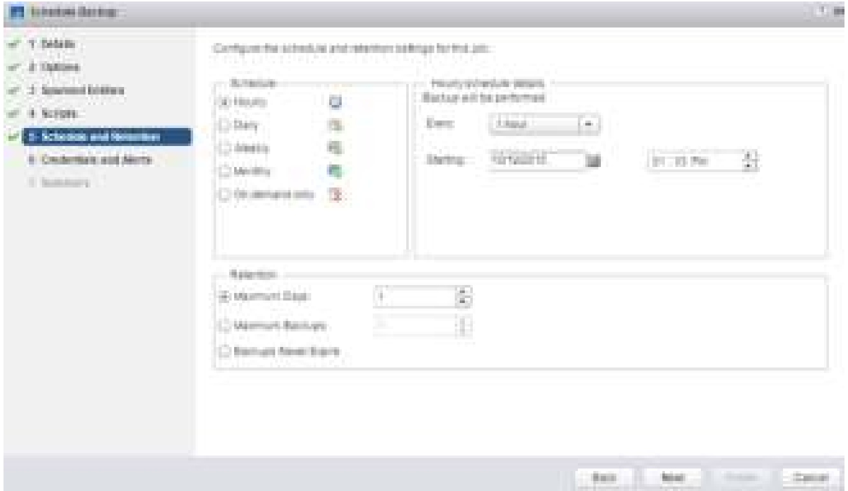

Schedule Backup

Configure the schedule and retention settings for this job:

Schedule

● Hourly

○ Daily

○ Weekly

○ Monthly

○ On demand only

Hourly schedule details

Backup will be performed

Every: 1 hour

Starting: 10/12/2015 01 : 33 PM

Retention

● Maximum Days: 1

○ Maximum Backups:

○ Backups Never Expire

11. Use the default vCenter credentials or type the user name and password for the vCenter server and click Next.

12. Specify backup retention details as per requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate them. Click Next.

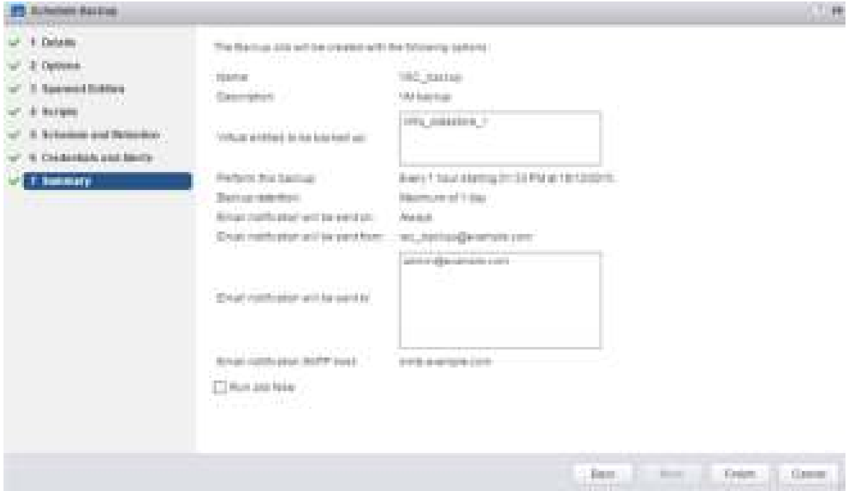

Schedule Backup

The Backup Job will be created with the following options:

Name: VSC_backup

Description: VM backup

Virtual entities to be backed up: infra_datastore_1

Perform the backup: Every 1 hour starting 01:33 PM at 10/12/2015.

Backup retention: Maximum of 1 day

Email notification will be send on: Always

Email notification will be send from: vsc_backup@example.com

Email notification will be send to: admin@example.com

Email notification SMTP host: smtp.example.com

□ Run Job Now

13. Review the summary page and click Finish. If you want to run the job immediately, select the Run Job Now option and then click Finish.

14. Click OK.

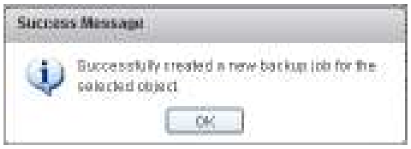

Success Message

Successfully created a new backup job for the selected object

15. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by entering the following command:

```
volume modify –volume infra_datastore_1 –snapshot-policy none
```

16. Also, to delete any existing automatic Snapshot copies that have been created on the volume, enter the following command:

```
volume snapshot show –volume infra_datastore_1
volume snapshot delete –volume infra_datastore_1 –vserver Infra-SVM –snapshot <snapshot name>
```

The wildcard character (*) can be used in snapshot names in the previous command.

OnCommand Unified Manager 6.3P2

OnCommand Unified Manager OVF Deployment

To install the OnCommand Unified Manager, complete the following steps:

Download and review the OnCommand Unified Manager Installation and Setup Guide.

1. Download OnCommand Unified Manager version 6.3P2 (OnCommandUnifiedManager-6.3P2.ova) from the OnCommand download site.
2. Log in to the vSphere web client. Go to vCenter > VMs and Templates.
3. At the top of the center pane, select Actions > Deploy OVF Template.

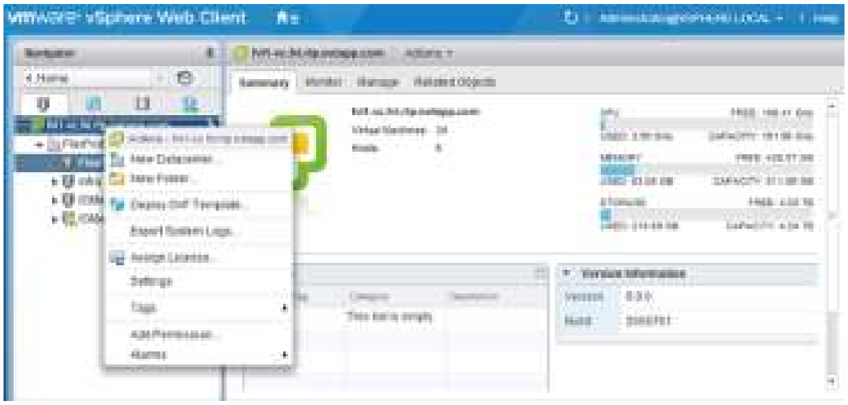

4. Browse the .ova file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.
5. Select Accept Extra Configuration Options, and click Next.
6. Review the deployment details, and click Next.

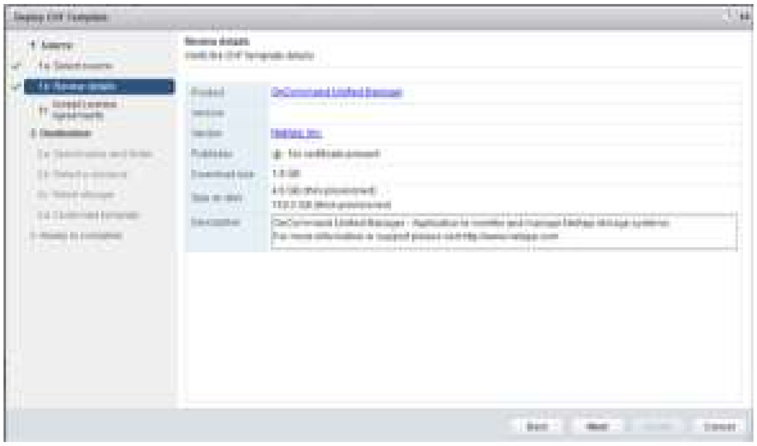

7. Read the end-user license agreement (EULA), and then click Accept to accept the agreement. Click Next.

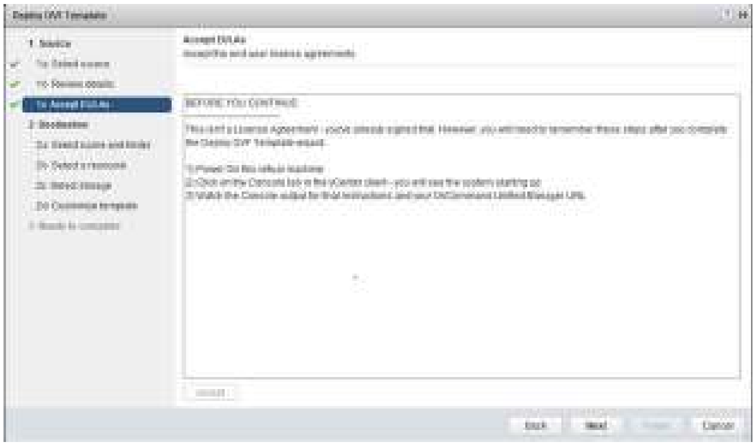

Deploy OVF Template

1 Source
1a Select source
1b Review details
1c Accept EULAs
2 Destination
2a Select name and folder
2b Select a resource
2c Select storage
2d Customize template
3 Ready to complete

Accept EULAs
Accept the end user license agreements

BEFORE YOU CONTINUE:

This isn't a License Agreement - you've already signed that. However, you will need to remember these steps after you complete the Deploy OVF Template wizard.

1) Power On this virtual machine
2) Click on the Console tab in the vCenter client - you will see the system starting up
3) Watch the Console output for final instructions and your OnCommand Unified Manager URL

8. Enter the name of the VM and select the FlexPod_DC folder to hold it. Click Next to continue.

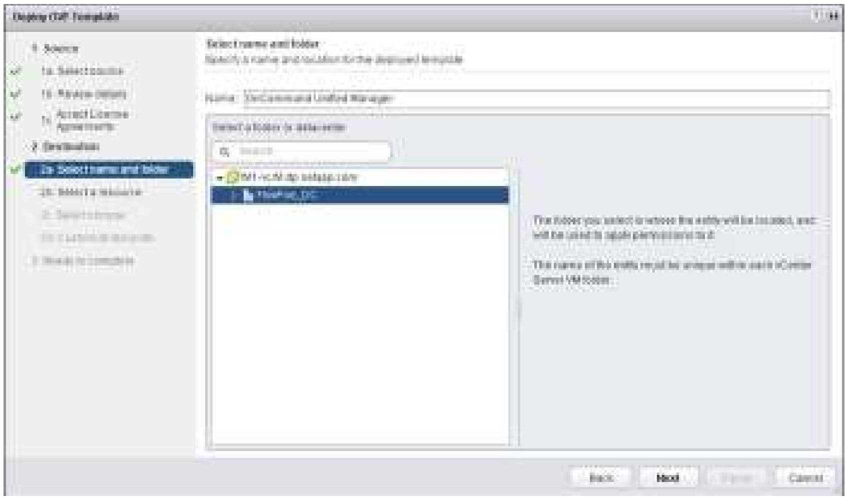

Deploy OVF Template

1 Source
1a Select source
1b Review details
1c Accept License Agreements
2 Destination
2a Select name and folder
2b Select a resource
2c Select storage
2d Customize template
3 Ready to complete

Select name and folder
Specify a name and location for the deployed template

Name: OnCommand Unified Manager

Select a folder or datacenter

[illegible]
FlexPod_DC

The folder you select is where the entity will be located, and will be used to apply permissions to it.

The name of the entity must be unique within each vCenter Server VM folder.

9. Select FlexPod_Management within the FlexPod_DC datacenter as the destination compute resource pool to host the VM. Click Next.

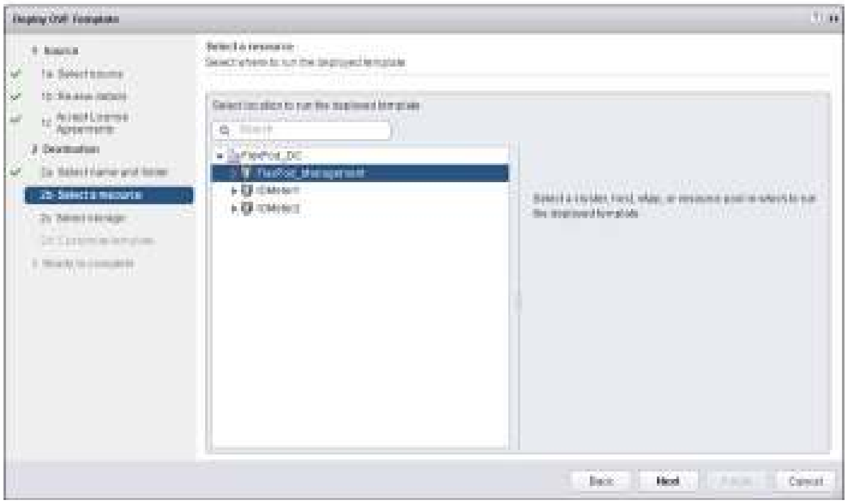

Deploy OVF Template

1 Source
1a Select source
1b Review details
1c Accept License Agreements
2 Destination
2a Select name and folder
2b Select a resource
2c Select storage
2d Customize template
3 Ready to complete

Select a resource
Select where to run the deployed template

Select location to run the deployed template

FlexPod_DC
FlexPod_Management
[illegible]
[illegible]

Select a cluster, host, vApp, or resource pool in which to run the deployed template.

10. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the virtual disk format. Click Next.

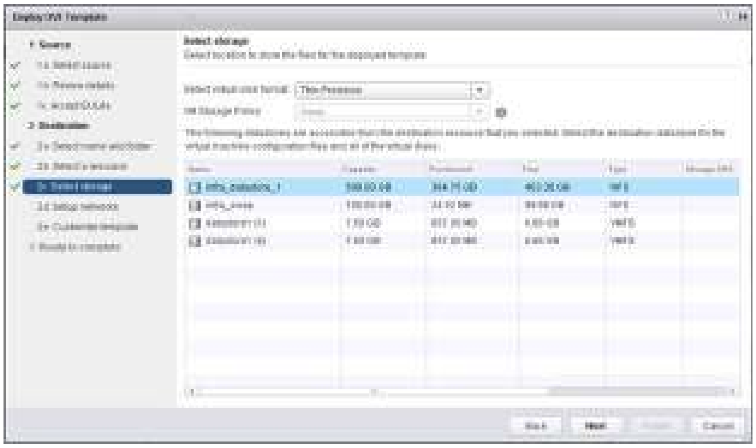

11. Select IB-MGMT-VLAN as the destination network for the nat source network. Click Next.

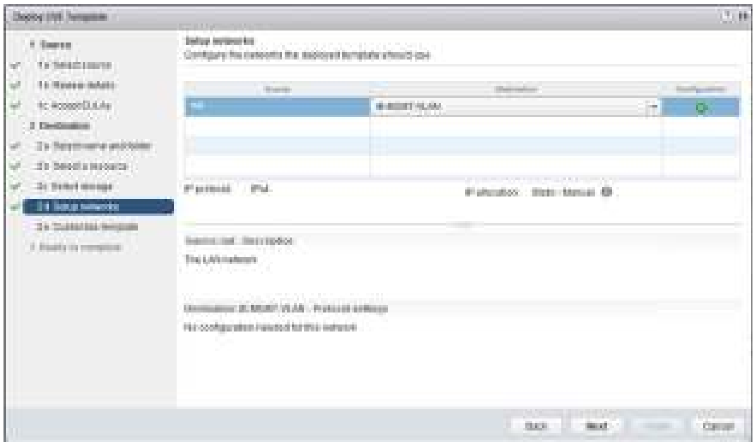

12. Complete the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS fields. Click Next.

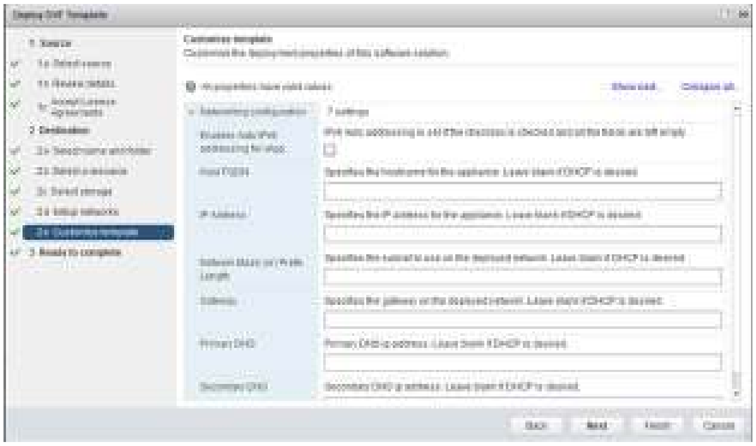

13. Clear the Power On After Deployment checkbox.
14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration information.
15. In the left pane, select vCenter -> Virtual Machines. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.

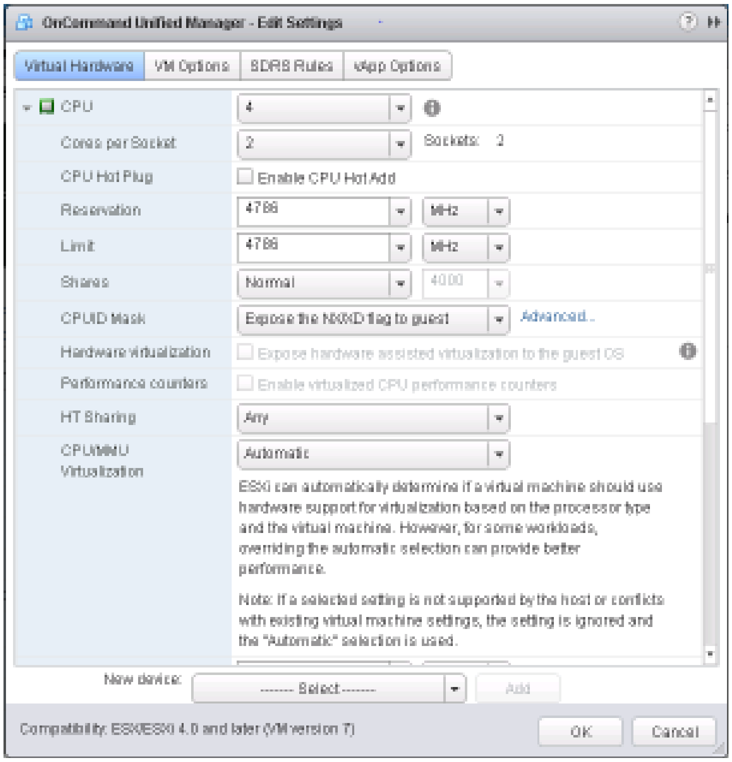

OnCommand Unified Manager - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

CPU: 4

Cores per Socket: 2 Sockets: 2

CPU Hot Plug: □ Enable CPU Hot Add

Reservation: 4786 MHz

Limit: 4786 MHz

Shares: Normal 4000

CPUID Mask: Expose the NX/XD flag to guest Advanced...

Hardware virtualization: □ Expose hardware assisted virtualization to the guest OS

Performance counters: □ Enable virtualized CPU performance counters

HT Sharing: Any

CPU/MMU Virtualization: Automatic

ESXi can automatically determine if a virtual machine should use hardware support for virtualization based on the processor type and the virtual machine. However, for some workloads, overriding the automatic selection can provide better performance.

Note: If a selected setting is not supported by the host or conflicts with existing virtual machine settings, the setting is ignored and the "Automatic" selection is used.

New device: ------- Select ------- Add

Compatibility: ESX/ESXi 4.0 and later (VM version 7)

OK Cancel

16. Size the VM's CPU and memory parameters according to the OnCommand Unified Manager 6.3 Installation and Setup Guide.
17. Click OK to accept the changes.
18. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Unified Manager Basic Setup

To setup the OnCommand Unified Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.
2. In the VMRC window, select VMRC > Manage > Install VMware Tools. VMware Tools will install in the VM.

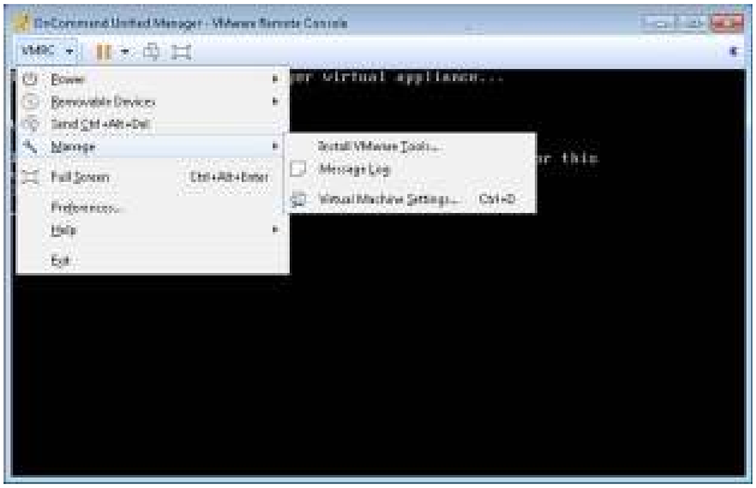

3. Set up OnCommand Unified Manager by answering the following questions in the console window:

```
Geographic area: <<Enter the number corresponding to your time zone>>
```

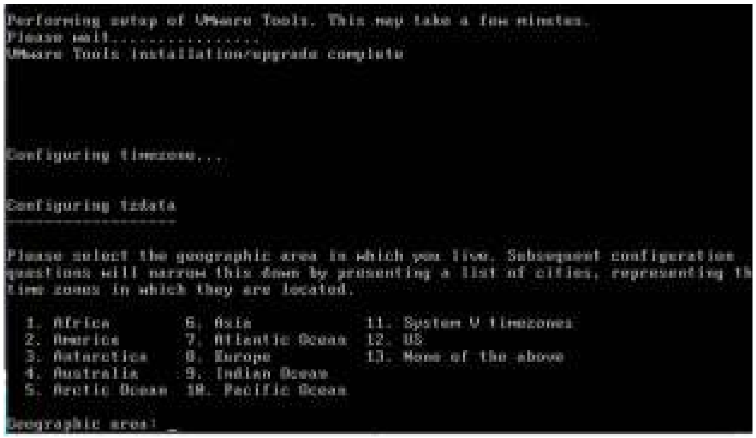

```
Performing setup of VMware Tools. This may take a few minutes.
Please wait.................
VMware Tools installation/upgrade complete




Configuring timezone...


Configuring tzdata
------------------

Please select the geographic area in which you live. Subsequent configuration
questions will narrow this down by presenting a list of cities, representing th
time zones in which they are located.

  1. Africa        6. Asia            11. System V timezones
  2. America       7. Atlantic Ocean  12. US
  3. Antarctica    8. Europe          13. None of the above
  4. Australia     9. Indian Ocean
  5. Arctic Ocean  10. Pacific Ocean

Geographic area: _
```

These commands complete the network configuration checks, generate SSL certificates for HTTPS and start the OnCommand Unified Manager services.

4. To Create a Maintenance User account, run the following commands:

> The maintenance user manages and maintains the settings on the OnCommand Unified Manager virtual appliance.

```
Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>
```

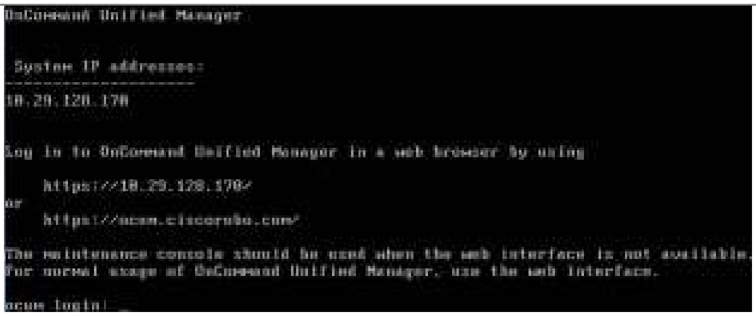

```
OnCommand Unified Manager


 System IP addresses:
---------------------
10.29.128.170


Log in to OnCommand Unified Manager in a web browser by using

    https://10.29.128.170/
or
    https://ocum.ciscorubu.com/

The maintenance console should be used when the web interface is not available.
For normal usage of OnCommand Unified Manager, use the web interface.

ocum login: _
```

5. With a web browser, navigate to the OnCommand Unified Manager using URL: https://<<var_oncommand_server_ip>>.

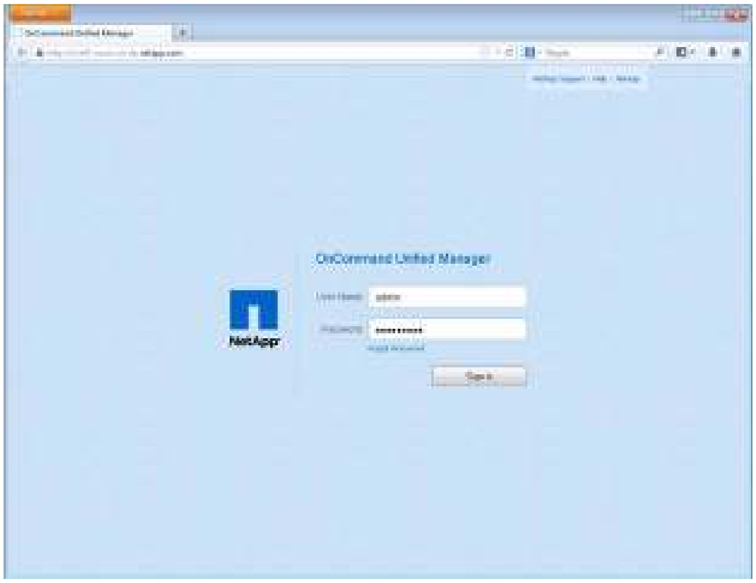

User Name: admin

Password: ••••••••••

6. Log in using the maintenance user account credentials.

7. Select Yes to enable AutoSupport capabilities.

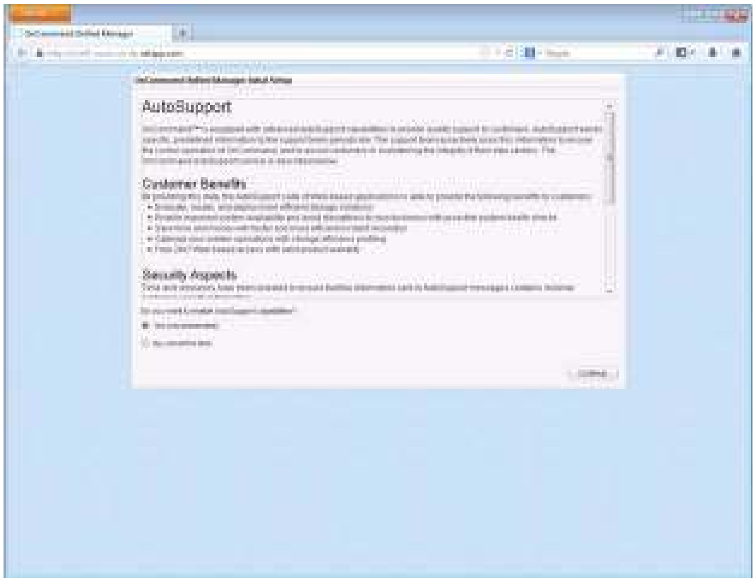

OnCommand Unified Manager Initial Setup

AutoSupport

Customer Benefits

Security Aspects

8. Click Continue.
9. Enter the NTP server IP address <<var_switch_a_ntp_ip>>.
10. Enter the storage admin e-mail address <<var_storage_admin_email>>.
11. Enter the SMTP server host name.

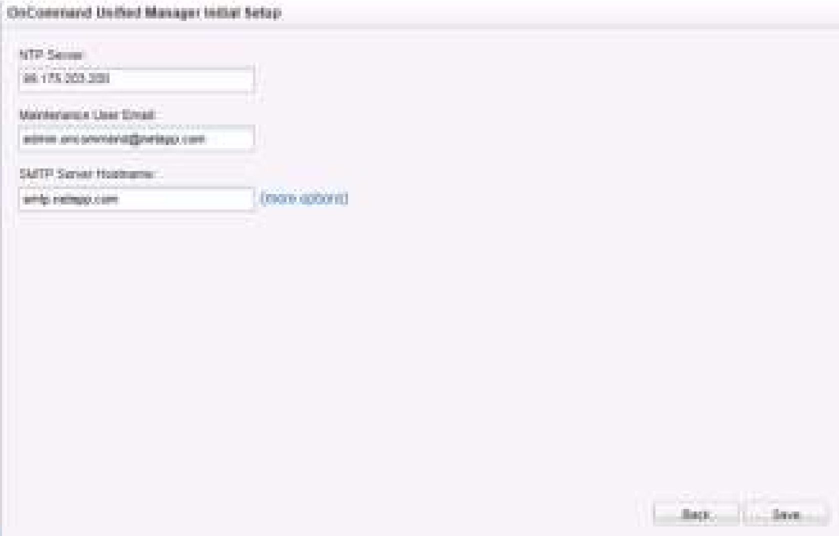

OnCommand Unified Manager Initial Setup

NTP Server: 66.175.203.200

Maintenance User Email: admin.oncommand@netapp.com

SMTP Server Hostname: smtp.netapp.com

12. Click Save.
13. Click Add Cluster.

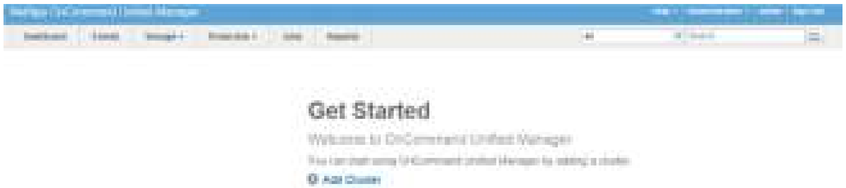

Get Started

Welcome to OnCommand Unified Manager

14. Provide the cluster management IP address, user name, password, protocol, and port.

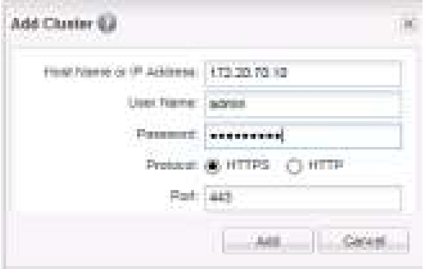

Add Cluster

Host Name or IP Address: 172.20.70.10

User Name: admin

Password: ••••••••••

Protocol: ● HTTPS ○ HTTP

Port: 443

15. Click Add.

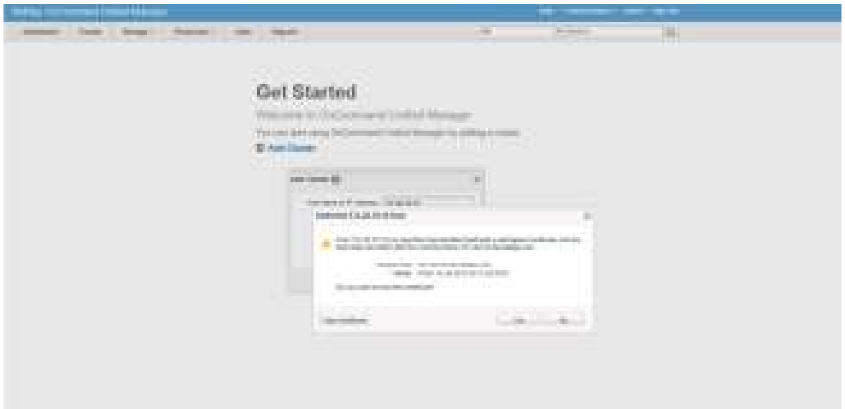

16. Click Yes to trust the certificate from the controller.

> The Cluster Add operation might take a couple of minutes.

17. After the cluster is added, it can be accessed by clicking on the Storage tab and selecting Clusters.

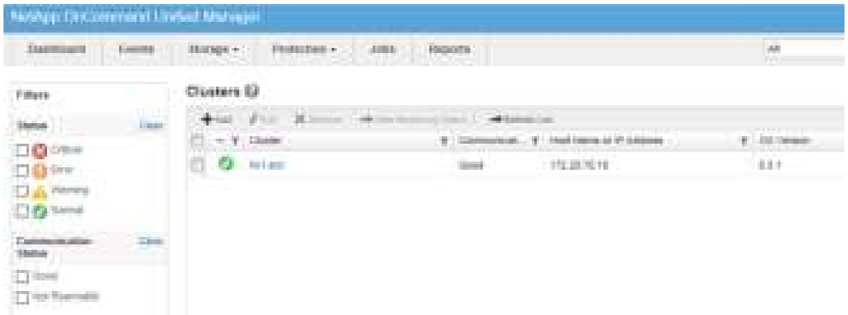

OnCommand Performance Manager 2.0

OnCommand Performance Manager OVF Deployment

To install OnCommand Performance Manager, complete the following steps:

1. Download and review the OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances.
2. Download OnCommand Performance Manager version 2.0 (OnCommandPerformanceManager-netapp-2.0.0.ova).
3. Log in to the vSphere Web Client. Select Home > VMs and Templates.
4. At the top of the center pane, click Actions > Deploy OVF Template.

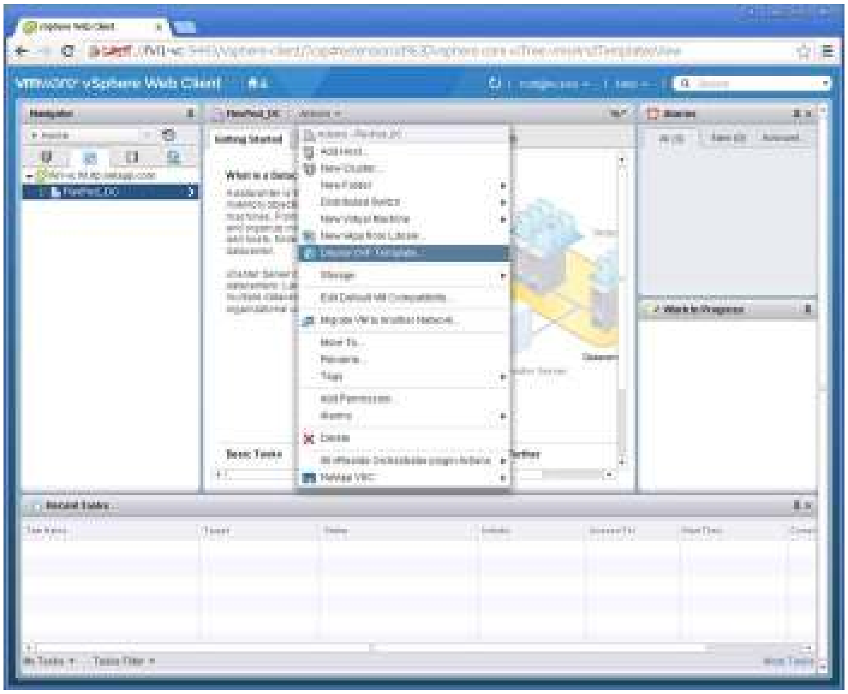

5. Browse to the OnCommandPerformanceManager-netapp-2.0.0.ova file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.

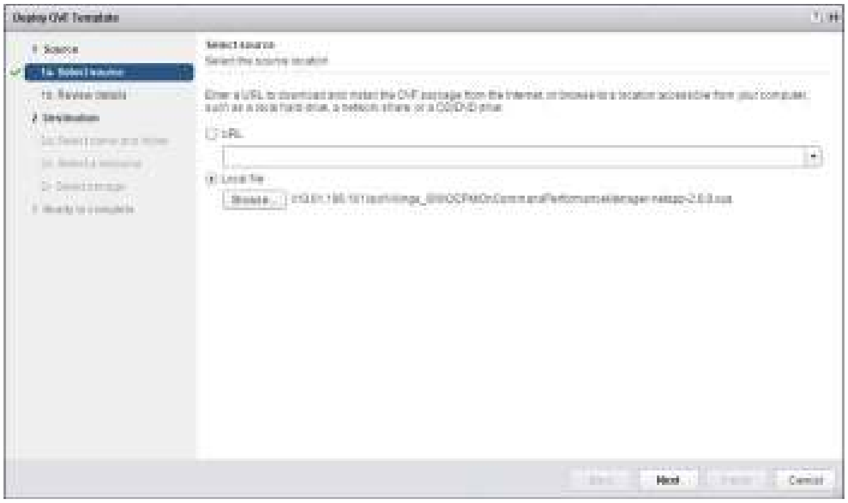

6. Review the details and click Next.

7. Read the EULA and click the Accept button to accept the agreement. Click Next.

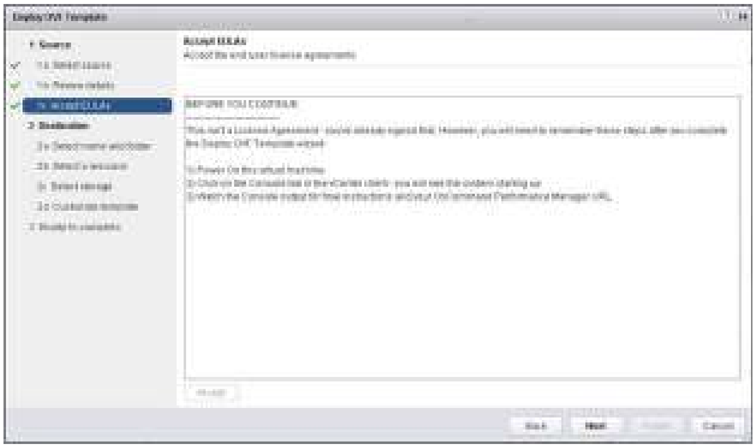

8. Enter the name of the VM and select the FlexPod_DC folder to hold the VM. Click Next.

9. Select FlexPod_Management within the FlexPod_DC datacenter as the destination compute resource pool to host the VM. Click Next.

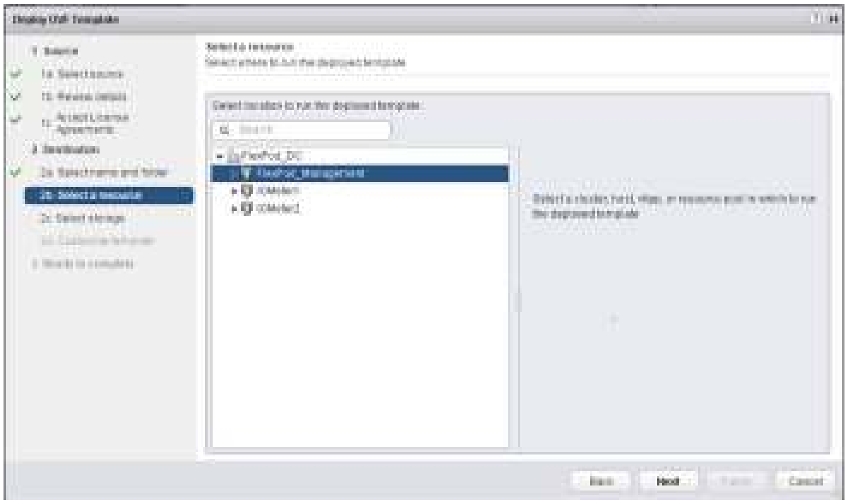

10. Select infra_datastore_1 as the storage target for the VM, and select Thin Provision as the virtual disk format. Click Next.

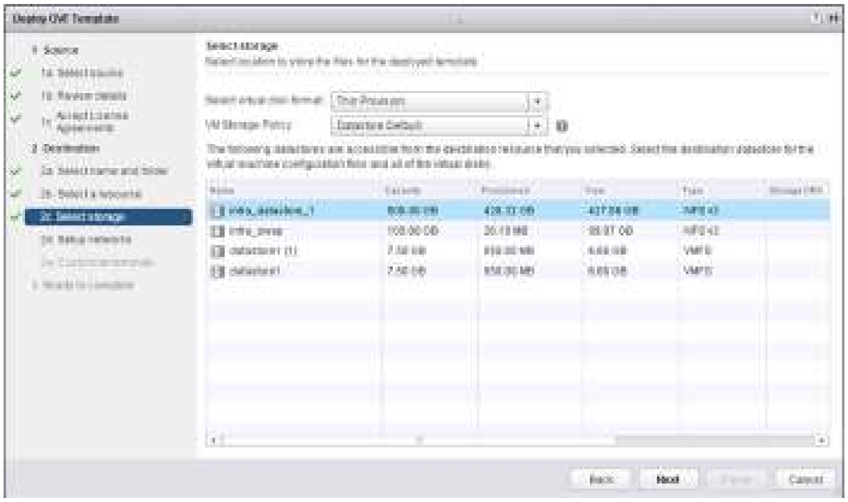

Deploy OVF Template

1 Source
1a Select source
1b Review details
1c Accept License Agreements
2 Destination
2a Select name and folder
2b Select a resource
2c Select storage
2d Setup networks
2e Customize template
3 Ready to complete

Select storage
Select location to store the files for the deployed template

Select virtual disk format: Thin Provision
VM Storage Policy: Datastore Default

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
infra_datastore_1	[illegible]	[illegible]	[illegible]	NFS v3	
infra_swap	100.00 GB	[illegible]	99.97 GB	NFS v3	
datastore1 (1)	7.50 GB	[illegible]	[illegible]	VMFS	
datastore1	7.50 GB	[illegible]	[illegible]	VMFS	

11. Select IB–MGMT-VLAN as the destination network for the nat source network. Click Next.

12. Enter the host name, IP address, network mask, gateway, primary DNS, and secondary DNS. Click Next.

13. Deselect Power On After Deployment.

14. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.

15. In the left pane, navigate to Home -> Hosts and Clusters. Expand the FlexPod_Management cluster and select the newly created OnCommand Performance Manager VM. After OVF deployment is complete, right-click the newly created VM and select Edit Settings.

16. Expand the CPU options.

a. The minimum required CPU reservation is 9572MHz. Determine the CPU frequency of the host.

b. Set the required number of CPUs (9572 / CPU frequency of the host).

c. Set the number of cores per socket where the socket number on the right matches the number of CPU sockets in the host. For example, if a host has two CPUs operating at a speed of 1999MHz, then the VM requires six virtual CPUs (9572 / 1999 = 4.79 - rounded to 6 virtual CPUs). If the host has two physical CPU sockets, then set three cores per socket.

For detailed information, see OnCommand Performance Manager 2.0 Installation and Administration Guide for VMware Virtual Appliances.

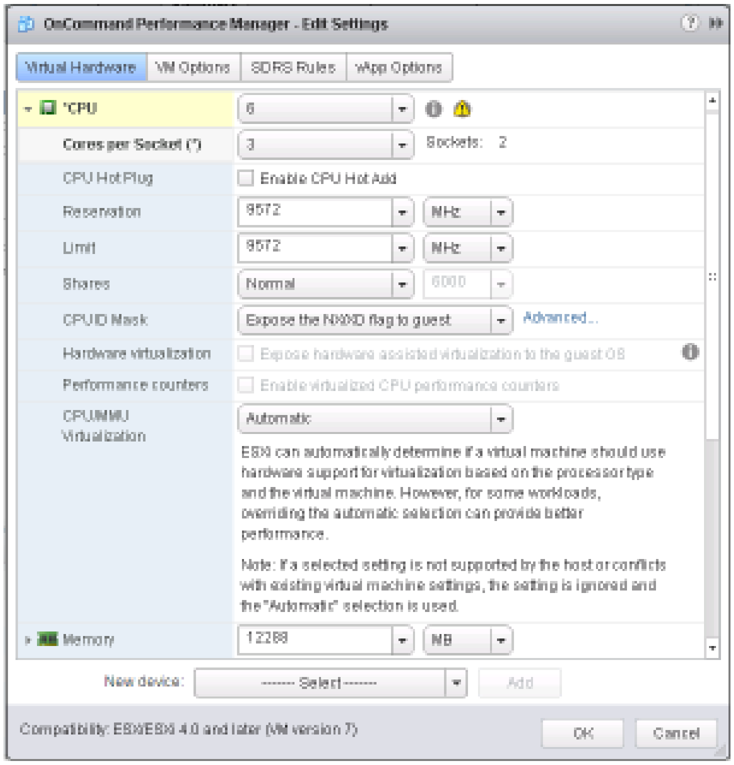

OnCommand Performance Manager - Edit Settings

*CPU	6	
Cores per Socket (*)	3	Sockets: 2
CPU Hot Plug	□ Enable CPU Hot Add	
Reservation	9572	MHz
Limit	9572	MHz
Shares	Normal	6000
CPUID Mask	Expose the NX/XD flag to guest	Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS	
Performance counters	□ Enable virtualized CPU performance counters	
CPU/MMU Virtualization	Automatic ESXi can automatically determine if a virtual machine should use hardware support for virtualization based on the processor type and the virtual machine. However, for some workloads, overriding the automatic selection can provide better performance. Note: If a selected setting is not supported by the host or conflicts with existing virtual machine settings, the setting is ignored and the "Automatic" selection is used.	
Memory	12288	MB

New device: ------- Select ------- Add

Compatibility: ESX/ESXi 4.0 and later (VM version 7)

17. Click OK to accept the changes.

18. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Performance Manager Basic Setup

To setup the OnCommand Performance Manager, complete the following steps:

1. Select the VM in the left-hand pane. In the center pane, select Launch Remote Console.
2. In the VMware Remote Console window, select VMRC > Manage > Install VMware Tools. VMware Tools installs in the VM.

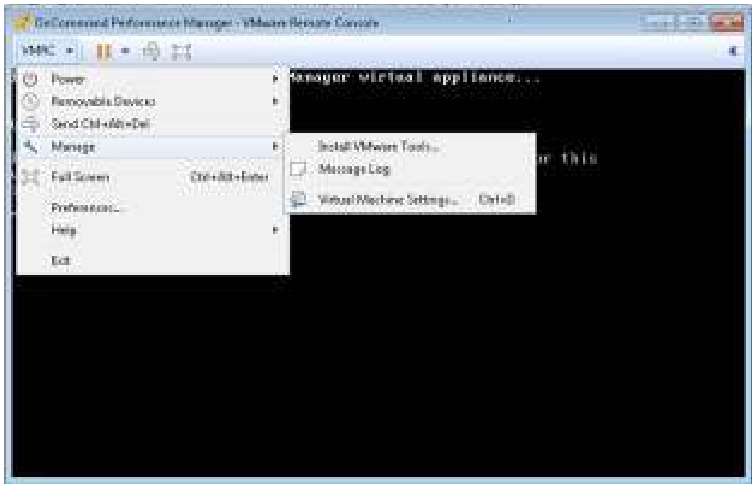
OnCommand Performance Manager - VMware Remote Console
VMRC
Power
Removable Devices
Send Ctrl+Alt+Del
Manage
Full Screen Ctrl+Alt+Enter
Preferences...
Help
Exit
Install VMware Tools...
Message Log
Virtual Machine Settings... Ctrl+D
Manager virtual appliance...
r this

3. Set up OnCommand Performance Manager by answering the following questions in the console window:

```
Geographic area: <<Enter your geographic location>>
Time zone: <<Select the city or region corresponding to your time zone>>
```

These commands complete the network configuration checks, generate SSL certificates, and start OnCommand Performance Manager services.

4. To create a maintenance user account, run the following commands:

The maintenance user manages and maintains the settings on the OnCommand Performance Manager virtual appliance.

```
Username : admin
Enter new UNIX password: <<var_password>>
Retype new UNIX password: <<var_password>>
```

5. Using a web browser, navigate to OnCommand Performance Manager using the URL https:// <<var_oncommand_pm_ip>>.
6. Log in using the maintenance user account (admin) credentials.
7. Enter a maintenance user e-mail address, SMTP mail server information, and the NTP server IP address. Click Save.
8. Select the Yes option to enable AutoSupport capabilities. Click Save.
9. Click Save and go to next step to not change the admin password.
10. Enter the storage cluster host name or IP address, the storage cluster admin user name, and the storage cluster admin password. Click Add Cluster, and then click Save and Complete Configuration. It can take up to 15 minutes for the cluster to be visible in OnCommand Performance Manager.

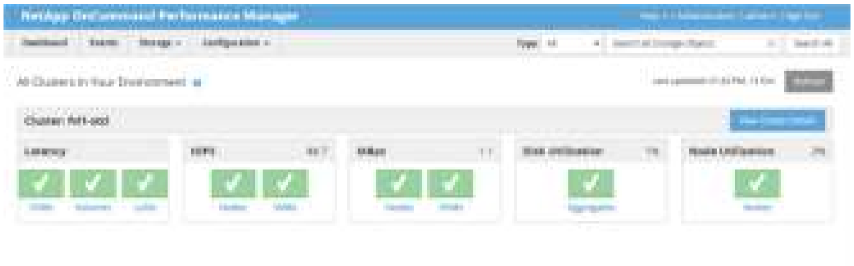

11. After the cluster is added, it can be accessed by clicking on Administration > Manage Data Sources.

Link OnCommand Performance Manager to OnCommand Unified Manager

To link OnCommand Performance Manager to the OnCommand Unified Manager, complete the following steps:

1. Using a web browser, navigate to OnCommand Unified Manager using the URL https:// <<var_oncommand_server_ip>>. Log in with the maintenance user ID and password set up earlier.
2. In the OnCommand Unified Manager web interface, select Administration > Manage Users to set up an Event Publication user.
3. Click Add to add a user.
4. Leave the Type set to Local User. Use eventpub as the name and enter and confirm a password. Enter an e-mail address for this user and set the Role to Event Publisher. Click Add.

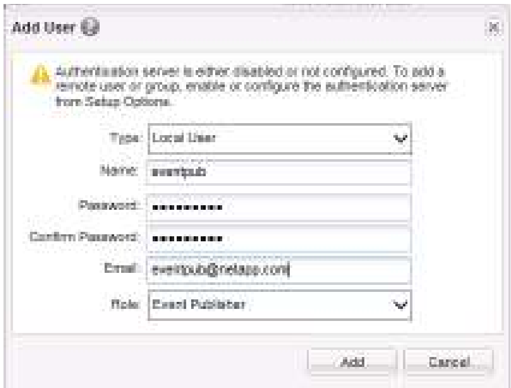

Add User

Authentication server is either disabled or not configured. To add a remote user or group, enable or configure the authentication server from Setup Options.

Type: Local User

Name: eventpub

Password: ••••••••••

Confirm Password: ••••••••••

Email: eventpub@netapp.com

Role: Event Publisher

5. In the OnCommand Performance Manager console window, log into the Command Line Interface with the maintenance user (admin) defined earlier.
6. Enter 5 to select Unified Manager Connection.

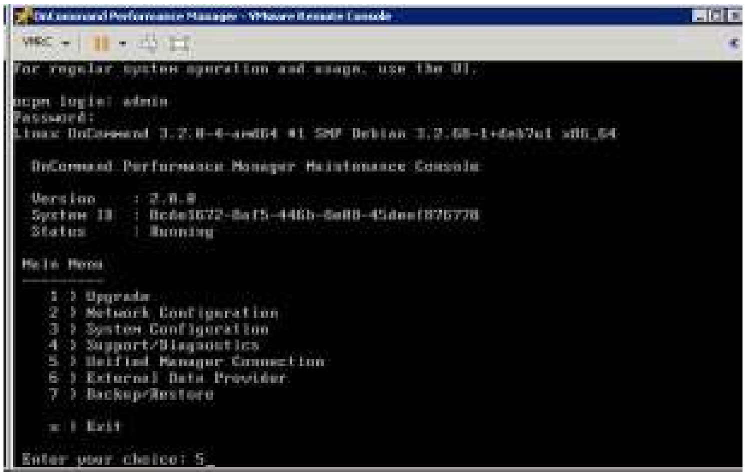

```
For regular system operation and usage, use the UI.

opm login: admin
Password:
Linux OnCommand 3.2.0-4-amd64 #1 SMP Debian 3.2.68-1+deb7u1 x86_64

  OnCommand Performance Manager Maintenance Console

  Version     : 2.0.0
  System ID   : 0cde1672-0af5-446b-0a00-45deef876770
  Status      : Running

 Main Menu
 ---------
    1 ) Upgrade
    2 ) Network Configuration
    3 ) System Configuration
    4 ) Support/Diagnostics
    5 ) Unified Manager Connection
    6 ) External Data Provider
    7 ) Backup/Restore

    x ) Exit

 Enter your choice: 5_
```

7. Enter 2 to Add / Modify Unified Manager Server Connection.
8. Enter y to continue.
9. Enter the OnCommand Unified Manager FQDN or IP address.
10. Click Enter to accept the default port 443.
11. Enter y to accept the Unified Manager security certificate.
12. Enter eventpub for the Event Publisher User Name.
13. Enter the eventpub password.
14. Enter y to accept the entered settings.
15. Press any key to continue.
16. Exit the OnCommand Performance Manager console. OnCommand Performance Manager events now appear in the OnCommand Unified Manager Dashboard.

NetApp NFS Plug-In 1.1.0 for VMware **VAAI**

Enable VMware vStorage for NFS in Clustered Data ONTAP

To enable VMware vStorage for NFS in ONTAP, complete the following steps:

1. From an SSH session to the storage cluster management address, log in with the admin user name and password.
2. Enable vStorage on the SVM.

```
vserver nfs modify –vserver Infra-SVM –vstorage enabled
```

3. Verify that the export policy rules are set up correctly.

```
vserver export-policy rule show –vserver Infra-SVM
```

The following text provides a sample output:

```
NetApp::> vserver export-policy rule show -vserver Infra-SVM
```

```
           Policy        Rule    Access   Client              RO
Vserver      Name            Index   Protocol Match               Rule
------------ --------------- ------  --------- --------------------- ---------
Infra-SVM    default        1        nfs       192.168.170.61        sys
Infra-SVM    default        2        nfs       192.168.170.60         sys
Infra-SVM    default        3        nfs       192.168.170.58         sys
Infra-SVM    default        4        nfs       192.168.170.59         sys
Infra-SVM    default        5        nfs       192.168.170.62         sys
Infra-SVM    default        6        nfs       192.168.170.63         sys
6 entries were displayed.
```

4. The access protocol for the FlexPod policy name should be nfs. If the access protocol is not nfs for a given rule index, run the following command to set nfs as the access protocol:

```
vserver export-policy rule modify –vserver Infra-SVM –policyname default –ruleindex <<var_rule_index>> -protocol nfs
```

Install NetApp NFS Plug-In for VMware VAAI

To install the NetApp NFS plug-in for VMware vStorage APIs for Array Integration (VAAI), complete the following steps:

1. From a console interface on the NetApp VSC VM, go to the Software Downloads page in the NetApp Support site.
2. Scroll down to locate the NetApp NFS Plug-in for VMware VAAI, select the ESXi6.0 platform, and click Go.
3. Click View & Download.
4. Click CONTINUE.
5. Click Accept.
6. Download the .vib file of the most recent plug-in version to the VSC VM Desktop as NetAppNasPlugin.vib.

It is important that the file be saved as NetAppNasPlugin.vib.

7. On the VSC VM desktop, move the NetAppNasPlugin.vib file to the C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web folder.

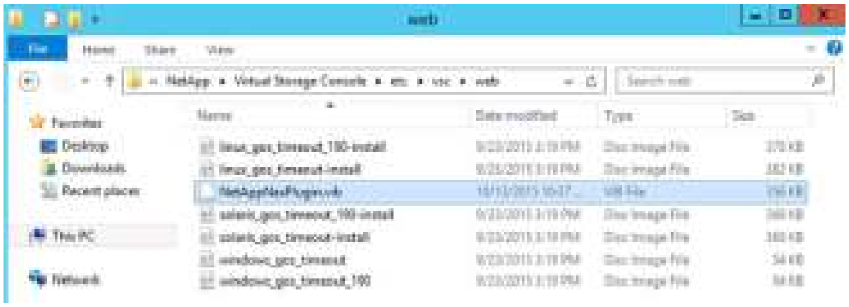

8. Go to the VMware vSphere Web Client and select VSC. Click NFS VAAI Tools. Make sure NFS Plug-in for VMware VAAI Version: 1.1.0-0 is shown.

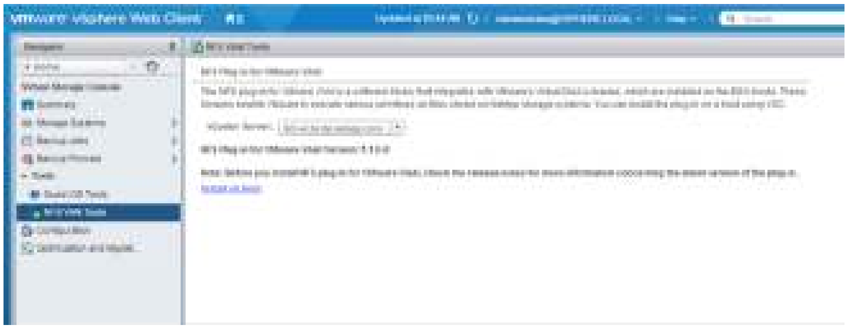

9. Click Install on Host. Select all of the hosts on which you want to install the plug-in.

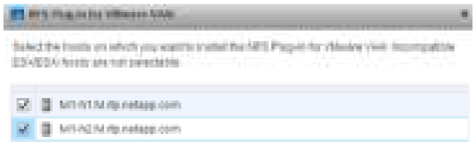

10. Click Install and then click OK.

11. One at a time, put each ESXi host into maintenance mode, reboot the host, and then exit maintenance mode. It might be necessary to manually migrate VMs to the other host to allow the host to enter maintenance mode.

12. When the reboots have completed, click Storage in the vSphere web client from the Home page. Then select the infra_datastore_1 datastore. Select Settings under the Manage tab in the center pane. Hardware Acceleration should now read Supported on All Hosts, as is shown in the following screenshot. All NFS datastores should now support hardware acceleration.

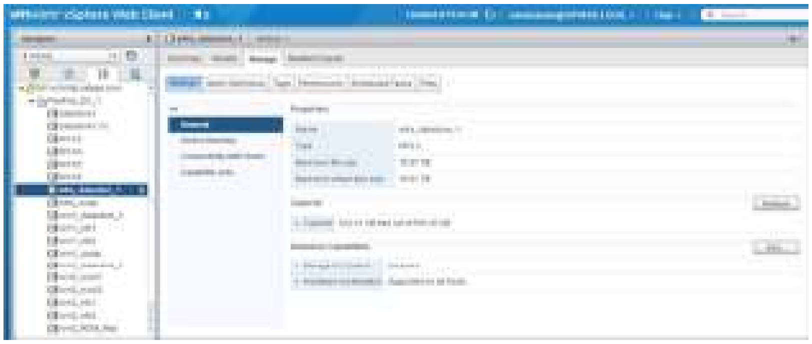

About the Authors

Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.

Ramesh Isaac is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Group. Ramesh has worked in data center and mixed-use lab settings since 1995. He started in information technology supporting UNIX environments and focused on designing and implementing multi-tenant virtualization solutions in Cisco labs over the last couple of years. Ramesh holds certifications from Cisco, VMware, and Red Hat.

Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Masters of Science in Information Security from East Carolina University.

Dave Derry, Technical Marketing Engineer, Infrastructure and Cloud Engineering, NetApp

Dave Derry is a Technical Marketing Engineer in the Converged Infrastructure Engineering team at NetApp. He has been with NetApp since 2012, serving in a variety of engineering roles. Prior to that, he was an engineer at Cisco Systems for over ten years, in a variety of development and solution test roles.

Acknowledgements

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

- John George, Cisco Systems, Inc.
- Haseeb Niazi, Cisco Systems, Inc.
- Melissa Palmer, NetApp

About NetApp

Throughout the world, leading organizations count on NetApp for software, systems, and services to store, manage, protect, and retain one of their most precious assets: their data. We enable enterprises, service providers and partners to envision, deploy, and evolve their IT environments. Customers benefit from our open collaboration with other technology leaders to create the specific solutions they need. We were incorporated in 1992 and created the world's first networked storage appliance.

Today, we offer a portfolio of products and services that satisfy a broad range of customer workloads across different data types and deployment models. NetApp's cloud-connected flash solutions, an element of a Data Fabric strategy, provide the simplicity, operational efficiency and protection needed to support innovation, add unprecedented performance and power the most demanding data sets and technologies, such as artificial intelligence or 5G networks. Only NetApp delivers everything that companies need to build their own unique Data Fabric that spans public cloud, private cloud and on-premises environments. With NetApp solutions, this Data Fabric can deliver applications that engage users and can provide analytics that turn insights into a competitive advantage.

The industry has awarded NetApp with hundreds of honors for our innovation, leadership, and culture.

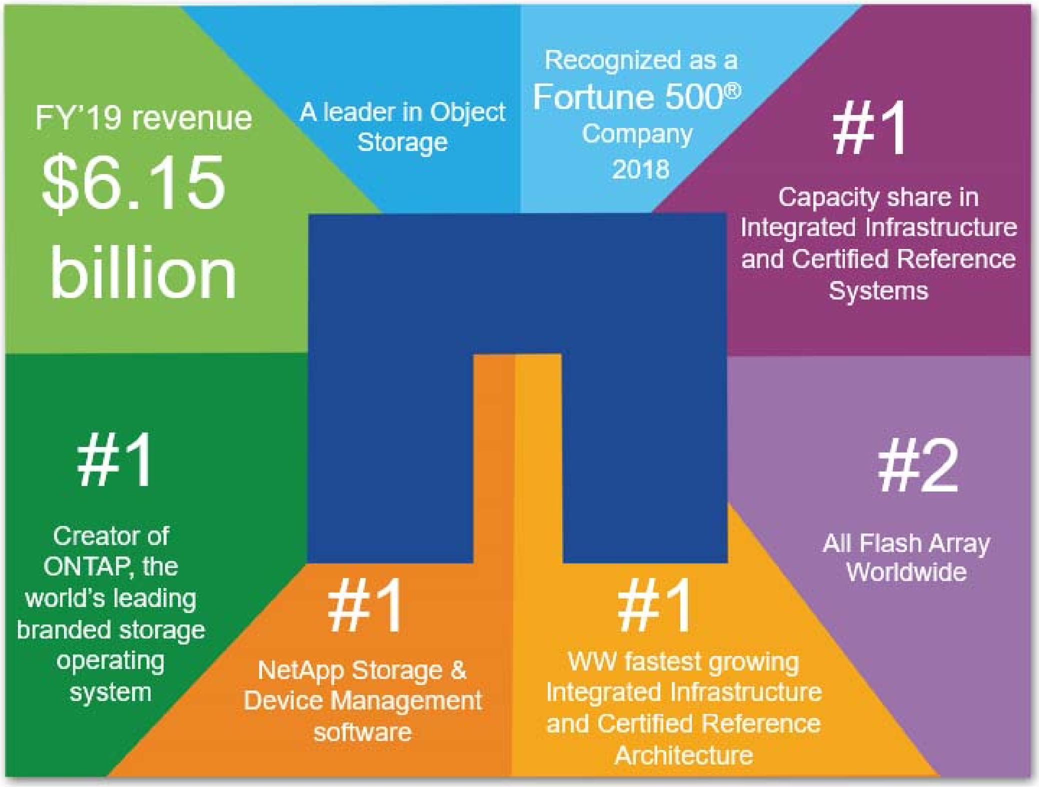
FY'19 revenue
$6.15 billion
A leader in Object Storage
Recognized as a Fortune 500® Company 2018
#1
Capacity share in Integrated Infrastructure and Certified Reference Systems
#1
Creator of ONTAP, the world's leading branded storage operating system
#1
NetApp Storage & Device Management software
#1
WW fastest growing Integrated Infrastructure and Certified Reference Architecture
#2
All Flash Array Worldwide

Figure 1: Successful customers drive NetApp Industry leadership.[1]

[1] http://fortune.com/fortune500/list/.

Industry Leadership

NetApp continues to lead the market with our integrated portfolio of platforms and solutions. We set the standards for the industry in Flash, OS, capacity, OpenStack, and Object Storage, to name a few. Our innovations help organizations to accelerate critical applications, increase data visibility, streamline data protection, and boost operational agility.

- #1 Integrated Infrastructure and Certified Reference Systems in Capacity Shipped[2]
- Leader in two Gartner Magic Quadrants:
  - General-Purpose Disk Arrays, November 8, 2018
  - Solid-State Arrays, July 23, 2018
- #1 in Commercial Storage for OpenStack[3]
- #1 WW NetApp Storage & Device Management software[4]
- #1 Creator of ONTAP, the world's leading open networked branded storage OS[5]
- Fastest growing Converged Systems vendor (YoY revenue growth) [2]
- Fastest growing Storage for Converged Systems vendor (YoY revenue growth)[2]
- #2 All Flash Array (AFA) Worldwide[5]
- Market-leading growth for six successive quarters with NetApp FlexPod[2]
- A Key Player in Data Services for Hybrid Cloud[6]

Financial Information

NetApp is a publicly held FORTUNE 500® [1] company with over $6 billion in revenue, and more than 10,000 employees in 149 offices worldwide. We are a member of the S&P 500 and NASDAQ 100 and our stock symbol is NTAP. Our annual reports can be found at http://investors.netapp.com/annuals.cfm.

The following table summarizes NetApp's financial information over the last five years. All amounts are in U.S. dollars.

Table 1: NetApp revenue and financial information for the last five years.

	FY 2019	FY 2018	FY 2017	FY 2016	FY 2015
(In millions, except per-share amounts)					
Total revenues	$6,146	$5,919	$5,491	$5,546	$6,123
Total cost of revenue	$2,201	$2,210	$2,127	$2,173	$2,290
Net income	$1,169	$116	$481	$229	$560
Current assets	$5,610	$6,952	$6,198	$6,448	$6,773

[2] **IDC Worldwide Quarterly Converged Systems Tracker - 2018Q4, April 2, 2019.**

[3] **OpenStack User Survey:** https://www.openstack.org/analytics.

[4] **IDC Worldwide Storage Software and Cloud Services Overview, 2019Q1, June 6, 2019 (#1 Storage & Device Management vendor – revenue share).**

[5] **IDC Worldwide Quarterly Enterprise Storage Systems Tracker, 2019Q1, June 6, 2019.**

[6] **IDC, Worldwide Data Services for Hybrid Cloud – Key Players Portfolio Analysis, IDC #US44266318, September 2018 A Key Player in Data Services for Hybrid Cloud**

	FY 2019	FY 2018	FY 2017	FY 2016	FY 2015
Total assets	$9,865	$9,493	$10,037	$9,401	$9,214
Shareholder's equity	$2,067	$2,780	$2,881	$3,414	$3,787

Our Customers

NetApp delivers data management solutions to over 30,000 businesses worldwide to help them unleash the full potential of their data. In a world transformed by digital technology, businesses cannot move ahead with data in silos. They need fast insights, seamless access, and the peace of mind of knowing their data is protected and within their control.

Figure 2: Businesses around the world and in every vertical market trust NetApp.

Strategic Partnerships

To empower our customers' continuing growth and success, we form partnerships with the industry's best reseller, application, infrastructure, consulting, and cloud service provider partners including:

- 50+ distributors globally
- 3500 active resellers
- 200+ Technology Alliance and Global System Integrator partners
- Hyperscale Service Providers

The following graphic shows a sample of our strategic alliances and partnerships.

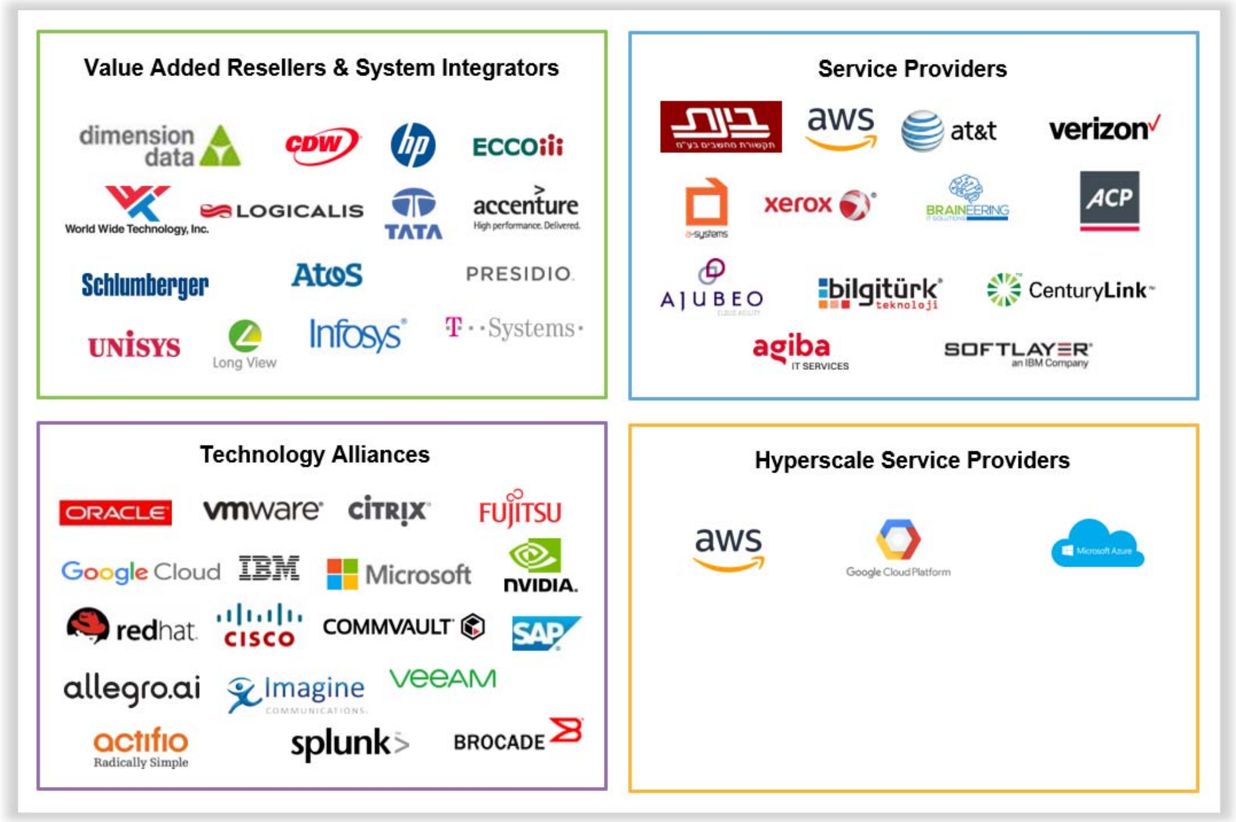

Figure 3: NetApp's partner ecosystem – *Developing and delivering integrated solutions to customers around the world.*

Research and Development

NetApp invests heavily in R&D to ensure that we provide the best value to our customers and maintain our market leadership. Our global R&D activities represent 14% to 16% of NetApp's net revenues.

One recent success story includes enabling control and choice in hybrid cloud. NetApp was the first data management vendor to build capabilities that support the cloud service provider community of 100s of xSPs, offering over 500 cloud-based services.

NetApp continues to advance its data replication, disaster recovery and data archive portfolio through building upon integration with SAP, Oracle, OpenStack, industry ISVs. Our acquisition of SolidFire helps accelerate our customers' transition to next generation data centers and web-scaled cloud applications.

As of June 2019, NetApp holds over 2,900 patent assets worldwide. This includes over 2,400 issued patents worldwide and hundreds of U.S. and international patent applications. In 2015, we held the second highest patent filings in our history, with 300 new patent applications filed and over 200 patents related to our Flash portfolio.

NetApp also supports innovative research in the academic community as part of its innovation strategy. The NetApp Advanced Technology Group is focused on exploring and evaluating the ideas, technologies, and trends that define the future of data management.

This group is a current member of the following university-industry consortiums:

- CERES: Center for Unstoppable Computing (University of Chicago)
- Parallel Data Laboratory (PDL), Carnegie Mellon University
- MIT CSAIL Alliance Program (Massachusetts Institute of Technology)
- Storage Systems Research Center (University of California, Santa Cruz)
- Wisconsin Institute on Software-defined Datacenters of Madison (University of Wisconsin –Madison)

The NetApp Advanced Technology Group has also sponsored research projects of professors and students at over twenty renowned IT universities across America, Canada, Italy and India.

NetApp's Global Support Capabilities

NetApp offers a broad portfolio of products and services including support services, professional services, and education. Businesses can rely on the NetApp support organization to resolve issues quickly, on a 24/7 basis. NetApp support is organized on a regional, country and city basis with overlapping service and parts coverage. Over 250,000 customer systems rely on our support worldwide.

The NetApp support organization manages:

- 400 depots
- 120 dispatchers with local language capability and security
- 3 test and repair centers
- 13 Technical Support locations: Amsterdam, Netherlands; Costa Rica; Dalian, China; Newcastle, U.K.; Lake Mary, Florida; Raleigh, North Carolina; Wichita, Kansas; Tokyo, Japan; Bangalore, India; Boulder, Colorado; Rochester, NY; Bogota, Colombia; Sofia, Bulgaria

NetApp Team

Our Team Is Data Driven and supports your vision for digital transformation. The strength of NetApp culture and vision attracts the industry's best, a fact reflected in our place on the FORTUNE "100 Best Companies to Work For®" list for 13 consecutive years. [7]

NetApp is committed to achieving market leadership by living our values and embracing strong principles. The NetApp executive team combines experience with innovation to help deliver our company strategies.

[7] Fortune Magazine.

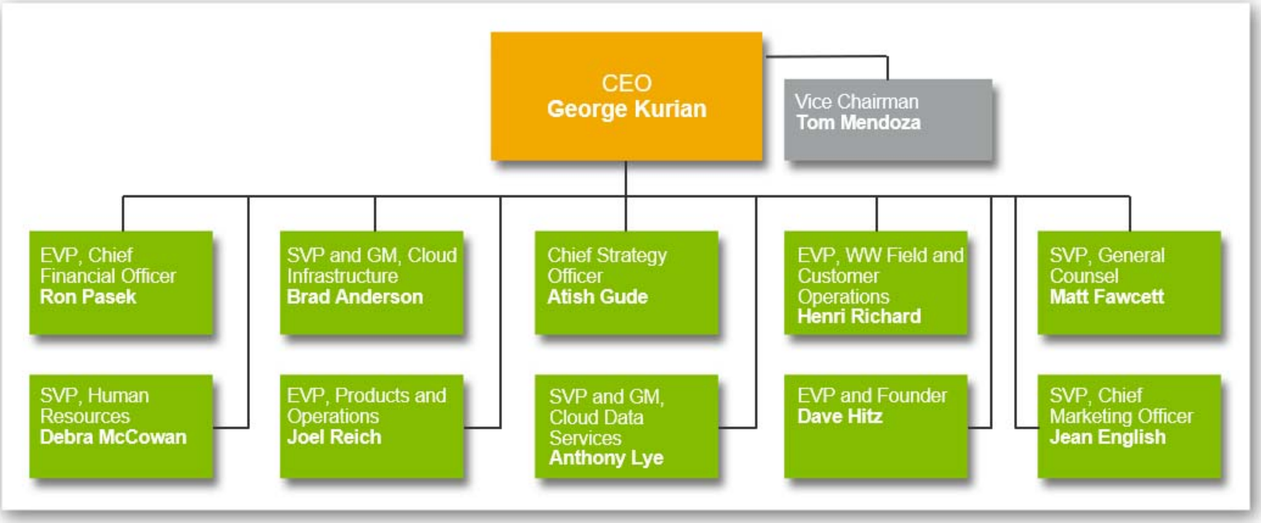

Figure 4: NetApp's executive team.

ONTAP 9: Harness the Power of the Hybrid Cloud

NetApp ONTAP simplifies data management for any application, anywhere; accelerate and protect data across the hybrid cloud; and future-proof your data infrastructure. ONTAP 9 is the next generation combining new levels of simplicity and flexibility with powerful data management capabilities and storage efficiencies. The latest version, ONTAP 9.5, enables additional integration of modern and traditional technologies—across flash, cloud, and software-defined architectures—to build a foundation for the data fabric.

WV Oasis's transformation into a digital business brings with it complexities in the short term. Your new priorities might require adding all-flash arrays for business-critical workloads, while integrating new applications into your existing environment, and managing data on premises as well as in the cloud—yet operations must be simplified, costs reduced, and budgets stretched.

NetApp® ONTAP 9® provides unmatched versatility, comprehensive data protection, and leading storage efficiency—delivering next-generation data management capabilities and efficiencies, fueled by simplicity and flexibility. ONTAP 9 enables you to deploy storage across your choice of architectures: engineered systems, software-defined storage (SDS), and the cloud, while unifying data management across each of them.

Leverage ONTAP 9 to:

- Simplify deployment and data management
- Adapt to changing business needs
- Power your enterprise applications

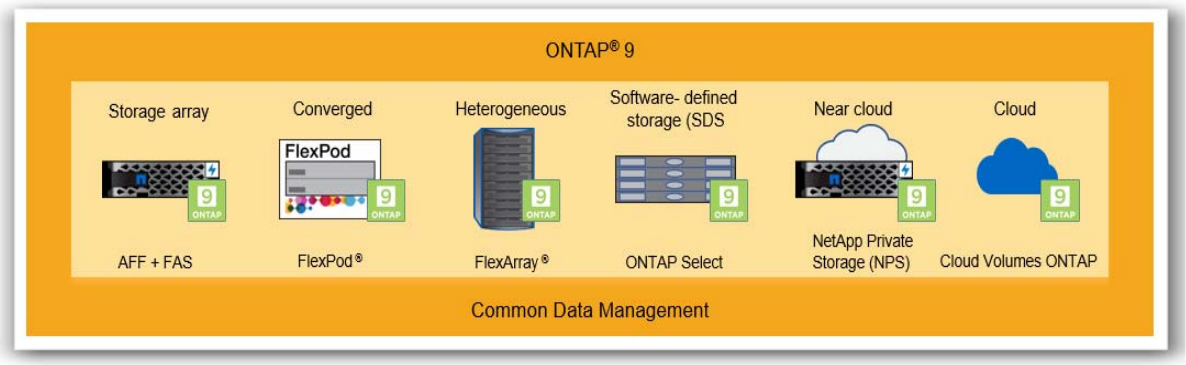

Figure 1: Standardize data management across architectures with a rich set of enterprise data services.

Simplify Deployment and Management

Although storage might double in size, it no longer means there is twice as much work to manage. ONTAP has a common set of features across deployment architectures that simplify complex tasks so your staff can be more productive.

Deploy New Workloads in Less Than 10 Minutes

New, fast provisioning workflows enable the deployment of key workloads such as Oracle, SQL Server, SAP-HANA, VDI, and VMware in less than 10 minutes from power-on to serving data. Years of NetApp experience and best practices are integrated into the system manager wizard and factory configurations, enabling you to quickly set up new configurations by answering a few questions. As new workloads are deployed, ONTAP 9 gives you the visibility to know which node has the most performance capacity available for optimal deployment.

Unified Data Management

Simplify your operations by unifying data management across a hybrid cloud that can span flash, disk, and cloud running SAN and NAS workloads. Increase the efficiency of your staff and easily move data between nodes to where it is most needed. ONTAP is the foundation for a Data Fabric that gives freedom, choice, and control across your storage environment.

Simplified, Powerful Management Capabilities

The NetApp OnCommand® software portfolio includes management products that manage virtualized private and hybrid cloud environments. Centrally monitor capacity, availability, performance, and data protection. You can take advantage of storage service analytics to make better informed decisions about your storage.

OnCommand management platform automates your data management processes by integrating into your data center orchestration platform for end-to-end service delivery for your private and hybrid cloud services.

Adapt to Changing Business Needs

ONTAP 9 provides the flexibility you need to design and deploy your storage environment across the widest range of architectures, so you can match the approach that is best for your evolving business needs:

- NetApp arrays: All Flash FAS (AFF) systems and hybrid-flash FAS systems
- Converged infrastructure: FlexPod®
- On commodity servers as SDS: ONTAP Select
- In front of third-party arrays: NetApp FlexArray® software
- Next to the cloud: NetApp Private Storage (NPS) for Cloud
- In the cloud: Cloud Volumes ONTAP

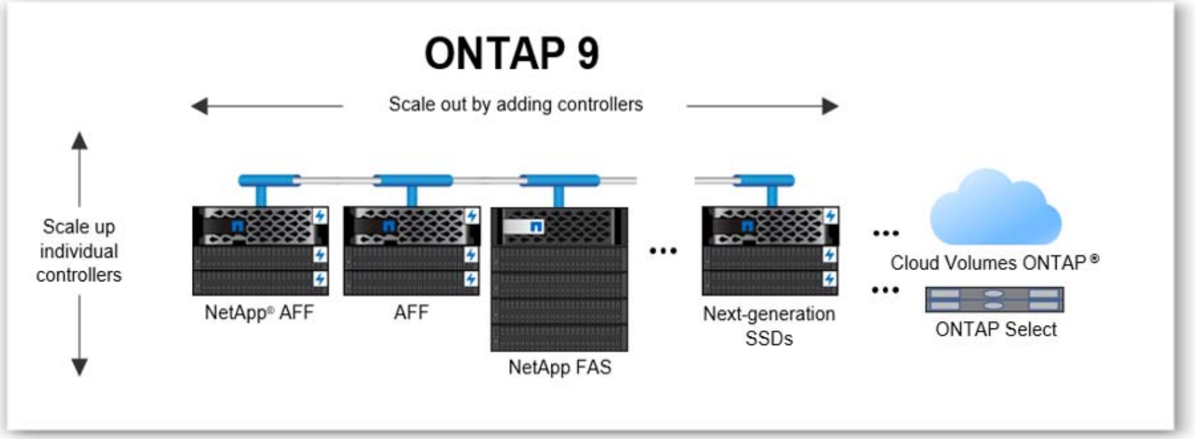

Figure 2: Scale seamlessly – *Scale out by intermixing your choice of flash and hybrid-flash nodes, upgrade hardware/software or scale up without disrupting users, incorporate software-defined, cloud, and future-generation flash.*

Flexibly consolidate both NAS and SAN workloads onto any ONTAP environment while delivering consistent data services. You can also seamlessly move your data between each deployment model to get your data onto the optimal environment for performance, capacity, and cost efficiency.

Add capacity as your business grows across both SAN and NAS environments. You can combine all-flash and hybrid-flash storage nodes into a larger storage cluster and connect them to the cloud. And ONTAP FabricPool technology can deliver up to 50% storage TCO savings by automatically tiering cold data from AFF/FAS systems, ONTAP Cloud, and ONTAP Select to the cloud, including Azure, AWS, and StorageGRID.

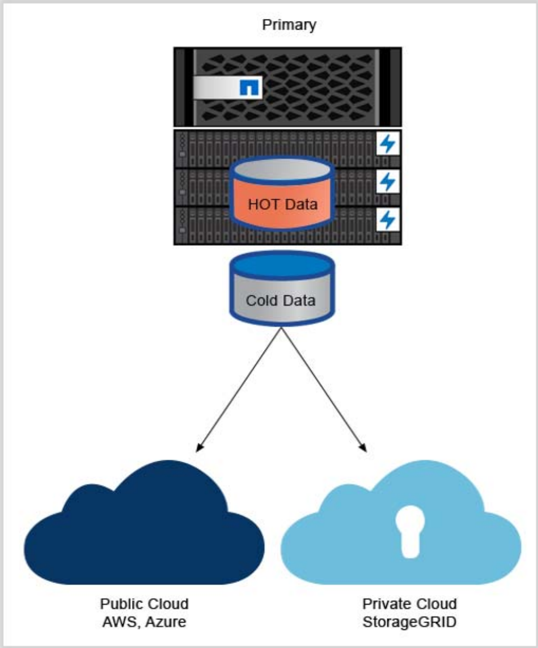

Figure 3: Automated cloud tiering of cold data.

Power Enterprise Applications

To support your critical applications, you need a storage environment that cost-effectively delivers high performance and availability that can also scale with business growth and protect your valuable data. ONTAP 9 delivers on all these requirements with highly efficient flash performance for scalable, nondisruptive operations.

Optimized for Flash

ONTAP 9 delivers the horsepower that critical applications require without compromising on rich data services. AFF systems running ONTAP 9 are optimized specifically for flash, including new NVMe technologies, providing up to twice the performance compared to the same workloads running on recent ONTAP 8 releases, while still delivering consistent submillisecond latency.

ONTAP 9 also enables FAS hybrid-flash storage systems to deliver flash-accelerated performance that is balanced with hard disk drives (HDD) economies. Hot data is automatically cached in flash to accelerate application performance.

Nondisruptive Operations

ONTAP gives you the ability to perform critical tasks without interrupting your business by dynamically assigning, promoting, and retiring storage resources without downtime over the lifecycle of an application. Data can be moved between controllers without application interruption. Storage controllers and disk shelves can be replaced without disruption, and with ONTAP you can mix models and generations of hardware to extend the life of existing investments.

NetApp MetroCluster™ technology delivers business continuity by synchronously mirroring between locations for continuous data availability. A MetroCluster storage array, leveraging FC or IP connectivity, can be deployed at a single site, across a metropolitan area, or in different cities.

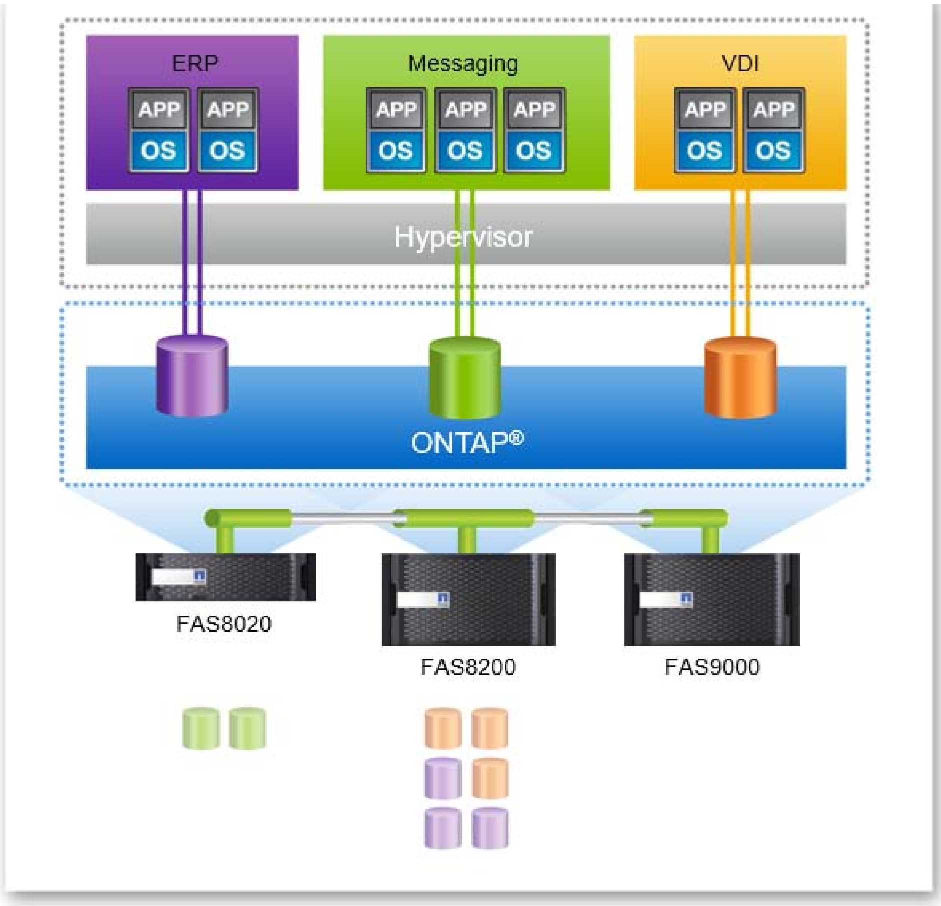

Figure 4: Nondisruptive operations – *Move data to available nodes and retire existing hardware.*

> *We can now provide access to our systems 24 hours a day. That's important to us and for our patients needing immediate care."*
>
> *— Tony Beaird, Director, Infrastructure and Security, DuPage Medical Group*

Integrated Data Protection

NetApp offers a complete suite of Integrated Data Protection (IDP) to safeguard your operations and keep them running smoothly. Meet your requirements for local backup with near-instant recovery by using space-efficient NetApp Snapshot™ copies. Achieve remote backup/recovery and disaster recovery with SnapMirror® asynchronous replication. Get zero data loss protection (RPO=0) with SnapMirror Synchronous replication.

NetApp also provides superior integration with enterprise backup vendors and leading applications. Our IDP solutions include integrated and unified disk-to-disk backup and disaster recovery in a single process for VMware and Microsoft virtualization.

Security and Compliance

Simplify and strengthen your security posture by integrating data security throughout your hybrid cloud. You can help meet governance, risk, and compliance (GRC) requirements such as HIPAA, PCI-DSS, and GDPR and cost effectively secure your NetApp ONTAP environment by incorporating industry-standard, built-in security that meets FIPS 140-2 compliance.

You> can easily and efficiently protect at-rest data with NetApp Storage Encryption (NSE)—that uses self-encrypting drives. Or encrypt any volume across FAS, AFF, and ONTAP Select deployments with NetApp Volume Encryption (NVE)—that does not require special encrypting disks. Key management can be delivered in a self-contained encryption solution using Onboard Key Manager (OKM), included with ONTAP, or with external key management solutions that provide separation of duties and a centralized key repository.

To meet stringent compliance and data retention policies, NetApp SnapLock® software enables write once, ready many (WORM) protected data for your ONTAP environment.

Superior Storage Efficiency

With ONTAP, you can reduce costs with one of the most comprehensive storage efficiency offerings in the industry. You get NetApp Snapshot copies, thin provisioning, and replication and cloning technologies. You also get inline data compression, inline deduplication, and inline compaction that work together to reduce data management costs and maximize effective capacity. In addition, FabricPool automates the cost-efficient tiering of cold data to both public and private clouds.

Maximized Shared Storage Investments

ONTAP gives you the ability to save time and money by sharing the same consolidated infrastructure for workloads or tenants that have different performance, capacity, and security requirements without fear that the activity in one tenant partition will affect another. With multitenancy, a storage cluster can be subdivided into secure partitions governed by rights and permissions. And quality of service (QoS) workload management allows you to control the resources that each workload can consume, to better manage performance spikes and improve customer satisfaction. Adaptive QoS can be used to set both maximum and minimum resource levels, which is especially important for business-critical workloads, and it automatically adjusts storage resource levels to respond to changes in workloads and deliver consistent performance.

Seamless Scalability

Storage systems that run ONTAP can transparently scale from a few terabytes up to 172PB. Scale up by adding solid-state drive (SSD) and HDD capacity. Or scale out by adding additional storage controllers to seamlessly expand your cluster up to 24 nodes as your business needs grow. Rebalance capacity to improve service levels by redeploying workloads dynamically and avoiding hot spots. You also have the ability to isolate workloads and offer levels of service by using different controller technologies, storage tiers, and QoS policies.

In addition, ONTAP supports massive NAS containers that are easy to manage. With FlexGroup, a single namespace can grow to 20PB and 400 billion files while maintaining consistent high performance and resiliency.

Maximize Investment Protection

ONTAP gives you the flexibility to create an integrated, scalable storage environment by clustering storage controllers from different families—AFF and FAS—as well as from different generations. Grow with the latest hardware and continue to use your older hardware. When it is time to retire a storage system, simply upgrade the controllers and keep data in place on the existing disk shelves. Get more value from existing investments in third-party arrays by virtualizing them with NetApp FlexArray virtualization and using the storage capacity in your ONTAP environment.

Simple, Straightforward Transition to ONTAP 9

No matter what your starting point, NetApp streamlines your move to ONTAP 9. You can:

- Upgrade from ONTAP 8.3 with a simple update of your ONTAP software—no disruption and zero downtime.
- Make a smooth transition from ONTAP 7-Mode with proven tools and best practices, including 7-Mode Transition Tool (7MTT) and Copy Free Transition (CFT).
- Use straightforward import processes from third-party storage to ONTAP 9.

Consult our experts to plan and implement your transition and gain the latest ONTAP advantages from day one. You can use either NetApp Services or NetApp Certified Services Partners, do it yourself using our proven tools and processes, or choose a combination of approaches.

ONTAP Technical Highlights

The building blocks for ONTAP scale-out storage configurations are high-availability (HA) pairs in which two storage controllers are interconnected to the same set of disks. If one controller fails, the other takes over its storage and continues serving data.

With ONTAP, each storage controller is referred to as a cluster node. Nodes can be different models and sizes of AFF and FAS systems. In addition, nodes can be FAS systems running FlexArray storage virtualization, leveraging third-party and NetApp E-Series arrays as the storage capacity behind the FAS system. Disks are made into aggregates, which are groups of disks of a particular type that are composed of one or more RAID groups protected by using NetApp RAID DP® and RAID TEC technology.

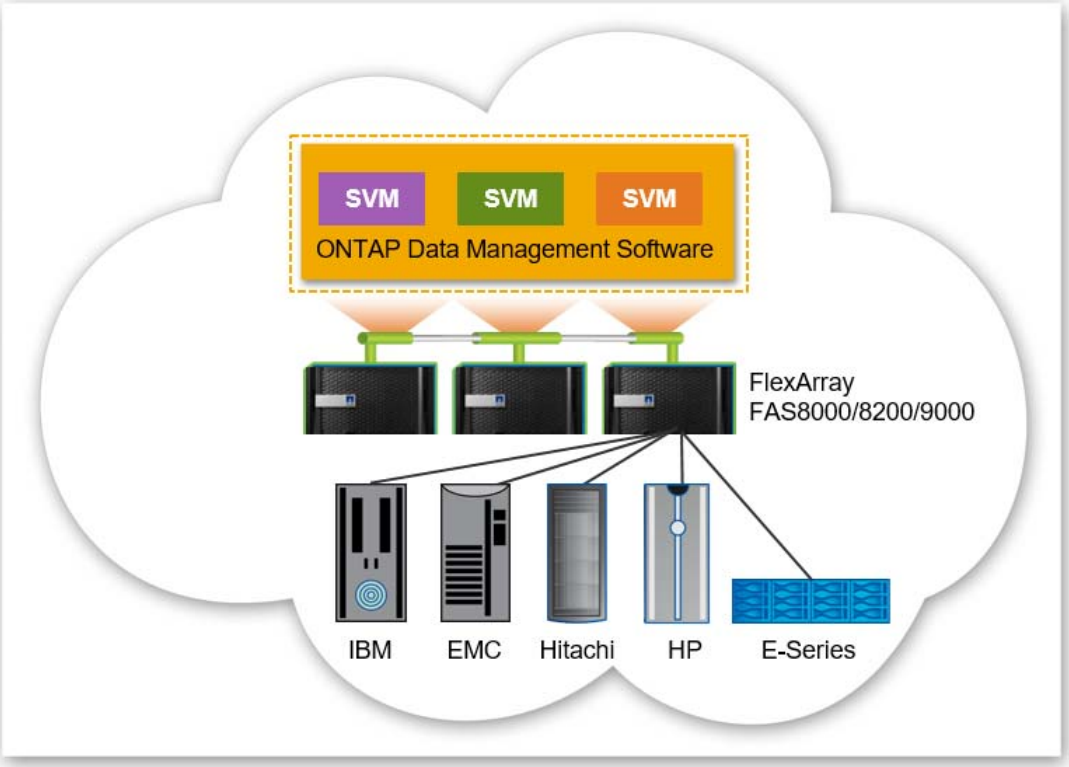

Figure 5: Investment protection – *Integrate your existing storage arrays into your private cloud with ONTAP and FlexArray.*

A key differentiator in an ONTAP environment is that numerous HA pairs are combined into a cluster to form a shared pool of physical resources that are available to applications, SAN hosts, and NAS clients. The shared pool appears as a single system image for management purposes. This means that there is a single common point of management, whether through the graphical user interface or command-line interface tools, for the entire cluster.

Although the members of each HA pair must be the same controller type, the cluster can consist of heterogeneous HA pairs of AFF all-flash arrays as well as FAS hybrid-flash arrays. Over time, as the cluster grows, and new controllers are released, it is likely to evolve into a combination of several different node types. All cluster capabilities are supported, regardless of the underlying controllers in the cluster.

To improve data access in NAS applications, NetApp virtualizes storage at the file-system level. This enables all client nodes to mount a single file system, access all stored data, and automatically accommodate physical storage changes that are fully transparent to the clients. Each client or server can access a huge pool of data residing across the ONTAP system through a single mount point.

Meet High-Availability Requirements

The proven reliability features in NetApp hardware and software result in data availability of more than 99.9999% as measured across the NetApp installed base. Backup and replication technologies integrated in the NetApp ONTAP data management software help keep your applications and data continuously available to users.

Nondisruptive Operations to Eliminate Downtime

Nondisruptive operations (NDO) are fundamental to the superior scale-out architecture of NetApp ONTAP. NDO is achieved as the storage infrastructure remains up and serving data throughout the execution of hardware and software maintenance operations as well as during other IT lifecycle operations. The goal of NDO is to eliminate downtime—whether it is preventable, planned, or unplanned—and to allow changes to your systems to occur at any time.

ONTAP allows you to transparently move data and network connections anywhere within the storage cluster. The capability to move individual data volumes or LUNs allows you to redistribute across a cluster at any time and for any reason. It’s transparent and nondisruptive to NAS and SAN hosts, and it enables the storage infrastructure to continue to serve data throughout these changes. This is helpful to rebalance capacity usage, to optimize for changing performance requirements, or to isolate one or more controllers or storage components when it becomes necessary to execute maintenance or lifecycle operations.

Table 1: Hardware and software maintenance operations can be performed nondisruptively with ONTAP.

Operation	Details
Upgrade software	Upgrade from one version of ONTAP to another
Upgrade firmware	System, disk, switch firmware upgrade
Replace failed controller or component within a controller	Network interface cards (NICs), host bus adapters (HBAs), and power supplies
Replace failed storage components	Cables, drives, shelves, and I/O modules

Table 2: Lifecycle operations can be performed nondisruptively with ONTAP.

Operation	Details
Scale storage	Add storage (shelves or controllers) to a cluster and redistribute volumes for future growth
Scale hardware	Add hardware to controllers to increase scalability, performance, or capability (HBAs, NICs, NetApp Flash Cache™ or Flash Pool™ caching)
Refresh technology	Upgrade storage shelves, storage controllers, back-end switch
Rebalance controller performance and storage utilization	Redistribute data across controllers to improve performance
Rebalance capacity	Redistribute data across controllers to account for future capacity growth
Rebalance disk performance and utilization	Redistribute data across storage tiers within a cluster to optimize disk performance

On-Demand Scalability—Expand as you Build

The ONTAP architecture is key to delivering maximum on-demand scalability for your shared IT infrastructure, offering performance, price, and capacity options.

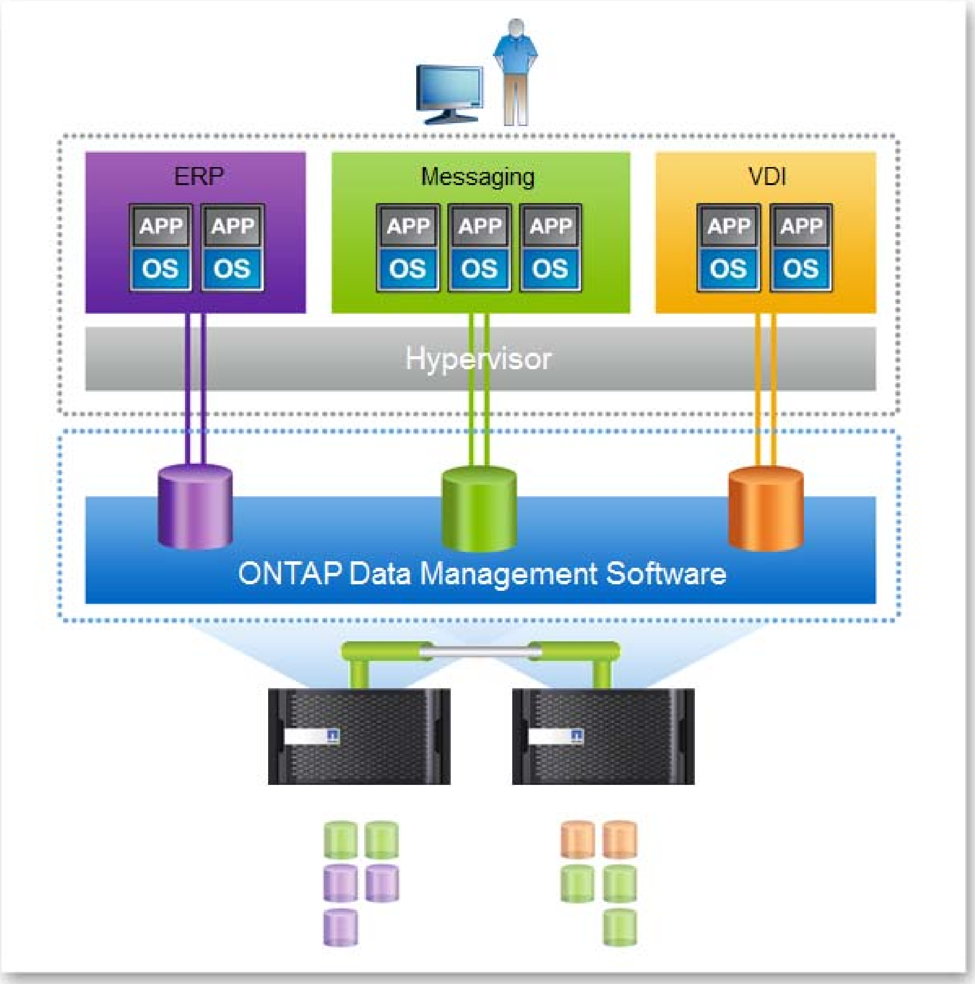

Figure 6: Expand as you build – *Start with a two-node cluster and expand controllers and capacity when you need to, nondisruptively.*

There are several approaches for leveraging flash in NetApp FAS hybrid-flash systems to accelerate workloads and reduce latency. Flash Cache can increase read performance for frequently accessed data. Plus, Flash Pool aggregates combine SSDs with traditional hard drives for delivering optimal performance and efficiency.

NetApp AFF all-flash systems offer the advantage of scalable performance with consistent low latency for SAN and NAS workloads. Customers can start with deploying AFF in an HA pair configuration to deliver enterprise-grade data management and high performance for a dedicated workload. If additional performance is required, AFF can scale out in a cluster—up to 24 nodes, delivering millions of IOPS at submillisecond latency and a total of over 88PB of SSD capacity.

The extra value of AFF shines when it is used as a high-performance node combined with hybrid-flash FAS systems in an ONTAP environment. This becomes a single storage repository for all workloads. And it enables nondisruptive movement of workloads to the node that best meets your performance and price/performance requirements at different points in time.

Multiprotocol Unified Architecture

A multiprotocol unified architecture provides the capability to support several data access protocols concurrently in the same overall storage system across a range of controller and disk storage types. ONTAP protocol support includes:

- CIFS/SMB
- NFS, pNFS
- iSCSI
- FC
- FCoE
- NVMe over FC (NVMe/FC)

Data replication and storage efficiency features in ONTAP are seamlessly supported across all protocols.

SAN Data Services

With the supported SAN protocols (FC, FCoE, iSCSI, and NVMe/FC), ONTAP provides LUN services. This is the capability to create LUNs and make them available to attached hosts. Because the cluster consists of numerous controllers, there are several logical paths to any individual LUN. A best practice is to configure at least one path per node in the cluster. Asymmetric Logical Unit Access is used on the hosts so that the optimized path to a LUN is selected and made active for data transfer. Support for multipath I/O is also available from leading OS and third-party driver vendors.

NAS Data Services

ONTAP can provide a single namespace with the supported NAS protocols such as SMB [CIFS] and NFS (NAS clients can access a very large data container by using a single NFS mount point or CIFS share). Each client, therefore, needs only to mount a single NFS file system mount point or access a single CIFS share, requiring only the standard NFS and CIFS client code for each operating system.

The namespace of ONTAP is composed of potentially thousands of volumes joined together by the cluster administrator. To the NAS clients, each volume appears as a folder or subdirectory, nested off the root of the NFS file system mount point or CIFS share. Volumes can be added at any time and are immediately available to the clients, with no remount required for visibility to the new storage.

The clients have no awareness that they are crossing volume boundaries as they move about in the file system, because the underlying structure is completely transparent.

ONTAP can be architected to provide a single namespace, yet it also supports the concept of several securely partitioned namespaces, called Storage Virtual Machines or SVMs. This accommodates the requirement for multi-tenancy or isolation of particular sets of clients or applications.

Opex and Capex Efficiency—Grow Your Business, Not IT Expense

NetApp storage solutions operating with ONTAP 9 deliver the industry's leading storage efficiency capabilities with features such as inline compression, inline deduplication, inline data compaction, thin provisioning, and thin clones. With these features coupled with space-efficient NetApp Snapshot copies, RAID DP, and RAID TEC, you can enjoy significant reductions in required disk capacity (varies by workload) when compared with traditional storage technologies.

Table 3: ONTAP 9 offers a robust set of standard and optional features.

NetApp Software and Features		
	Function	**Benefits**
Balance placement	Automates loading of new workloads onto a cluster	Increases cluster utilization and performance by adding a new workload to the optimal node
Data compaction	Packs more data into each storage block for greater data reduction	Works with compression to reduce the amount of storage that you need to purchase and operate
Data compression	Provides transparent inline and postprocess data compression for data reduction	Reduces the amount of storage that you need to purchase and maintain
Deduplication	Performs general-purpose deduplication for removal of redundant data	Reduces the amount of storage that you need to purchase and maintain
FabricPool	Automates data tiering to the cloud (public and private)	Decreases storage costs for cold data
Flash Pool™ Caching	Creates a mixed-media storage pool by using SSDs and HDDs	Increases the performance and efficiency of HDD pools with flash acceleration
FlexCache®	Caches datasets within a cluster and at remote sites	Accelerates read performance for hot datasets
FlexClone®	Instantaneously creates file, LUN, and volume clones without requiring additional storage	Saves you time in testing and development and increases your storage capacity
FlexGroup	Enables a single namespace to scale up to 20PB and 400 billion files	Supports compute-intensive workloads and data repositories that require a massive NAS container while maintaining consistent high performance and resiliency
FlexVol®	Creates flexibly sized volumes across a large pool of disks and one or more RAID groups	Enables storage systems to be used at maximum efficiency and reduces hardware investment

NetApp Software and Features		
	Function	**Benefits**
Headroom	Provides visibility of performance capacity that is available for deploying new workloads on storage nodes	Simplifies management and enables more effective provisioning of new workloads to the optimal node
MetroCluster	Combines array-based clustering with synchronous mirroring to deliver continuous availability and zero data loss; up to 300km distance between nodes	Maintains business continuity for critical enterprise applications and workloads if a data center disaster occurs
QoS (adaptive)	Simplifies setup of QoS policies and automatically adjusts storage resources to respond to work-load changes (number of TB of data, priority of the workload, etc.)	Simplifies operations and maintains consistent workload performance within your prescribed minimum and maximum IOPS boundaries
RAID-TEC™ and RAID DP® technologies	Provides triple parity or double-parity RAID 6 implementation that prevents data loss when three or two drives fail	Protect your data without the performance impact of other RAID implementations; reduce risks during long rebuilds of large-capacity HDDs
SnapCenter®	Provides host-based data management of NetApp storage for databases and business applications	Offers application-aware backup and clone management; automates error-free data restores
SnapLock	Provides WORM file-level locking	Supports regulatory compliance and organizational data retention requirements
SnapMirror	Provides integrated remote backup/recovery and disaster recovery with incremental asynchronous data replication; preserves storage efficiency savings during and after data transfer	Provides flexibility and efficiency when replicating data to support remote backup/recovery, disaster recovery, and data distribution
SnapMirror Synchronous	Delivers incremental, volume-granular, synchronous data replication; preserves storage efficiency savings during and after data transfer	Achieve zero data loss protection (RPO=0)
SnapRestore®	Rapidly restores single files, directories, or entire LUNs and volumes from any Snapshot copy backup	Instantaneously recovers files, databases, and complete volumes from your backup

NetApp Software and Features		
	Function	**Benefits**
Snapshot	Makes incremental data-in-place, point-in-time copies of a LUN or a volume with minimal performance impact	Enables you to create frequent space-efficient backups with no disruption to data traffic
Volume encryption	Provides data-at-rest encryption that is built into ONTAP	Let's you easily and efficiently protect your at-rest data by encrypting any volume on an AFF or FAS system; no special encrypting disks are required

NetApp OnCommand Data Management Software

The more information you have about your storage infrastructure, the better equipped you are to effectively manage it. NetApp OnCommand® management software can help you to improve storage and service efficiency. It offers functions that help you control, automate, and analyze your shared storage infrastructure. OnCommand tools offer simplified, effective, cost-efficient management of your shared storage infrastructure so that you can optimize utilization, meet SLAs, reduce risk, and boost performance.

Table 4: OnCommand management software product portfolio.

Product Name	Description
System Manager	Provides device-level management of NetApp storage systems. Ideal for one-off and nonrepeatable management and configuration tasks.
Unified Manager	Monitors the availability, capacity, performance, and protection of NetApp FAS and All Flash FAS resources. Unified Manager provides a single view of NetApp storage health. It also collects, retains, and analyzes NetApp storage performance statistics so users can troubleshoot and resolve issues quickly.
NetApp Service Level Manager and API Services	NetApp Service Level Manager simplifies storage consumption and enables delivery of predictable performance for your workloads. Provides a set of REST APIs to integrate with your automation and orchestrators including industry leading ITSM solutions. Also provides monitoring and conformance checking APIs.
Workflow Automation	Customized automation and delegation of repeatable storage management and storage service tasks. It facilitates your specific needs for storage infrastructure management via customized workflows based on NetApp best practices. It also integrates with orchestrators for end-to-end automated service delivery.
Cloud Manager	Deploys NetApp ONTAP® management software on AWS and Azure cloud storage in minutes. Manage and track cloud resources with ease. Cloud Manager is for ONTAP Cloud and private storage environments and supports CloudSync.

Product Name	Description
Insight	OnCommand Insight is an open data center management platform that provides operational intelligence, business insight, and IT ecosystem integration within complex enterprise IT environments.

All Flash FAS

NetApp® All Flash FAS (AFF) is an all-flash array that delivers high performance, flexibility, low latency, and superior data management without sacrificing enterprise capabilities. AFF enables a smooth transition to flash for your data center, built on NetApp ONTAP® data management software.

As businesses go through digital transformation, they must modernize their IT infrastructure to improve speed and responsiveness to support critical business operations. Although all-flash storage systems have been widely adopted to accelerate typical enterprise applications, newer workloads such as data analytics, artificial intelligence (AI), and deep learning—demand higher performance that first-generation flash systems cannot deliver.

In addition, as more organizations adopt a “cloud first” strategy, it is critical to offer enterprise-grade data management capabilities for a shared environment across on-premise data centers and the cloud. Many all-flash array solutions available today lack robust data management, integrated data protection, seamless scalability, new levels of performance, deep application, and cloud integration.

Cloud-Connected All-Flash Storage Powered by ONTAP

NetApp® All Flash FAS (AFF) is a robust scale-out platform built for virtualized environments, combining low-latency performance with comprehensive data management, built-in efficiencies, integrated data protection, multiprotocol support, and nondisruptive operations.

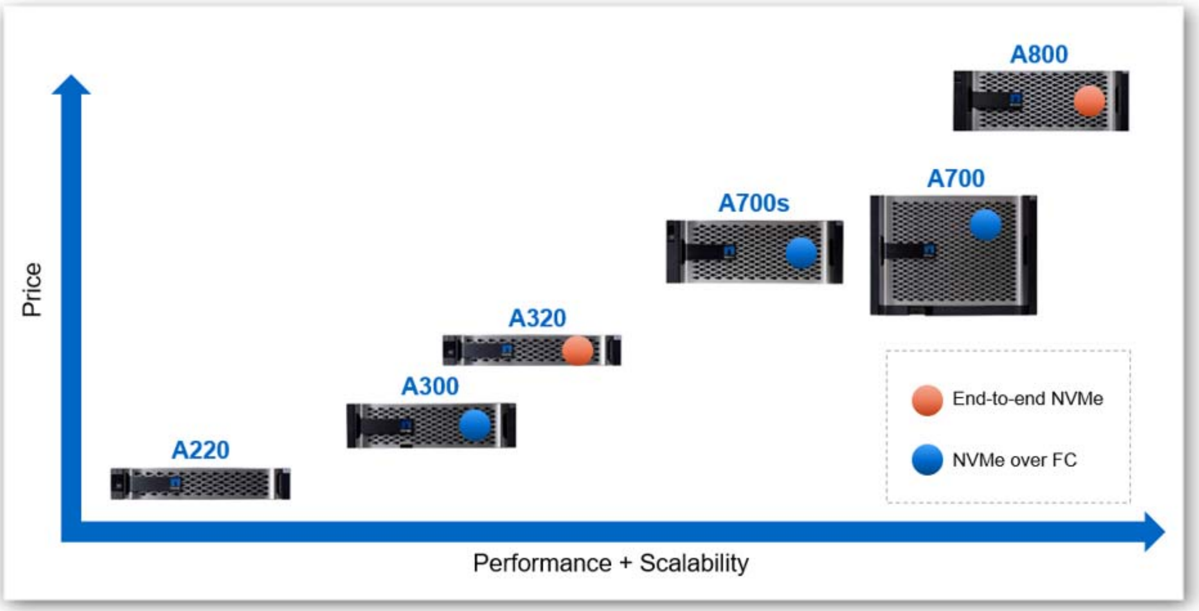

Figure 1: AFF A-Series portfolio—from enterprise to mid-size business.

“We're able to fit a whole lot more in a smaller amount of space and still provide more performance than we had before.”

— CI Engineer, financial services firm

NetApp AFF A-Series systems are designed to help businesses accelerate infrastructure transformation and fuel data-driven strategies. Powered by NetApp ONTAP® data management software, AFF systems accelerate, manage, and protect business-critical data and give you an easy and smooth transition to flash for your digital transformation in the hybrid cloud. With AFF systems, you can:

- Increase operational efficiency
- Accelerate applications and future-proof your infrastructure
- Keep business-critical data available, protected, and secure.

Increase Operational Efficiency

AFF offers the broadest application ecosystem integration for enterprise application, such as virtual desktop infrastructure (VDI), database, and server virtualization—supporting Oracle, Microsoft SQL Server, VMware, SAP, MySQL, and more. Infrastructure management tools simplify and automate common storage tasks so that you can:

- Provision and rebalance workloads by monitoring clusters and nodes
- Use one-click automation and self-service for provisioning and data protection
- Import LUNs from third-party storage arrays directly into an AFF system to seamlessly migrate data

In addition, with the NetApp Active IQ® intelligence engine you can optimize your NetApp systems with predictive analytics and proactive support tool, provide real-time insights and recommendations to prevent problems and optimize your data infrastructure.

> *"With the NetApp solution, we can slash the time needed to create an environment from 6 hours to 5 minutes regardless of scale, while provisioning additional environments simultaneously. That translates to a time savings of 70% for each product line."*
>
> *— Sandrine Kalk |Director of Global DevOps and Operations, Verint*

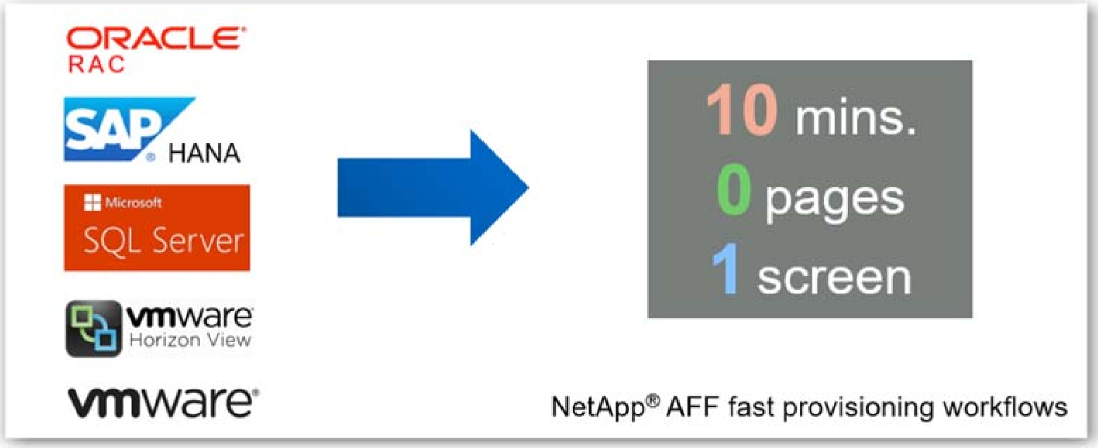

Figure 2: Application-aware data management – *Deploy key workloads in less than 10 minutes with OnCommand System Manager.*

Achieve Storage Savings, Backed by the Industry’s Most Effective Guarantee

With AFF, reduce your data center costs with the best effective capacity for any workload, backed by the industry’s most effective guarantee. We guarantee in writing:

- 3:1 guarantee across all workloads
- 4:1 for VVOL and 8:1 for VDI
- Use snapshots and get 10x higher efficiency

AFF system’s support for solid state drives (SSDs) with multistream write technology, combined with advanced SSD partitioning, provides maximum usable capacity, regardless of the type of data that you store. Thin provisioning; NetApp Snapshot™ copies; and inline data reduction features, such as deduplication, compression, and compaction, provide additional space savings—without affecting performance—so you can purchase the least amount of storage capacity possible.

Build your Hybrid Cloud with Ease

The NetApp Data Fabric helps you simplify and integrate data management across cloud and on-premises to meet business demands and gain a competitive edge. With AFF, you connect to more clouds for more data services, data tiering, caching, and disaster recovery. FabricPool gives you the ability to move data automatically between AFF and the cloud storage tiers to maximize performance and reduce overall data management cost. Simplify hybrid cloud backup and recovery with cloud-resident NetApp Data Availability Services and accelerate read performance for data that is shared throughout your organization and across hybrid cloud deployments.

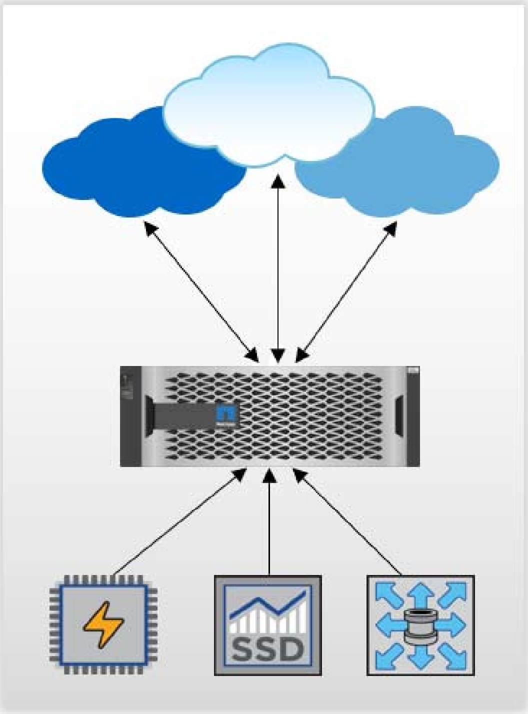

Figure 3: Future-proof your infrastructure with the most cloud-connected all-flash array – *Designed for the cloud era to connect to more clouds, in more ways, and to more services—to virtually any service provider or private cloud.*

Accelerate Applications and Future-Proof Your Infrastructure

NetApp AFF systems deliver industry-leading performance proven by SPC-1[1] and SPEC SFS industry benchmarks, making them ideal for demanding, highly transactional applications such as Oracle, Microsoft SQL Server, MongoDB databases, VDI, and server virtualization. The AFF A800 system achieved:

- 2,401,000 SPC-1 IOPS at 0.590 SPC-1 IOPS Response Time in a new SPC-1v3 result
- Lowest latency and $/GB among the top 5 results
- Predictable and consistent latency
  - ~0.4ms latency @ 80% load
  - 0.351ms SPC-1 Overall Response Time
- Highest storage capacity utilization
  - 66% versus ~30% from most others

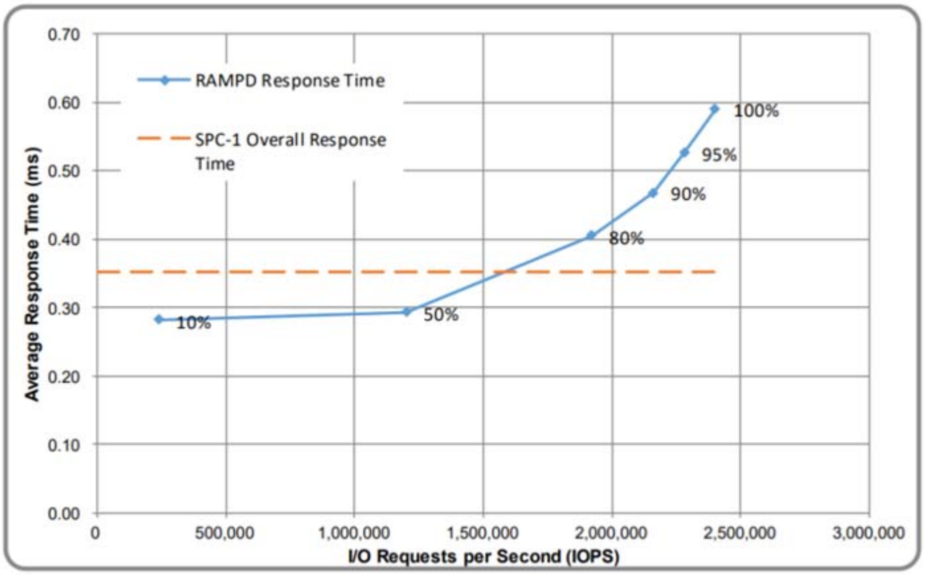

Figure 4: AFF A800 Places in the Top 4 of SPC-1v3 – *Best performance and value among major vendors who publish benchmarks.*

Accelerate Demanding Workloads

Accelerate the most demanding workloads with an AFF A800 and AFF A320 system. The AFF A800 combines NVMe SSDs and NVMe/FC connectivity to provide an ultrafast end-to-end data path to your applications. The midrange AFF A320 system supports NVMe/RoCE connectivity on the backend to the NVMe drive shelf and NVMe/FC on the front-end to the host. The AFF A320 leads the market with the best combination of NVMe-oF technologies.

Consolidate all workloads on AFF systems, which deliver up to 11.4 million IOPS at lms latency in a cluster with a truly unified scale-out architecture. You can manage a scalable NAS container of up to 20PB and 400 billion files with a single namespace by using NetApp FlexGroup volumes, while maintaining consistent high performance with adaptive quality of service (QoS) and resiliency. NetApp FlexCache® software improves the speed and productivity

[1] Link to SPC-1 report: http://spcresults.org/benchmarks/results/spc1-spc1e#A32007.

of collaboration across multiple locations and increases data throughput for read-intensive applications.

> *The NVMe-ready AFF A800s awarded the Product of the Year award for Enterprise Storage from CRN.*

Modernize with Advanced NVMe

Designed specifically for flash, the AFF A-Series all-flash systems deliver industry-leading performance, capacity density, scalability, security, and network connectivity in dense form factors. AFF A-Series systems support NVMe/FC host connectivity, so you can gain twice the IOPS and cut application response time in half compared with traditional FC. These systems support a range of ecosystems, including VMware, Microsoft Windows 10, and Linux, with storage path failover. For most customers, integrating NVMe/FC into an existing SAN is a simple, nondisruptive software upgrade.

In addition, integrate new technologies and private or public cloud into your infrastructure nondisruptively. AFF is the only all-flash array where you can combine different controllers, SSD sizes, and new technologies—protecting your investment.

Keep Important Data Available, Protected, and Secure

Support backup and disaster recovery needs through our complete suite of integrated data protection and replication features. NetApp Integrated Data Protection technologies protect data and accelerate recovery; for easier management they integrate with leading backup applications. Benefit from features and capabilities such as NetApp Snapshot™ copies, cloning, encryption, and both synchronous and asynchronous replication for backup and disaster recovery. Key capabilities and benefits include:

- Reduced data management costs with native space efficiency with cloning and NetApp Snapshot copies. Up to 1,023 copies are supported.
- Unified, scalable platform and plug-in suite for application-consistent data protection and clone management with NetApp SnapCenter®.
- Reduced overall system costs with NetApp SnapMirror® replication software, which replicates to any type of FAS/AFF system: all-flash, hybrid, or HDD, on the premises or in the cloud.
- Synchronous replication with NetApp MetroCluster™ software, a capability in the all-flash-array market that delivers zero RPO and low to zero RTO for mission-critical workloads.
- Regulatory compliance with NetApp SnapLock® technology, which is enabled with Integrated Data Protection and storage efficiency.

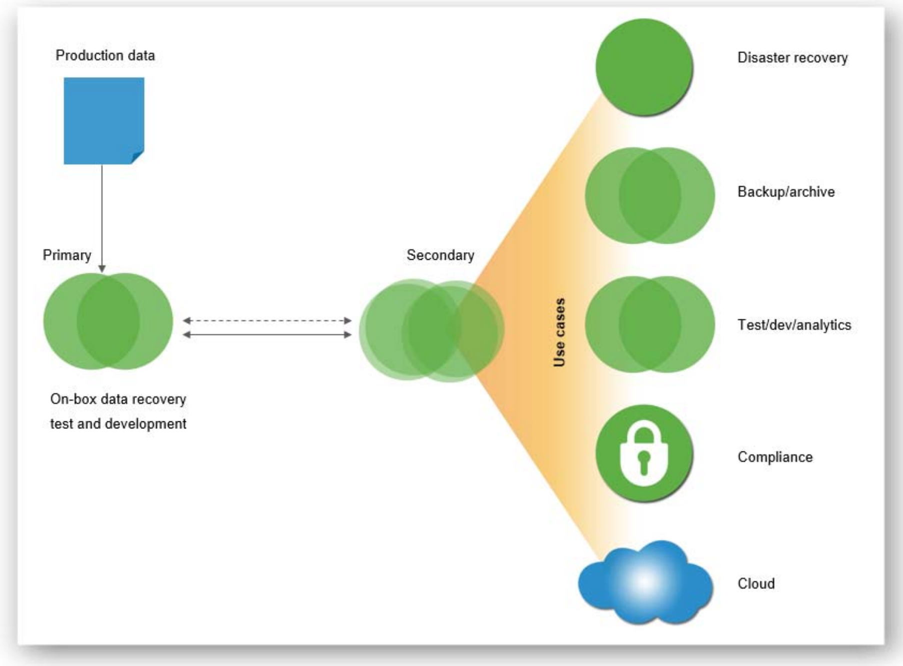

Figure 5: NetApp integrated data protection – *Offers one data management flexible platform that provides data availability to keep applications running, mitigate risk, control costs, and improve data protection processes.*

In addition, flexible encryption and key management help guard sensitive data on the premises, in the cloud, and in transit. With the simple and efficient security solutions, you can:

- Achieve FIPS 140-2 compliance (Level 1 and Level 2) with self-encrypting drives and use any type of drives with software-based encryption.
- Meet governance, risk, and compliance requirements with security features such as secure purge; logging and auditing monitors; and write once, read many (WORM) file locking.
- Protect against threats with multifactor authentication, role-based access control, secure multitenancy, and storage-level file security.

> *“NetApp's multiprotocol capability was a major draw for our colleges. With NetApp, we can enable our colleges to retain their skillsets. They don't have to learn something new or put in a mix of products just to accommodate their protocols.”*
>
> *— Daniel Black |Director of Engineering, Technical College System of Georgia*

Future-Proof Your Investment with Maximum Flexibility

NetApp solutions establish a seamless, well-integrated hybrid cloud architecture or Data Fabric that easily ties together private cloud, service providers, and hyperscale cloud providers along with their data management environments. This Data Fabric gives you the ability to implement the hybrid cloud on its own terms. Move data and applications to an AFF system, on commodity hardware with software-defined storage, or in the cloud. The Data Fabric offers a broad set of application ecosystem integration for database, VDI, and server virtualization.

With AFF, which is Data Fabric ready, your investment is protected as performance and capacity needs change or your cloud strategy evolves:

- AFF systems eliminate performance silos. Seamless integration with hybrid FAS systems means that workloads can transparently move between high-performance tiers and low-cost capacity tiers.
- Seamlessly adapt to changing needs with the only all-flash array that offers the ability to intermix different controllers, SSD sizes, and next-generation technologies.
- AFF is data fabric ready, with proven cloud connectivity. FabricPool enables you to move data automatically between AFF and the cloud storage tiers to maximize performance and reduce overall data management cost.
- Optimize data management for enterprise workload environments with leading application integration with Oracle, Microsoft, VMware, SAP, OpenStack, and many more.

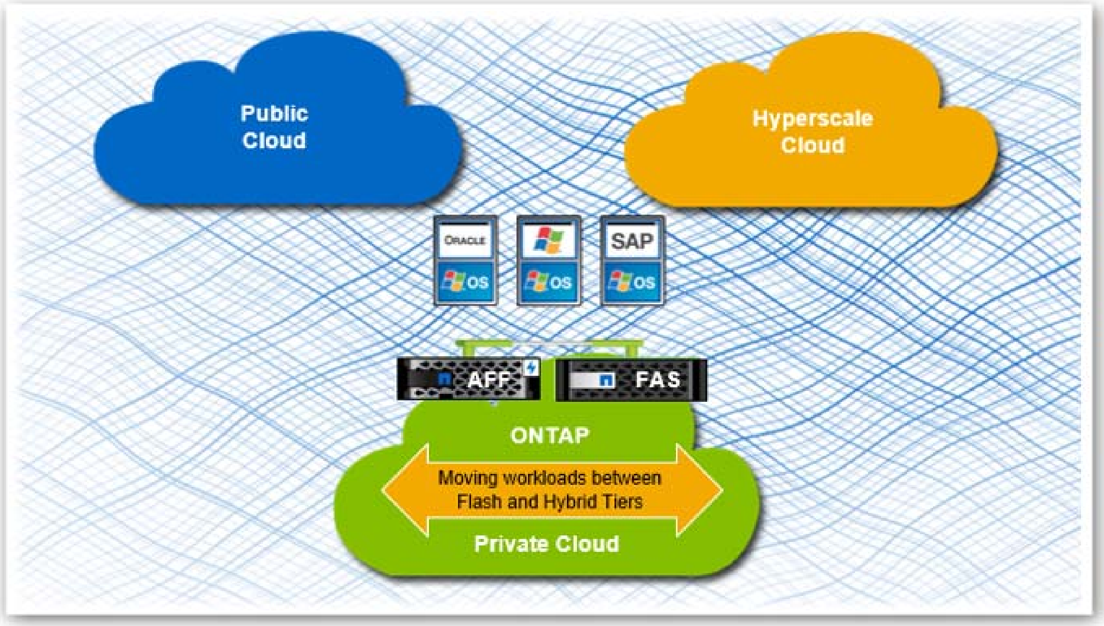

Figure 6: AFF is Data Fabric ready—moving data between tiers and different clouds.

> *"With NetApp All Flash FAS, we can improve the quality of healthcare in our own hospitals and others throughout the region by offering high-performing electronic patient records and virtual desktops to healthcare providers."*
>
> *— Reinoud Reynders, IT Manager, Infrastructure and Operations at UZ Leuven*

All-Flash Performance Powered by End-to-End NVMe Technology

AFF systems are excellent for performance-demanding applications and mixed-workload environments that consist of, for example, Oracle, Microsoft SQL Server, MongoDB databases, VDI, and server virtualization. With NVMe-based AFF A800, AFF is also a great choice for AI and deep-learning environments:

- Combined with ONTAP cloud integration and software-defined capabilities, AFF enables the full range of the data pipeline that spans the edge, the core, and the cloud for AI and deep learning, leveraging the same ONTAP data management.
- The end-to-end NVMe-based AFF A800 delivers 1.3 million IOPS at below 500µs latency.
- Built-in adaptive QoS safeguards SLAs in multiworkload and multitenant environments. It optimizes performance control dynamically with superior scalability of up to 40,000 workloads per cluster at LUN, file, and VVol levels.
- With the latest ONTAP release, AFF delivers up to 90% performance increase for Microsoft SQL Server with multichannel SMB.

Storage Efficiency Technologies

NetApp is known for its superior storage efficiency technologies, such as inline deduplication, inline compression, thin provisioning, and space-efficient Snapshot copies. These technologies apply to AFF systems and further reduce your total cost of ownership by lowering cost per effective gigabyte of storage:

- Performance-efficient inline data reduction technologies provide an average of 5 to 10 times space savings for a typical use case.
- Space-saving inline data compaction technology places multiple logical data blocks from the same volume into a single 4KB block. Space savings as high as 67:1 from this feature have been observed when using inline data compaction and inline compression with an Oracle database.
- There is a near-zero performance impact with inline compression. Incompressible data detection eliminates wasted cycles.
- You can increase space savings by eliminating redundant blocks using inline deduplication—effective for operations such as VDI OS patches in which this deduplication can achieve 70:1 reduction rates.
- As the first all-flash array to support SSDs with MSW technology, and combined with advanced SSD partitioning in ONTAP, AFF further increases usable capacity by up to 42%.

NetApp OnCommand Simplifies Management

NetApp OnCommand® management software provides automated tools to further simplify management of storage operations:

- Set up and configure AFF quick and easy with preconfigured systems for SAN and NAS deployments. It takes less than 10 minutes with OnCommand System Manager.
- OnCommand Workflow Automation automates common storage tasks such as provisioning and data protection. It provides fast one-click automation and self-service.
- To optimize storage for peak performance and to keep everything running smoothly, OnCommand Performance Manager provisions and rebalances workloads by monitoring clusters and nodes to assure performance headroom.
- Import LUNs from storage arrays that are not based on ONTAP software directly into an AFF system to seamlessly migrate data from older storage arrays.

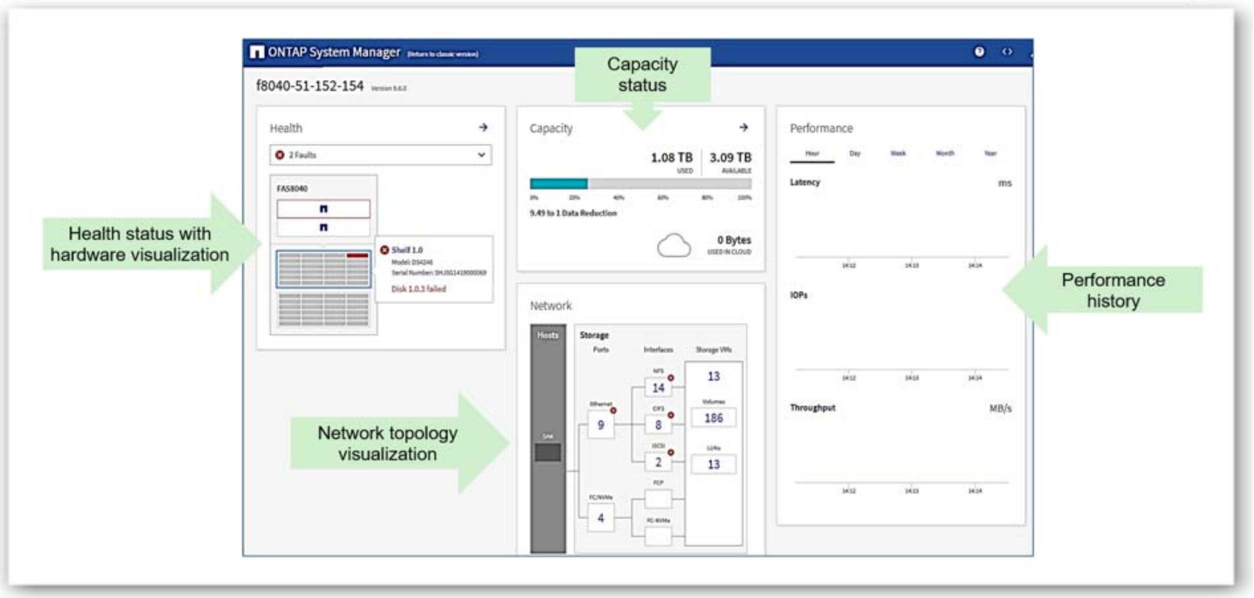

Figure 7: **Intuitive ONTAP System Manager** *–Based on REST APIs, the new System Manager dashboard is more intuitive and displays richer information in a more actionable view.*

Get More Business Value with Services

To help you fully realize the benefits of NetApp solutions, NetApp Services and our NetApp certified services partners will collaborate with you through a full portfolio of services that covers the company’s IT lifecycle. NetApp offers:

- Assessment services to evaluate the performance and efficiency of workloads across heterogeneous environments
- Advisory services to determine the best workload candidates to move to flash
- Deploy and optimize services to prepare your environment and deliver continuous operation of AFF systems
- Managed upgrade services to secure your storage environment and to protect your investment by ensuring your ONTAP software is the most current version.

NetApp Support offerings, such as the NetApp Active IQ® cloud-based predictive cloud-based analytics and proactive support tool, provide real-time insights and recommendations to prevent problems and optimize your data infrastructure. Learn more at netapp.com/services.

AFF A-Series Systems

NetApp AFF systems help you meet your enterprise storage requirements with the following AFF A-Series Systems:

AFF A800

The AFF A800 is designed for the most demanding workloads requiring ultra-low latency and is the first flash array on the market to support NVMe SSDs and NVMe over Fabrics (NVMe-oF). It provides end-to-end NVMe connectivity between storage arrays and host servers for maximum bandwidth, high IOPS, and the lowest possible latency. Each 4U chassis accommodates dual controllers for high availability (HA) and includes 48 slots for NVMe SSDs. In addition to 32Gb and 16Gb FC, network options include the storage industry's first 100GbE connectivity, as well as 40GbE and 10GbE. An NVMe-powered SAN scale-out cluster supports up to 12 nodes (6 HA pairs) with 1,440 drives and nearly 160PB of effective capacity. NAS scale-out clusters support up to 24 nodes (12 HA pairs). The AFF A800 future-proofs your data infrastructure with NetApp ONTAP 9 the industry's leading data management software.

> *"NetApp once again hits it out of the park with the enterprise focused A800. The performance profile is very strong, taking its position at the top of the ONTAP family."*
>
> *— StorageReview Editors' Choice, May 2019*

AFF A700

The AFF A700 is a high-end NetApp storage controller designed for performance-driven workloads and data centers requiring a modular design. The AFF A700 can dramatically enhance performance and high-performance I/O density in a new 8U HA form factor and it includes options for 40GbE and 32Gb FC along with the latest in SAS connectivity, the SAS 3.0 standard with 12Gb speeds. This controller also provides the most versatile I/O interface available, the UTA2 connections that support 10GbE and 16Gb FC and that can be easily changed between these two protocols in the field. AFF A700 controllers support up to 12 nodes for SAN deployments and up to 24 nodes in NAS deployments.

AFF A700s

The AFF A700s is an integrated high-end all-flash array and best for performance-driven workloads and data centers requiring a small footprint. The AFF A700s comes in a compact form factor with dual controllers and 24 internal SSDs in a single 4U chassis. A700s provides data center efficiencies and excellent performance with reduced power and cooling. AFF A700s performance is comparable to that of AFF A700; however, they offer different connectivity and capacity options to address different solutions and customer requirements.

AFF A320

The new AFF A320 midrange end-to-end NVMe NetApp AFF storage controller is a modern NVMe Flash storage system. It provides application performance improvements with lower low latencies compared to the AFF A300. For enterprise applications that require the best performance at value, the AFF A320 includes dense 2U form factor with two HA controllers, extreme bandwidth with 16 onboard 100GbE ports and four expansion slots in an HA pair, adapter support includes100GbE, 32Gb FC, 25GbE, and 10GbE support, NVDIMMs for persistent write cache of data received but not yet committed to flash media, and host-side NVMe/FC support for low-latency, high-performance remote direct memory access (RDMA) connectivity to the NVMe SSDs.

AFF A300

The A300 firmly targets enterprise applications that require best balance of performance and cost. It is more powerful than the AFA A220 for users that need additional capacity and performance The AFF A300 is easy to set up and runs the latest version of ONTAP and supports SSDs up to 30TB. It requires just 12 SSDs to start but scales to over 140PB raw

(560PB effective) in NAS config and 70PB raw (280PB effective) as SAN. The A300 supports 10GbE, 40GbE as well as Fibre Channel up to 32Gb and NVMe/FC with the 32Gb FC adapter.

> *The midrange AFF A300 recently won the Editor’s Choice Award from StorageReview, which bestows this award for “performance in excess of competitive offerings, a feature set that is innovative and sets a new bar for competitive offerings or for defining a new category or space within enterprise IT”. Through Storage Review’s independent testing with Oracle, SQL, VDI workloads, AFF A300 stands out with its impressive performance and feature set.*
>
> *— StorageReview Editors' Choice, November 2018*

AFF A220

The AFF A220 is the entry system to the NetApp all-flash array series. It is ideal for mid-size business and small enterprises that require simplicity and best value. With the AFF A220 you can accelerate business insights and demanding workloads. This 2U array enables enhanced storage efficiency based on the types of workloads. With a potential maximum raw capacity of up to 48.3 PB and maximum memory of 768 GB, NetApp ensures the effectiveness of its inline data reduction technologies, including compression, deduplication and data compaction. It offers 4x 10 GbE cluster interconnect channels for distribution of the processing across an array of nodes in the clusters, and high-data rate and low-latency communication between node processes.

> *“In addition to accelerating every application without disruption, the NetApp AFF A200 dramatically improves data center economics and enables data-driven enterprises to modernize their infrastructures with confidence. Editor’s Choice award for the NetApp AFF A200 for phenomenal performance at sub-millisecond latencies.”*
>
> *— StorageReview Editors' Choice, November 2017*

Table 1: All Flash FAS technical specifications.

AFF Technical Specifications						
	AFF A800	**AFF A700s**	**AFF A700**	**AFF A320**	**AFF A300**	**AFF A220**
Maximum scale-out	2–24 nodes (12 HA pairs)					
Maximum SSD	2,880	2,529	5,760	576	4,608	1,728
Max effective capacity[2]	316.3PB	316.3PB	702.7PB	35PB	562.2PB	193.3PB
Per-System Specifications (Active-Active Dual Controller)						
Controller form factor	4U with 48 SSD slots	4U with 48 SSD slots	8U	2U	3U	2U with two 24 SSD slots

[2] Effective capacity is based on 5:1 storage efficiency ratios with the maximum number of SSDs installed. The actual ratio can be higher depending on workloads and use cases.

Table 2: AFF A-Series software.

AFF A-Series Software	
Features and software Included with ONTAP software	**Efficiency**: NetApp FlexVol®, inline deduplication, inline compression, inline compaction, and thin provisioning **Availability**: Multipath I/O and active-active HA pair **Data protection**: NetApp RAID DP®, NetApp RAID TEC®, and Snapshot technology **Whole cluster synchronous replication**: MetroCluster **Performance control:** Adaptive QoS and balanced replacement **Management**: OnCommand Workflow Automation, ONTAP System Manager, and Active IQ Unified Manager (formerly OnCommand Unified Manager) **Scalable NAS container**: NetApp ONTAP FlexGroup **Storage protocols supported**: NVMe/FC, FC, FCoE, iSCSI, NFS, pNFS, and SMB
Flash bundle	• NetApp **SnapRestore®** software: Restore entire Snapshot copies in seconds • NetApp **SnapMirror** software: Simple, flexible backup and replication for disaster recovery • NetApp **FlexClone®** technology: Instant virtual copies of files, LUNs, and volumes • NetApp **SnapCenter®**: Unified, scalable platform and plug-in suite for application-consistent data protection and clone management • NetApp **SnapManager** software: Application-consistent backup/recovery for enterprise applications Go to NetApp.com for information on additional software available from NetApp.
Extended-value software (optional)	**NetApp OnCommand Insight**: Flexible, efficient resource management for heterogeneous environments **NetApp SnapLock**: Compliance software for write once, read many (WORM) protected data **NetApp Volume Encryption (free license)**: Granular, volume-level, data-at-rest encryption **NetApp FabricPool feature**: Automatic data tiering to the cloud **SnapMirror Synchronous**: Synchronous data replication with zero recovery point objective **NetApp Data Availability Services**: Cloud native backup solution for NetApp ONTAP storage **NetApp FlexCache**: Acceleration for data access for single or multisite deployment

FlexPod Converged Infrastructure Platform

FlexPod converged infrastructure platform integrates NetApp storage systems, Cisco UCS servers, and Cisco Nexus fabric switches into a validated enterprise-class IT platform. It delivers cloud services faster and with minimal business disruptions. Through the ability to run multiple enterprise workloads on a single, flexible architecture, you can speed the deployment of cloud-based data center infrastructures and business-critical applications while reducing costs, complexity, and risk.

Supporting business objectives from an IT perspective is getting more complex. In the past, data was hidden in on-premises locations, organized and collected in a steady and predictable manner. Data is no longer locked away and hidden behind firewalls but distributed—organizations can store data in multiple on-premises and cloud locations and move it between. Data needs to be delivered to the places where it can be most effective. The velocity of information from social media streams, customer attributes, and demographics is constantly fluctuating. And data is diverse—acquired from multiple sources, suppliers, and the Internet of Things (IoT). However, to digitally transform, you need to blend structured, unstructured, and streaming data with machine-learning data to gain holistic views, create revenue streams, and optimize operations.

IT is facing a transformation in data management and needs to adapt with simpler approaches to deploy and manage its infrastructure to enable greater agility by optimizing costs and existing staff resources, while also maintaining data security and adherence to compliance requirements. In anticipation of growing businesses' requests such as expansion to new markets and support for mobile users—as well as opportunities in analytics and hybrid cloud, IT also needs to modernize the data center to become an agile, innovative service provider.

FlexPod Simplifies IT and Accelerates Applications

Built on NetApp® and Cisco technology, the FlexPod® converged infrastructure platform meets and exceeds these IT challenges. A leader in converged infrastructure solutions, FlexPod is trusted by thousands of customers worldwide. Composed of prevalidated storage, networking, and server technologies, FlexPod increases IT responsiveness to organizational requirements and reduces the cost of computing with maximum uptime and minimal risk.

FlexPod integrates components that can scale, compute, and store independently to meet changing, granular business requirements. It provides a proven solution that is agile and flexible for a broad set of applications and demands. The prevalidated FlexPod architecture delivers application performance, and agility that drive higher productivity, faster decision making, and greater opportunities for growth.

> *"Among the biggest benefits of FlexPod are integrated components that help enable us to centrally manage all our data center requirements."*
>
> *— Darrell Williams, Director of Information Systems, Katz, Sapper & Miller*

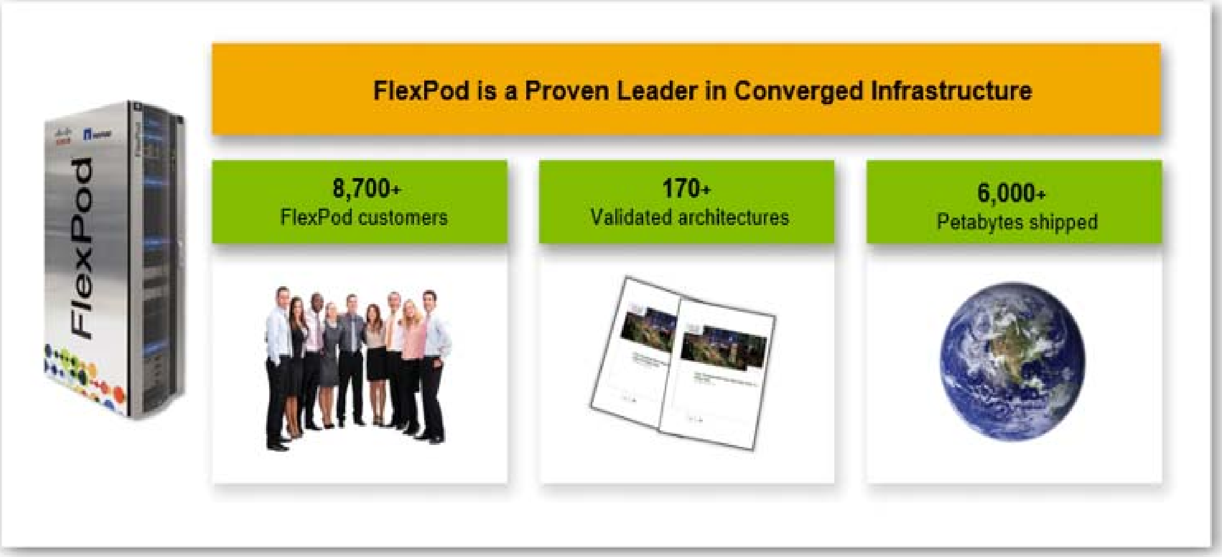

FlexPod is a Proven Leader in Converged Infrastructure

8,700+
FlexPod customers

170+
Validated architectures

6,000+
Petabytes shipped

Figure 1: FlexPod is leading the integrated infrastructure field.

> *According to the IDC Worldwide Quarterly Converged Systems Tracker, FlexPod is #1 in Capacity Share and #2 in Revenue Share.*[1]

WV OASIS can accelerate any application deployment project or facilitate a transition to the cloud with FlexPod. This converged infrastructure solution integrates compute, storage, and network components into a single architecture that scales to fit a variety of virtualized and nonvirtualized environments. Using Cisco Validated Designs (CVDs) and NetApp Validated Designs (NVAs) accelerates predeployment planning, reduces risk, and provides a platform for effective shared virtualization and private cloud.

The FlexPod architecture delivers three core benefits:

- Application optimization
- Hybrid cloud flexibility
- Simplified IT infrastructure

Application Optimization

FlexPod is tested, prevalidated and preconfigured to run hundreds of applications and workloads with reliable performance that scales along with business demands. Host multiple instances of mixed applications, consolidated on a shared infrastructure with centralized, simplified management. With All Flash FAS (AFF) on FlexPod, you speed up application performance by up to 20x. As your workloads increase, FlexPod can be modified to add storage, networking, and compute layers without losing the design validation. In addition, FlexPod can help customers achieve software-defined, policy-based automated deployment of enterprise applications with Cisco Application-Centric Infrastructure (ACI). FlexPod gives you the ability to:

[1] IDC, Worldwide Quarterly Converged Systems Tracker - Q3 2018, December 18, 2018 (Fastest growing Integrated Infrastructure & Certified Reference Systems Year over Year revenue growth).

- Decrease deployment risk and time to solution with fully tested and validated designs
- Focus IT resources on high-value activities by using pre-validated converged infrastructure
- Improve operational efficiency through simplified management and automation

A single All Flash FAS system delivered 239,000 submillisecond IOPS on an Oracle RAC validation. Another test showed support for 2,500 virtual desktops with near millisecond average response times for read and write operations. For AI and DL Scenarios on FlexPod AI, you can achieve over a million IOPs with submillisecond latency for the highest Machine Learning demands.

For more information on configurations and test details, see ESG Lab Report: "FlexPod Datacenter with NetApp All-Flash FAS, Converged Infrastructure for High-performance Latency-sensitive Databases and Virtual Desktop Solutions."

Hybrid Cloud Flexibility

IT administrators work with multiple clouds. Rapid shifts in technology are placing new and evolving applications in a combination of public, private, and hybrid clouds. As more of these applications become business critical, it is crucial to deliver trusted platforms with resilience and scale to meet increasingly stringent performance and availability requirements. IT administrators also need to consider new ways to address data control, elastic capacity, and location flexibility.

FlexPod makes working in the cloud easier with innovative technologies by NetApp and Cisco. FlexPod is ready for the cloud, regardless of where you are in your cloud strategy. Operate across hybrid cloud resources with the software-defined capabilities of NetApp Data Fabric, while maintaining security, control, and workload portability with Cisco CloudCenter. The Data Fabric simplifies data management by seamlessly integrating enterprise IT, private data management, and clouds of all shapes and sizes to accelerate digital transformation. NetApp has achieved this capability by bringing industry-leading products to market and enabling them for managing data across multiple environments and providers. You can:

- Manage data from flash to disk to cloud with the simplicity of a single set of tools
- Optionally use OpenStack software on FlexPod to create a private or hybrid cloud

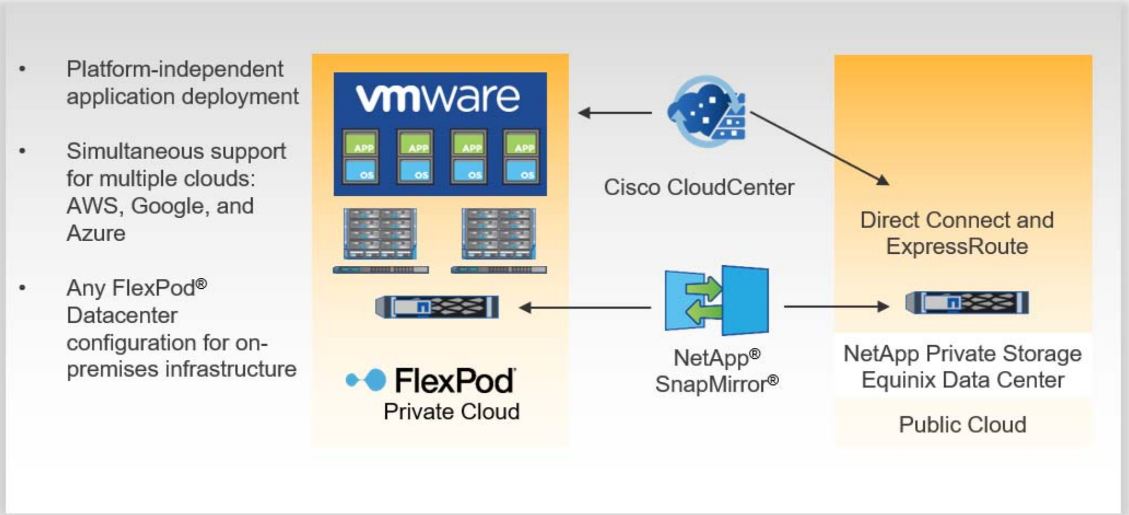
• Platform-independent application deployment
• Simultaneous support for multiple clouds: AWS, Google, and Azure
• Any FlexPod® Datacenter configuration for on-premises infrastructure
vmware
APP
OS
FlexPod
Private Cloud
Cisco CloudCenter
NetApp® SnapMirror®
Direct Connect and ExpressRoute
NetApp Private Storage Equinix Data Center
Public Cloud

Figure 2: Hybrid cloud agility with FlexPod.

> *The FlexPod Datacenter solution from NetApp and Cisco continues as the market leader in converged infrastructure through numerous innovations like cloud agility.*

Simplify IT Infrastructure

FlexPod optimizes the data center for the most demanding applications so you can increase productivity and continue innovation across your enterprise. With a single view and centralized automation and orchestration, Cisco Unified Computing System (UCS) Director provides end-to-end management of your FlexPod, freeing up valuable IT time. The validated designs help you deploy FlexPod platforms in a wide range of operating environments with less risk and accelerated ROI. Our cooperative or solutions support models are designed to simplify and streamline support for your FlexPod converged infrastructure.

> *“Our executives loved the simplicity and power of the integrated stack in FlexPod. And for IT, the prevalidated architecture with prescriptive sizing and design guides reduced our risk.”*
>
> *— Wojciech Biernacki, IT Systems Administrator, University of Tennessee*

Choose Your FlexPod Configuration

FlexPod has a range of configuration options that are designed to meet your specific capacity and performance requirements:

- For large enterprises and cloud service providers that have mature IT processes and rapid growth expectations and want to deploy a highly scalable shared infrastructure for multiple critical applications
- For midsized organizations and branch departments, as a cost-effective starting point for infrastructure consolidation and virtualization solutions
- For workloads that require high-performance computing or very large data capacity environments such as databases, artificial intelligence (AI), machine learning (ML), big data analytics, and dedicated application optimization
- For private or hybrid cloud deployments that need a consistent set of data management tools for edge, private, and public clouds

All FlexPod solutions can be scaled up or out and duplicated in a modular fashion to accommodate your future growth. They can also scale to a larger FlexPod configuration with a clearly defined upgrade path that leverages all existing components and management processes.

> *Our efforts were recently recognized by TechTarget, which made FlexPod the winner of the 2017 TechTarget Impact Award for Best Converged/Hyper-Converged Infrastructure.*

Build Your Own FlexPod

Configure and build your own FlexPod. You can plan the power, floor space, usable capacity, performance, and cost of each FlexPod deployment. Go to Buildaflexpod.com and choose your network, compute, and storage gear. See what your FlexPod will look like as each component is added. Then, save and share your design.

Supports a Broad Range of Environments

FlexPod continues to add business-critical workload support, expanding the broad range of validations available. FlexPod supports the broadest range of hypervisors, operating systems, applications, and infrastructure software. Through extensive lab validation, NetApp and Cisco ensure that these applications are supported.

Table 1: Applications supported by NetApp and Cisco.

NetApp and Cisco Supported Applications	
SAP and SAP Hana VMware vSphere®	Microsoft Private Cloud Docker & Container Solutions
VMware View	Cloudera's Distribution, including Apache Hadoop
Citrix XenDesktop, VMWare Horizon Suit	NetApp SnapProtect® technology
Red Hat Enterprise Linux	Cisco Nexus data center switches
Red Hat Enterprise Linux OpenStack Platform	Hortonworks Data Platform
AI / ML Frameworks with Nvidia V100 Analytic GPU (Graphical Processing Unit)	Oracle (RAC, JD Edwards, Oracle Linux, Oracle VM Server)
Microsoft Exchange, SQL Server, and SharePoint	
Epic and MEDITECH healthcare solutions	

FlexPod maintains a consistent advantage because all designs share the same storage vendor, compute, network architecture, and support, as well as a common flexible approach to a validated design.

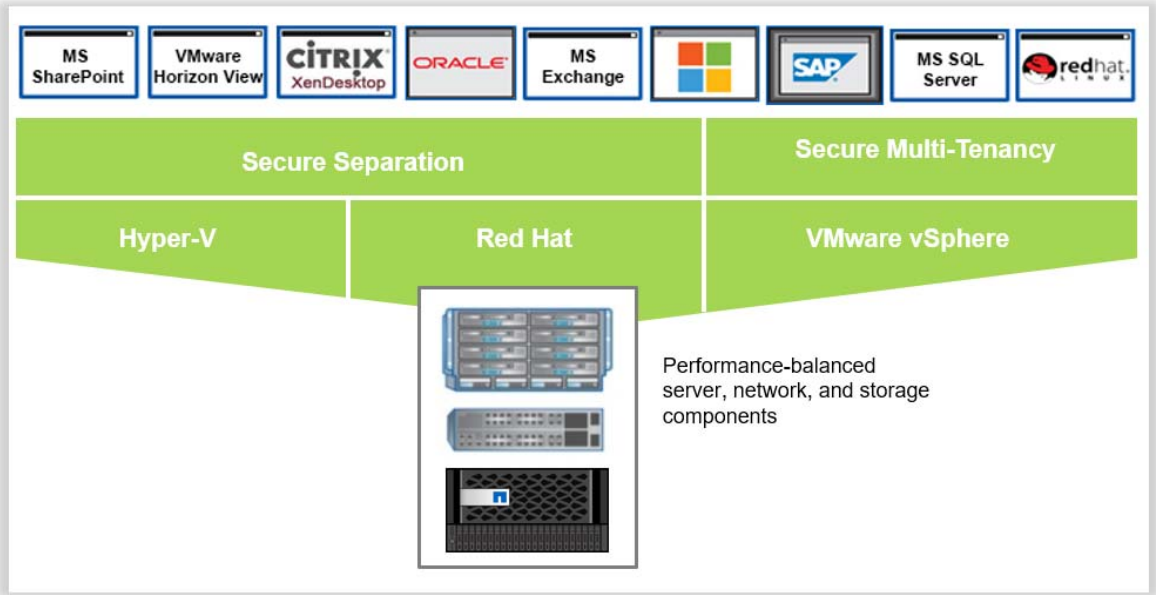

Figure 3: FlexPod validated application stacks – *Expanding business-critical workload support.*

Jointly Developed Reference Architectures

FlexPod is a flexible platform that is designed to handle mixed workloads in many environments. Its reference architectures can be sized for small or large enterprises, and the architecture remains consistent with the same capabilities, technology, and management tools. Scale from a small environment to a very large environment without having to make major upgrades in technology or change management tools or methodologies. To help you integrate and flex the solution to best meet specific requirements, NetApp and Cisco have developed reference architectures for critical environments.

Table 2: Reference architectures developed NetApp and Cisco.

Reference Architectures	
Workload Consolidation	With FlexPod you can consolidate and virtualize your business applications onto less hardware. Along with improved hardware utilization, this approach frees up data center space and reduces power and cooling requirements, enabling you to lower your infrastructure costs by up to 50%.
Virtual Desktop Infrastructure	FlexPod is a self-contained virtual desktop solution in a rack. Its modular design facilitates rapid, repeatable deployment of thousands of virtual desktops. It is optimized for VMware View and Citrix Virtual Apps and Desktops. You can gain efficiencies by deduplicating up to 90% of redundant user and OS data. I/O performance can be accelerated by up to 50% with the NetApp Virtual Storage Tier.
Development and Test	FlexPod enables rapid provisioning and deprovisioning of virtual resources, making it ideal for development and test environments. NetApp FlexClone® software facilitates rapid development/test setup with cloning technology that gives you the ability to deploy thousands of space-efficient VMs for new projects in minutes, accelerating time to market. Clones can also be redeployed to secondary sites, reducing preparation time for initiatives such as disaster recovery testing.
Business and Disaster Recovery	FlexPod can be configured with integrated data protection software to provide fast recovery from system, site, and regional outages for business continuity. The combination of NetApp MetroCluster™ and SnapMirror® technologies with Cisco UCS Manager offers automated monitoring and failover, as well as cost-effective replication to a secondary site for continuous protection against unplanned downtime. Move virtual server and storage resources and data nondisruptively across hardware to eliminate planned downtime.
Secure Multi-tenancy and Secure Separation	FlexPod leverages Cisco and NetApp technologies to deliver secure multi-tenancy with solutions such as Cisco Secure Enclaves. Resources and data for each tenant—application, business unit, or customer—are securely isolated within the FlexPod environment. This combines the data separation and service-level guarantees offered by application silos with the efficiencies of a converged, virtualized infrastructure.

> *"One of the reasons we are so excited about the FlexPod solution is that it offers us an extreme amount of flexibility in terms of workloads that we run to support our colleges. "It also provides each and every one of our colleges with the flexibility to scale their solution up or down to meet their unique needs."*
>
> *— Daniel Black, Director of Engineering, Technical College System of Georgia's System Office*

Enhanced Data Center Efficiency

FlexPod components are integrated in a standardized configuration that scales from entry-level designs to high-performance large database workloads. This integrated approach can significantly reduce capital and operating expenses through end-to-end virtualization and higher efficiencies at each layer.

Cisco Unified Computing System

Cisco UCS is a data center platform that eliminates time-consuming manual configuration, reduces TCO, and increases business agility. It combines compute and network resources, storage access, and virtualization into a scalable, modular system that is easily managed as a single entity by Cisco UCS Manager. Service profile templates enable automatic, policy-based hardware configuration and deployment for large, stateless computing environments.

Cisco Nexus Data Center Switches

Cisco Nexus switches use award- winning unified fabric technology to identify and consolidate all network traffic onto a single simplified, cost- effective architecture based on Fibre Channel over Ethernet. The switches offer "zero-touch" installation, automatic configuration, enterprise-class scalability, and nondisruptive in-service upgrades. A single point of policy management also increases efficiency, availability, and security.

The option of Cisco Nexus 7000 Series switches provides greater networking scale, throughput, availability, and advanced features for data center interconnect requirements. Cisco Nexus 9000 switches lay the foundation for software-defined innovations such as Cisco Application Centric Infrastructure, allowing intelligent software to automate hardware resources across next-generation data centers.

NetApp FAS Storage

NetApp FAS storage systems reduce the cost and complexity for virtualized infrastructures by meeting your storage requirements with a single, highly scalable solution. NetApp's unified storage platform supports all protocols, so you no longer need to purchase separate systems to accommodate different storage needs. You can lower capacity use by up to 50% with built-in deduplication, thin provisioning, and space-efficient backup and cloning.

NetApp systems enhance operational efficiency with automated data management, data protection, and security. The ONTAP operating system brings a new level of nondisruptive operations, scalability, and efficiency to enterprise storage. Performance is optimized with innovative flash technologies and 10GbE and FCoE support. With NetApp data management, you can deploy the exact proportion of flash to spinning media for your environment. And for extreme performance for dedicated workloads.

NetApp All Flash FAS

Like all the FlexPod Systems, the FlexPod Datacenter with NetApp All Flash FAS (AFF) is comprised of compute (database, virtualization, application and management servers from Cisco), network (three-layer network and SAN technologies from Cisco), and storage (NetApp AFF storage systems). Individually and together, the components of the FlexPod solution are designed to keep applications available. This converged infrastructure solution integrates high-availability and disaster recovery capabilities that give you the ability to achieve "always on" performance with the agility to meet the rapidly changing demands of today's business environment.

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

NetApp Converged Systems Advisor

NetApp Converged Systems Advisor (CSA) is cloud-based management software for FlexPod. It supports FlexPod throughout its lifecycle with configuration, health, monitoring, and reporting capabilities. NetApp CSA combines an on-premises agent with a cloud-based portal to validate, monitor, and advise you on how to optimize your FlexPod system.

NetApp CSA performs the following tasks to validate your FlexPod implementation:

- Hardware and firmware assessment and inventory
- Verification of compliance with FlexPod best practices
- Verification of system configuration
- Comprehensive map diagram of system cabling, connectivity, components, and topology

NetApp CSA automates tasks and provides the following information to better monitor your FlexPod configuration:

- Scheduled system health checks for continuous validation
- Health check archive (configuration Snapshot copy) for configuration, change control, and status history
- Change control historical analysis
- E-mail notification for compliance failures

NetApp CSA helps optimize your FlexPod system with the following capabilities:

- Identification of bottlenecks or no resilient design that can affect performance
- Notifications for availability of hardware and firmware updates
- Recommendations for optimal hardware and firmware compliance

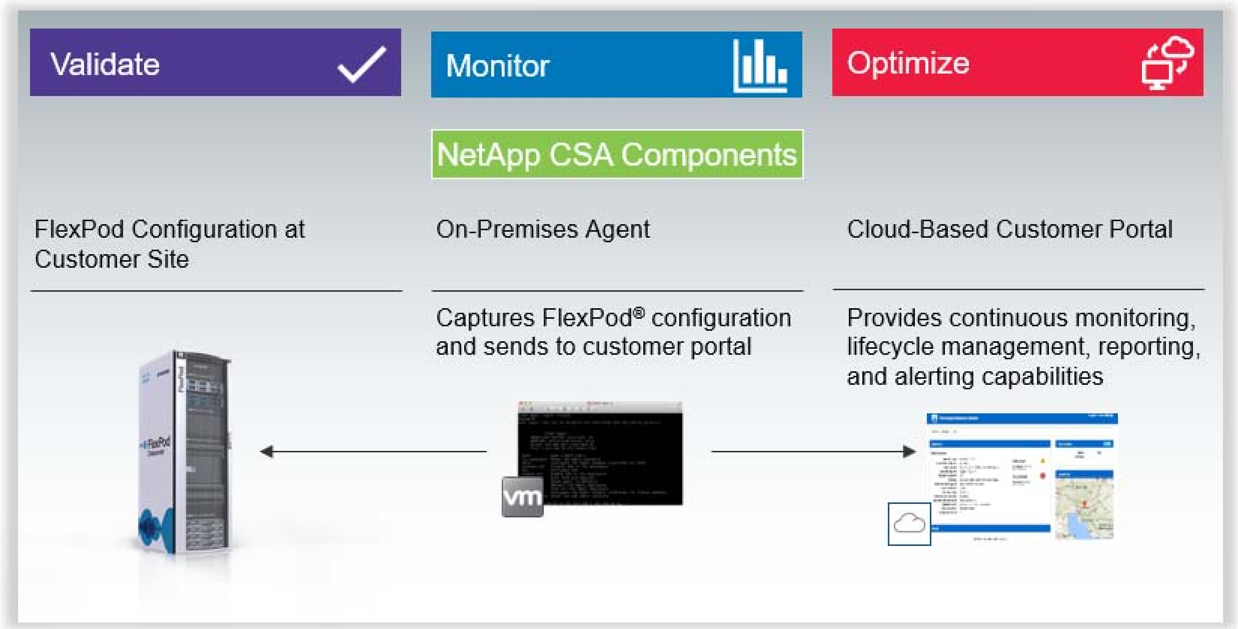

Validate ✓

Monitor

Optimize

NetApp CSA Components

FlexPod Configuration at Customer Site

On-Premises Agent

Captures FlexPod® configuration and sends to customer portal

Cloud-Based Customer Portal

Provides continuous monitoring, lifecycle management, reporting, and alerting capabilities

Figure 4: Lifecycle management automation – *NetApp Converged Systems Advisor automates key operations tasks and provides remote monitoring capabilities.*

FlexPod Global Partner Network

FlexPod is a meet-in-the-channel solution, with trained resellers delivering FlexPod solutions. Our reseller partners have NetApp channel certification and Cisco UCT, DCA, and ATP certifications. We now have more than 1,100 joint partners worldwide who deliver FlexPod, including 13 out of 14 of the world's largest system integrators.

Our network also includes 188 FlexPod Premium partners. This elite group of resellers has the highest level of training and certification and has made a significant investment in a FlexPod services practice.

You can choose from a global network of FlexPod Premium Partners and other highly qualified solution delivery partners to implement your FlexPod solution. These partners understand your unique requirements and are certified on NetApp, Cisco, and complementary technologies to deliver a complete cloud solution.

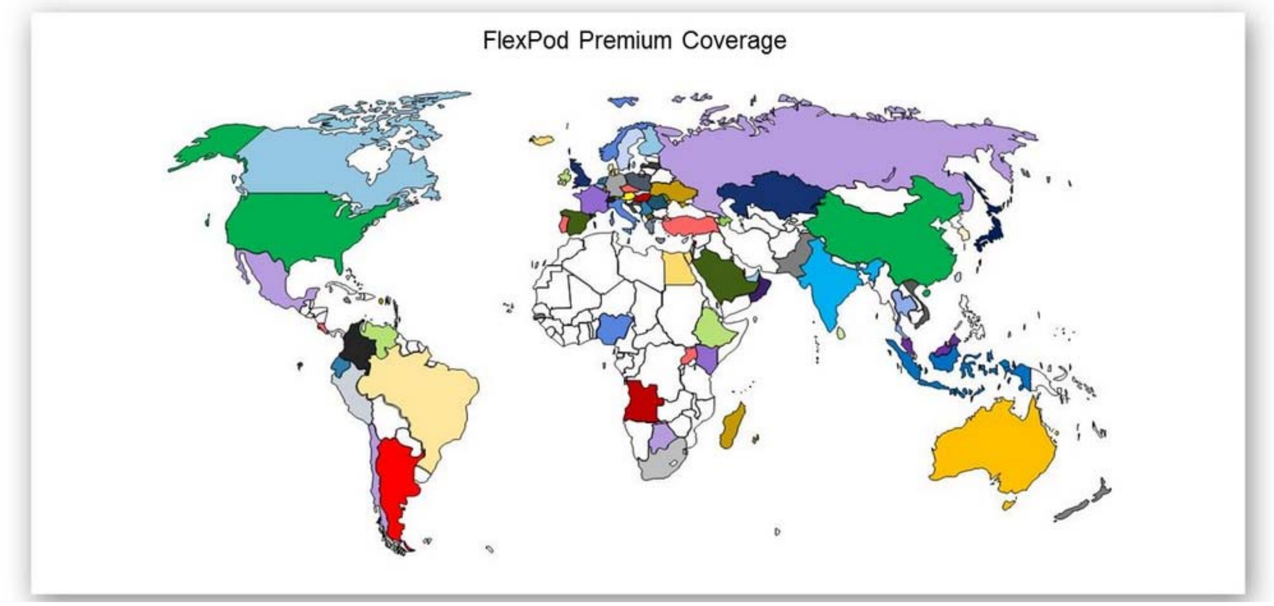

Figure 5: **Global FlexPod partner network** – *Assembling and delivering FlexPod to partners.*

FlexPod Solutions Support

FlexPod offers support that best fits your needs. You can choose to support your FlexPod solutions using Cooperative Support, Solution Support from either NetApp or Cisco, or support from your partners, if the partner provides support as part of their offering.

FlexPod Cooperative Support

FlexPod Cooperative Support is a multivendor program that includes NetApp, Cisco, VMware, Microsoft, Red Hat, and Citrix.

Your IT staff chooses which vendor to call based on your initial assessment of the problem's origin. Knowledgeable FlexPod engineers work to resolve issues quickly using shared communications, expertise gained through ongoing joint training, and a formal escalation process. This support model includes every area of support including people, processes, and technologies across companies. The result is a rapid resolution to your technical issues.

FlexPod Cooperative Support provides:

- Direct Access. No barrier between you and the experts, whenever you need them.
- Coordinated Support. Formal communications, training, and escalation processes across vendors and partners.
- FlexPod Expertise. A multivendor cooperative support lab to simulate and pretest support solutions.

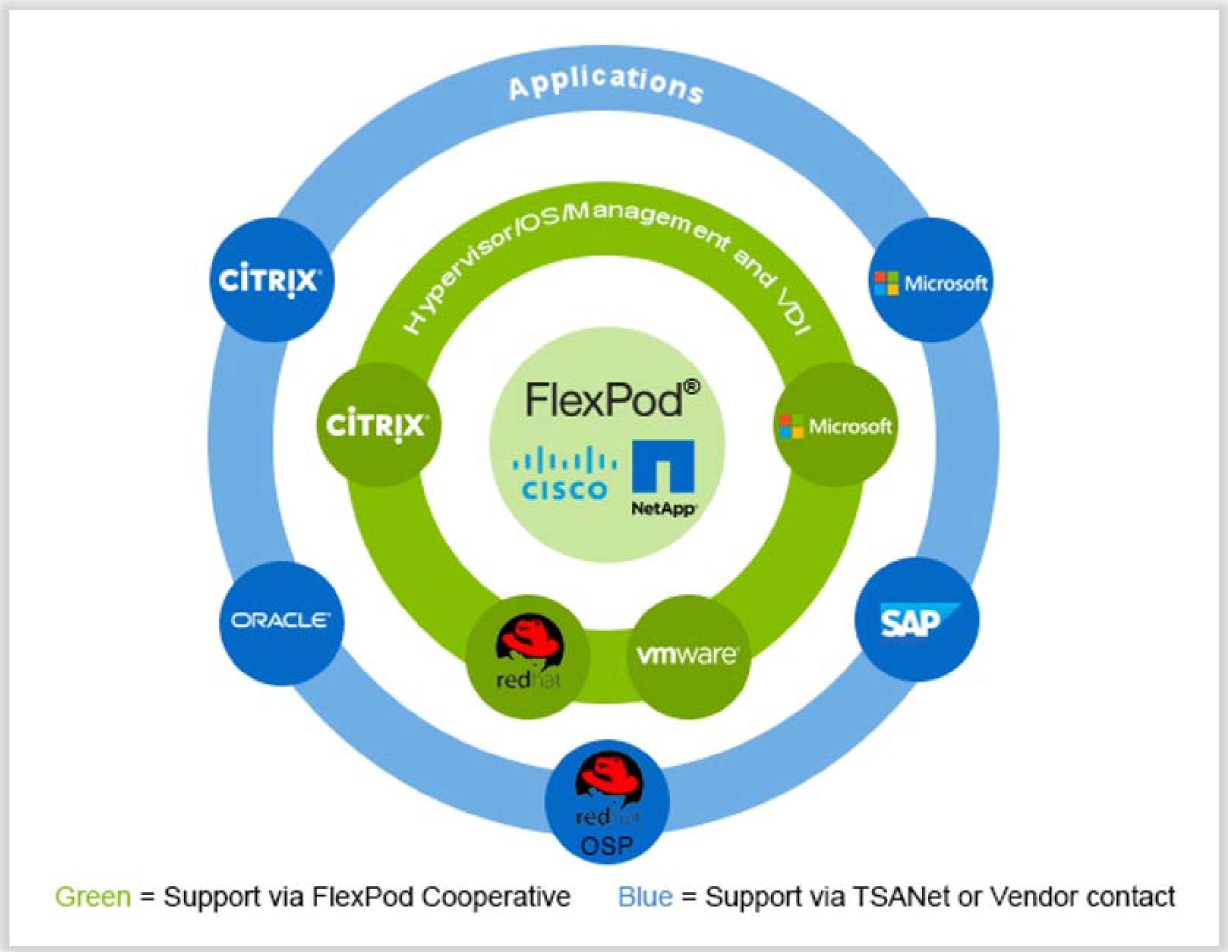
Applications
Hypervisor/OS/Management and VDI
FlexPod®
CISCO
NetApp
CITRIX
Microsoft
CITRIX
Microsoft
ORACLE
SAP
red hat
vmware
red hat OSP
Green = Support via FlexPod Cooperative
Blue = Support via TSANet or Vendor contact

Figure 6: The FlexPod Cooperative Support Model is an ecosystem of technology partners.

> *"To the Walz Group, the partnership among Cisco, NetApp, and VMware with a mutual commitment to support has meant time-saving technology integration, simpler update processes, and more open-minded, collaborative service."*
>
> *— Walz Group*

FlexPod Solution Support

Solution Support is premium-level support for the FlexPod solution that provides a primary point of contact offering FlexPod and data center expertise—including solution partner applications—for complete coverage in the stack.

NetApp Solution Support for FlexPod (NSS-FP) is a support offering that is layered on top of the required support contracts from Cisco, VMware, and NetApp. It offers full-stack, single-point-contact support for all the key components, including storage, storage network, compute, compute OS, LAN networking, and hypervisor. NetApp Support will help you resolve any problem with any component of your system, and will facilitate a resolution to any problem, even if that solution involves another vendor's product. NSS-FP provides the following services:

- A single contact point for any problem relating to your FlexPod infrastructure
- Access to experts in the storage, compute, networking, hypervisor, and many of the layered products typically deployed on a FlexPod
- Collaboration with other vendor’s support teams
- End-to-end case management

Note: NSS-FP includes a license for CSA.

Datasheet

NetApp Storage Encryption

Full disk encryption that protects data at rest with no operational impact

The NetApp® ONTAP® storage management solution continues to evolve, with security as an integral part of the solution. With NetApp Storage Encryption (NSE), full disk encryption is available using self-encrypting drives (SEDs) that are FIPS 140-2 level 2 certified. In addition, the strength of the portfolio and ONTAP solution continues with the arrival of NVMe SEDs, available in ONTAP 9.6 (not FIPS 140-2 certified); NetApp Volume Encryption (NVE), available in ONTAP 9.1; and NetApp Aggregate Encryption (NAE), available in ONTAP 9.6. NVE and NAE let you encrypt the data at a volume level, making the solution agnostic of the physical drive. In addition, the ability to take advantage of both hardware and software encryption options, providing double encryption at rest, is an industry first.

The Challenge

Encrypt your data without getting in the way

You work for a government, financial, or healthcare entity and are subject to regulations surrounding data protection. The requirement to keep all the personally identifiable information, personal healthcare information, and customer information protected within your storage infrastructure becomes a challenge when you repurpose drives, return defective drives, or upgrade to larger drives by selling them or trading them in. Wouldn't it be nice if there were a way for all of your data to be encrypted all the time without affecting everyday operations?

The Solution

NetApp Storage Encryption

NSE uses FIPS 140-2 level 2 SEDs to facilitate compliance and spares return by enabling the protection of data at rest, through AES 256-bit transparent disk encryption. The drives perform all the data encryption operations internally, as depicted in Figure 1, including encryption key generation. To prevent unauthorized access to the data, the storage system must authenticate itself with the drive by using an authentication key that is established the first time the drive is used. This can be done with either the onboard key manager (OKM) or an external key manager.

NSE can use the OKM to set and store the authentication keys for NSE drives. When the system uses an external key manager, the authentication keys are backed up externally using the industry-standard OASIS Key Management Interoperability Protocol (KMIP). With the external key manager, only the storage system, drive, and key manager have access to the key, and the drive cannot be unlocked if it is moved outside the security domain, thus preventing data leakage.

Key Benefits

Gain Full Disk Encryption

SEDs prevent data access until the drive's encryption key is unlocked by an authorized administrator.

Perform Mandatory Data Encryption

NSE and NVMe SEDs are file system and network independent: No action is required by the operator when aggregates, volumes, shares, or LUNs are created or deleted, and your data is always protected.

Enhance Data Confidentiality and Integrity

Use NSE or NVMe SEDs along with NVE and NAE to take advantage of the industry's first double encryption solution using two distinct layers. This combination provides a more robust data encryption solution.

Maintain Storage Efficiencies

By using NSE, NVMe SEDs, NVE, or NAE, you can encrypt your data while maintaining NetApp storage efficiencies such as deduplication and compression.

Satisfy Governance and Compliance Requirements

Use established security best practices to adhere to and support industry regulation and security compliance, including FIPS 140-2 level 2 with NSE.

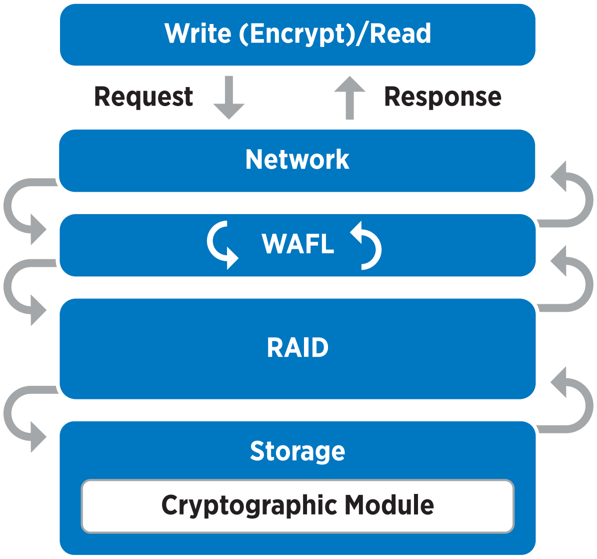

Figure 1) NSE cryptographic function.

Do I Need NSE?

Here are some questions to ask yourself:

- Must physical media be encrypted?
- Do I have any physical drive encryption requirements—for example, tamper evidence solutions?
- Must root aggregate volumes and storage virtual machine volumes be encrypted?
- Do I require ubiquitous disk encryption of all data?

If the answer to any of these questions is yes, then NSE is a viable solution.

Combine encryption for double encryption (layered defense)

If you need to segregate access to data as well as make sure that data is protected all the time, NSE can be combined with network- or fabric-level encryption. NSE can act like a backstop if an administrator forgets to configure or misconfigures higher-level encryption. For two distinct layers of encryption, you can combine NSE drives with NVE.

NSE Supported Storage Architectures

- NetApp AFF A-Series
- NetApp FAS9000 series
- NetApp FAS8200 series
- NetApp FAS2650 series
- NetApp FAS2620 series

Contact your account team to find out more about how the NSE solution can solidify your organization's needs. The following table lists some of the basics of NSE.

Entire disk encrypted
Hardware based
AES 256 encryption
NSE SEDs required
Onboard or external key management for the authentication key
FIPS 140-2 validated when used with external key manager; FIPS level depends on key manager use and implementation
All drives (including HA pairing) must be NSE drives; no mixing NSE and non-NSE drives

Resources

- NetApp Storage Encryption, NVMe Self-Encrypting Drives, NetApp Volume Encryption, and NetApp Aggregate Encryption datasheet
- NetApp Volume Encryption and NetApp Aggregate Encryption datasheet

About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit www.netapp.com. #DataDriven