

FlexPod Datacenter for SAP Solution with Cisco UCS Manager 3.2 and Cisco Application Centric Infrastructure (ACI)

Design and Deployment Guide for FlexPod Datacenter for SAP Solution with IP-Based Storage using NetApp AFF A-Series, Cisco UCS Manager 3.2, and Cisco Application Centric Infrastructure 3.2

Last Updated: September 11, 2018



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, see

<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 NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

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

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

Table of Contents

Executive Summary	11
Solution Overview	12
Introduction	12
Audience	12
Purpose of this Document.....	12
Technology Overview	14
Cisco Unified Computing System.....	14
Cisco UCS Manager	14
Cisco UCS Fabric Interconnect	15
Cisco UCS 6332UP Fabric Interconnect.....	16
Cisco UCS 2304XP Fabric Extender	16
Cisco UCS Blade Chassis	16
Cisco UCS B480 M5 Blade Server	17
Cisco UCS C480 M5 Rack Servers	18
Cisco UCS C240 M5 Rack Servers	19
Cisco I/O Adapters for Blade and Rack-Mount Servers	20
Cisco VIC 1340 Virtual Interface Card.....	20
Cisco VIC 1380 Virtual Interface Card.....	20
Cisco VIC 1385 Virtual Interface Card.....	21
Cisco UCS Differentiators	22
Cisco Application Centric Infrastructure (ACI)	23
ACI Components	25
Cisco Application Policy Infrastructure Controller.....	25
Leaf Switches.....	25
Spine Switches.....	26
NetApp All Flash FAS and ONTAP	26
FlexPod with Cisco ACI-Components	27
Validated System Hardware Components.....	30
Solution Architecture	31
Hardware and Software Components	33
Operating System Provisioning	34
SAP HANA Database Volumes.....	35
Management Pod	37

SAP HANA Solution Implementations.....	39
SAP HANA System on a Single Host - Scale-Up	39
SAP HANA System on Multiple Hosts Scale-Out	39
Infrastructure Requirements for the SAP HANA Database	40
CPU	40
Memory	40
CPU and Memory Combinations	40
Network.....	41
Storage.....	43
Filesystem Layout	44
Operating System	45
High Availability	45
Software Revisions	46
Configuration Guidelines	47
Device Cabling.....	54
Management Pod Cabling.....	62
Management Pod Installation	64
Network Configuration for Management Pod	64
Dual-Homed FEX Topology (Active/Active FEX Topology)	64
Cisco Nexus 9000 Series Switches–Network Initial Configuration Setup	64
Enable Appropriate Cisco Nexus 9000 Series Switches - Features and Settings	67
Create VLANs for Management Traffic.....	68
Configure Virtual Port Channel Domain	68
Configure Network Interfaces for the VPC Peer Links	69
Configure Network Interfaces to Cisco UCS C220 Management Server.....	70
Direct Connection of Management Pod to FlexPod Infrastructure	71
Uplink into Existing Network Infrastructure	73
Management Server Installation	73
Server Configuration.....	73
CIMC Configuration	73
Storage Configuration.....	75
VMware ESXi Installation	79
Install ESXi.....	80
Set Up Management Networking for ESXi Hosts	81
VMware ESXi Host ESXi-Mgmt-01	82

FlexPod Cisco ACI Network Configuration for SAP HANA	86
Cisco APIC Initial Configuration.....	86
Cisco IMC Configuration	86
Configure Cisco APIC Power State	87
Configure Cisco IMC NTP	88
Cisco APIC Initial Configuration Setup	88
Cisco ACI Fabric Discovery.....	92
Configure NTP for Cisco ACI Fabric	94
POD Policy Group.....	96
Configuring Access Policies.....	98
CDP Interface Policies	98
LLDP Interface Policies	100
LACP Interface Policies	102
Link Level Policy	103
VLAN Pool for SAP HANA.....	104
Domain Configuration	106
Attachable Access Entity Profile Configuration.....	108
Leaf Interface Policy Groups	110
Interface Profile Configuration	116
VPC Configuration for NetApp Storage.....	121
Switch Policy Configuration	129
Switch Profile Configuration.....	130
Tenant Configuration	133
Tenant Creation	136
Application Profiles for HANA	137
Bridge Domains	138
Application EPGs	142
Contracts.....	150
Cisco UCS Solution for SAP HANA TDI	156
Cisco UCS Server Configuration	157
Initial Setup of Cisco UCS 6332 Fabric Interconnect.....	157
Cisco UCS 6332 Fabric Interconnect A.....	157
Cisco UCS 6332 Fabric Interconnect B.....	158
Log in to Cisco UCS Manager	158
Upgrade Cisco UCS Manager Software to Version 3.2(3d)	159

Add Block of IP Addresses for KVM Access	159
Synchronize Cisco UCS to NTP.....	160
Cisco UCS Blade Chassis Connection Options.....	160
Edit Chassis Discovery Policy	160
Enable Server and Uplink Ports.....	161
Acknowledge Cisco UCS Chassis and Rack-Mount Servers	162
Create Uplink Port Channels to Cisco ACI Leaf Switches	162
Create New Organization	166
Create MAC Address Pools	167
Create UUID Suffix Pool.....	170
Power Policy.....	171
Power Control Policy	171
Create Host Firmware Package.....	172
Create Local Disk Configuration Policy (Optional)	173
Create Server BIOS Policy	174
Create Serial Over LAN Policy	177
Update Default Maintenance Policy.....	178
IPMI Access Profiles	179
Network Configuration	180
Set Jumbo Frames in Cisco UCS Fabric.....	180
Update Default Network Control Policy to Enable CDP.....	181
LAN Tab Configurations.....	181
Create VLANs.....	182
Assigning Port Channels to VLANs	184
Create QoS Policies.....	185
Create vNIC Template	186
Create vNIC/vHBA Placement Policy	200
Create PXE Boot Policies	201
Create Service Profile Templates SAP HANA Scale-Out.....	202
Create Service Profile from the Template	214
Create Service Profile Templates SAP HANA iSCSI	216
Create Service Profile Templates for SAP HANA Scale-Up.....	227
Create Service Profiles	230
Storage Configuration	231
Complete Configuration Worksheet	231

Configure ONTAP Nodes	231
Configure Node 01	231
Configure Node 02	233
Set Up Node 01	234
Set Up Node 02	238
Log into the Cluster	239
Set Auto-Revert on Cluster Management	239
Set Up Management Broadcast Domain	240
Set Up Service Processor Network Interface	240
Create Aggregates	240
Verify Storage Failover.....	241
Disable Flow Control on 40GbE Ports	242
Configure Network Time Protocol	242
Configure Simple Network Management Protocol.....	243
Configure AutoSupport	243
Enable Cisco Discovery Protocol	244
Create Broadcast Domains	244
Create Interface Groups	245
Create VLANs.....	245
Configure HTTPS Access	246
Upgrade to ONTAP 9.2/9.3	248
Configure SVM for the Infrastructure.....	254
Create SVM for the Infrastructure	255
Create Load-Sharing Mirrors	256
Create Export Policies for the Root Volumes.....	256
Add Infrastructure SVM Administrator.....	256
Create Export Policies for the Infrastructure SVM.....	257
Create iSCSI LIFs.....	257
Create PXE LIFs.....	257
Create Block Protocol (iSCSI) Service.....	258
Create FlexVol Volumes.....	258
Configure LUNs for iSCSI Boot	258
Configure SVM for HANA.....	259
Create SVM for SAP HANA.....	260
Create Load-Sharing Mirrors	260

Create Export Policies for the Root Volumes.....	261
Add HANA SVM Administrator.....	261
Create Export Policies for the HANA SVM.....	261
Create NFS LIF for SAP HANA Data.....	261
Create NFS LIF for SAP HANA Log.....	262
HANA Node Preparation.....	263
Preparation of PXE Boot Environment.....	263
Installing SLES 12 SP3 based PXE Server VM on the Management Servers.....	263
Customize PXE Server.....	265
Mount Volume for PXE Boot Configuration.....	266
Update Packages for PXE Server VM.....	267
PXE Services Configuration.....	268
Installing RHEL 7.4 based PXE Server VM on the Management Servers.....	272
Customize PXE Server.....	274
Update Packages for PXE Server VM.....	275
Mount Volume for PXE Boot Configuration.....	276
PXE Services Configuration.....	276
Operating System Installation SUSE SLES12SP3.....	279
PXE Boot Preparation for SUSE OS Installation.....	279
PXE booting with SUSE Linux Enterprise Server 12 SP3.....	280
Create Swap Partition in a File.....	290
Update OS Master.....	291
Install Cisco enic Driver.....	293
Operating System Optimization for SAP HANA.....	295
Cloning OS Volumes.....	296
Post Installation OS Customization.....	299
Operating System Installation Red Hat Enterprise Linux 7.4.....	302
Post Installation Tasks.....	309
Install Cisco enic Driver.....	312
Prepare NFS Root Volume.....	313
Cloning OS Volumes.....	315
Post Installation OS Customization.....	317
Storage Provisioning for SAP HANA.....	321
Configuring SAP HANA Single-Host Systems.....	322
Configuration Example for a SAP HANA Single-Host System.....	323

Create Data Volume and Adjust Volume Options	324
Create a Log Volume and Adjust the Volume Options.....	324
Create a HANA Shared Volume and Qtrees and Adjust the Volume Options	324
Update the Load-Sharing Mirror Relation.....	325
Create Mount Points	325
Configuration for SAP HANA Multiple-Host Systems	326
Configuration Example for a SAP HANA Multiple-Host Systems	328
Create Data Volumes and Adjust Volume Options.....	329
Create Log Volume and Adjust Volume Options.....	329
Create HANA Shared Volume and Qtrees and Adjust Volume Options.....	330
Update Load-Sharing Mirror Relation	330
Create Mount Points	330
Mount File Systems	331
SAP HANA Installation	333
Important SAP Notes	333
SAP HANA IMDB Related Notes	333
Linux Related Notes.....	333
SAP Application Related Notes.....	334
Third Party Software	334
High-Availability (HA) Configuration for Scale-Out.....	335
High-Availability Configuration.....	335
Enable the SAP HANA Storage Connector API.....	337
Test the IPMI Connectivity	337
SAP HANA Data Protection	338
SAP HANA Backup	338
SAP HANA Disaster Recovery with Asynchronous Storage Replication.....	339
References	341
About the Authors.....	342
Acknowledgements	342



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 infrastructure.

This document describes the architecture and deployment procedures for SAP HANA Tailored DataCenter Integration option on FlexPod infrastructure composed of Cisco® compute and switching products that leverage Cisco Application Centric Infrastructure [ACI] - the industry-leading software-defined networking solution (SDN) along with NetApp® A-series AFF arrays. The intent of this document is to show the design principles with the detailed configuration steps for SAP HANA deployment.

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers deploying the FlexPod Datacenter Solution for SAP HANA with NetApp clustered Data ONTAP®. External references are provided wherever applicable, but readers are expected to be familiar with the technology, infrastructure, and database security policies of the customer installation.

Solution Overview

Introduction

FlexPod is a defined set of hardware and software that serves as an integrated foundation for virtualized and non-virtualized data center solutions. It provides a pre-validated, ready-to-deploy infrastructure that reduces the time and complexity involved in configuring and validating a traditional data center deployment. The FlexPod Datacenter solution for SAP HANA includes NetApp storage, NetApp ONTAP, 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 and can be deployed according to a customer's data center design. A key 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. The solution is designed to host scalable SAP HANA workloads.

SAP HANA is SAP SE's implementation of in-memory database technology. A SAP HANA database takes advantage of low cost main memory (RAM), the data-processing capabilities of multicore processors, and faster data access to provide better performance for analytical and transactional applications. SAP HANA offers a multi-engine query-processing environment that supports relational data with both row-oriented and column-oriented physical representations in a hybrid engine. It also offers graph and text processing for semi-structured and unstructured data management within the same system.

With the introduction of SAP HANA TDI for shared infrastructure, the FlexPod solution provides you the advantage of having the compute, storage, and network stack integrated with the programmability of the Cisco UCS. SAP HANA TDI enables organizations to run multiple SAP HANA production systems in one FlexPod solution. It also enables customers to run the SAP applications servers and the SAP HANA database on the same infrastructure.

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers deploying the FlexPod Datacenter Solution for SAP HANA with NetApp clustered Data ONTAP®. External references are provided wherever applicable, but readers are expected to be familiar with the technology, infrastructure, and database security policies of the customer installation.

Purpose of this Document

This document describes the steps required to deploy and configure a FlexPod Datacenter Solution for SAP HANA with Cisco **ACI**. **Cisco's validation provides further confirmation with regard to component**

compatibility, connectivity and correct operation of the entire integrated stack. This document showcases one of the variants of cloud architecture for SAP HANA. While readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to the deployment of this solution are provided in this CVD.

Technology Overview

The FlexPod Datacenter Solution for SAP HANA with Cisco ACI is composed of Cisco UCS servers, Cisco Nexus switches and NetApp AFF storage. This section describes the main features of these different solution components.

Cisco Unified Computing System

Cisco Unified Computing System™ (Cisco UCS®) is an integrated computing infrastructure with embedded management to automate and accelerate deployment of all your applications, including virtualization and cloud computing, scale-out and bare-metal workloads, and in-memory analytics, as well as edge computing that supports remote and branch locations and massive amounts of data from the Internet of Things (IoT). The main components of Cisco UCS: unified fabric, unified management, and unified computing resources.

The Cisco Unified Computing System is the first integrated data center platform that combines industry standard, x86-architecture servers with networking and storage access into a single unified system. The system is smart infrastructure that uses integrated, model-based management to simplify and accelerate deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud **computing environments. Employing Cisco's innovative SingleConnect technology, the system's unified I/O** infrastructure uses a unified fabric to support both network and storage I/O. The Cisco fabric extender architecture extends the fabric directly to servers and virtual machines for increased performance, security, and manageability. Cisco UCS helps change the way that IT organizations do business, including the following:

- Increased IT staff productivity and business agility through just-in-time provisioning and equal support for both virtualized and bare-metal environments
- Reduced TCO at the platform, site, and organization levels through infrastructure consolidation
- A unified, integrated system that is managed, serviced, and tested as a whole
- Scalability through a design for up to 160 discrete servers and thousands of virtual machines, the capability to scale I/O bandwidth to match demand, the low infrastructure cost per server, and the capability to manage up to 6000 servers with Cisco UCS Central Software
- Open industry standards supported by a partner ecosystem of industry leaders
- A system that scales to meet future data center needs for computing power, memory footprint, and I/O bandwidth; it is poised to help you move to 40 Gigabit Ethernet with the new Cisco UCS 6300 Series Fabric Interconnects

Cisco UCS Manager

Cisco UCS Manager (UCSM) provides unified, embedded management of all software and hardware **components of the Cisco Unified Computing System™ (Cisco UCS) and Cisco HyperFlex™ Systems across** multiple chassis and rack servers and thousands of virtual machines. It supports all Cisco UCS product models, including Cisco UCS B-Series Blade Servers and C-Series Rack Servers, Cisco UCS Mini, and Cisco HyperFlex hyperconverged infrastructure, as well as the associated storage resources and networks. Cisco

UCS Manager is embedded on a pair of Cisco UCS 6300 or 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The manager participates in server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection.

An instance of Cisco UCS Manager with all Cisco UCS components managed by it forms a Cisco UCS domain, which can include up to 160 servers. In addition to provisioning Cisco UCS resources, this infrastructure management software provides a model-based foundation for simplifying the day-to-day processes of updating, monitoring, and managing computing resources, local storage, storage connections, and network connections. By enabling better automation of processes, Cisco UCS Manager allows IT organizations to achieve greater agility and scale in their infrastructure operations while reducing complexity and risk. The manager provides flexible role- and policy-based management using service profiles and templates.

Cisco UCS Manager manages Cisco UCS systems through an intuitive HTML 5 or Java user interface and a command-line interface (CLI). It can register with Cisco UCS Central Software in a multi-domain Cisco UCS environment, enabling centralized management of distributed systems scaling to thousands of servers. The manager can be integrated with Cisco UCS Director to facilitate orchestration and to provide support for converged infrastructure and Infrastructure as a Service (IaaS).

The Cisco UCS API provides comprehensive access to all Cisco UCS Manager functions. The unified API provides Cisco UCS system visibility to higher-level systems management tools from independent software vendors (ISVs) such as VMware, Microsoft, and Splunk as well as tools from BMC, CA, HP, IBM, and others. ISVs and in-house developers can use the API to enhance the value of the Cisco UCS platform according to their unique requirements. Cisco UCS PowerTool for Cisco UCS Manager and the Python Software Development Kit (SDK) help automate and manage configurations in Cisco UCS Manager.

Cisco UCS Fabric Interconnect

The Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6300 Series offers line-rate, low-latency, lossless 10 and 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions. The Cisco UCS 6300 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers, 5100 Series Blade Server Chassis, and C-Series Rack Servers managed by Cisco UCS. All servers attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6300 Series provides both LAN and SAN connectivity for all servers within its domain.

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

Cisco UCS 6332UP Fabric Interconnect

The Cisco UCS 6332 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. All servers attached to 6332 Fabric Interconnects become part of one highly available management domain. The Cisco UCS 6332UP 32-Port Fabric Interconnect is a 1-rack-unit 40 Gigabit Ethernet, FCoE and Fibre Channel switch offering up to 2.56 Tbps throughput and up to 32 ports. The switch has 32 fixed 40-Gbps Ethernet and FCoE ports. Cisco UCS 6332UP 32-Port Fabric Interconnect have ports that can be configured for the breakout feature that supports connectivity between 40 Gigabit Ethernet ports and 10 Gigabit Ethernet ports. This feature provides backward compatibility to existing hardware that supports 10 Gigabit Ethernet. A 40 Gigabit Ethernet port can be used as four 10 Gigabit Ethernet ports. Using a 40 Gigabit Ethernet SFP, these ports on a Cisco UCS 6300 Series Fabric Interconnect can connect to another fabric interconnect that has four 10 Gigabit Ethernet SFPs.

Figure 1 Cisco UCS 6332 UP Fabric Interconnect



Cisco UCS 2304XP Fabric Extender

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

Figure 2 Cisco UCS 2304 XP



Cisco UCS Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server.

The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors. Four hot-swappable power supplies are accessible from the front of the chassis, and single-phase AC, -48V DC, and 200 to 380V DC power supplies and chassis are available. These power supplies are up to 94 percent efficient and meet the requirements for the 80 Plus Platinum rating. The power subsystem can be configured to support nonredundant, N+1 redundant, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays that can support either Cisco UCS 2000 Series Fabric Extenders or the Cisco UCS 6324 Fabric Interconnect. A passive

midplane provides up to 80 Gbps of I/O bandwidth per server slot and up to 160 Gbps of I/O bandwidth for two slots.

The Cisco UCS Blade Server Chassis is shown in Error! Reference source not found..

Figure 3 Cisco Blade Server Chassis (front and back view)



Cisco UCS B480 M5 Blade Server

The enterprise-class Cisco UCS B480 M5 Blade Server delivers market-leading performance, versatility, and density without compromise for memory-intensive mission-critical enterprise applications and virtualized workloads, among others. With the Cisco UCS B480 M5, you can quickly deploy stateless physical and virtual workloads with the programmability that Cisco UCS Manager and Cisco® SingleConnect technology enable.

The Cisco UCS B480 M5 is a full-width blade server supported by the Cisco UCS 5108 Blade Server Chassis. The Cisco UCS 5108 chassis and the Cisco UCS B-Series Blade Servers provide inherent architectural advantages:

- Through Cisco UCS, gives you the architectural advantage of not having to power, cool, manage, and purchase excess switches (management, storage, and networking), Host Bus Adapters (HBAs), and Network Interface Cards (NICs) in each blade chassis
- Reduces the Total Cost of Ownership (TCO) by removing management modules from the chassis, making the chassis stateless
- **Provides a single, highly available Cisco Unified Computing System™ (Cisco UCS) management domain** for all system chassis and rack servers, reducing administrative tasks

The Cisco UCS B480 M5 Blade Server offers:

- Four Intel® Xeon® Scalable CPUs (up to 28 cores per socket)
- 2666-MHz DDR4 memory and 48 DIMM slots with up to 6 TB using 128-GB DIMMs
- Cisco FlexStorage® storage subsystem
- Five mezzanine adapters and support for up to four GPUs
- Cisco UCS Virtual Interface Card (VIC) 1340 modular LAN on Motherboard (mLOM) and upcoming fourth-generation VIC mLOM
- Internal Secure Digital (SD) and M.2 boot options

Figure 4 Cisco UCS B480 M5 Blade Server



Cisco UCS C480 M5 Rack Servers

The Cisco UCS C480 M5 Rack Server is a storage and I/O optimized enterprise-class rack server that delivers industry-leading performance for in-memory databases, big data analytics, virtualization, Virtual Desktop Infrastructure (VDI), and bare-metal applications. The Cisco UCS C480 M5 delivers outstanding **levels of expandability and performance for standalone or Cisco Unified Computing System™ (Cisco UCS)** managed environments in a 4RU form-factor, and because of its modular design, you pay for only what you need. It offers these capabilities:

- Latest Intel® Xeon® Scalable processors with up to 28 cores per socket and support for two or four processor configurations
- 2666-MHz DDR4 memory and 48 DIMM slots for up to 6 TeraBytes (TB) of total memory
- 12 PCI Express (PCIe) 3.0 slots
 - Six x8 full-height, full length slots
 - Six x16 full-height, full length slots
- Flexible storage options with support up to 24 Small-Form-Factor (SFF) 2.5-inch, SAS, SATA, and PCIe NVMe disk drives
- Cisco® 12-Gbps SAS Modular RAID Controller in a dedicated slot
- Internal Secure Digital (SD) and M.2 boot options
- Dual embedded 10 Gigabit Ethernet LAN-On-Motherboard (LOM) ports

Cisco UCS C480 M5 servers can be deployed as standalone servers or in a Cisco UCS managed environment. When used in combination with Cisco UCS Manager, the Cisco UCS C480 M5 brings the power and automation of unified computing to enterprise applications, including Cisco® SingleConnect technology, drastically reducing switching and cabling requirements. Cisco UCS Manager uses service profiles, templates, and policy-based management to enable rapid deployment and help ensure deployment consistency. It also enables end-to-end server visibility, management, and control in both virtualized and bare-metal environments.

Figure 5 Cisco UCS C480 M5 Rack Server



Cisco UCS C240 M5 Rack Servers

The Cisco UCS C240 M5 Rack Server is a 2-socket, 2-Rack-Unit (2RU) rack server offering industry-leading performance and expandability. It supports a wide range of storage and I/O-intensive infrastructure workloads, from big data and analytics to collaboration. Cisco UCS C-Series Rack Servers can be deployed **as standalone servers or as part of a Cisco Unified Computing System™ (Cisco UCS) managed environment** to take advantage of **Cisco's standards-based unified computing innovations that help reduce customers' Total Cost of Ownership (TCO) and increase their business agility.**

In response to ever-increasing computing and data-intensive real-time workloads, the enterprise-class Cisco UCS C240 M5 server extends the capabilities of the Cisco UCS portfolio in a 2RU form factor. It incorporates the Intel® Xeon® Scalable processors, supporting up to 20 percent more cores per socket, twice the memory capacity, and five times more Non-Volatile Memory Express (NVMe) PCI Express (PCIe) Solid-State Disks (SSDs) compared to the previous generation of servers. These improvements deliver significant performance and efficiency gains that will improve your application performance. The Cisco UCS C240 M5 delivers outstanding levels of storage expandability with exceptional performance, along with the following:

- Latest Intel Xeon Scalable CPUs with up to 28 cores per socket
- Up to 24 DDR4 DIMMs for improved performance
- Up to 26 hot-swappable Small-Form-Factor (SFF) 2.5-inch drives, including 2 rear hot-swappable SFF drives (up to 10 support NVMe PCIe SSDs on the NVMe-optimized chassis version), or 12 Large-Form-Factor (LFF) 3.5-inch drives plus 2 rear hot-swappable SFF drives
- Support for 12-Gbps SAS modular RAID controller in a dedicated slot, leaving the remaining PCIe Generation 3.0 slots available for other expansion cards
- Modular LAN-On-Motherboard (mLOM) slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot, supporting dual 10- or 40-Gbps network connectivity
- Dual embedded Intel x550 10GBASE-T LAN-On-Motherboard (LOM) ports
- Modular M.2 or Secure Digital (SD) cards that can be used for boot

Cisco UCS C240 M5 servers can be deployed as standalone servers or in a Cisco UCS managed environment. When used in combination with Cisco UCS Manager, the Cisco UCS C240 M5 brings the

power and automation of unified computing to enterprise applications, including Cisco® SingleConnect technology, drastically reducing switching and cabling requirements. Cisco UCS Manager uses service profiles, templates, and policy-based management to enable rapid deployment and help ensure deployment consistency. It also enables end-to-end server visibility, management, and control in both virtualized and bare-metal environments.

Figure 6 Cisco UCS C240 M5 Rack Server



Cisco I/O Adapters for Blade and Rack-Mount Servers


This section discusses the Cisco I/O Adapters used in this solution.

Cisco VIC 1340 Virtual Interface Card

The Cisco UCS blade server has various Converged Network Adapters (CNA) options.

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

Figure 7 Cisco UCS 1340 VIC Card



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

Cisco VIC 1380 Virtual Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1380 is a dual-port 40-Gbps Ethernet, or dual 4 x 10 Fibre Channel over Ethernet (FCoE)-capable mezzanine card designed exclusively for the M5 generation of Cisco

UCS B-Series Blade Servers. The card enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1380 supports Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.


Figure 8 Cisco UCS 1380 VIC Card



Cisco VIC 1385 Virtual Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1385 is a Cisco® innovation. It provides a policy-based, stateless, agile server infrastructure for your data center. This dual-port Enhanced Quad Small Form-Factor Pluggable (QSFP) half-height PCI Express (PCIe) card is designed exclusively for Cisco UCS C-Series Rack Servers. The card supports 40 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). It incorporates **Cisco's next**-generation converged network adapter (CNA) technology and offers a comprehensive feature set, providing investment protection for future feature software releases. The card can present more than 256 PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the VIC supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology. This technology extends the Cisco UCS Fabric Interconnect ports to virtual machines, simplifying server virtualization deployment.

Figure 9 Cisco UCS 1385 VIC Card



Cisco UCS Differentiators

Cisco Unified Computing System is revolutionizing the way servers are managed in data-center. The following are the unique differentiators of Cisco Unified Computing System and Cisco UCS Manager:

- **Embedded management:** In Cisco Unified Computing System, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers. Also, a pair of FIs can manage up to 40 chassis, each containing 8 blade servers. This gives enormous scaling on management plane.
- **Unified fabric:** In Cisco Unified Computing System, from blade server chassis or rack server fabric extender to FI, there is a single Ethernet cable used for LAN, SAN and management traffic. This converged I/O, results in reduced cables, SFPs and adapters – reducing capital and operational expenses of overall solution.
- **Auto discovery:** By simply inserting the blade server in the chassis or connecting rack server to the fabric extender, discovery and inventory of compute resource occurs automatically without any management intervention. Combination of unified fabric and auto-discovery enables wire-once architecture of Cisco Unified Computing System, where compute capability of Cisco Unified Computing System can extend easily, while keeping the existing external connectivity to LAN, SAN and management networks.
- **Policy based resource classification:** When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD focuses on the policy-based resource classification of Cisco UCS Manager.
- **Combined Rack and Blade server management:** Cisco UCS Manager can manage Cisco UCS B-Series Blade Servers and Cisco UCS C-Series Rack Servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic. This CVD focuses on the combination of B-Series and C-Series Servers to demonstrate stateless and form factor independent computing work load.

- Model-based management architecture: Cisco UCS Manager Architecture and management database is model based and data driven. Open, standard based XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management system, such as VMware vCloud director, Microsoft system center, and Citrix CloudPlatform.
- Policies, Pools, Templates: Management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- Loose referential integrity: In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each-other. This provides great flexibilities where different experts from different domains, such as network, storage, security, server and virtualization work together to accomplish a complex task.
- Policy resolution: In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to other policy by name is resolved in the org hierarchy with closest policy match. **If no policy with specific name is found in the hierarchy till root org, then special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibilities to owners of different orgs.**
- Service profiles and stateless computing: Service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- Built-in multi-tenancy support: Combination of policies, pools and templates, loose referential integrity, policy resolution in org hierarchy and service profile based approach to compute resources make Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
- Virtualization aware network: VM-FEX technology makes access layer of network aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when virtual network is managed by port-**profiles defined by the network administrators’ team.** VM-FEX also offloads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.
- Simplified QoS: Even though fibre-channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

Cisco Application Centric Infrastructure (ACI)

The Cisco Nexus 9000 family of switches supports two modes of operation: NxOS standalone mode and Application Centric Infrastructure (ACI) fabric mode. In standalone mode, the switch performs as a typical

Nexus switch with increased port density, low latency and 40Gb connectivity. In fabric mode, the administrator can take advantage of Cisco ACI. Cisco Nexus 9000 based FlexPod design with Cisco ACI consists of Cisco Nexus 9500 and 9300 based spine/leaf switching architecture controlled using a cluster of three Application Policy Infrastructure Controllers (APICs).

Cisco ACI delivers a resilient fabric to satisfy today's dynamic applications. ACI leverages a network fabric that employs industry proven protocols coupled with innovative technologies to create a flexible, scalable, and highly available architecture of low-latency, high-bandwidth links. This fabric delivers application instantiations through the use of profiles that house the requisite characteristics to enable end-to-end connectivity.


The ACI fabric is designed to support the industry trends of management automation, programmatic policies, and dynamic workload provisioning. The ACI fabric accomplishes this with a combination of hardware, policy-based control systems, and closely coupled software to provide advantages not possible in other architectures.

The Cisco ACI fabric consists of three major components:

- The Application Policy Infrastructure Controller (APIC)
- Spine switches
- Leaf switches

The ACI switching architecture uses a leaf-and-spine topology, in which each leaf switch is connected to every spine switch in the network, with no interconnection between leaf switches or spine switches. Each leaf and spine switch is connected with one or more 40 Gigabit Ethernet links or with 100 Gigabit links. Each APIC appliance should connect to two leaf switches for resiliency purpose.

Figure 10 Cisco ACI Fabric Architecture



ACI Components

Cisco Application Policy Infrastructure Controller

The Cisco Application Policy Infrastructure Controller (APIC) is the unifying point of automation and management for the ACI fabric. The Cisco APIC provides centralized access to all fabric information, optimizes the application lifecycle for scale and performance, and supports flexible application provisioning across physical and virtual resources. Some of the key benefits of Cisco APIC are:

- Centralized application-level policy engine for physical, virtual, and cloud infrastructures
- Detailed visibility, telemetry, and health scores by application and by tenant
- Designed around open standards and open APIs
- Robust implementation of multi-tenant security, quality of service (QoS), and high availability
- Integration with management systems such as VMware, Microsoft, and OpenStack

The software controller, APIC, is delivered as an appliance and three or more such appliances form a cluster for high availability and enhanced performance. The controller is a physical appliance based on a Cisco UCS® rack server with two 10 Gigabit Ethernet interfaces for connectivity to the leaf switches. The APIC is also equipped with 1 Gigabit Ethernet interfaces for out-of-band management. Controllers can be configured with 10GBASE-T or SFP+ Network Interface Cards (NICs), and this configuration must match the physical format supported by the leaf. In other words, if controllers are configured with 10GBASE-T, they have to be connected to a Cisco ACI leaf with 10GBASE-T ports.

APIC is responsible for all tasks enabling traffic transport including:

- Fabric activation
- Switch firmware management
- Network policy configuration and instantiation

Although the APIC acts as the centralized point of configuration for policy and network connectivity, it is never in line with the data path or the forwarding topology. The fabric can still forward traffic even when communication with the APIC is lost.

APIC provides both a command-line interface (CLI) and graphical-user interface (GUI) to configure and control the ACI fabric. APIC also provides a northbound API through XML and JavaScript Object Notation (JSON) and an open source southbound API

For more information on Cisco APIC, refer to:

<http://www.cisco.com/c/en/us/products/cloud-systems-management/application-policy-infrastructure-controller-apic/index.html>

Leaf Switches

In Cisco ACI, all workloads connect to leaf switches. Leaf switch, typically can be a fixed form Nexus 9300 series or a modular Nexus 9500 series switch that provides physical server and storage connectivity as well

as enforces ACI policies. The latest Cisco ACI fixed form factor leaf nodes allow connectivity up to 25 and 40 Gbps to the server and uplinks of 100 Gbps to the spine. There are a number of leaf switch choices that differ based on functions like port speed, medium type, multicast routing support, scale of endpoints etc.

For a summary of leaf switch options available refer the Cisco ACI Best Practices Guide https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_0111.html

Spine Switches

The Cisco ACI fabric forwards traffic primarily based on host lookups. A mapping database stores the information about the leaf switch on which each IP address resides. This information is stored in the fabric cards of the spine switches. All known endpoints in the fabric are programmed in the spine switches. The spine models also differ in the number of endpoints supported in the mapping database, which depends on the type and number of fabric modules installed

For a summary of spine switch options available refer the Cisco ACI Best Practices Guide: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/ACI_Best_Practices/b_ACI_Best_Practices/b_ACI_Best_Practices_chapter_0111.html

NetApp All Flash FAS and ONTAP

NetApp All Flash FAS (AFF) systems address enterprise storage requirements with high performance, superior flexibility, and best-in-class data management. Built on NetApp ONTAP data management software, AFF systems speed up business without compromising on the efficiency, reliability, or flexibility of IT operations. As an enterprise-grade all-flash array, AFF accelerates, manages, and protects business-critical data and enables an easy and risk-free transition to flash for your data center.

Designed specifically for flash, the NetApp AFF A series all-flash systems deliver industry-leading performance, capacity density, scalability, security, and network connectivity in dense form factors. At up to 7M IOPS per cluster with submillisecond latency, they are the fastest all-flash arrays built on a true unified scale-**out architecture**. **As the industry's first** all-flash arrays to provide both 40 Gigabit Ethernet (40GbE) and 32Gb Fibre Channel connectivity, AFF A series systems eliminate the bandwidth bottlenecks that are increasingly moved to the network from storage as flash becomes faster and faster.

AFF comes with a full suite of acclaimed NetApp integrated data protection software. Key capabilities and benefits include the following:

- Native space efficiency with cloning and NetApp Snapshot® copies, which reduces storage costs and minimizes performance effects.
- Application-consistent backup and recovery, which simplifies application management.
- NetApp SnapMirror® replication software, which replicates to any type of FAS/AFF system—all flash, hybrid, or HDD and on the premises or in the cloud— and reduces overall system costs.

AFF systems are built with innovative inline data reduction technologies:

- Inline data compaction technology uses an innovative approach to place multiple logical data blocks from the same volume into a single 4KB block.


- Inline compression has a near-zero performance effect. Incompressible data detection eliminates wasted cycles.
- Enhanced inline deduplication increases space savings by eliminating redundant blocks.

This version of FlexPod introduces the NetApp AFF A300 series unified scale-out storage system. This controller provides the high-performance benefits of 40GbE and all flash SSDs and occupies only 3U of rack space. Combined with a disk shelf containing 3.8TB disks, this solution provides ample horsepower and over 90TB of raw capacity while taking up only 5U of valuable rack space. The AFF A300 features a multiprocessor Intel chipset and leverages high-performance memory modules, NVRAM to accelerate and optimize writes, and an I/O-tuned PCIe gen3 architecture that maximizes application throughput. The AFF A300 series comes with integrated unified target adapter (UTA2) ports that support 16Gb Fibre Channel, 10GBE, and FCoE. In addition 40GBE add-on cards are available.

FlexPod with Cisco ACI—Components

FlexPod with ACI is designed to be fully redundant in the compute, network, and storage layers. There is no single point of failure from a device or traffic path perspective. 0 shows how the various elements are connected together.


Figure 11 FlexPod Design with Cisco ACI



Fabric: As in the previous designs of FlexPod, link aggregation technologies play an important role in FlexPod with ACI providing improved aggregate bandwidth and link resiliency across the solution stack. The NetApp storage controllers, Cisco Unified Computing System, and Cisco Nexus 9000 platforms support active port channeling using 802.3ad standard Link Aggregation Control Protocol (LACP). In addition, the Cisco Nexus 9000 series features virtual Port Channel (vPC) capabilities. vPC allows links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single "logical" port channel to a third device, essentially offering device fault tolerance. Note in the Figure above that vPC peer links are no longer needed for leaf switches. The peer link is handled in the leaf to spine connections and any two leaves in an ACI fabric can be paired in a vPC. The Cisco UCS Fabric Interconnects and NetApp FAS controllers benefit from the Cisco Nexus vPC abstraction, gaining link and device resiliency as well as full utilization of a non-blocking Ethernet fabric.


Compute: Each Fabric Interconnect (FI) is connected to both the leaf switches and the links provide a robust 40GbE connection between the Cisco Unified Computing System and ACI fabric. Figure 12 illustrates the use of vPC enabled 40GbE uplinks between the Cisco Nexus 9000 leaf switches and Cisco UCS 3rd Gen FIs. Additional ports can be easily added to the design for increased bandwidth as needed. Each Cisco UCS 5108 chassis is connected to the FIs using a pair of ports from each IO Module for a combined 40G uplink. FlexPod design supports direct attaching the Cisco UCS C-Series servers into the FIs. FlexPod designs mandate Cisco UCS C-Series management using Cisco UCS Manager to provide a uniform look and feel across blade and standalone servers.

Figure 12 Compute Connectivity



Storage: The ACI-based FlexPod design is an end-to-end IP-based storage solution that supports SAN access by using iSCSI. The solution provides a 40GbE fabric that is defined by Ethernet uplinks from NetApp storage controllers connected to the Cisco Nexus switches.

Figure 13 Storage Connectivity to ACI Leaf Switches



0 shows the initial storage configuration of this solution as a two-node high availability (HA) pair running clustered Data ONTAP in a switchless cluster configuration. Storage system scalability is easily achieved by adding storage capacity (disks and shelves) to an existing HA pair, or by adding more HA pairs to the cluster or storage domain.

Validated System Hardware Components

The following components were used to validate this Cisco Nexus 9000 ACI design:

- Cisco Unified Computing System 3.2 (3d)
- Cisco Nexus 2304 Fabric Extenders
- Cisco UCS B480M5 servers
- Cisco UCS C220M5 servers
- Cisco Nexus 93180LC-EX Series Leaf Switch
- Cisco Nexus 9336PQ Spine Switch
- Cisco Application Policy Infrastructure Controller (APIC) 2.2(3s)
- NetApp All-Flash FAS Unified Storage

Solution Architecture

The FlexPod Datacenter ACI solution for SAP HANA with NetApp All Flash FAS storage provides an end-to-end architecture with Cisco and NetApp technologies that demonstrate support for multiple SAP and SAP HANA workloads with high availability and server redundancy. The architecture uses Cisco UCS Manager with combined Cisco UCS B-Series blade servers and C-Series rack servers with NetApp AFF A300 series storage attached to the Cisco Nexus 93180LC-EX ACI leaf switches. The Cisco UCS C-Series Rack Servers are connected directly to Cisco UCS Fabric Interconnect with single-wire management feature. This infrastructure provides PXE and iSCSI boot options for hosts with file-level and block-level access to shared storage. **The reference architecture reinforces the “wire-once” strategy, because when the additional storage is added to the architecture, no re-cabling is required from hosts to the Cisco UCS Fabric Interconnect.**

Figure 14 shows the FlexPod Datacenter reference architecture for SAP HANA workload, described in this Cisco Validation Design. It highlights the FlexPod hardware components and the network connections for a configuration with IP-based storage.

Figure 14 FlexPod ACI Solution for SAP HANA - Reference Architecture

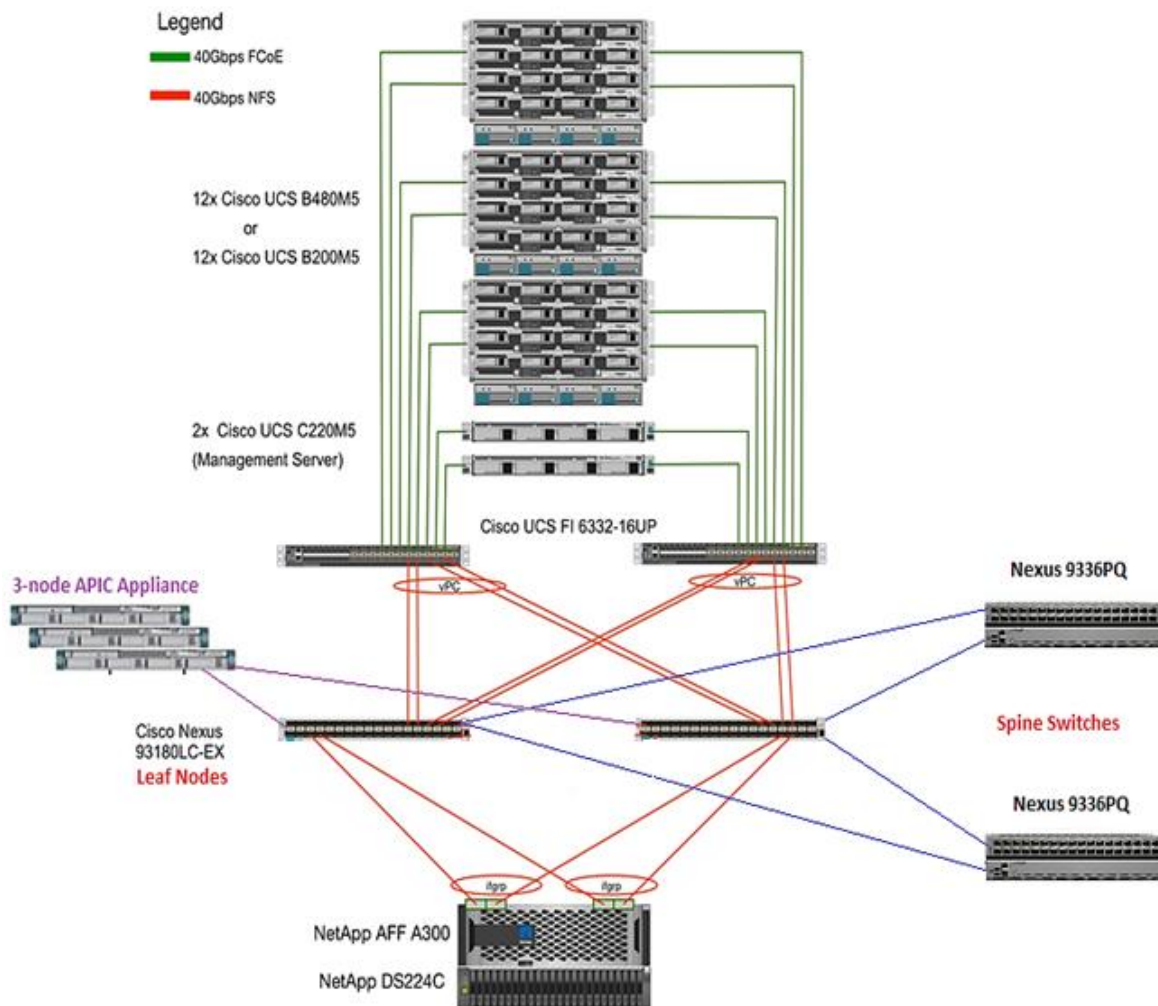


Figure 14 includes the following:

- Cisco Unified Computing System
 - 2 x Cisco UCS 6332 32 x 40Gb/s
 - 3 x Cisco UCS 5108 Blade Chassis with 2 x Cisco UCS 2304 Fabric Extenders with 4x 40 Gigabit Ethernet interfaces
 - 12 x Cisco UCS B480 M5 High-Performance Blade Servers with 2x Cisco UCS Virtual Interface Card (VIC) 1380 and 2x Cisco UCS Virtual Interface Card (VIC) 1340

Or

 - 12 x Cisco UCS C480 M5 High-Performance Rack-Mount Servers with 2x Cisco UCS Virtual Interface Card (VIC) 1385.

Or

 - 12 x Cisco UCS B200 M5 High-Performance Blade Servers with Cisco UCS Virtual Interface Card (VIC) 1340
 - 2 x Cisco UCS C220 M5 High-Performance Rack Servers with Cisco UCS Virtual Interface Card (VIC) 1385
- Cisco ACI
 - 2 x Cisco Nexus 93180LC-EX Switch for 40/100 Gigabit Ethernet for the ACI Leafs
 - 2 x Cisco Nexus 9336PQ Switch for ACI Spines
 - 3 node APIC cluster appliance
- NetApp AFF A300 Storage
 - NetApp AFF A300 Storage system using ONTAP 9.3
 - 1 x NetApp Disk Shelf DS224C with 24x 3.8TB SSD

Although this is the base design, each of the components can be scaled easily to support specific business requirements. Additional servers or even blade chassis can be deployed to increase compute capacity without additional Network components. Two Cisco UCS 6332 Fabric interconnect can support up to:

- 20 Cisco UCS B-Series B480 M5
- 20 Cisco UCS C480 M5 Server
- 40 Cisco UCS C220 M5/C240 M5 Server

For every twelve Cisco UCS Servers, one NetApp AFF A300 HA paired with ONTAP is required to meet the SAP HANA storage performance. While adding compute and storage for scaling, it is required to increase the network bandwidth between Cisco UCS Fabric Interconnect and Cisco Nexus 9000 switch. Each NetApp Storage requires an additional two 40 GbE connectivity from each Cisco UCS Fabric Interconnect to Cisco Nexus 9000 switches.

The number of Cisco UCS C-Series or Cisco UCS B-Series Servers and the NetApp FAS storage type depends on the number of SAP HANA instances. SAP specifies the storage performance for SAP HANA, based on a per server rule independent of the server size. In other words, the maximum number of servers per storage remains the same if you use Cisco UCS B200 M5 with 192GB physical memory or Cisco UCS B480 M5 with 6TB physical memory.

Hardware and Software Components

This architecture is based on and supports the following hardware and software components:

- SAP HANA
 - SAP Business Suite on HANA or SAP Business Warehouse on HANA
 - S/4HANA or BW/4HANA
 - SAP HANA single-host or multiple-host configurations
- Operating System
 - SUSE Linux Enterprise for SAP (SLES for SAP)
 - Red Hat Enterprise Linux
- Cisco UCS Server
 - Bare Metal
- Network
 - 40GbE end-to-end
 - NFS for SAP HANA data access
 - PXE NFS or iSCSI for OS boot
- Storage
 - NetApp AFF A-series array

Figure 15 shows an overview of the hardware and software components.

Figure 15 Hardware and Software Component Overview



SAP HANA
SLES and RHEL



Cisco UCS



NFS and iSCSI with
40Gb Ethernet



NetApp All Flash FAS


Operating System Provisioning

All operating system images are provisioned from the external NetApp storage, leveraging either:

- PXE boot and NFS root file system
- iSCSI boot

Figure 16 shows an overview of the different operating system provisioning methods.

Figure 16 Overview Operating System Provisioning



SAP HANA Database Volumes

All SAP HANA database volumes, data, log, and the shared volumes are mounted with NFS from the central storage.

Figure 17 shows an overview of the SAP HANA database volumes.


Figure 17 SAP HANA Database Volumes



Figure 18 shows a block diagram of a complete SAP Landscape built using the FlexPod architecture. It is comprised of multiple SAP HANA systems and SAP applications with shared infrastructure as illustrated in the figure. The FlexPod Datacenter reference architecture for SAP solutions supports SAP HANA system in both Scale-Up mode and Scale-Out mode with multiple servers with the shared infrastructures.

The FlexPod datacenter solution manages the communication between the application server and the SAP HANA database. This approach enhances system performance by improving bandwidth and latency. It also improves system reliability by including the application server in the disaster-tolerance solution with the SAP HANA database.

Figure 18 Shared Infrastructure Block Diagram



The FlexPod architecture for SAP HANA TDI can run other workloads on the same infrastructure, as long as the rules for workload isolation are considered.

You can run the following workloads on the FlexPod architecture:

1. Production SAP HANA databases
2. SAP application servers

3. Non-production SAP HANA databases
4. Production and non-production SAP systems on traditional databases
5. Non-SAP workloads

In order to make sure that the storage KPIs for SAP HANA production databases are fulfilled, the SAP HANA production databases must have dedicated storage controller of a NetApp FAS Storage HA pair. SAP application servers could share the same storage controller with the production SAP HANA databases.

This document describes in detail the procedure for the reference design and outlines the network, compute and storage configurations and deployment process for running SAP HANA on FlexPod platform.



This document does not describe the procedure for deploying SAP applications.

Management Pod

Comprehensive management is an important element for a FlexPod environment running SAP HANA, especially in a system involving multiple FlexPod platforms; Management pod was built to handle this efficiently. It is optional to build a dedicated Management environment; you can use your existing Management environment for the same functionality. Management Pod includes (but is not limited to) a pair of Cisco Nexus 9000 Series switches (readily available Cisco Nexus 9396PX switches were used in the validation PoD) in standalone mode and a pair of Cisco UCS C220 M5 Rack-Mount Servers. The Cisco Nexus 9000 series switches provide the out-of-band management network through a dual homed Nexus 2K switch providing 1GbE connections. The Cisco UCS C220 M5 Rack-Mount Servers will run ESXi with PXE boot server, vCenter with additional management, and monitor virtual machines.




It is recommended to use additional NetApp FAS Storage in the Management Pod for redundancy and failure scenarios.

Management Pod switches can connect directly to FlexPod switches or your existing network infrastructure. If your existing network infrastructure is used, the uplink from FlexPod switches are connected same pair of switch as uplink from Management Pod switches as shown in Figure 19.




The LAN switch must allow all the necessary VLANs for managing the FlexPod environment.

Figure 19 Management Pod Using Customer Existing Network



The dedicated Management Pod can directly connect to each FlexPod environment as shown in Figure 20. In this topology, the switches are configured as port-channels for unified management. This CVD describes the procedure for the direct connection option.

Figure 20 Direct Connection of Management Pod to FlexPod



SAP HANA Solution Implementations

This section describes the various implementation options and their requirements for a SAP HANA system.

SAP HANA System on a Single Host - Scale-Up

A single-host system is the simplest of the installation types. It is possible to run an SAP HANA system entirely on one host and then scale the system up as needed. All data and processes are located on the same server and can be accessed locally. The network requirements for this option minimum one 1-Gb Ethernet (access) and one 10/40-Gb Ethernet storage networks are sufficient to run SAP HANA scale-up.

With the SAP HANA TDI option, multiple SAP HANA scale-up systems can be built on a shared infrastructure.

SAP HANA System on Multiple Hosts Scale-Out

SAP HANA Scale-Out option is used if the SAP HANA system does not fit into the main memory of a single server based on the rules defined by SAP. In this method, multiple independent servers are combined to form one system and the load is distributed among multiple servers. In a distributed system, each index server is usually assigned to its own host to achieve maximum performance. It is possible to assign different tables to different hosts (partitioning the database), or a single table can be split across hosts (partitioning of tables). SAP HANA Scale-Out supports failover scenarios and high availability. Individual hosts in a distributed system have different roles master, worker, slave, standby depending on the task.



Some use cases are not supported on SAP HANA Scale-Out configuration and it is recommended to check with SAP whether a use case can be deployed as a Scale-Out solution.

The network requirements for this option are higher than for Scale-Up systems. In addition to the client and application access and storage access network, a node-to-node network is necessary. One 10 Gigabit Ethernet (access) and one 10 Gigabit Ethernet (node-to-node) and one 10 Gigabit Ethernet storage networks are required to run SAP HANA Scale-Out system. Additional network bandwidth is required to support system replication or backup capability.

Based on the SAP HANA TDI option for shared storage and shared network, multiple SAP HANA Scale-Out systems can be built on a shared infrastructure.

Infrastructure Requirements for the SAP HANA Database

There are hardware and software requirements defined by SAP to run SAP HANA systems in Tailored Datacenter Integration (TDI) option. This Cisco Validated Design uses guidelines provided by SAP.

Additional information is available at: <http://saphana.com>.



This document does not cover the updated information published by SAP after Q1/2017.

CPU

SAP HANA 2.0 (TDI) supports servers equipped with Intel Xeon processor E7-8880v3, E7-8890v3, E7-8880v4, E7-8890v4 **and all Skylake CPU's > 8 cores**. In addition, the Intel Xeon processor E5-26xx v4 is supported for scale-up systems with the SAP HANA TDI option.

Memory

SAP HANA is supported in the following memory configurations:

- Homogenous symmetric assembly of dual in-line memory modules (DIMMs) for example, DIMM size or speed should not be mixed
- Maximum use of all available memory channels
- SAP HANA 2.0 Memory per socket up to 1024 GB for SAP NetWeaver Business Warehouse (BW) and DataMart
- SAP HANA 2.0 Memory per socket up to 1536 GB for SAP Business Suite on SAP HANA (SoH) on 2- or 4-socket server

CPU and Memory Combinations

SAP HANA allows for a specific set of CPU and memory combinations. Table 1 describes the list of certified Cisco UCS servers for SAP HANA with supported Memory and CPU configuration for different use cases.

Table 1 List of Cisco UCS Servers Defined in FlexPod Datacenter Solution for SAP

Cisco UCS Server	CPU	Supported Memory	Scale UP/Suite on HANA	Scale-Out
Cisco UCS B200 M5	2 x Intel Xeon	128 GB to 2 TB BW 128 GB to 3 TB for SoH	Supported	Not supported
Cisco UCS C220 M5	2 x Intel Xeon	128 GB to 2 TB BW 128 GB to 3 TB for SoH	Supported	Not supported
Cisco UCS	2 x Intel	128 GB to 2 TB BW	Supported	Not

Cisco UCS Server	CPU	Supported Memory	Scale UP/Suite on HANA	Scale-Out
C240 M5	Xeon	128 GB to 3 TB for SoH		supported
Cisco UCS B480 M5	4 x Intel Xeon	256 GB to 4 TB for BW 256 GB to 6 TB for SoH	Supported	Supported
Cisco UCS C480 M5	4 x Intel Xeon	256 GB to 4 TB for BW 256 GB to 6 TB for SoH	Supported	Supported
Cisco C880 M5	8x Intel Xeon	2TB - 6TB for BW 2TB - 12TB for SoH	Supported	Supported

Network

A SAP HANA data center deployment can range from a database running on a single host to a complex distributed system. Distributed systems can get complex with multiple hosts located at a primary site having one or more secondary sites; supporting a distributed multi-terabyte database with full fault and disaster recovery.

SAP HANA has different types of network communication channels to support the different SAP HANA scenarios and setups:

- Client zone: Channels used for external access to SAP HANA functions by end-user clients, administration clients, and application servers, and for data provisioning through SQL or HTTP
- Internal zone: Channels used for SAP HANA internal communication within the database or, in a distributed scenario, for communication between hosts
- Storage zone: Channels used for storage access (data persistence) and for backup and restore procedures

Table 2 lists all the networks defined by SAP or Cisco or requested by customers.

Table 2 List of Known Networks

Name	Use Case	Solutions	Bandwidth requirements
Client Zone Networks			
Application Server Network	SAP Application Server to DB communication	All	10 or 40 GbE
Client Network	User / Client Application to DB communication	All	10 or 40 GbE

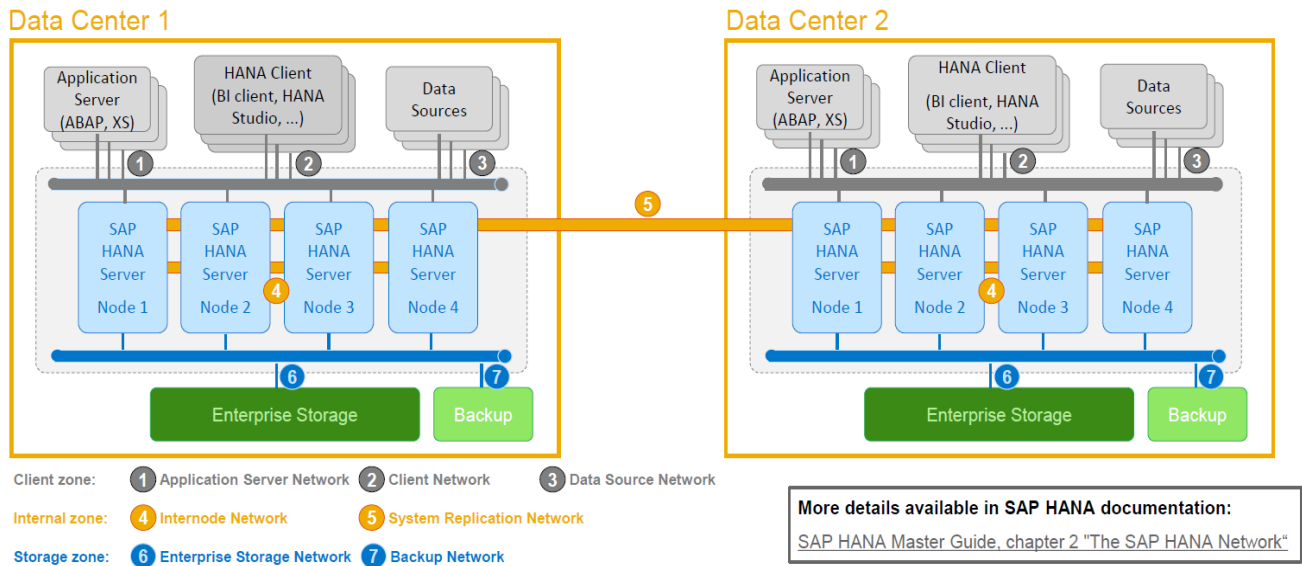
Name	Use Case	Solutions	Bandwidth requirements
Data Source Network	Data import and external data integration	Optional for all SAP HANA systems	10 or 40 GbE
Internal Zone Networks			
Inter-Node Network	Node to node communication within a scale-out configuration	Scale-Out	40 GbE
System Replication Network		For SAP HANA Disaster Tolerance	TBD with Customer
Storage Zone Networks			
Backup Network	Data Backup	Optional for all SAP HANA systems	10 or 40 GbE
Storage Network	Node to Storage communication	All	40 GbE
Infrastructure Related Networks			
Administration Network	Infrastructure and SAP HANA administration	Optional for all SAP HANA systems	1 GbE
Boot Network	Boot the Operating Systems via PXE/NFS or iSCSI	Optional for all SAP HANA systems	40 GbE

Details about the network requirements for SAP HANA are available in the white paper from SAP SE at: <http://www.saphana.com/docs/DOC-4805>.

The network needs to be properly segmented and must be connected to the same core/ backbone switch as shown in Figure 21 based on your **customer’s high**-availability and redundancy requirements for different SAP HANA network segments.

Figure 21 High-Level SAP HANA Network Overview

High-Level SAP HANA Network Overview



Based on the listed network requirements, every server must be equipped with 2x 10 Gigabit Ethernet for scale-up systems to establish the communication with the application or user (Client Zone) and a 10 GbE Interface for Storage access.

For Scale-Out solutions, an additional redundant network for SAP HANA node to node communication with 10 GbE is required (Internal Zone).



For more information on SAP HANA Network security, refer to the [SAP HANA Security Guide](#).

Storage

As an in-memory database, SAP HANA uses storage devices to save a copy of the data, for the purpose of startup and fault recovery without data loss. The choice of the specific storage technology is driven by various requirements like size, performance and high availability. To use Storage system in the Tailored Datacenter Integration option, the storage must be certified for SAP HANA TDI option at:

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/enterprise-storage.html>.

All relevant information about storage requirements is documented in this white paper:

<https://www.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html>.



SAP can only support performance related SAP HANA topics if the installed solution has passed the validation test successfully.


Refer to the [SAP HANA Administration Guide](#) section 2.8 Hardware Checks for Tailored Datacenter Integration for Hardware check test tool and the related documentation.

Filesystem Layout

Figure 22 shows the file system layout and the required storage sizes to install and operate SAP HANA. For the Linux OS installation (/root) 10 GB of disk size is recommended. Additionally, 50 GB must be provided for the /usr/sap since the volume used for SAP software that supports SAP HANA.

While installing SAP HANA on a host, specify the mount point for the installation binaries (/hana/shared/<sid>), data files (/hana/data/<sid>) and log files (/hana/log/<sid>), where sid is the instance identifier of the SAP HANA installation.

Figure 22 File System Layout for 2 Node Scale-Out System



The storage sizing for filesystem is based on the amount of memory equipped on the SAP HANA host.

Below is a sample filesystem size for a single system appliance configuration:

Root-FS: 10 GB

/usr/sap: 50 GB

/hana/shared: 1x RAM or 1TB whichever is less

/hana/data: 1 x RAM

/hana/log: ½ of the RAM size for systems ≤ 256GB RAM and min ½ TB for all other systems

With a distributed installation of SAP HANA Scale-Out, each server will have the following:

Root-FS: 10 GB

/usr/sap: 50 GB

The installation binaries, trace and configuration files are stored on a shared filesystem, which should be accessible for all hosts in the distributed installation. The size of shared filesystem should be 1 X RAM of a worker node for each 4 nodes in the cluster. For example, in a distributed installation with three hosts with 512 GB of memory each, shared file system should be 1 x 512 GB = 512 GB, for 5 hosts with 512 GB of memory each, shared file system should be 2 x 512 GB = 1024GB.

For each SAP HANA host there should be a mount point for data and log volume. The size of the file system for data volume with TDI option is one times the host memory:

/hana/data/<sid>/mntXXXXX: 1x RAM

For solutions based on Intel Skylake 81XX CPU the size of the Log volume must be as follows:

- Half of the server RAM size for systems with \leq 512 GB RAM
- 512 GB for systems with > 512 GB RAM

Operating System

The supported operating systems for SAP HANA are as follows:

- SUSE Linux Enterprise Server for SAP Applications
- Red Hat Enterprise Linux for SAP HANA

High Availability

The infrastructure for a SAP HANA solution must not have single point of failure. To support high-availability, the hardware and software requirements are:

- Internal storage: A RAID-based configuration is preferred
- External storage: Redundant data paths, dual controllers, and a RAID-based configuration are required
- Ethernet switches: Two or more independent switches should be used

SAP HANA Scale-Out comes with in integrated high-availability function. If a SAP HANA system is configured with a stand-by node, a failed part of SAP HANA will start on the stand-by node automatically. For automatic host failover, storage connector API must be properly configured for the implementation and operation of the SAP HANA.

For detailed information from SAP see: <http://saphana.com> or <http://service.sap.com/notes>.

Software Revisions

Table 3 details the software revisions used for validating various components of the FlexPod Datacenter Reference Architecture for SAP HANA.

Table 3 Hardware and Software Components of the FlexPod Datacenter Reference Architecture for SAP Solutions

Component		Software version	Count
Network	Nexus 93180LC-EX <i>[ACI Leaf]</i>	3.2(2l) [build 13.2(2l)]	2
	Nexus 9336PQ <i>[ACI Spine]</i>	3.2(2l) [build 13.2(2l)]	2
	3 node Cluster APIC appliance	3.2(2l)	1 unit
Compute	Cisco UCS Fabric Interconnect 6332 - 16UP	3.2(3d)	2
	Cisco UCS B480 M5 w/1340/1380 VIC	3.3(3d)B	12
	Cisco UCS C220 M5 w/1385 VIC <i>[Management Servers]</i>	3.1(3g)	2
Storage	NetApp AFF A300	ONTAP 9.3	1
	NetApp DS224C Flash disk shelf		1
Software OS	SLES4SAP	12 SP3	
	Red Hat	7.4	

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document and Cisco Nexus A and 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 are identified sequentially: HANA-Server01, HANA-Server02, and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. Review 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, etc. 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	Description	Customer Implementation Value
<<var_nexus_mgmt_A_hostname>>	Cisco Nexus Management A host name	
<<var_nexus_mgmt_A_mgmt0_ip>>	Out-of-band Cisco Nexus Management A management IP address	
<<var_nexus_mgmt_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_mgmt_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_mgmt_B_hostname>>	Cisco Nexus Management B host name	

Variable	Description	Customer Implementation Value
<<var_nexus_mgmt_B_mgmt0_ip>>	Out-of-band Cisco Nexus Management B management IP address	
<<var_nexus_mgmt_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_mgmt_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_oob_vlan_id>>	Out-of-band management network VLAN ID	
<<var_admin_vlan_id_mgmt>> <<var_admin_vlan_id>>	Mgmt PoD - Admin Network VLAN Admin network VLAN ID - UCS	
<<var_boot_vlan_id>> <<var_boot_vlan_id_aci>> <<var_boot_vlan_id_aci_storage>>	PXE boot network VLAN ID -UCS PXE boot network VLAN ID - PXEserver Mgmt Pod PXE boot network VLAN ID Storage	
<<var_nexus_vpc_domain_mgmt_id>>	Unique Cisco Nexus switch VPC domain ID for Management Switch	
<<var_nexus_vpc_domain_id>>	Unique Cisco Nexus switch VPC domain ID for Management Switch	
<<var_vm_host_mgmt_01_ip>>	ESXi Server 01 for Management Server IP Address	
<<var_vm_host_mgmt_02_ip>>	ESXi Server 02 for Management Server IP Address	
<<var_nexus_A_hostname>>	Cisco Nexus Mgmt-A host name	
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus Mgmt-A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_B_hostname>>	Cisco Nexus Mgmt-B host name	

Variable	Description	Customer Implementation Value
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus Mgmt-B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_storage_data_vlan_id>> <<var_storage_data_vlan_id_aci>> <<var_storage_log_vlan_id>> <<var_storage_log_vlan_id_aci>>	Storage network for HANA Data VLAN ID - UCS Storage HANA Data Network VLAN ID - ACI Storage network for HANA Log VLAN ID - UCS Storage HANA Log Network VLAN ID - ACI	
<<var_internal_vlan_id>>	Node to Node Network for HANA Data/log VLAN ID	
<<var_backup_vlan_id>>	Backup Network for HANA Data/log VLAN ID	
<<var_client_vlan_id>>	Client Network for HANA Data/log VLAN ID	
<<var_appserver_vlan_id>>	Application Server Network for HANA Data/log VLAN ID	
<<var_datasource_vlan_id>>	Data source Network for HANA Data/log VLAN ID	
<<var_replication_vlan_id>>	Replication Network for HANA Data/log VLAN ID	
<<iSCSI_vlan_id_A>>	iSCSI-A VLAN ID initiator UCS	
<<iSCSI_vlan_id_B>>	iSCSI-B VLAN ID initiator UCS	
<<iSCSI_vlan_id_A_aci>>	iSCSI-A VLAN ID Target ACI	
<<iSCSI_vlan_id_B_aci>>	iSCSI-B VLAN ID Target ACI	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address	

Variable	Description	Customer Implementation Value
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address	
<<var_cimc_gateway>>	Out-of-band management network default gateway	
<<var_ib-mgmt_vlan_id>>	In-band management network VLAN ID	
<<var_node01_mgmt_ip>>	Out-of-band management IP for storage cluster node 01	
<<var_node01_mgmt_mask>>	Out-of-band management network netmask	
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_url_boot_software>>	Data ONTAP 9.x URL; format: http://	
<<var_number_of_disks>>	Number of disks to assign to each storage controller	
<<var_node02_mgmt_ip>>	Out-of-band management IP for storage cluster node 02	
<<var_node02_mgmt_mask>>	Out-of-band management network netmask	
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_clustername>>	Storage cluster host name	
<<var_cluster_base_license_key>>	Cluster base license key	
<<var_nfs_license>>	NFS protocol license key	
<<var_iscsi_license>>	iSCSI protocol license key	
<<var_flexclone_license>>	FlexClone license key	
<<var_password>>	Global default administrative password	
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster	
<<var_clustermgmt_mask>>	Out-of-band management network netmask	

Variable	Description	Customer Implementation Value
<<var_clustermgmt_gateway>>	Out-of-band management network default gateway	
<<var_dns_domain_name>>	DNS domain name	
<<var_nameserver_ip>>	DNS server IP(s)	
<<var_node_location>>	Node location string for each node	
<<var_node01>>	Storage Cluster node 01 host name	
<<var_node02>>	Storage Cluster node 02 host name	
<<var_node01_sp_ip>>	Out-of-band Storage cluster node 01 service processor management IP	
<<var_node01_sp_mask>>	Out-of-band management network netmask	
<<var_node01_sp_gateway>>	Out-of-band management network default gateway	
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP	
<<var_node02_sp_mask>>	Out-of-band management network netmask	
<<var_node02_sp_gateway>>	Out-of-band management network default gateway	
<<var_timezone>>	FlexPod time zone (for example, America/New_York)	
<<var_snmp_contact>>	Administrator e-mail address	
<<var_snmp_location>>	Cluster location string	
<<var_oncommand_server_fqdn>>	VSC or OnCommand virtual machine fully qualified domain name (FQDN)	
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name	
<<var_mailhost>>	Mail server host name	
<<var_storage_admin_email>>	Storage Administrator e-mail address	
<<var_security_cert_vserver_common_name>>	Infrastructure Vserver FQDN	
<<var_country_code>>	Two-letter country code	
<<var_state>>	State or province name	

Variable	Description	Customer Implementation Value
<<var_city>>	City name	
<<var_org>>	Organization or company name	
<<var_unit>>	Organizational unit name	
<<var_security_cert_cluster_comm on_name>>	Storage cluster FQDN	
<<var_security_cert_node01_commo n_name>>	Cluster node 01 FQDN	
<<var_security_cert_node02_commo n_name>>	Cluster node 02 FQDN	
<<var_clustermgmt_port>>	Port for cluster management	
<<var_vsadmin_password>>	Password for VS admin account	
<<var_vserver_mgmt_ip>>	Management IP address for Vserver	
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver	
<<var_node01_boot_lif_ip>>	Storage Cluster node 01 NFS Boot VLAN IP address	
<<var_node01_boot_lif_mask>>	Storage Cluster node 01 NFS Boot VLAN netmask	
<<var_node02_boot_lif_ip>>	Storage Cluster node 02 NFS Boot IP address	
<<var_node02_boot_lif_mask>>	Storage Cluster node 02 NFS Boot VLAN netmask	
<<var_node01_storage_data_lif_ip >>	Cluster node 01 Storage for HANA Data VLAN IP address	
<<var_node01_storage_log_lif_ip> >	Cluster node 01 Storage for HANA Log VLAN IP address	
<<var_node01_storage_data_lif_ma sk>>	Node 01 Storage for HANA Data VLAN netmask	
<<var_node01_storage_log_lif_mas k>>	Node 01 Storage for HANA Log VLAN netmask	
<<var_node02_storage_data_lif_ip >>	Cluster node 02 Storage for HANA Data VLAN IP address	
<<var_node02_storage_data_lif_ip >>	Cluster node 02 Storage for HANA Log VLAN IP address	

Variable	Description	Customer Implementation Value
<<var_node02_storage_data_lif_mask>>	Node 02 Storage for HANA Data VLAN netmask	
<<var_node02_storage_log_lif_mask>>	Node 02 Storage for HANA Data VLAN netmask	
<< var_node01_iscsi_A_IP>>	Cluster node 01 iSCSI A VLAN IP address	
<< var_node01_iscsi_B_IP>>	Cluster node 01 iSCSI B VLAN IP address	
<< var_node02_iscsi_A_IP>>	Cluster node 02 iSCSI A VLAN IP address	
<< var_node02_iscsi_B_IP>>	Cluster node 02 iSCSI B VLAN IP address	
<<var_backup_node01>>	NetApp Storage 01 for Backup	
<<var_backup_node02>>	NetApp Storage 02 for Backup	
<<var_host_boot_subnet>>	Boot VLAN IP range	
<<var_rule_index>>	Rule index number	
<<var_ftp_server>>	IP address for FTP server	
<<var_pxe_oob_IP>>	Out-of-band IP address for PXE boot Server	
<<var_pxe_oob_subnet>>	Out-of-band netmask for PXE boot Server	
<<var_pxe_boot_IP>>	Boot VLAN IP address for PXE boot Server	
<<var_pxe_boot_subnet>>	Boot VLAN netmask for PXE boot Server	
<<var_pxe_admin_IP>>	Admin Network IP address for PXE boot Server	
<<var_pxe_admin_subnet>>	Admin VLAN netmask for PXE boot Server	

Device Cabling

The information in this section is provided as a reference for cabling the network and storage components. The tables in this section contain details for the prescribed and supported configuration of the NetApp AFF A300 running NetApp ONTAP 9.3. For any modifications of this prescribed architecture, consult the NetApp [Interoperability Matrix Tool](#) (IMT). To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables show the out-of-band management ports connectivity into Management Pod Cisco Nexus 9000 Series Switches. To utilize a preexisting management infrastructure, the Management Ports cabling needs to be adjusted accordingly. These Management interfaces will be used in various configuration steps



In addition to the NetApp AFF A300 configurations listed in the tables below, other configurations can be used so long as the configurations match the descriptions given in the tables and diagrams in this section.

Figure 23 shows a cabling diagram for a FlexPod configuration using the Cisco ACI and NetApp AFF A300 storage systems with NetApp ONTAP. The NetApp Storage Controller and disk shelves are connected according to best practices for the specific storage controller and disk shelves as shown. For disk shelf cabling, refer to the [Universal SAS and ACP Cabling Guide](#).

Figure 23 FlexPod ACI for SAP HANA - Cable Connection Diagram




Table 5 through Table 15 provides the details of all the connections.

Table 5 Cisco Nexus 93180LC-EX Leaf A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180LC-EX Leaf A	Eth1/1	40GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	40GbE	Uplink to Customer Data Switch B	Any
	Eth1/3*	40GbE	Cisco APIC 1 via QSA adapter	VIC1225 - Port0
	Eth1/5*	40GbE	Cisco APIC 2 via QSA adapter	VIC1225-Port0
	Eth1/7*	40GbE	Cisco APIC 3 via QSA adapter	VIC1225-Port0
	Eth1/9	40GbE	NetApp controller A	e1a
	Eth1/10	40GbE	NetApp controller B	e1a
	Eth1/11	40GbE	Cisco UCS fabric interconnect A	Eth 1/29
	Eth1/12	40GbE	Cisco UCS fabric interconnect A	Eth 1/30
	Eth1/13	40GbE	Cisco UCS fabric interconnect B	Eth 1/29
	Eth1/14	40GbE	Cisco UCS fabric interconnect B	Eth1/30
	Eth1/15	40GbE	Cisco UCS fabric interconnect A	Eth1/33
	Eth1/16	40GbE	Cisco UCS fabric interconnect B	Eth1/33
	Eth1/21	40GbE	Mgmt Pod Cisco Nexus 9000 A	Eth2/1
	Eth1/22	40GbE	Mgmt PoD Cisco Nexus 9000 B	Eth2/1
	Eth1/29	40GbE	Cisco Nexus 9336PQ Spine Switch A	Eth1/4

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/31	40GbE	Cisco Nexus 9336PQ Spine Switch B	Eth1/3
	MGMT0	GbE	Mgmt PoD Nexus 9K's N2K FEX	Eth1/25



*For devices requiring 10GbE connectivity, use compliant QSA adapters

Table 6 Cisco Nexus 93180LC-EX Leaf B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180LC-EX Leaf B	Eth1/1	40GbE	Uplink to Customer Data Switch A	Any
	Eth1/2	40GbE	Uplink to Customer Data Switch B	Any
	Eth1/3	40GbE	Cisco APIC 1 via QSA adapter	VIC1225-Port1
	Eth1/5	40GbE	Cisco APIC 2 via QSA adapter	VIC1225-Port1
	Eth1/7	40GbE	Cisco APIC 3 via QSA adapter	VIC1225-Port1
	Eth1/8	40GbE	NetApp controller A	e1e
	Eth1/10	40GbE	NetApp controller B	e1e
	Eth1/11	40GbE	Cisco UCS fabric interconnect A	Eth1/31
	Eth1/12	40GbE	Cisco UCS fabric interconnect A	Eth1/32
	Eth1/13	40GbE	Cisco UCS fabric interconnect B	Eth1/31
	Eth1/14	40GbE	Cisco UCS fabric interconnect B	Eth1/32
	Eth1/15	40GbE	Cisco UCS fabric interconnect A	Eth1/34
	Eth1/16	40GbE	Cisco UCS fabric interconnect B	Eth1/34
	Eth1/21	40GbE	Mgmt Pod Cisco Nexus 9000 A	Eth2/2

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/22	40GbE	Mgmt PoD Cisco Nexus 9000 B	Eth2/2
	Eth1/29	40GbE	Cisco Nexus 9336PQ Spine Switch A	Eth1/3
	Eth1/31	40GbE	Cisco Nexus 9336PQ Spine Switch B	Eth1/4
	MGMT0	GbE	Mgmt PoD Nexus 9K's N2K FEX	Eth1/26



For devices requiring 10GbE connectivity, use compliant QSA adapters.

Table 7 Cisco ACI APIC-1 Controller Cabling Information

Local Device	VIC1225 Port	Connection	Remote Device	Remote Port
Cisco APIC 1	Port 0	10GbE	Cisco Nexus 93180LC-EX Leaf A	Eth1/3
	Port 1	10GbE	Cisco Nexus 93180LC-EX Leaf B	Eth1/3
	MGMT0	GBE	Mgmt PoD Nexus9K > N2k FEX	Eth1/13

Table 8 Cisco ACI APIC-2 Controller Cabling Information

Local Device	VIC1225 Port	Connection	Remote Device	Remote Port
Cisco APIC 2	Port 0	10GbE	Cisco Nexus 93180LC-EX Leaf A	Eth1/5
	Port 1	10GbE	Cisco Nexus 93180LC-EX Leaf B	Eth1/5
	MGMT0	GBE	Mgmt PoD Nexus9K > N2k FEX	Eth1/14

Table 9 Cisco ACI APIC-3 Controller Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco APIC 3	Port 0	10GbE	Cisco Nexus 93180LC-EX Leaf A	Eth1/7
	Port 1	10GbE	Cisco Nexus 93180LC-EX Leaf B	Eth1/7
	MGMT0	GBE	Mgmt PoD Nexus9K > N2k FEX	Eth1/15

Table 10 Cisco Nexus 9336PQ ACI Spine A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336PQ Spine A	Eth1/4	40GbE	Cisco Nexus 93180LC-EX Leaf A	Eth1/31
	Eth1/3	40GbE	Cisco Nexus 93180LC-EX Leaf B	Eth1/29
	MGMT0	GBE	Mgmt PoD Nexus9K > N2k FEX	Eth1/16

Table 11 Cisco Nexus 9336PQ ACI Spine B Cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336PQ Spine B	Eth1/3	40GbE	Cisco Nexus 93180LC-EX Leaf A	Eth1/29
	Eth1/4	40GbE	Cisco Nexus 93180LC-EX Leaf B	Eth1/31
	MGMT0	GBE	Mgmt PoD Nexus9K > N2k FEX	Eth1/17

Table 12 NetApp Controller-1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1	e0M	GbE	Mgmt PoD Nexus9K > N2k FEX	Eth1/21
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp controller 2	e0a
	e0c	10GbE	NetApp controller 2	e0c
	e1a	40GbE	Cisco ACI Leaf A	Eth1/9
	e1e	40GbE	Cisco ACI Leaf B	Eth1/9



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 13 NetApp Controller-2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2	e0M	GbE	Mgmt PoD Nexus9K > N2k FEX	Eth1/22
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp controller 1	e0a

Local Device	Local Port	Connection	Remote Device	Remote Port
	e0c	10GbE	NetApp controller 1	e0c
	e1a	40GbE	Cisco ACI Leaf A	Eth1/10
	e1e	40GbE	Cisco ACI Leaf B	Eth1/10



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 14 Cisco UCS Fabric Interconnect A - Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/17	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) A	IOM1/1
	Eth1/19	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) A	IOM 1/2
	Eth1/21	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) A	IOM 1/1
	Eth1/23	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) A	IOM 1/2
	Eth1/25	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) A	IOM 1/1
	Eth1/27	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) A	IOM 1/2
	Eth1/29	40GbE	Cisco ACI Leaf A	eth1/11
	Eth1/30	40GbE	Cisco ACI Leaf A	eth1/12
	Eth1/31	40GbE	Cisco ACI Leaf B	Eth1/11
	Eth1/32	40GbE	Cisco ACI Leaf B	Eth1/12
	Eth1/33	40GbE	Cisco ACI Leaf A	Eth1/15
	Eth1/34	40GbE	Cisco ACI Leaf B	Eth1/15
	MGMT0	GbE	Mgmt PoD N9k > N2k FEX	ETH1/1
	L1	GbE	Cisco UCS Fabric Interconnect B	L1
	L2	GbE	Cisco UCS Fabric Interconnect B	L2

Table 15 Cisco UCS Fabric Interconnect B - Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/17	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) B	IOM 1/1
	Eth1/19	40GbE	Cisco UCS Chassis 1 Fabric Extender (FEX) B	IOM 1/2
	Eth1/21	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) B	IOM 1/1
	Eth1/23	40GbE	Cisco UCS Chassis 2 Fabric Extender (FEX) B	IOM 1/2
	Eth1/25	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) B	IOM 1/1
	Eth1/27	40GbE	Cisco UCS Chassis 3 Fabric Extender (FEX) B	IOM 1/2
	Eth1/29	40GbE	Cisco ACI Leaf A	eth1/13
	Eth1/30	40GbE	Cisco ACI Leaf A	eth1/14
	Eth1/31	40GbE	Cisco ACI Leaf B	Eth1/13
	Eth1/32	40GbE	Cisco ACI Leaf B	Eth1/14
	Eth1/33	40GbE	Cisco ACI Leaf A	Eth1/16
	Eth1/34	40GbE	Cisco ACI Leaf B	Eth1/16
	MGMT0	GbE	Mgmt PoD N9k > N2k FEX	ETH1/2
	L1	GbE	Cisco UCS Fabric Interconnect B	L1
	L2	GbE	Cisco UCS Fabric Interconnect B	L2

Management Pod Cabling

Figure 24 FlexPod ACI – Management POD Cable Connection Diagram




Table 16 through Table 19 provides the details of the connections used for Management Pod. As described earlier, in this reference design the Management Pod is directly connected to FlexPod as shown in Figure 24.

Table 16 Cisco Nexus 9000-A Management Pod Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9000** Mgmt A	Eth2/1	40GbE	Cisco ACI Leaf A	Eth1/21
	Eth2/2	40GbE	Cisco ACI Leaf B	Eth1/21
	Eth1/5	10GbE	Cisco UCS C-220-A	Port 0
	Eth1/7	10GbE	Cisco UCS C-220-B	Port 0
	Eth1/9-12*	10GbE	Cisco Nexus 9000 Mgmt B - vPC Peer link	Eth1/9-12*

* The ports ETH1/9-12 can be replaced with E2/11 and E2/12 for 40G connectivity.

Table 17 Cisco Nexus 9000-B Management Pod Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9000** Mgmt B	Eth2/1	40GbE	Cisco ACI Leaf A	Eth1/22
	Eth2/2	40GbE	Cisco ACI Leaf B	Eth1/22
	Eth1/5	10GbE	Cisco UCS C-220-A	Port 1
	Eth1/7	10GbE	Cisco UCS C-220-B	Port 1
	Eth1/9-12*	10GbE	Cisco Nexus 9000 Mgmt A - vPC Peer link	Eth1/9-12*

* The ports ETH1/9-12 can be replaced with E2/11 and E2/12 for 40G connectivity.

Table 18 Cisco UCS C-Series Server-A*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-220-A	CIMC Port M	1GbE	Cisco Nexus 9000 Management A	Eth 1/17
	Port 0	10GbE	Cisco Nexus 9000 Management A	Eth 1/5
	Port 1	10GbE	Cisco Nexus 9000 Management B	Eth 1/5

Table 19 Cisco UCS C-Series Server-B

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-220-B	CIMC Port M	1GbE	Cisco Nexus 9000 Management B	Eth 1/17
	Port 0	10GbE	Cisco Nexus 9000 Management A	Eth 1/7
	Port 1	10GbE	Cisco Nexus 9000 Management B	Eth 1/7

*Validation setup had one management server. **Cisco Nexus 9396 switches were used in the Management Pod of the validation setup.

Management Pod Installation

This section describes the configuration of the Management Pod to manage the multiple FlexPod environments for SAP HANA. In this reference architecture, the Management Pod includes a pair of Cisco Nexus 9000 Switches in standalone mode for out of band management network and a pair of Cisco UCS C220 M5 Rack-Mount Servers. The rack-mount servers for management are built on VMware ESXi. ESXi hosts will run PXE boot server, VMware vCenter and Windows Jump Host for Management. The next sections outline the configurations of each component in the Management Pod.

Network Configuration for Management Pod


The following section provides a detailed procedure for configuring the Cisco Nexus 9000 Series Switches for the Management Pod. It is based on cabling plan described in the Device Cabling section. If the systems connected on different ports, configure the switches accordingly following the guidelines described in this section.



The configuration steps detailed in this section provides guidance for configuring the Cisco Nexus 9000 running release 6.1(2) within a multi-VDC environment.

Dual-Homed FEX Topology (Active/Active FEX Topology)

The dual-homed FEX (Active/Active) topology is supported with NX-OS 7.0(3)I5(2) and later using Cisco Nexus 9300 and Nexus 9300-EX Series switches. The following topology shows that each FEX is dual-homed with two Cisco Nexus 9300 Series switches. The FEX-fabric interfaces for each FEX are configured as a vPC on both peer switches. The host interfaces on the FEX appear on both peer switches.



Cisco Nexus 9000 Series Switches—Network Initial Configuration Setup

This section provides the steps for the initial Cisco Nexus 9000 Series Switch setup.

Cisco Nexus 9000 A

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



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.

```

----- Basic System Configuration Dialog VDC: 1 -----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

*Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

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

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

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

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

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

Enter the switch name : <<var_nexus_mgmt_A_hostname>>

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

Mgmt0 IPv4 address : <<var_nexus_mgmt_A_mgmt0_ip>>

Mgmt0 IPv4 netmask : <<var_nexus_mgmt_A_mgmt0_netmask>>

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

IPv4 address of the default gateway : <<var_nexus_mgmt_A_mgmt0_gw>>

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

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

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

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

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

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

NTP server IPv4 address : <<var_global_ntp_server_ip>>

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

The following configuration will be applied:
password strength-check
switchname <<var_nexus_mgmt_A_hostname>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_mgmt_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 2048 force
feature ssh

```

```

ntp server <<var_global_ntp_server_ip>>
copp profile strict
interface mgmt0
ip address <<var_nexus_mgmt_A_mgmt0_ip>> <<var_nexus_mgmt_A_mgmt0_netmask>>
no shutdown

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

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

[#####] 100%
Copy complete.

```

Cisco Nexus 9000 B

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



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.

```

----- Basic System Configuration Dialog VDC: 1 -----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

*Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

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

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

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

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

Enter the switch name : <<var_nexus_mgmt_B_hostname>>

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

Mgmt0 IPv4 address : <<var_nexus_mgmt_B_mgmt0_ip>>

Mgmt0 IPv4 netmask : <<var_nexus_mgmt_B_mgmt0_netmask>>

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

IPv4 address of the default gateway : <<var_nexus_mgmt_B_mgmt0_gw>>

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

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

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

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

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

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

NTP server IPv4 address : <<var_global_ntp_server_ip>>

```

```

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

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

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

The following configuration will be applied:
 password strength-check
 switchname <<var_nexus_mgmt_B_hostname>>
 vrf context management
 ip route 0.0.0.0/0 <<var_nexus_mgmt_B_mgmt0_gw>>
 exit
 no feature telnet
 ssh key rsa 2048 force
 feature ssh
 ntp server <<var_global_ntp_server_ip>>
 copp profile strict
 interface mgmt0
 ip address <<var_nexus_mgmt_B_mgmt0_ip>> <<var_nexus_mgmt_B_mgmt0_netmask>>
 no shutdown

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

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

[#####] 100%
Copy complete.

```

Enable Appropriate Cisco Nexus 9000 Series Switches - Features and Settings

Cisco Nexus 9000 A and Cisco Nexus 9000 B

To enable the IP switching feature and set default spanning tree behaviors, complete the following steps:

1. On each Nexus 9000, enter configuration mode:

```
config terminal
```

2. Use the following commands to enable the necessary features:

```
feature udld
feature lacp
feature vpc
feature interface-vlan
feature lldp
```

3. Configure spanning tree defaults:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
```

4. Save the running configuration to start-up:

```
copy run start
```

Create VLANs for Management Traffic

Cisco Nexus 9000 A and Cisco Nexus 9000 B

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

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

```
vlan <<var_oob_vlan_id>>
name OOB-Mgmt

vlan <<var_admin_vlan_id>>
name HANA-Admin

vlan <<var_boot_vlan_id>>
name HANA-Boot
```

Configure Virtual Port Channel Domain

Cisco Nexus 9000 A

To configure virtual port channels (vPCs) for switch A, complete the following steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain <<var_nexus_vpc_domain_mgmt_id>>
```

2. Make Nexus 9000A the primary vPC peer by defining a low priority value:

```
role priority 10
```

3. Use the management interfaces on the supervisors of the Nexus 9000s to establish a keepalive link:

```
peer-keepalive destination <<var_nexus_mgmt_B_mgmt0_ip>> source <<var_nexus_mgmt_A_mgmt0_ip>>
```

4. Enable following features for this vPC domain:

```
peer-switch
delay restore 150
peer-gateway
auto-recovery
```

Cisco Nexus 9000 B

To configure vPCs for switch B, complete the following steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain <<var_nexus_vpc_domain_mgmt_id>>
```

2. Make Cisco Nexus 9000 B the secondary vPC peer by defining a higher priority value than that of the Nexus 9000 A:

```
role priority 20
```

- Use the management interfaces on the supervisors of the Cisco Nexus 9000s to establish a keepalive link:

```
peer-keepalive destination <<var_nexus_mgmt_A_mgmt0_ip>> source <<var_nexus_mgmt_B_mgmt0_ip>>
```

- Enable following features for this vPC domain:

```
peer-switch
delay restore 150
peer-gateway
auto-recovery
```

Configure Network Interfaces for the VPC Peer Links

Cisco Nexus 9000 A

- Define a port description for the interfaces connecting to VPC Peer <<var_nexus_mgmt_B_hostname>>.

```
interface Eth1/9
description VPC Peer <<var_nexus_mgmt_B_hostname>>:1/9

interface Eth1/10
description VPC Peer <<var_nexus_mgmt_B_hostname>>:1/10

interface Eth1/11
description VPC Peer <<var_nexus_mgmt_B_hostname>>:1/11

interface Eth1/12
description VPC Peer <<var_nexus_mgmt_B_hostname>>:1/12
```

- Apply a port channel to both VPC Peer links and bring up the interfaces.

```
interface Eth1/9-12
channel-group 1 mode active
no shutdown
```

- Define a description for the port-channel connecting to <<var_nexus_mgmt_B_hostname>>.

```
interface Po1
description vPC peer-link
```

- Make the port-channel a switchport, and configure a trunk to allow Management VLANs.

```
switchport
switchport mode trunk
switchport trunk allowed vlan <<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>
```

- Make this port-channel the VPC peer link and bring it up.

```
spanning-tree port type network
vpc peer-link
no shutdown
```

Cisco Nexus 9000 B

- Define a port description for the interfaces connecting to VPC peer <<var_nexus_A_hostname>>.

```
interface Eth1/9
```

```

description VPC Peer <<var_nexus_mgmt_A_hostname>>:1/9

interface Eth1/10
description VPC Peer <<var_nexus_mgmt_A_hostname>>:1/10

interface Eth1/11
description VPC Peer <<var_nexus_mgmt_A_hostname>>:1/11

interface Eth1/12
description VPC Peer <<var_nexus_mgmt_A_hostname>>:1/12

```

2. Apply a port channel to both VPC peer links and bring up the interfaces.

```

interface Eth1/9-12
channel-group 1 mode active
no shutdown

```

3. Define a description for the port-channel connecting to <<var_nexus_A_hostname>>.

```

interface Po1
description vPC peer-link

```

4. Make the port-channel a switchport, and configure a trunk to allow Management VLANs.

```

switchport
switchport mode trunk
switchport trunk allowed vlan <<var_admin_vlan_id>>,<<var_boot_vlan_id>>

```

5. Make this port-channel the VPC peer link and bring it up.

```

spanning-tree port type network
vpc peer-link
no shutdown

```

6. Save the running configuration to start-up in both Nexus 9000s.

```

copy run start

```

Configure Network Interfaces to Cisco UCS C220 Management Server

Cisco Nexus 9000 A

1. Define a port description for the interface connecting to <<var_c220>>-A and <<var_c220>>-B.

```

interface Eth1/5
description << var_c220>>-A:P1

interface Eth1/7
description << var_c220>>-B:P1

```

2. Make a switchport, and configure a trunk to allow NFS, PXE, Management, VM traffic VLANs.

```

interface Eth1/5
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>,<<var_pxe_vlan_id>>

interface Eth1/7

```

```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>,<<var_pxe_vlan_id>>
```

Cisco Nexus 9000 B

1. Define a port description for the interface connecting to <<var_c220>>-A and <<var_c220>>-B.

```
interface Eth1/5
description << var_C220>>-A:P2

interface Eth1/7
description << var_C220>>-B:P2
```

2. Make a switchport, and configure a trunk to allow NFS, PXE, Management, VM traffic VLANs.

```
interface Eth1/5
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>,<<var_pxe_vlan_id>>

interface Eth1/7
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>,<<var_pxe_vlan_id>>
```

Direct Connection of Management Pod to FlexPod Infrastructure

This section describes the configuration steps for Cisco Nexus 9000 switches in the Management Pod **connected to each FlexPod instance's Leaf Switches.**

Cisco Nexus 9000 A

1. Define a port description for the interface connecting to <<var_nexus_A_hostname>>.

```
interface eth2/1
description <<var_nexus_A_hostname>>:1/21
```

2. Apply it to a port channel and bring up the interface.

```
interface eth2/1
channel-group 6 mode active
no shutdown
```

3. Define a port description for the interface connecting to <<var_nexus_B_hostname>>.

```
interface eth2/2
description <<var_nexus_B_hostname>>:1/21
```

4. Apply it to a port channel and bring up the interface.

```
interface eth2/2
channel-group 6 mode active
no shutdown
```

5. Define a description for the port-channel connecting to FlexPod Switch.

```
interface Po6
description <<var_nexus_A_hostname>>
```

6. Make the port-channel a switchport, and configure a trunk to allow all Management VLANs.

```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>,<<var_pxe_vlan_id>>
```

7. Make the port channel and associated interfaces spanning tree network ports.

```
spanning-tree port type network
```

8. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

9. Make this a VPC port-channel and bring it up.

```
vpc 6
no shutdown
```

10. Save the running configuration to start-up.

```
copy run start
```

Cisco Nexus 9000 B

1. Define a port description for the interface connecting to <<var_nexus_A_hostname>>.

```
interface eth2/1
description <<var_nexus_A_hostname>>:eth1/22
```

2. Apply it to a port channel and bring up the interface.

```
interface eth2/1
channel-group 6 mode active
no shutdown
```

3. Define a port description for the interface connecting to <<var_nexus_B_hostname>>.

```
interface eth2/2
description <<var_nexus_B_hostname>>:eth1/22
```

4. Apply it to a port channel and bring up the interface.

```
interface eth2/2
channel-group 6 mode active
no shutdown
```

5. Define a description for the port-channel connecting to FlexPod Switch.

```
interface Po6
description <<var_nexus_A_hostname>>
```

6. Make the port-channel a switchport, and configure a trunk to allow all Management VLANs.


```
switchport
switchport mode trunk
switchport trunk allowed vlan
<<var_admin_vlan_id>>,<<var_boot_vlan_id>>,<<var_oob_vlan_id>>,<<var_pxe_vlan_id>>
```

7. Make the port channel and associated interfaces spanning tree network ports.

```
spanning-tree port type network
```

8. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

9. Make this a VPC port-channel and bring it up.

```
vpc 6
no shutdown
```

10. Save the running configuration to start-up.

```
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink from the Management environment to connect to FlexPod SAP HANA environment. If an existing Cisco Nexus environment is present, Cisco recommends using vPCs to uplink the Cisco Nexus 9000 switches in the Management environment to the FlexPod SAP HANA environment. 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.

Management Server Installation

The Cisco UCS C220 M5 Server acts as a management server for this solution. It requires VMware ESXi 6.5 for the Cisco UCS C220 M5 Servers and for the PXE boot tasks, either a SLES or Red Hat Server. Windows based system can also be considered (optional) for these management servers.

Server Configuration

The Cisco UCS C220 M5 Rack-Mount Servers are recommended for use as management servers in the FlexPod environment.

Cisco Integrated Management Controller (CIMC) of Cisco UCS C220 M5 Servers and both the Cisco UCS VIC card ports must be connected to Cisco Nexus 9000 Series Switches in the management network, as defined in the Cabling Section. Three IP addresses are necessary for each of the server; one each for the CIMC, ESXi console and PXE boot VM networks.

CIMC Configuration

To configure the IP-Address on the CIMC, complete the following steps:

1. With a direct attached monitor and keyboard press F8 when the following screen appears:



2. Configure the CIMC as required to be accessible from the Management LAN.

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode                               NIC redundancy
Dedicated:      [X]                   None:           [X]
Shared LOM:     [ ]                   Active-standby: [ ]
Cisco Card:     [ ]                   Active-active:  [ ]
  Riser1:       [ ]                   VLAN (Advanced)
  Riser2:       [ ]                   VLAN enabled:   [ ]
  MLom:         [ ]                   VLAN ID:        1
Shared LOM Ext: [ ]                   Priority:        0
IP (Basic)
IPV4:           [X]                   IPV6:           [ ]
DHCP enabled    [ ]
CIMC IP:        192.168.76.91
Prefix/Subnet:  255.255.255.0
Gateway:        192.168.76.1
Pref DNS Server: 0.0.0.0
Smart Access USB
Enabled         [ ]
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F1>Additional settings
    
```

3. When connecting the CIMC to Management Switch, complete the following steps:
 - a. Choose Dedicated under NIC mode
 - b. Enter the IP address for CIMC which is accessible from the Management Network
 - c. Enter the Subnet mask for CIMC network
 - d. Enter the Default Gateway for CIMC network
 - e. Choose NIC redundancy as None
 - f. Enter the Default password for admin user under Default User (Basic) and Reenter password


Storage Configuration

To create a redundant virtual drive (RAID 1) on the internal disks to host ESXi and VMs, complete the following steps:




RAID1 for two internal disks in the Management server can be set up from the CIMC web Browser by completing the following steps:

1. Open a web browser and navigate to the Cisco C220-M5 CIMC IP address.
2. Enter admin as the user name and enter the administrative password, which was previously set.

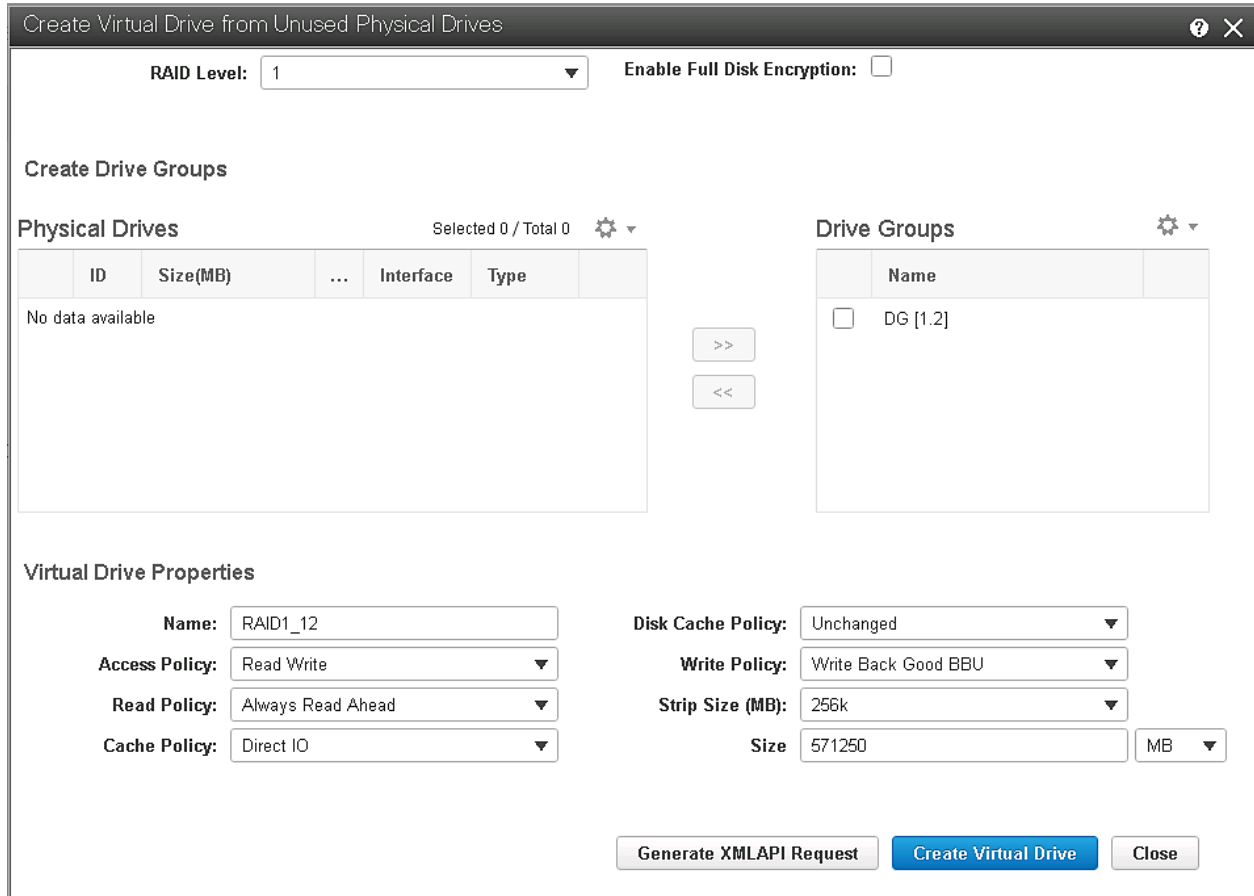


3. Click Login to log in to CIMC.
4. On the Navigation Pane click the Storage tab. Select Cisco 12G Modular Raid Controller.




5. Click Create Virtual Drive from Unused Physical Drives.

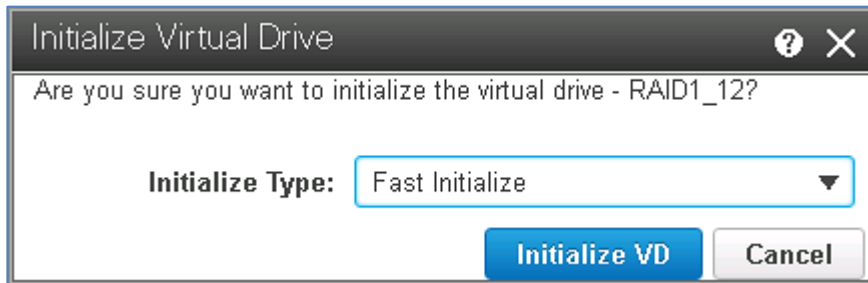
6. Choose RAID Level 1 and Select the Disks and click >> to add them in the Drive Groups.




7. Click Create Virtual Drive to create the virtual drive.
8. Click the Virtual Drive Info tab.
9. Select the Virtual Drive created and Click Initialize.



10. Click Initialize VD.



11. As a pre-requisite for ESXi installation, under Compute BIOS setting's Security sub-tab, make sure the Intel Trusted Execution Technology Support is Enabled.



VMware ESXi Installation

Install VMware ESXi 6.5d on the Cisco UCS M5 C-Series server and configure both Cisco UCS VIC interfaces as the ESX Management Network by completing the following steps.

Download Cisco Custom Image for ESXi 6.5a

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 the following link [Cisco ESXi 6.5U2 GA Install CD Download](#).
4. Click Download.
5. Save it to your destination folder.

VMware ESXi Hosts ESXi-Mgmt-01 and ESXi-Mgmt-02

To prepare the server for the OS installation, complete the following steps on each ESXi host:


1. On your Browser go to IP address Set for CIMC.
2. In the Navigation Pane Server > Summary.
3. Click Launch KVM Console.
4. Open with Java JRE installed.
5. Click the Virtual Media tab.
6. Click Map CD/DVD.
7. Browse to the ESXi installer ISO image file and click Open.
8. Select the Mapped checkbox to map the newly added image.
9. Under Power tab select Power Cycle System to reboot the server..

Install ESXi

Management Server ESXi-Mgmt-01 and ESXi-Mgmt-02

To install VMware ESXi on the local disk, 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 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 local disk which was previously created 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 existing partitions will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, clear the Mapped checkbox (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.
9. The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.
10. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Click Yes to unmap the image.
11. 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.

Configure Management Access

To configure the ESXi-Mgmt-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 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_oob_vlan_id>>` and press Enter.
6. From the Configure Management Network menu, select IP Configuration and press Enter.
7. Select the Set Static IP Address and Network Configuration option by using the space bar.
8. Enter the IP address for managing the first ESXi host: `<<var_vm_host_mgmt_01_ip>>`.

9. Enter the subnet mask for the first ESXi host.
10. Enter the default gateway for the first ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Select the IPv6 Configuration option and press Enter.
13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.
14. Select the DNS Configuration option and press Enter.
15. 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.



Repeat the above steps to configure the ESXi-Mgmt-02 ESXi host.

VMware ESXi Host ESXi-Mgmt-01

Set Up VMkernel Ports and Virtual Switch



Repeat the steps in this section for all the ESXi Hosts.

To set up the VMkernel ports and the virtual switches on the ESXi-Mgmt-01 ESXi host, complete the following steps:

1. From each Web client, select the host in the inventory.

2. Click the Networking in the main pane.
3. Click virtual switches on the folder tap.
4. Select the Add standard virtual switch configuration and click Edit.
5. Specify the Name - FlexPod, MTU - 9000 and up-link 1 (select the first enic interface) and click OK to finalize the setup for VM Network.

Add standard virtual switch - FlexPod	
Add uplink	
vSwitch Name	FlexPod
MTU	9000
Uplink 1	vmnic2 - Up, 40000 mbps
▼ Link discovery	
Mode	Listen
Protocol	Cisco discovery protocol (CDP)
▼ Security	
Promiscuous mode	<input type="radio"/> Accept <input checked="" type="radio"/> Reject
MAC address changes	<input type="radio"/> Accept <input checked="" type="radio"/> Reject
Forged transmits	<input type="radio"/> Accept <input checked="" type="radio"/> Reject

6. On the left, click Add Uplink.
7. Similarly add vmnic3 to the vSwitch and click Save.
8. Configure additional port groups on this new vSwitch.
9. Select Networking in the main pane.
10. Select Port groups in the Navigation tab.
11. Select Add port group.
12. For Network Label enter HANA-Boot.
13. Enter VLAN ID for PXE Boot.
14. Click Finish.

Add port group - HANA-Boot	
Name	HANA-Boot
VLAN ID	277
Virtual switch	FlexPod ▼
▶ Security	Click to expand

15. Add additional port groups for the Management network as well to the vSwitch.

16. Repeat the last section for the Mgmt network.

Add port group - Mgmt	
Name	Mgmt
VLAN ID	76
Virtual switch	FlexPod ▼
▶ Security	Click to expand

17. Click Add

Mount Required Datastores

For VMware ESXi Hosts ESXi-Mgmt-01 and ESXi-Mgmt-02, it is recommended to use Additional NetApp Storage for Management Pod for redundancy and failure scenarios. If you have NetApp storage for Management, then create a volume for datastores, create a VM Kernel port for storage, assign IP address and complete the following steps on each of the ESXi hosts:

1. From each vSphere Client, select the host in the inventory.
2. Click Storage in the Navigation pane.
3. From the Datastores area, click New datastore to open the New Datastore wizard.
4. Select Mount NFS datastore and click Next.
5. The wizard prompts for the location of the NFS server. Enter the IP address for NFS Storage Device.
6. Enter Volume path for the NFS share.

7. Enter mgmt_datastore_01 as the datastore name.
8. Click Next to continue with the NFS datastore creation.
9. Click Finish to finalize the creation of the NFS datastore.

Configure NTP on ESXi Hosts

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From each vSphere Client, select the host in the inventory.
2. Click the Configuration tab to enable configurations.
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 Options dialog box, complete the following steps:
 - a. Click General in the left pane, 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_global_ntp_server_ip>> as the IP address of the NTP server and click OK.
8. In the NTP Daemon Options dialog box, select the Restart NTP Service to Apply Changes checkbox and click OK.
9. 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.
10. The NTP server time may vary slightly from the host time.

FlexPod Cisco ACI Network Configuration for SAP HANA

The following sections provide a detailed procedures to configure the Cisco Application Centric Infrastructure (ACI) for SAP HANA environment. The network configuration in this section is based on the cabling details described in the Device Cabling section. Follow the guidelines provided in this section to configure the switches for FlexPod Cisco ACI for SAP HANA solution.

Cisco APIC Initial Configuration

This section provides the required configuration details for setting up Cisco ACI.

Cisco IMC Configuration

To configure the IP address on the Cisco IMC, complete the following steps:

1. With a direct attached monitor and keyboard press F8 when the following screen appears.

Figure 25 Cisco IMC Configuration



2. Configure the Cisco IMC as required to be accessible from the management LAN.

Figure 26 Cisco IMC Configuration Utility

```

CIMC Configuration Utility  Version 1.7  Cisco Systems, Inc.
*****
NIC Properties
NIC mode                               NIC redundancy
Dedicated:      [X]                   None:          [X]
Shared LOM:     [ ]                   Active-standby:[ ]
Cisco Card:     [ ]                   Active-active: [ ]
Shared LOM Ext: [ ]

IPV4 (Basic)                            Factory Defaults
DHCP enabled:   [ ]                   CIMC Factory Default:[ ]
CIMC IP:        192.168.76.104        Default User (Basic)
Subnetmask:     255.255.255.0         Default password:
Gateway:        192.168.76.1         Reenter password:

VLAN (Advanced)                         Port Profile
VLAN enabled:   [ ]                   Reset:         [ ]
VLAN ID:        1                     Name:
Priority:        0

Port Properties
Auto Negotiation: [X]
Speed[1000/100 Mbps]: 100
Duplex mode[half/full]: full
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit

```

3. When connecting the Cisco IMC to Management Switch, complete the following:
 - a. Choose Dedicated under NIC model.
 - b. Enter the IP address for CIMC which is accessible from the Management Network.
 - c. Enter the Subnet mask for CIMC network.
 - d. Enter the Default Gateway for CIMC network. Choose NIC redundancy as None.
 - e. Enter the Default password for admin user under Default User (Basic) and Reenter password.
 - f. Press F10 to save your configuration.
 - g. Press F5 to refresh your settings after 45 seconds to see the changes.
 - h. Press ESC to exit and continue booting with the new settings.


Configure Cisco APIC Power State

To power on the APIC controller to restore power, complete the following steps:

1. Browse to https://<cimc_ip_address>.
2. Log in using admin as the username and use the password defined during Cisco IMC setup.
3. Click on the Server tab on the left pane of the Cisco UCS Manager.

4. Click the Power Policy.
5. Select Restore Last State from the drop down list of Power Restore Policy.
6. Click Save Changes.

Figure 27 APIC Cisco IMC Power Policy



Configure Cisco IMC NTP

1. Browse to https://<cimc_ip_address>.
2. Log in using admin as username and use the password defined during CIMC setup.
3. Click the Admin tab.
4. Click Network.
5. Click NTP Settings.
6. Check Enable NTP.
7. Enter the IP address of one or more NTP servers.
8. Click Save Changes.

Cisco APIC Initial Configuration Setup

1. Log into the APIC CIMC using a web browser and launch the KVM.
2. Browse to https://<cimc_ip_address>.
3. Log in using admin as username and use the password defined during Cisco IMC setup.
4. From the Server tab, select Summary and click Launch KVM Console.
5. KVM application will be launched and initial APIC setup screen should be visible.

Figure 28 APIC Configuration

```

09 blocks
dev/mapper/vg_ifc0-data2: 84/11927552 files (0.8% non-contiguous), 824678/47718
08 blocks
[ OK ]
remounting root filesystem in read-write mode: [ OK ]
mounting local filesystems: [ OK ]
enabling /etc/fstab swaps: [ OK ]

Press any key to continue...
Starting Setup Utility

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.

Press Enter at anytime to assume the default values. Use ctrl-c
at anytime to restart from the beginning.

Cluster configuration ...
Enter the fabric name [ACI Fabric]: _

```

6. Press return to select the default value for Enter the fabric name. This value can be changed if desired.
7. Press return to select the default value of three for Enter the number of controllers in the fabric. While the fabric can operate with a single APIC, 3 APICs are recommended for redundancy. While a minimum of 3 nodes are a minimum for production use-cases, up to 5 nodes may be needed in large installations.
8. Press return to select the default value for the POD ID as 1.
9. Press return to select the default value NO for the standby controller question.
10. Enter the controller number currently being set up under Enter the controller ID (1-3). Enter the value 1 and press enter.
11. Press return to retain the default value for controller name apic1. This can be changed if desired.
12. Press return to select the default pool under Enter the address pool for TEP addresses. If the network is already in use, please choose a different range.
13. Enter a unique unreserved VLAN for Enter the VLAN id for infra network.
14. Press return to select the default range for Enter address pool for BD multicast addresses.
15. Enter appropriate values for the out of band management network configuration. The out of band management IP address will be used to access the APIC from client browsers.
16. Press Y to enforce a strong password.
17. Enter the admin password (controller 1 only).
18. Press return to accept the configuration without changes.
19. Let the APIC complete its boot process.

Figure 29 APIC Initial Setup

```

Cluster configuration ...
  Enter the fabric name [ACI Fabric1]: flexpod-hana
  Enter the fabric ID (1-128) [1]:
  Enter the number of active controllers in the fabric (1-9) [3]:
  Enter the POD ID (1-9) [1]:
  Is this a standby controller? [NO]:
  Enter the controller ID (1-3) [1]:
  Enter the controller name [apic1]:
  Enter address pool for TEP addresses [10.0.0.0/16]:
  Note: The infra VLAN ID should not be used elsewhere in your environment
        and should not overlap with any other reserved VLANs on other platforms.
  Enter the VLAN ID for infra network (1-4094): 4093
  Enter address pool for BD multicast addresses (GIPO) [225.0.0.0/15]:

Out-of-band management configuration ...
  Enable IPv6 for Out of Band Mgmt Interface? [N]:
  Enter the IPv4 address [192.168.10.1/24]: 192.168.76.50/24
  Enter the IPv4 address of the default gateway [None]: 192.168.76.1
  Enter the interface speed/duplex mode [auto]:

admin user configuration ...
  Enable strong passwords? [Y]:

```



When APIC-1 boots up for the first time, it might take up to 5 minutes to allow login using the admin password set during the setup procedure. If something went wrong during the setup, APIC does allow login using a special user called **rescue-user**. If admin password was never set or was not setup properly, rescue-user will allow access to APIC without any password. If an admin password was set previously, use rescue-user with the admin password.

-
20. Repeat steps 1 through 18 for APIC controllers APIC 2 and APIC 3. Make sure to select the controller IDs and controller name accordingly for them.



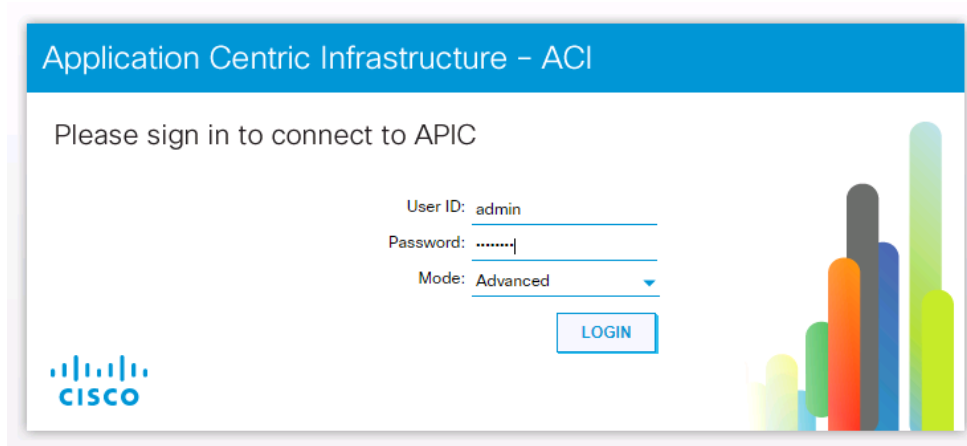
Only APIC 1 is configured with a password. APIC 2 and 3 will get the password from APIC 1 once the Cisco ACI fabric is configured.

Cisco ACI Fabric Discovery

The APIC is responsible for fabric discovery. It manages the device addressing. The fabric discovery happens via Link Layer Discovery Protocol [LLDP].


1. Log in to the Cisco APIC GUI using a web browser:
2. Browse to <https://<Out of Band IP address of APIC 1>>.
3. Log in using admin as username and use the password defined during initial setup.

Figure 30 APIC Login Screen



4. Click FABRIC from the top bar. Under INVENTORY, expand Fabric Membership.
5. At least one of the leaves should be visible. The Leaf node connected to active port of bonded interface of the APIC controller-1 is the first to be discovered via LLDP.

Figure 31 APIC Fabric Membership



6. Log in to the leaf using console connection (admin/<no password needed>) and use the serial number to identify discovered leaf (Leaf-1 or Leaf-2 in the physical setup).


```
switch# show inventory

NAME: "Chassis", DESCR: "Nexus C93180LC-EX Chassis"

PID: N9K-C93180LC-EX , VID: V01, SN: FDO204917D8
```


7. Double click the identified leaf description on the right hand side and assign 101 as NODE ID value and NODE NAME <device name>. Click UPDATE.

Figure 32 Update Device Parameter



8. As the fabric discovery continues, the leaf discovers the spine switches it is connected to and they are populated under the Fabric Membership window. Repeat Step 6 to assign the NODE ID and NODE NAME to these spine switches.
9. When the spine switches are configured, they discover the remaining leaf switches they are connected to the information is populated under the Fabric Membership window. Repeat Step 6 to assign the NODE ID and NODE NAME to these leaf switches [in our case the remaining leaf switch].
10. When the NODE ID and NODE NAME values are assigned, APIC assigns IP addresses from TEP, the pool defined during initial setup.
11. Fabric now self assembles and the discover process is complete.

Figure 33 Fabric Membership



Configure NTP for Cisco ACI Fabric

To configure NTP server for Cisco ACI fabric, complete the following steps:

1. Click FABRIC and select FABRIC POLICIES under the sub-menu.
2. Expand Pod Policies in the left pane and then expand Policies.
3. Right-click Date and Time Policy and select Create Date and Time Policy.
4. In the menu box, enter NTP as the policy name; select the Administrative State as enabled and the Authentication State as disabled.
5. Click NEXT.

Figure 34 ACI Fabric NTP Configuration

Create Date And Time Policy

STEP 1 > Identity

1. Identity 2. NTP Servers

Specify the information about the Date/Time Policy

Name: NTP

Description: optional

Administrative State: disabled enabled

Authentication State: disabled enabled

PREVIOUS NEXT CANCEL

6. On the NTP Servers window, Click + to add NTP servers.
7. Enter the IP address of the NTP server in the name field and choose the option default (Out-of-Band) for Management EPG.
8. Click OK.

Figure 35 NTP Server Configuration

Create Providers

Specify the information about the NTP Server

Name: 192.168.76.6

Description: optional

Preferred:

Minimum Polling Interval: 4

Maximum Polling Interval: 6

Management EPG: select an option

OK CANCEL

9. Click FINISH.

POD Policy Group

To create a POD Policy, complete the following steps:

1. Click FABRIC and select FABRIC POLICIES under the sub-menu.
2. Expand POD Policies in the left pane.
3. Right-click the Policy Groups and then select Create POD Policy Group.
4. Enter HANA-PoD as the policy name.
5. Select the option NTP from the Date Time Policy drop-down list.
6. Select default for BGP Route Reflector Policy from the drop-down list.
7. Click SUBMIT.

Figure 36 Create POD Policy

Create Pod Policy Group

Specify the Policy Group properties

Name: HANA-PoD

Description: optional

Date Time Policy: NTP

ISIS Policy: select a value

COOP Group Policy: select a value

BGP Route Reflector Policy: default

Management Access Policy: select a value

SNMP Policy: select a value

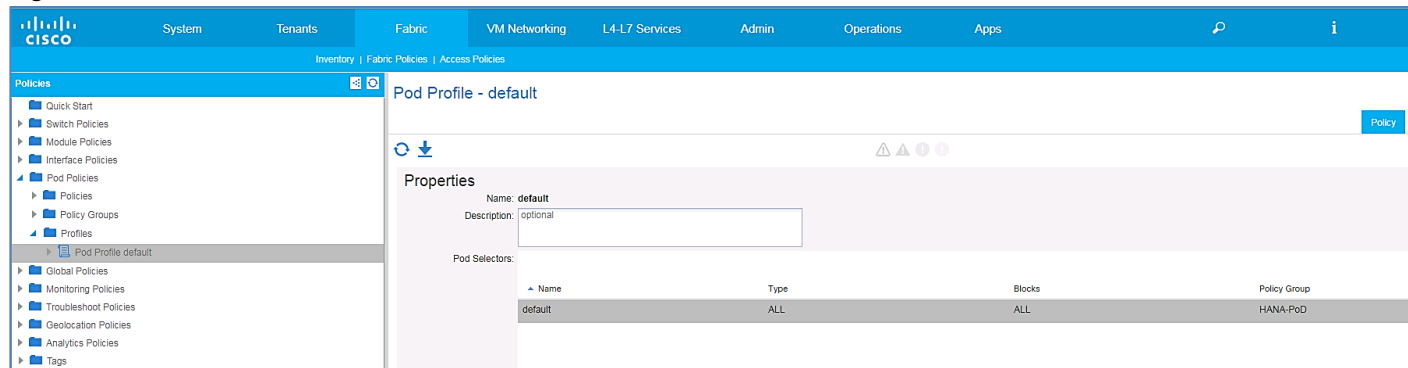
SUBMIT CANCEL

To set default POD for the configuration, complete the following steps:

1. Click FABRIC and select FABRIC POLICIES under the sub-menu.
2. Expand the POD Policies and then expand Profiles.
3. Click Pod Profile default.

4. Select HANA-PoD for the Fabric Policy Group from the drop-down list.
5. Click SUBMIT.

Figure 37 POD Profile – default



Configuring Access Policies

Access policies configure external-facing interfaces that connect to devices such as hosts, network attached storage or layer 2 switches for example, Cisco UCS Fabric Interconnects. Access policies enable the configuration of port channels and virtual port channels, protocols such as Link Layer Discovery Protocol (LLDP), Cisco Discovery Protocol (CDP), or Link Aggregation Control Protocol (LACP).

As observed in the reference architecture, all the devices namely Cisco UCS Fabric Interconnects, NetApp controllers and management PoD Nexus switches connect to leaf switches using vPC. The leaf switches do not have any direct connected end devices [either in access mode or in Port Channel configuration].

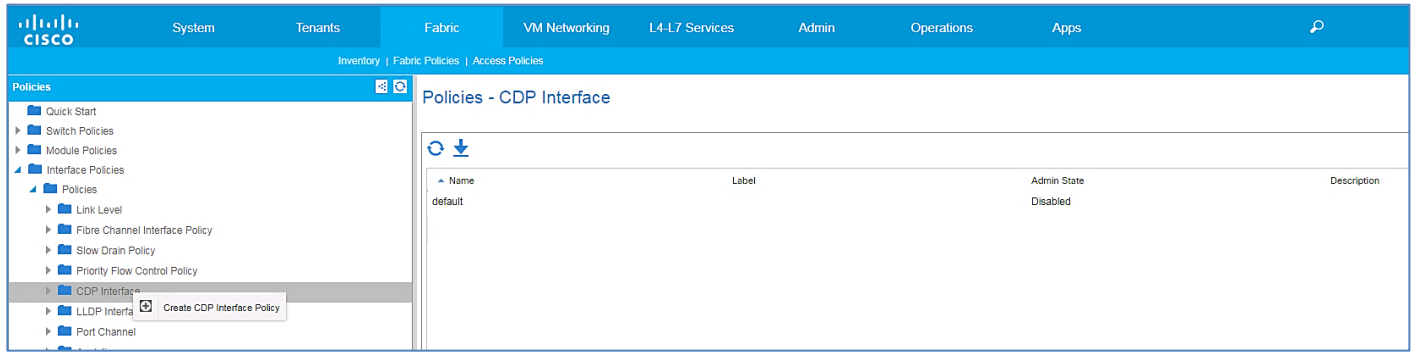
In this section, initially Interface Policies that configure various protocol options such as Link Level, CDP, LLDP and Port Channel are defined. The interface profiles [representing the switch interfaces] along with Interface Policies together map to Interface Policy Groups [IPGs] to define the vPC configurations that leaf switches use to connect to end devices.

Further VLAN pools – policies defining ID ranges used for VLAN encapsulation, along with connectivity options such as Physical Domain – single management domain that is the scope for policy enforcement, and Attachable Access Entity Profile [AAEP] – a template to deploy policies on a designated set of leaf ports are defined.

CDP Interface Policies

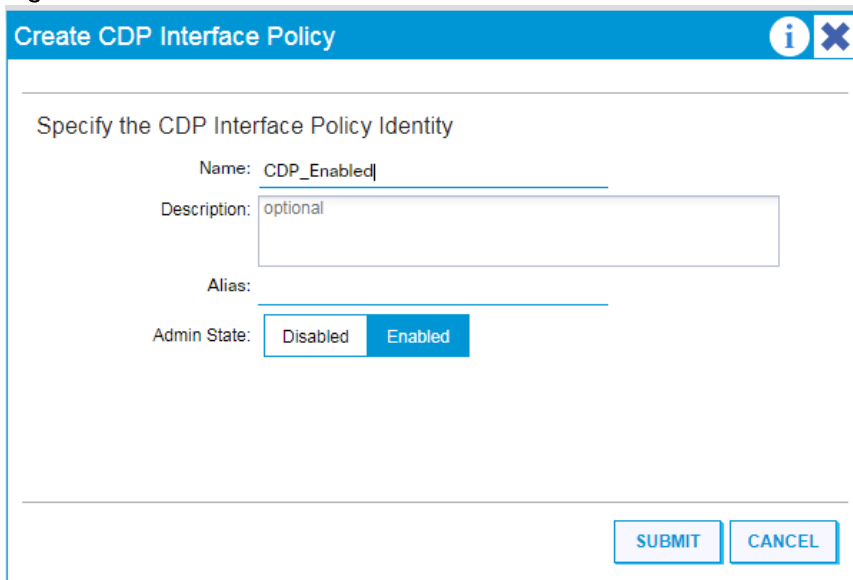
1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. From the left menu bar, expand Interface Policies.
3. Expand Policies. A **'default' policy with CDP disabled already exists**.
4. Right-click the CDP Interface and select Create CDP Interface Policy.

Figure 38 Access Policies - CDP Policy Creation



5. In the menu box, enter CDP_Enabled as the policy name and select Admin State as Enabled.

Figure 39 Access Policies – CDP_Enabled



6. Click SUBMIT.

7. Right-click and select Create CDP Interface Policy again

8. In the menu box, enter CDP_Disabled as the policy name and select Admin State as Disabled.

Figure 40 Access Policies – CDP Disabled

Create CDP Interface Policy

Specify the CDP Interface Policy Identity

Name: CDP_Disabled

Description: optional

Alias:

Admin State: Disabled Enabled

SUBMIT CANCEL

9. Click SUBMIT.

10. In summary – we created two policies one with CDP enabled and another disabling it.

Inventory | Fabric Policies | Access Policies

Policies - CDP Interface

Name	Label	Admin State	Description
CDP_Disabled		Disabled	
CDP_Enabled		Enabled	
default		Disabled	

LLDP Interface Policies

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. From the left menu bar, expand Interface Policies.
3. Expand Policies. A **'default' policy with LLDP enabled for both Receive and Transmit states** already exists.
4. Right-click LLDP Interface and select Create LLDP Interface Policy.

5. In the menu box, enter LLDP_Enabled as the policy name and set both Transmit State and Receive State Enabled.

Figure 41 Access Policies - LLDP Enabled

The screenshot shows a dialog box titled "Create LLDP Interface Policy". The main heading is "Specify the LLDP Interface Policy Properties". The fields are as follows:

- Name: LLDP_Enabled
- Description: optional
- Alias: (empty)
- Receive State: Enabled (selected)
- Transmit State: Enabled (selected)

Buttons: SUBMIT, CANCEL

6. Click SUBMIT.
7. Similarly create LLDP_Disabled with both Transmit and Receive state set to disabled as shown below.

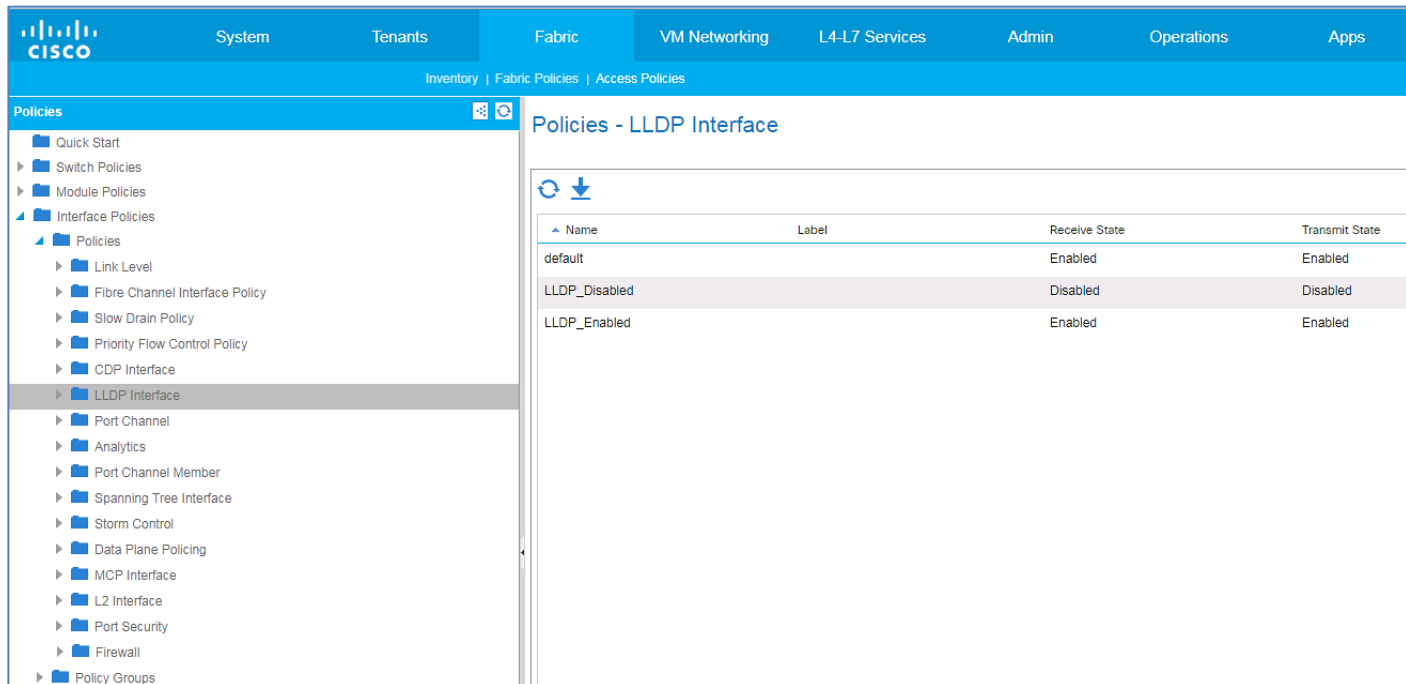
Figure 42 Access Policies - LLDP Disabled

The screenshot shows a dialog box titled "Create LLDP Interface Policy". The main heading is "Specify the LLDP Interface Policy Properties". The fields are as follows:

- Name: LLDP_Disabled
- Description: optional
- Alias: (empty)
- Receive State: Disabled (selected)
- Transmit State: Disabled (selected)

Buttons: SUBMIT, CANCEL

8. In summary we created two LLDP policies one with states enabled and another set to disabled.



LACP Interface Policies

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. From the left menu bar, expand Interface Policies.
3. Expand Policies. A 'default' policy with Mode set to 'Static Channel - Mode On' already exists.
4. Right-click Port Channel Policies and select Create LACP Policy.
5. In the menu box, enter LACP_Active as the policy name and select LACP Active for Mode. Leave remaining options as default.

Figure 43 Access Policies – LACP Active

6. Click SUBMIT.

You now have a LACP_Active policy.

Name	Label	Control	Mode	Minimum Links	Maximum Links
default			Static Channel - Mode On	1	16
LACP_Active		Fast Select Hot Standby Ports Graceful Convergence Suspend Individual Port	LACP Active	1	16

Link Level Policy

To create link level policies for the interfaces, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. From the left menu bar, expand Interface Policies.

3. Expand Policies. A **'default'** policy with speed set to **'inherit'** already exists.
4. Right-click Link Level and select Create Link Level Policy.
5. Enter 40GB as the name of Link Level Policy.
6. Select the speed to be 40 Gbps. Leave remaining options as default.

7. Click SUBMIT.



Depending on the supported bandwidth of end devices that connect to leaf switches and specific use-case, you may need to create Link Level policies with 10G or 25G link speed. However, the current setup being end to end 40G, a lone 40G Link Level policy suffices.

VLAN Pool for SAP HANA

To create VLANs in ACI Fabric for SAP HANA, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. From the left menu bar, expand Pools.
3. Right-click VLAN and select Create VLAN Pool.
4. Enter HANA-VLANs as the Pool name and select Static Allocation.
5. Under Encap Blocks click + to add a VLAN range.
6. Enter the VLAN range to be used for SAP HANA.

Figure 44 VLAN Pool Block Range

Specify the Encap Block Range

Type: VLAN

Range: VLAN 75 - VLAN 399

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

OK CANCEL

7. Click OK

Figure 45 VLAN Pool Identity

Specify the Pool identity

Name: HANA-VLANs

Description: optional

Allocation Mode: Dynamic Allocation Static Allocation

Encap Blocks:

VLAN Range	Allocation Mode
[75-399]	Inherit allocMode from parent

SUBMIT CANCEL

8. Click SUBMIT.

Domain Configuration

To create Physical Domain in the Cisco ACI Fabric for SAP HANA, complete the following steps:

1. Click FABRIC and select Access Policies under the sub-menu.
2. From the left pane, expand Physical and External Domains.
3. Right-click the Physical Domains and select Create Physical Domain.
4. Enter HANA for the physical domain name field.

5. Select HANA-VLANs for the VLAN Pool field from the drop-down list.

Figure 46 Physical Domain VLAN Properties

Specify the domain name and the VLAN Pool

Name: HANA

Associated Attachable Entity Profile: select a value

VLAN Pool: HANA-VLANs(static)

Security Domains:

Select	Name	Description
<input type="checkbox"/>		

SUBMIT CANCEL

6. Click SUBMIT



Associated Attachable Entity Profile will be added later and bound to the Physical Domain.

Attachable Access Entity Profile Configuration

To configure the Attachable Access Entity Profile, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand Global Policies on the left pane.
3. Right-click the Attachable Access Entity Profile and select Create Attachable Access Entity Profile.
4. Enter HANA-AAEP as the profile name.
5. Under Domains, click the + symbol to add Domain Profile.
6. From the Domain Profile drop-down list select HANA (Physical) and click Update.

Figure 47 Attachable Access Entity Profile Domains

Create Attachable Access Entity Profile

STEP 1 > Profile

1. Profile 2. Association To Interfaces

Specify the name, domains and infrastructure encaps

Name: HANA-AAEP

Description: optional

Enable Infrastructure VLAN:

Domains (VMM, Physical or External) To Be Associated To

Interfaces:

Domain Profile	Encapsulation
HANA (Physical)	

UPDATE CANCEL

EPG DEPLOYMENT (All Selected EPGs Will Be Deployed On All The Interfaces Associated.)

Application EPGs	Encap	Primary Encap	Mode
------------------	-------	---------------	------

PREVIOUS NEXT CANCEL

7. Click NEXT.

Figure 48 Attachable Access Entity Profile Interface

Create Attachable Access Entity Profile

STEP 2 > Association To Interfaces

1. Profile 2. Association To Interfaces

Select the interfaces

Interface Policy Group	Type	Associated Attachable Access Entity Profile	Switches / Fexes	Interfaces	Select Interfaces

PREVIOUS FINISH CANCEL

8. Click FINISH.



EPG information addition of step 1 and Interfaces association of step 2 to the AAEP will be done later.

Leaf Interface Policy Groups

This section describes the configuration of Interface Policy Groups [IPG] required for the vPC connections from Cisco ACI Leaf switches to Cisco UCS Fabric Interconnects, NetApp controllers and Management PoD Nexus switches.

VPC Configuration for Cisco UCS

This section describes the configuration of policies required for VPC connection from ACI Leaf switches to the Cisco UCS FIs.

Policy Group Configuration for VPC Connectivity Cisco UCS

To create VPC interface policies for the interfaces connecting to Cisco UCS Fabric Interconnect A, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand the Interface Policies in the left pane.

3. Expand Policy Groups and right-click the Leaf Policy Groups and select Create VPC Interface Policy Group.
4. Enter vPC-FiA-11 for the Name of Policy group.
5. Select 40GB for Link level Policy from the drop-down list.
6. Select CDP_Enabled for CDP Policy from the drop-down list.
7. Select LLDP_Disabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for PortChannel Policy from the drop-down list.
9. Select HANA-AAEP for Attached Entity Profile from the drop-down list.

Figure 49 VPC Interface Policy Group for Fabric Interconnect A

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: **vPC-FiA-11**

Description: optional

Link Aggregation Type: Port Channel **VPC**

Link Level Policy: 40GB

CDP Policy: CDP_Enabled

MCP Policy: select a value

LLDP Policy: LLDP_Disabled

STP Interface Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

Port Channel Policy: LACP_Active

Monitoring Policy: select a value

Storm Control Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Attached Entity Profile: HANA-AAEP

Connectivity Filters:

Switch IDs	Interfaces

SUBMIT CANCEL

10. Click SUBMIT.

To create VPC interface policies for the interfaces connecting to UCS Fabric Interconnect B, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand Interface Policies in the left pane.
3. Expand Policy Groups and right-click the Leaf Policy Groups and select Create VPC Interface Policy Group.
4. Enter vPC-FiB-12 for the Name of Policy group.
5. Select 40GB for Link level Policy from the drop-down list.
6. Select CDP_Enabled for CDP Policy from the drop-down list.
7. Select LLDP_Disabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for LACP Policy from the drop-down list.
9. Select HANA-AAEP for Attached Entity Profile from the drop-down list.
10. Click SUBMIT.

Figure 50 vPC Interface Policy Group for Fabric Interconnect B

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: **vPC-FiB-12**

Description: optional

Link Aggregation Type: **Port Channel** **VPC**

Link Level Policy: 40GB

CDP Policy: CDP_Enabled

MCP Policy: select a value

LLDP Policy: LLDP_Disabled

STP Interface Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

Port Channel Policy: LACP_Active

Monitoring Policy: select a value

Storm Control Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Attached Entity Profile: HANA-AAEP

Connectivity Filters:

Switch IDs	Interfaces

SUBMIT **CANCEL**

To create separate IPG for vPC connecting to UCS Fabric Interconnects for exclusive SAP HANA backup network usage, complete the following steps:

Cisco UCS Fabric Interconnect A

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand Interface Policies in the left pane.
3. Right-click the Policy Groups and select Create VPC Interface Policy Group.
4. Enter FI-A-PO21 for the Policy group name.
5. Select 10GB for Link level Policy from the drop-down list.

6. Select CDP_Enabled for CDP Policy from the drop-down list.
7. Select LLDP_Disabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for LACP Policy from the drop-down list.
9. Select HANA-Physical for Attached Entity Profile from the drop-down list.

Figure 51 IPG config for vPC - Exclusive Backup Network Usage FI-A

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: vPC-FiA-21

Description: optional

Link Level Policy: 40GB

CDP Policy: CDP_Enabled

MCP Policy: select a value

LLDP Policy: LLDP_Disabled

STP Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

Attached Entity Profile: HANA-AAEP

Connectivity Filters:

Switch IDs	Interfaces

Port Channel Policy: LACP_Active

Monitoring Policy: select a value

SUBMIT CANCEL

10. Click SUBMIT.

Cisco UCS Fabric Interconnect B

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand Interface Policies in the left pane.

3. Right-click the Policy Groups and select Create VPC Interface Policy Group.
4. Enter vPC-FiB-22 for the Policy group name.
5. Select 40GB for Link level Policy from the drop-down list.
6. Select CDP_Enabled for CDP Policy from the drop-down list.
7. Select LLDP_Disabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for LACP Policy from the drop-down list.
9. Select HANA-AAEP for Attached Entity Profile from the drop-down list.

Figure 52 IPG config for vPC - Exclusive Backup Network Usage FI-B

Create VPC Interface Policy Group
i X

Specify the Policy Group identity

Name: vPC-FiB-22

Description:

Link Level Policy: 40GB ▼ □

CDP Policy: CDP_Enabled ▼ □

MCP Policy: select a value ▼

LLDP Policy: LLDP_Disabled ▼ □

STP Interface Policy: select a value ▼

L2 Interface Policy: select a value ▼

Port Security Policy: select a value ▼

Egress Data Plane Policing Policy: select a value ▼

Ingress Data Plane Policing Policy: select a value ▼

Priority Flow Control Policy: select a value ▼

Fibre Channel Interface Policy: select a value ▼

Slow Drain Policy: select a value ▼

Attached Entity Profile: HANA-AAEP ▼ □

Connectivity Filters: x +

Switch IDs	Interfaces

Port Channel Policy: LACP_Active ▼ □

Monitoring Policy: select a value ▼

SUBMIT
CANCEL

10. Click SUBMIT.


Interface Profile Configuration

This section describes the procedure to create the interface profiles that will inherit the associated policy groups.

To create the interface profiles for ports connecting to Cisco UCS Fabric Interconnect A, complete the following steps:


1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand the Interface Policies in the left pane.
3. Expand Profiles and right-click Leaf Profiles and select Create Leaf Interface Profile.

Figure 53 Leaf Interface Profile



4. Enter FI-A-PO11 as the Interface Profile name.
5. For Interface Selectors, click the + symbol to add interfaces.

Figure 54 Interface Profile Identity



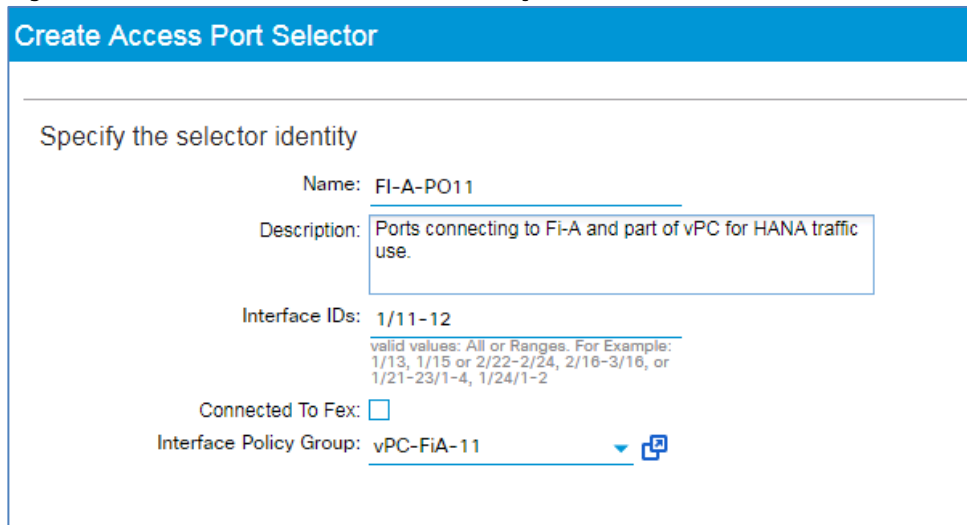
6. Enter FI-A-PO11 for port selector identity name and 1/11,1/12 as the Interface IDs.



These are ports used which is connected to UCS Fabric Interconnect A as per the Device Cabling section.

7. Select vPC-FiA-11 for Interface Policy Group from the drop-down list.

Figure 55 Interface Port Selector Identity

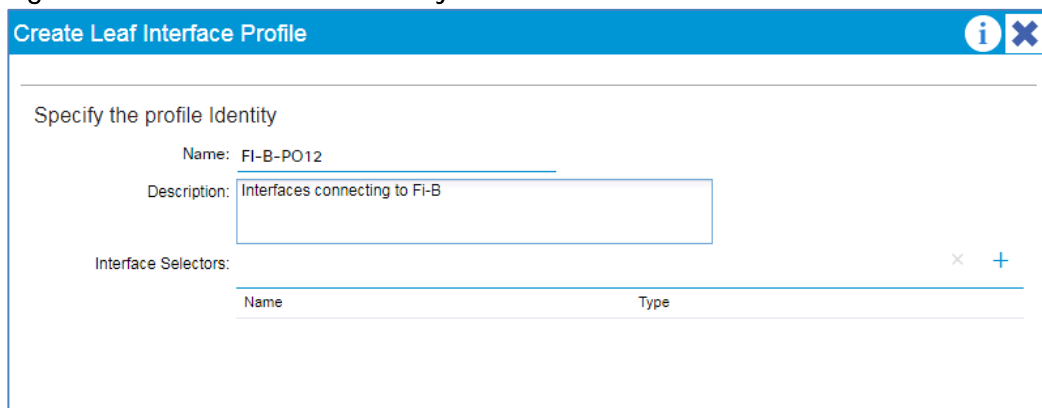


8. Click OK and then click SUBMIT.

Similarly, create interface profiles for ports connecting to Cisco UCS Fabric Interconnect B:

1. Right-click Leaf Profiles in the left pane and select Create Leaf Interface Profile.
2. Enter FI-B-PO12 as the Interface Profile name for Interface Selectors, click the + symbol to add interfaces.


Figure 56 Interface Profile Identity



3. Enter Fi-B-PO12 for port selector identity name and 1/13,1/14 for the Interface IDs.

4. For the Interface Policy Group select vPC-FiB-12 from the drop-down list.

Figure 57 Interface Port Selector Identity

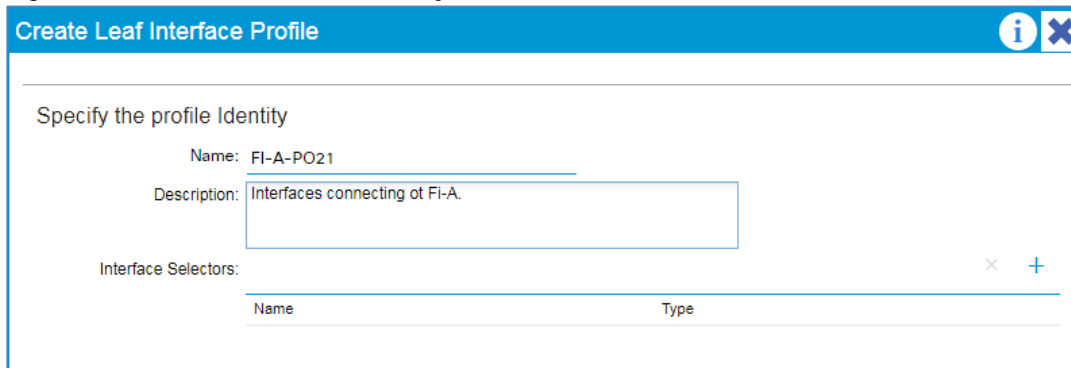


5. Click OK and then click SUBMIT.

To configure the interfaces planned for the SAP HANA backup network, complete the following steps:

1. Right-click the Leaf Profiles in the left pane and select Create Leaf Interface Profile.
2. Enter FI-A-PO21 as the Interface Profile name.

Figure 58 Interface Profile Identity



3. For Interface Selectors, click the + symbol to add interfaces.
4. Enter FI-A-PO21 for port selector identity name and 1/15 for the Interface ID.
5. Select vPC-FiA-21 for the Interface Policy Group from the drop-down list.

Figure 59 Interface Port Selector Identity

6. Click OK and then click SUBMIT.

Cisco UCS FI-B

1. Right-click the Profiles and select Create Interface Profile.
2. Enter FI-B-PO22 as the Interface Profile name.
3. For Interface Selectors, click the + symbol to add interfaces.

Figure 60 Interface Profile Identity

Name	Type

4. Enter FI-B-PO22 for port selector identity name and 1/16 as the Interface ID.
5. Select vPC-FiB-22 for the Interface Policy Group from the drop-down list.

Figure 61 Interface Port Selector Identity

Create Access Port Selector

Specify the selector identity

Name: FI-B-PO22

Description: Ports connecting to FI-B and part of vPC for HANA Backup traffic use

Interface IDs: 1/16
valid values: All or Ranges. For Example:
 1/13, 1/15 or 2/22-2/24, 2/16-3/16, or
 1/21-23/1-4, 1/24/1-2

Connected To Fex:

Interface Policy Group: vPC-FIB-22

6. Click OK and then click SUBMIT.

VPC Configuration for NetApp Storage

This section describes the configuration of policies required for VPC connection from ACI Leaf switches to the NetApp Storage Controllers.

Policy Group Configuration for VPC Connectivity to NetApp Storage

To create VPC interface policies for the interfaces connecting to the NetApp Storage Controller A, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand the Interface Policies in the left pane.
3. Expand Policy Groups and right-click the Leaf Policy Groups and select Create VPC Interface Policy Group.
4. Enter vPC-Netapp-cntrl-a for the Policy Group name.
5. Select 40GB for Link level Policy from the drop-down list.
6. Select CDP_Disabled for CDP Policy from the drop-down list.
7. Select LLDP_Enabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for LACP Policy from the drop-down list.
9. Select HANA-AAEP for Attached Entity Profile from the drop-down list.

Figure 62 Interface Policy Group for Netapp Controller-A

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: vPC-Netapp-cntrl-a

Description: optional

Link Level Policy: 40GB

CDP Policy: CDP_Disabled

MCP Policy: select a value

LLDP Policy: LLDP_Enabled

STP Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

Attached Entity Profile: HANA-AAEP

Connectivity Filters:

Switch IDs	Interfaces

Port Channel Policy: LACP_Active

Monitoring Policy: select a value

SUBMIT CANCEL

10. Click SUBMIT.

To create VPC interface policy for the interfaces connecting to the NetApp Storage Controller B, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand the Interface Policies in the left pane.
3. Expand Policy Groups and right-click the Leaf Policy Groups and select Create VPC Interface Policy Group.
4. Enter vPC-Netapp-cntrl-b for the Name of Policy group.
5. Select 40GB for Link level Policy from the drop-down list.
6. Select CDP_Disabled for CDP Policy from the drop-down list.

7. Select LLDP_Enabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for LACP Policy from the drop-down list.
9. Select HANA-AAEP for Attached Entity Profile from the drop-down list.

Figure 63 Interface Policy Group for Netapp Controller-B

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: vPC-Netapp-cntrl-b

Description: optional

Link Level Policy: 40GB

CDP Policy: CDP_Disabled

MCP Policy: select a value

LLDP Policy: LLDP_Enabled

STP Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

Attached Entity Profile: HANA-AAEP

Connectivity Filters: × +

Switch IDs	Interfaces

Port Channel Policy: LACP_Active

Monitoring Policy: select a value

SUBMIT CANCEL

10. Click SUBMIT.

Interface Profile Configuration

To create VPC interface policies for the interfaces connecting to NetApp Storage, complete the following steps:

1. Expand Interface Profiles.
2. Expand Profiles and right-click the Leaf Profiles and select Create Leaf Interface Profile.
3. Enter Netapp-A-PO40 as the Interface Profile name.

4. For Interface Selectors, click the + symbol to add interfaces.

Figure 64 Interface Profile Identity – Netapp Controller-A

Create Leaf Interface Profile

Specify the profile Identity

Name: Netapp-A-PO40

Description: Interfaces connecting to Netapp Controller A

Interface Selectors: x +

Name	Type

5. Enter Netapp-A-PO40 for port selector identity name and 1/9 for the Interface ID.
6. Select vPC-Netapp-cntrl-a for the Interface Policy Group from the drop-down list.

Figure 65 Interface Port Selector Identity – Netapp Controller-A

Create Access Port Selector

Specify the selector identity

Name: Netapp-A-PO40

Description: Ports connecting to Netapp Controller A

Interface IDs: 1/9
valid values: All or Ranges. For Example:
 1/13, 1/15 or 2/22-2/24, 2/16-3/16, or
 1/21-23/1-4, 1/24/1-2

Connected To Fex:

Interface Policy Group: vPC-Netapp-cntrl-a

7. Click OK and then click SUBMIT.

Controller-B

1. Right-click the Leaf Profiles and select Create Leaf Interface Profile.
2. Enter FI-B-PO41 as the Interface Profile name.
3. For Interface Selectors, click the + symbol to add interfaces.

Figure 66 Interface Profile Identity – Netapp Controller-B

Create Leaf Interface Profile

Specify the profile Identity

Name: Netapp-B-PO41

Description: Interfaces connecting to Netapp Controller B

Interface Selectors: × +

Name	Type
------	------

4. Enter Netapp-B-PO41 for the port selector identity name and 1/10 for the Interface ID.
5. Select vPC-Netapp-cntrl-b for the Interface Policy Group from the drop-down list.

Figure 67 Interface Port Selector Identity – Netapp Controller-A

Create Access Port Selector

Specify the selector identity

Name: Netapp-B-PO41

Description: Ports connecting to Netapp Controller B

Interface IDs: 1/10
valid values: All or Ranges. For Example: 1/13, 1/15 or 2/22-2/24, 2/16-3/16, or 1/21-23/1-4, 1/24/1-2

Connected To Fex:

Interface Policy Group: vPC-Netapp-cntrl-b

6. Click OK and then click SUBMIT.

VPC Configuration for Management PoD

This section describes the configuration of policies required for VPC connection from Cisco ACI Leaf switches to Management switch.

Policy Group Configuration for VPC Connectivity to Management PoD

To create VPC interface policies for the interfaces connecting to Nexus 9000 switches of the Management PoD, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand the Interface Policies in the left pane.
3. Expand Policy Groups and right-click the Leaf Policy Groups and select Create VPC Interface Policy Group.

4. Enter Mgmt-PoD for the Policy Group name.
5. Select 40GB for Link level Policy from the drop-down list.
6. Select CDP_Disabled for CDP Policy from the drop-down list.
7. Select LLDP_Enabled for LLDP Policy from the drop-down list.
8. Select LACP_Active for LACP Policy from the drop-down list.
9. Select HANA-AAEP for Attached Entity Profile from the drop-down list.

Figure 68 Interface Policy Group for Management PoD

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: vPC-Mgmt-PoD

Description: Interfaces connecting to Management PoD

Link Level Policy: 40GB

CDP Policy: CDP_Disabled

MCP Policy: select a value

LLDP Policy: LLDP_Enabled

STP Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

Attached Entity Profile: HANA-AAEP

Connectivity Filters:

Switch IDs	Interfaces

Port Channel Policy: LACP_Active

Monitoring Policy: select a value

SUBMIT CANCEL

10. Click SUBMIT.

Interface Profile Configuration

This section describes the procedure to create the interface profiles that will inherit the associated policy groups.

To create the interface profiles for ports connecting to Cisco Nexus 9000 switches of the Management PoD, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand the Interface Policies in the left pane.
3. Expand Profiles and right-click the Leaf Profiles and select Create Leaf Interface Profile.
4. Enter Mgmt-PoD as the Name of the Interface Profile.
5. For Interface Selectors, click the + symbol to add interfaces.

Figure 69 Interface Profile Identity – Mgmt PoD

6. Enter Mgmt-PoD for port selector identity name and 1/21-22 for the Interface IDs.
7. Select vPC-Mgmt-PoD for the Interface Policy Group from the drop-down list.

Figure 70 Interface Port Selector Identity – Mgmt PoD

8. Click OK and then click SUBMIT.

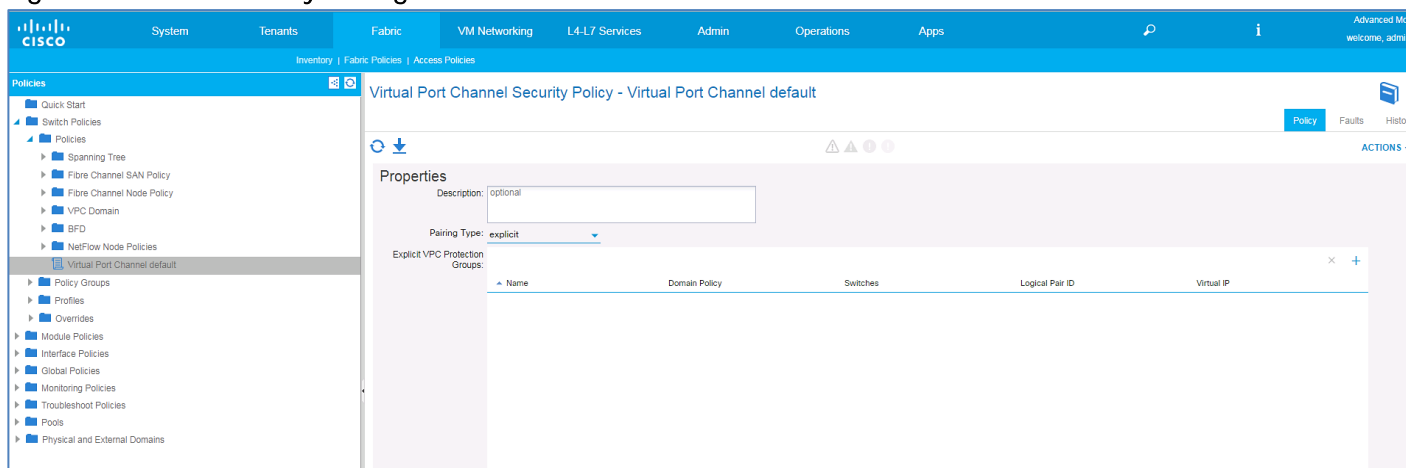
Switch Policy Configuration

This section describes the procedure to configure Switch Policy, which will be applied to the Cisco ACI Leaf switches.

To configure the Switch Policy, complete the following steps:

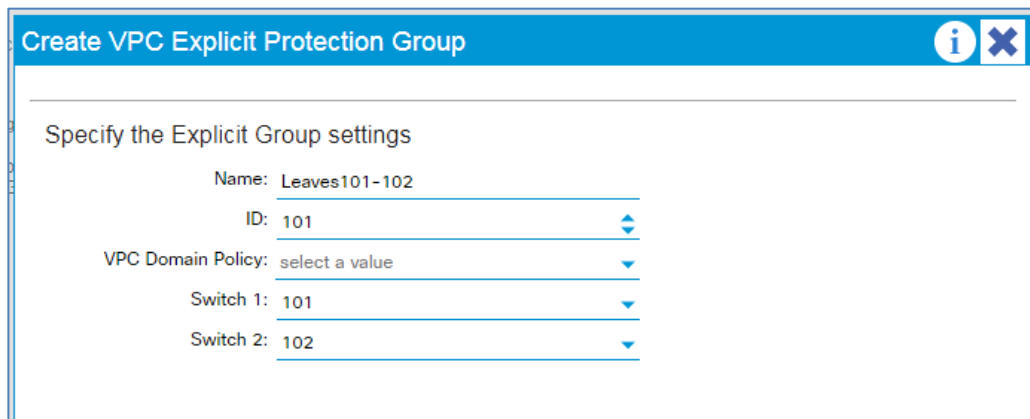
1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand Switch Policies.
3. Expand Policies.
4. Click Virtual Port Channel Default.
5. Under Explicit VPC Protection Groups, click the + symbol.

Figure 71 Switch Policy Configuration



6. Enter Leaves101-102 as the VPC Explicit Group name.
7. Enter 101 in the ID field.
8. Check if 101 is selected for Switch 1 and 102 is selected for Switch 2.
9. Click SUBMIT.

Figure 72 Switch Policy Group Setting

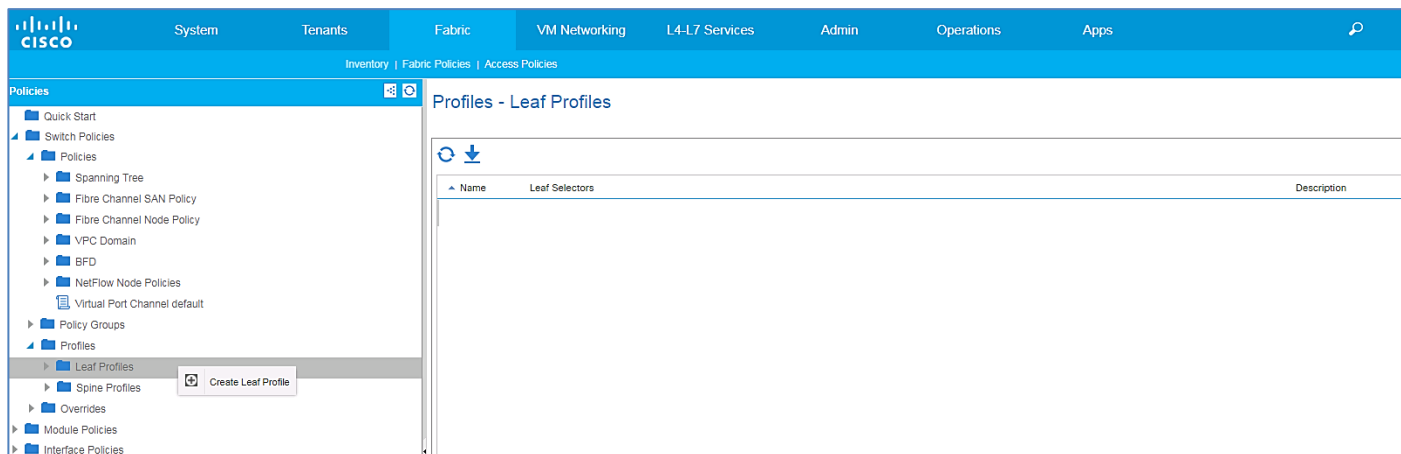


Switch Profile Configuration

This section describes the procedure to create switch profile in order to apply the switch configuration to the Leaf switches.


To configure ACI Leaf switches using Switch Profile, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.
2. Expand Switch Policies in the left pane.
3. Expand Profiles and right-click the Leaf Profiles and select Create Leaf Profile.



4. Enter Leaves101 -102 as the Name of the Switch Profile.
5. Under Switch Selectors, click the + symbol.

Figure 73 Switch Profile Configuration



6. Enter Leaf101 in the Name field.
7. Select 101 under BLOCKS.
8. Click UPDATE to save the configuration.
9. Click the + symbol again to add another switch.
10. Enter Leaf102 in the Name field.
11. Select 102 under BLOCKS.
12. Click UPDATE to save the configuration.

Figure 74 Switch Profile Configuration with Leaf Switches

Create Leaf Profile

STEP 1 > Profile

1. Profile 2. Associations

Specify the profile Identity

Name: Leaves101-102

Description: Profile with Leaf s/ws connecting to end devices via vPC to Fls, Netapp controllers and Mgmt Pod Nexus s/ws

Leaf Selectors:

Name	Blocks	Policy Group
Leaf101	101	
Leaf102	102	

PREVIOUS NEXT CANCEL

13. Click Next.

14. In the ASSOCIATIONS windows, select all the Interface Profiles created for Interface Selector Profiles.

Figure 75 Switch Profile Association of VPC

1. Profile 2. Associations

STEP 2 > Associations

Select the interface/module selector profiles to associate

Interface Selector Profiles:

Select	Name	Description
<input checked="" type="checkbox"/>	FI-A-PO11	Port connecting to FI-A. PC for HANA traffic usage.
<input checked="" type="checkbox"/>	FI-A-PO21	Ports connecting to FI-A. Backup network use
<input checked="" type="checkbox"/>	FI-B-PO12	Port connecting to FI-B. HANA traffic usage.
<input checked="" type="checkbox"/>	FI-B-PO22	Interfaces connecting to FI-B. Backup network use.

Module Selector Profiles:

Select	Name	Description
--------	------	-------------

PREVIOUS FINISH CANCEL

15. Click FINISH.

Tenant Configuration

While the Fabric and Access Policies dealt with physical aspects of the fabric setup, Tenant provides for a logical container or a folder for application policies. The Tenant can represent an actual tenant, a customer, an organization, or can just be used for the convenience of organizing information. A tenant represents a unit of isolation from a policy perspective.

All application configurations in Cisco ACI are part of a tenant. Within a tenant, you define one or more Layer 3 networks (VRF instances), one or more bridge domains per network, and EPGs to divide the bridge domains.



We define a SAP HANA customer T01 [**HANA-T01**] that wants to address a typical HANA Scale-Out system on FlexPod.

Application End Point Group (EPG): An End Point Group (EPG) is a collection of physical and/or virtual end points that require common services and policies. An End Point Group example is a set of servers or storage LIFs on a common VLAN providing a common application function or service. While the scope of an EPG definition is much wider, in the simplest terms an EPG can be defined on a per VLAN segment basis where all the servers or VMs on a common LAN segment become part of the same EPG.

You can define the EPGs based on the following:

- HANA nodes configured with various networks
 - HANA-T01-Boot-Nodes based on Boot network – nodes to boot from NFS root volumes on the NetApp array utilizing PXE services of PXE Boot Server in the Management PoD.
 - HANA-T01-Internode based on Inter-node network – used by HANA Nodes for inter-node communication
 - HANA-T01-Management – management interface communication and node administration
 - HANA-T01-Node-Data, HANA-T01-Node-Log, HANA-T01-NFS – networks used by nodes for NFS communication for HANA Data, HANA Log, and HANA Shared filesystems.
 - HANA-T01-Client – for communication to Client/User network.
 - HANA-T01-Backup – for communication between HANA system and backup storage network
 - HANA-T01-Replication-Site01 – for communication between HANA system and replication to another HANA system
 - HANA-T01-AppServer – for communication between HANA system and SAP Application Servers.
 - HANA-T01-DataSource – for communication between HANA system and data source network for Data uploads
 - HANA-T01-Access – for user access to HANA system
 - HANA-T01-iSCSI-InitiatorA and HANA-T01-iSCSI-InitiatorB – in case HANA nodes need to use iSCSI boot leveraging iSCSI LUNs from the NetApp array.
- NetApp storage controllers providing via designated networks
 - HANA-T01-Boot-Storage – Network configured on the array providing boot volumes
 - HANA-T01-HANA-Data, HANA-T01-HANA-Log and HANA-T01-NFS – Data, Log and Hana shared Filesystems provider network on NetApp array.
 - HANA-T01-iSCSI-TargetA and HANA-T01-iSCSI-TargetB – iSCSI service provider network on the array providing boot LUN access.
- Management PoD providing PXE services for HANA nodes.
 - HANA-T01-Mgmt-External – Management network access to Management PoD
 - HANA-T01-PXEServer – Management PoD node configured as PXE server providing TFTP/DHCP services to the HANA nodes for booting from NFs root volumes on NetApp array.

Application Profile: An application profile models application requirements and contains as many (or as few) End Point Groups (EPGs) as necessary that are logically related to providing the capabilities of an application.

You could define one Application Profile: HANA-T01-ScaleOut **representing a customer T01's solution** HANA Scale-Out system containing the EPGs as explained above.

Bridge Domain: A bridge domain represents a L2 forwarding construct within the fabric. One or more EPG can be associated with one bridge domain or subnet. A bridge domain can have one or more subnets associated with it. One or more bridge domains together form a tenant network.

You could make EPGs that need to talk to each other, part of the same bridge domain while you configure it per EPG basis for those that are isolated. To create Bridge Domains, complete the following steps:

1. HANA-T01-Boot – set of subnets associated with a common function, booting. HANA-T01-Boot-Nodes, HANA-T01-Boot-Storage and HANA-T01-PXEServer EPGs are associated with this bridge domain.
2. HANA-T01-Storage – set of subnets associated with HANA data filesystems provider and consumer functions. EPGs HANA-T01-HANA-Data and HANA-T01-Node-Data are associated with this domain.
3. HANA-T01-Log – set of networks associated with HANA log filesystem provider and consumer functions. EPGs HANA-T01-HANA-Log and HANA-T01-Node-Log are associated with this domain.
4. HANA-T01-iSCSIA – set of networks associated with iSCSI booting. EPGs HANA-T01-iSCSI-InitiatorA and HANA-T01-iSCSI-TargetA are associated with this domain. Similarly a bridge domain HANA-T01-iSCSIB having EPGs HANA-T01-iSCSI-InitiatorB and HANA-T01-iSCSI-TargetB associated with it is defined.
5. HANA-T01-Mgmt – set of networks associated with management access. EPGs HANAT01-Mgmt and HANA-T01-Mgmt-External are associated with this domain.
6. HANA-T01-Internal containing the EPG HANA-T01-Internode, associated with inter node communication in the SAP HANA Scale-out cluster.
7. HANA-T01-L3Out containing EPGs HANA-T01-Access, HANA-T01-Client, HANA-T01-Datasource, HANA-T01-AppServer and HANA-T01-Replication-Site01 are defined keeping in mind the requirement for HANA database access outside of the landscape network.

Contracts: A service contract can exist between two or more participating peer entities, such as two applications running and communicating with each other behind different endpoint groups, or between providers and consumers, such as a DNS contract between a provider entity and a consumer entity. Contracts utilize filters to limit the traffic between the applications to certain ports and protocols.

You need to create individual contracts between the EPGs that need to talk to each other leveraging the **‘default’ filter that allows all traffic between them.**



For a Multi-Tenancy environment, each Tenant can be configured with identical categories, port channels, etc. However, VLAN use within a tenant would need to be unique between tenants.



The common/default filter allows all communication between the corresponding endpoints. A filter can be created to limit the communication port and protocol between two endpoints; this configuration is beyond the scope of this document

Tenant Creation

To create tenant for SAP HANA, complete the following steps:

1. From the main menu, click TENANTS and from the sub-menu click ADD TENANT.
2. In the CREATE TENANT dialog box, type HANA-T01 as the name of the tenant.
3. For VRF Name enter HANA-T01-VRF.

Figure 76 Create Tenant

Create Tenant

Specify tenant details

Name: HANA-T01

Alias:

Description: optional

Tags:
 enter tags separated by comma

GUID:

Provider	GUID	Account Name

Monitoring Policy: select a value

Security Domains:

Name	Description

VRF Name: HANA-T01-VRF

Take me to this tenant when I click finish

SUBMIT CANCEL

4. Check the 'Take me to the tenant when I click finish' checkbox.
5. Click SUBMIT.

Application Profiles for HANA

To create HANA-T01-ScaleOut Application Profile, complete the following steps:

1. Click TENANTS and from the sub-menu click HANA-T01 tenant.
2. Expand Tenant HANA-T01 in the left pane.
3. Right-click the Application Profiles and click Create Application Profile.
4. In CREATE APPLICATION PROFILE dialog box, enter HANA-T01-ScaleOut in the Name field.

Figure 77 Application Profile HANA-T01-ScaleOut

Create Application Profile

Specify Tenant Application Profile

Name:

Alias:

Description:

Tags:

enter tags separated by comma

Monitoring Policy:

EPGs x +

Name	Alias	BD	Domain	Static Path	Static Path VLAN	Provided Contract	Consumed Contract

5. Click SUMBIT.

Bridge Domains

To create the planned Bridge Domains for use with HANA-T01-ScaleOut Application Profile, complete the following steps:

1. Click TENANTS and from the sub-menu click HANA-T01 tenant.
2. Expand Tenant HANA-T01 in the left pane.
3. Expand Networking and right-click the Bridge Domains and click Create Bridge Domain
4. In the STEP1>Main page, enter HANA-T01-Boot for name and select the previously created HANA-T01/HANA-T01-VRF and click NEXT.

Create Bridge Domain [i] [X]

STEP 1 > Main 1. Main 2. L3 Configurations 3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Type:


VRF: [copy]

Forwarding: [v]

End Point Retention Policy: [v]
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: [v]

5. In the Step2 > L3 Configurations page, additionally enable ARP Flooding. Leave the rest parameter selections unchanged. Click NEXT.



6. In the Step3 > Advanced/Troubleshooting page, click FINISH.
7. Under Bridge Domains, select the created HANA-T01-Boot on the left pane. On the right pane, on Properties page, change the L2 Unknown Unicast value from Hardware Proxy to Flood.

The screenshot displays the Cisco ACI configuration interface for a Bridge Domain named 'HANA-T01-Boot'. The left sidebar shows a navigation tree under 'Tenant HANA-T01' with 'Bridge Domains' selected. The main panel shows the configuration for 'Bridge Domain - HANA-T01-Boot' with the following properties:

- Name: HANA-T01-Boot
- Alias: (empty)
- Description: optional
- Type: fc, regular
- Global Alias: (empty)
- Legacy Mode: No
- VRF: HANA-T01/HANA-T01-
- Resolved VRF: HANA-T01/HANA-T01-VRF
- L2 Unknown Unicast: Flood, Hardware Proxy
- L3 Unknown Multicast Flooding: Flood, Optimized Flood
- Multi Destination Flooding: Flood in BD, Drop, Flood in Encapsulation
- PIM:
- IGMP Policy: select an option
- ARP Flooding:
- Endpoint Dataplane Learning:
- Clear Remote MAC Entries:
- Limit IP Learning To Subnet:
- End Point Retention Policy: select a value
- IGMP Snoop Policy: select a value

Follow the steps above from 1 through 7 to create the rest of the planned Bridge Domains: HANA-T01-Internal, HANA-T01-L3Out, HANA-T01-Log, HANA-T01-Mgmt, HANA-T01-Storage, HANA-T01-iSCSIA and HANA-T01-iSCSIB.



Make sure to use HANA-T01-HANA-T01-VRF value on Step1 page and **enable ARP Flooding** on Step2 page. Selecting the created Bridge Domain, change the **L2 Unknown Unicast** value from **Hardware Proxy** to **Flood**.

Figure 78 Summary of Created Bridge Domains

Name	Alias	Type	Unknown Unicast Traffic Class ID	Segment	Multicast Address	Custom MAC Address	L2 Unknown Unicast	ARP Flooding	Unicast Routing	Description
HANA-T01-Boot		regular	32770	14811121	225.0.166.192	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-Internal		regular	16388	14942176	225.0.169.32	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-L3Out		regular	16398	15269816	225.0.235.32	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-Log		regular	32782	15171525	225.0.137.16	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-ISCsIA		regular	49156	15564692	225.1.145.96	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-L3Out		regular	16393	14876648	225.1.185.0	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-Log		regular	16394	15531928	225.1.148.240	00:22:BD:F8:19:FF	Flood	True	True	
HANA-T01-Mgmt		regular	16390	15630221	225.0.154.32	00:22:BD:F8:19:FF	Flood	True	True	

Application EPGs

You need to create the EPGs based on the requirements, as detailed in the Tenant Creation section. The same information is tabulated below:

Application EPG	Bridge Domain	Physical Domain	Paths	Port Encap	Sample vLAN ID
HANA-T01-Boot-Nodes	HANA-T01/HANA-T01-Boot	HANA	Pod-1/Node 101-102/ vPC-FIA-11	vlan-<<var_boot_vlan_id>>	227
			Pod-1/Node 101-102/ vPC-FIB-12	vlan-<<var_boot_vlan_id>>	227
HANA-T01-Boot-Storage	HANA-T01/HANA-T01-Boot	HANA	Pod-1/Node 101-102/ vPC-Netapp-cntrl-a	vlan-<<var_boot_vlan_id_aci_storage>>	127
			Pod-1/Node 101-102/ vPC-Netapp-cntrl-b	vlan-<<var_boot_vlan_id_aci_storage>>	127
HANA-T01-PXEServer			Pod-1/Node 101-102/ vPC-Mgmt-PoD	vlan-<<var_boot_vlan_id_aci>>	277
HANA-T01-HANA-Data	HANA-T01/HANA-T01-Storage	HANA	Pod-1/Node 101-102/ vPC-Netapp-cntrl-a	vlan-<<var_storage_data_vlan_id_aci>>	201
			Pod-1/Node 101-102/ vPC-Netapp-cntrl-b	vlan-<<var_storage_data_vlan_id_aci>>	201
HANA-T01-Node-Data	HANA-T01/HANA-T01-Storage	HANA	Pod-1/Node 101-102/ vPC-FIA-11	vlan-<<var_storage_data_vlan_id>>	271
			Pod-1/Node 101-102/ vPC-FIB-12	vlan-<<var_storage_data_vlan_id>>	271
HANA-T01-HANA-Log	HANA-T01/HANA-T01-Storage	HANA	Pod-1/Node 101-102/ vPC-Netapp-cntrl-a	vlan-<<var_storage_log_vlan_id_aci>>	228
			Pod-1/Node 101-102/ vPC-Netapp-cntrl-b	vlan-<<var_storage_log_vlan_id_aci>>	228
HANA-T01-Node-Log	HANA-T01/HANA-T01-Storage	HANA	Pod-1/Node 101-102/ vPC-FIA-11	vlan-<<var_storage_log_vlan_id>>	278
			Pod-1/Node 101-102/ vPC-FIB-12	vlan-<<var_storage_log_vlan_id>>	278
HANA-T01-Mgmt	HANA-T01/HANA-T01-Mgmt	HANA	Pod-1/Node 101-102/ vPC-FIA-11	vlan-<<var_admin_vlan_id>>	176
			Pod-1/Node 101-102/ vPC-FIB-12	vlan-<<var_admin_vlan_id>>	176
HANA-T01--Mgmt-External	HANA-T01/HANA-T01-Mgmt	HANA	Pod-1/Node 101-102/ vPC-FIA-11	vlan-<<var_admin_vlan_id_mgmt>>	76
			Pod-1/Node 101-102/ vPC-FIB-12	vlan-<<var_admin_vlan_id_mgmt>>	76
HANA-T01-iSCSI-InitiatorA	HANA-T01/HANA-T01-iSCsIA	HANA	Pod-1/Node 101-102/ vPC-FIA-11	<<iSCSI_vlan_id_A>>	328
			Pod-1/Node 101-102/ vPC-FIB-12	<<iSCSI_vlan_id_A>>	328
HANA-T01-iSCSI-TargetA			Pod-1/Node 101-102/ vPC-Netapp-cntrl-a	<<iSCSI_vlan_id_A_aci>>	128
HANA-T01-iSCSI-InitiatorB	HANA-T01/HANA-T01-iSCsIB	HANA	Pod-1/Node 101-102/ vPC-FIA-11	<<iSCSI_vlan_id_B>>	329
			Pod-1/Node 101-102/ vPC-FIB-12	<<iSCSI_vlan_id_B>>	329
HANA-T01-iSCSI-TargetB			Pod-1/Node 101-102/ vPC-Netapp-cntrl-b	<<iSCSI_vlan_id_B_aci>>	129

Application EPG	Bridge Domain	Physical Domain	Paths	Port Encap	Sample vLAN ID
HANA-T01-Internode	HANA-T01/HANA-T01-Internal	HANA	Pod-1/Node 101-102/ vPC-FIA-11 Pod-1/Node 101-102/ vPC-FIB-12	<<var_internal_vlan_id>> <<var_internal_vlan_id>>	220 220
HANA-T01-Backup-Node	HANA-T01/HANA-T01-Backup	HANA	Pod-1/Node 101-102/ vPC-FIA-21 Pod-1/Node 101-102/ vPC-FIB-22	<<var_internal_vlan_id>> <<var_internal_vlan_id>>	224 224
HANA-T01-Access			Pod-1/Node 101-102/ vPC-FIA-11 Pod-1/Node 101-102/ vPC-FIB-12	<<var_internal_vlan_id>> <<var_internal_vlan_id>>	301 301
HANA-T01-AppServer	HANA-T01/HANA-T01-L3Out	HANA	Pod-1/Node 101-102/ vPC-FIA-11 Pod-1/Node 101-102/ vPC-FIB-12	<<var_appserver_vlan_id>> <<var_appserver_vlan_id>>	226 226
HANA-T01-DataSource			Pod-1/Node 101-102/ vPC-FIA-11 Pod-1/Node 101-102/ vPC-FIB-12	<<var_datasource_vlan_id>> <<var_datasource_vlan_id>>	225 225
HANA-T01-Replication	HANA-T01/HANA-T01-Replication	HANA	Pod-1/Node 101-102/ vPC-FIA-21 Pod-1/Node 101-102/ vPC-FIB-22	<<var_replication_vlan_id>> <<var_replication_vlan_id>>	300 300

The table serves as a ready reference for the values for key inputs during the creation of EPGs.

The steps to create HANA-T01-Boot-Nodes, HANA-T01-Boot-Storage and HANA-T01-PXEServer are provided below.

HANA-T01-Boot-Nodes EPG

To create the planned Application EPGs, complete the following steps:

1. Click TENANTS and from the sub-menu click HANA-T01 tenant.
2. Expand Tenant HANA-T01 in the left pane.
3. Expand Application Profiles, expand HANA-T01-ScaleOut, right-click the Application EPGs and click Create Application EPG.
4. In the STEP1>Main page, enter HANA-T01-Boot-Nodes for name and select the previously created Bridge Domain HANA-T01/HANA-T01-Boot.
5. **Enable** “Statically Link with Leaves/Paths” option. Click NEXT.

Create Application EPG
i ✕

STEP 1 > Identity

1. Identity
2. Leaves/Paths

Specify the EPG Identity

Name:

Alias:

Description:

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Intra EPG Isolation:

Preferred Group Member:

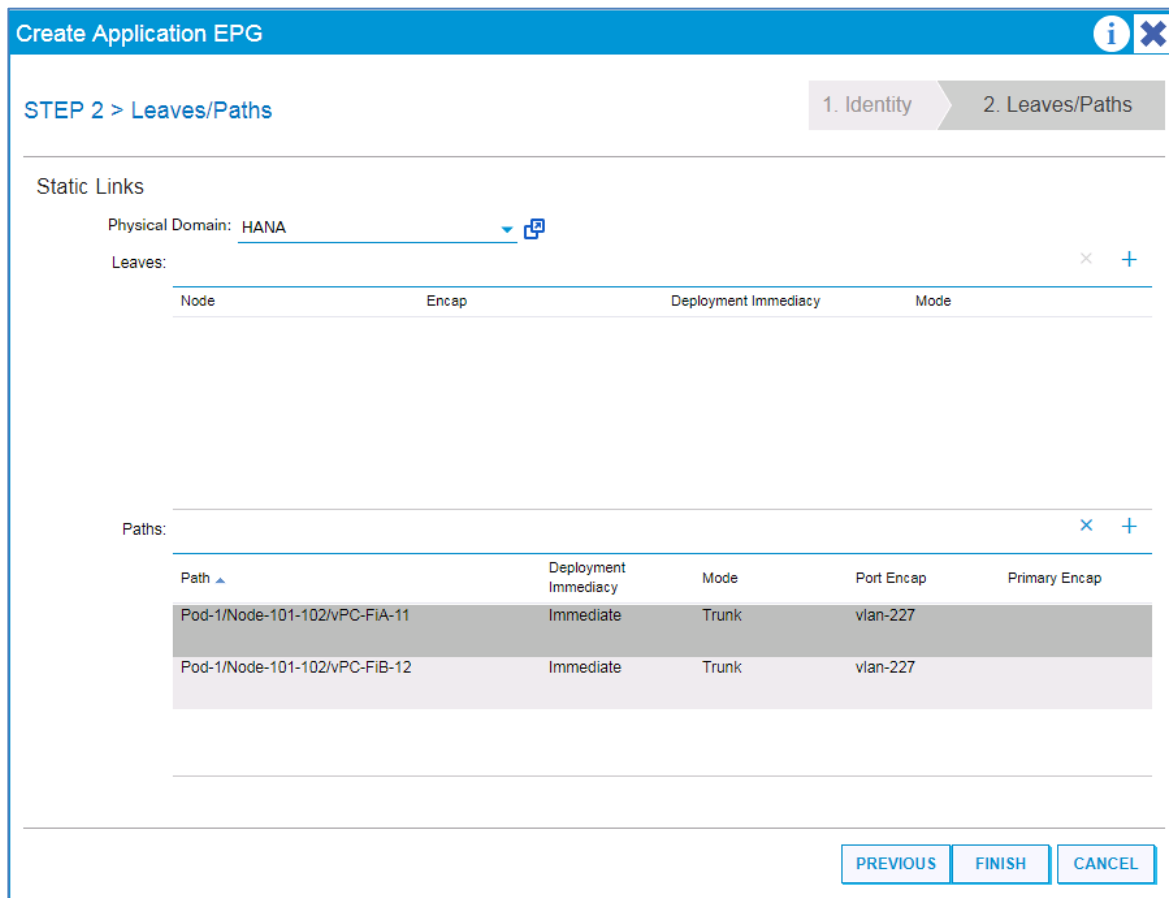
Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

6. In the STEP2 > Leaves/Paths window, for the Physical Domain, choose HANA from the drop-down list.
7. Under Paths, click + for adding Paths.
8. From the drop down list for Path, **select “Pod-1/Node 101-102/vPC-FiA-11” and enter** `vlan-
<<var_boot_vlan_id>>` **for “Encap” to associated** VLAN for HANA-T01-Boot-Node. Click UPDATE.
9. Again click + for adding Paths.
10. From the drop down list for Path, **select “Pod-1/Node 101-102/vPC-FiB-12” and enter** `vlan-
<<var_boot_vlan_id>>` **for “Encap” to associated** VLAN for HANA-T01-Boot-Node. Click UPDATE.



11. Click FINISH.

HANA-T01-Boot-Storage EPG

1. Click TENANTS and from the sub-menu click HANA-T01 tenant.
2. Expand Tenant HANA-T01 in the left pane.
3. Expand Application Profiles, expand HANA-T01-ScaleOut, right-click the Application EPGs and click Create Application EPG.
4. In the STEP1>Main page, enter HANA-T01-Boot-Storage for name and select the previously created Bridge Domain HANA-T01/HANA-T01-Boot.
5. **Enable** “Statically Link with Leaves/Paths” option. Click NEXT.

Create Application EPG
i
✕

STEP 1 > Identity

1. Identity

2. Leaves/Paths

Specify the EPG Identity

Name:

Alias:

Description:

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

6. In the STEP2 > Leaves/Paths window, for the Physical Domain, choose HANA from the drop down list.
7. Under Paths, click + for adding Paths.
8. From the drop down list for Path and **select** "Pod-1/Node 101-102/vPC-Netapp-cntrl-a" **and enter** `vlan- $\langle\langle$ var_boot_vlan_id_aci_storage $\rangle\rangle$ for "Encap" to associated vlan for HANA-T01-Boot-Node-Storage. Click UPDATE.`
9. In the Static Link section, click the + for adding Paths.
10. From the drop down list for Path and **select** "Pod-1/Node 101-102/vPC-Netapp-cntrl-b" **and enter** `vlan- $\langle\langle$ var_boot_vlan_id_aci_storage $\rangle\rangle$ for "Encap" to associated vlan for HANA-Boot-NFSStorage. Click UPDATE.`

i X
Create Application EPG

STEP 2 > Leaves/Paths

1. Identity
2. Leaves/Paths

Static Links

Physical Domain: HANA ▼ 📄

Leaves: ✕ +

Node	Encap	Deployment Immediacy	Mode

Paths: ✕ +

Path	Deployment Immediacy	Mode	Port Encap	Primary Encap
Pod-1/Node-101-102/vPC-Netapp-cntrl-a	Immediate	Trunk	vlan-127	
Pod-1/Node-101-102/vPC-Netapp-cntrl-b	Immediate	Trunk	vlan-127	

PREVIOUS
FINISH
CANCEL

11. Click FINISH.

HANA-T01-PXEServer EPG

1. Click TENANTS and from the sub-menu click HANA-T01 tenant.
2. Expand Tenant HANA-T01 in the left pane.
3. Expand Application Profiles, expand HANA-T01-ScaleOut, right-click the Application EPGs and click Create Application EPG.
4. In the STEP1>Main page, enter HANA-T01-PXEServer for name and select the previously created Bridge Domain HANA-T01/HANA-T01-Boot
5. **Enable “Statically Link with Leaves/Paths” option.** Click NEXT.

Create Application EPG [i] [X]

STEP 1 > Identity 1. Identity | 2. Leaves/Paths

Specify the EPG Identity

Name:

Alias:

Description:

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Intra EPG Isolation: Enforced Unenforced

Preferred Group Member: Exclude Include

Bridge Domain: [copy]

Monitoring Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

PREVIOUS | NEXT | CANCEL

6. In the STEP2 > Leaves/Paths window, for the Physical Domain, choose HANA from the drop-down list.
7. Under Paths, click + for adding Paths.
8. From the drop down list for Path, **select "Pod-1/Node 101-102/ vPC-Mgmt-PoD" and enter** `vlan-
<<var_boot_vlan_id_aci>>` **for "Encap" to associated VLAN** for HANA-T01-PXEServer. Click UPDATE.

Create Application EPG
i X

STEP 2 > Leaves/Paths

1. Identity
2. Leaves/Paths

Static Links

Physical Domain: HANA v +

Leaves: x +

Node	Encap	Deployment Immediacy	Mode

Paths: x +

Path	Deployment Immediacy	Mode	Port Encap	Primary Encap
Pod-1/Node-101-102/VPC-Mgmt-PoD	Immediate	Trunk	vlan-277	


PREVIOUS
FINISH
CANCEL

9. Click FINISH.



Repeat the sequence of steps substituting the input values from the reference table above to create rest of the EPGs.

In summary, at the end of this section, you will have created the following EPGs:



Contracts

This section describes the procedure to create and add contracts EPGs that need to communicate with each other.

To create required contract for NFS Boot access, complete the following steps:

1. From the main menu, click TENANTS and from the sub-menu click the HANA-T01 tenant.
2. Expand Tenant HANA-T01 in the left pane and right-click Contracts and select Create Contract.
3. In the Create Contract window, enter Boot-NFS for the Contract name.

Figure 79 Application Profile Contract Boot-NFS

Create Contract ⓘ ✕

Specify Identity Of Contract

Name:

Alias:

Scope: ▼

QoS Class: ▼

Target DSCP: ▼

Description:

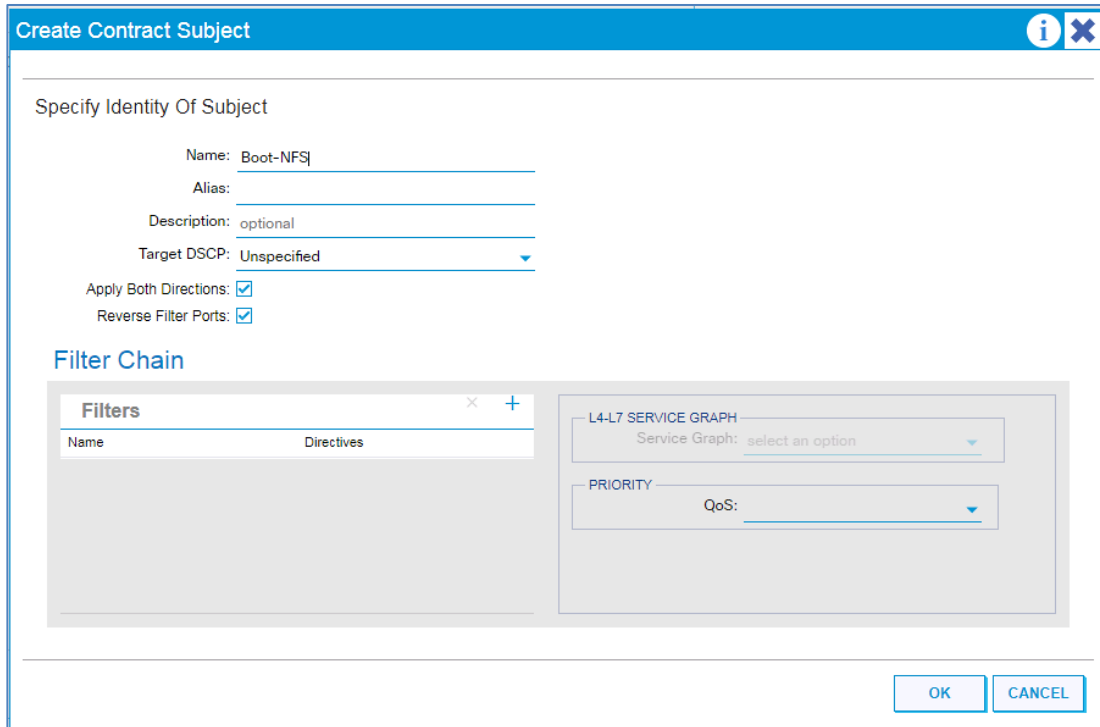
Tags: ▼
enter tags separated by comma

Subjects: ✕ +

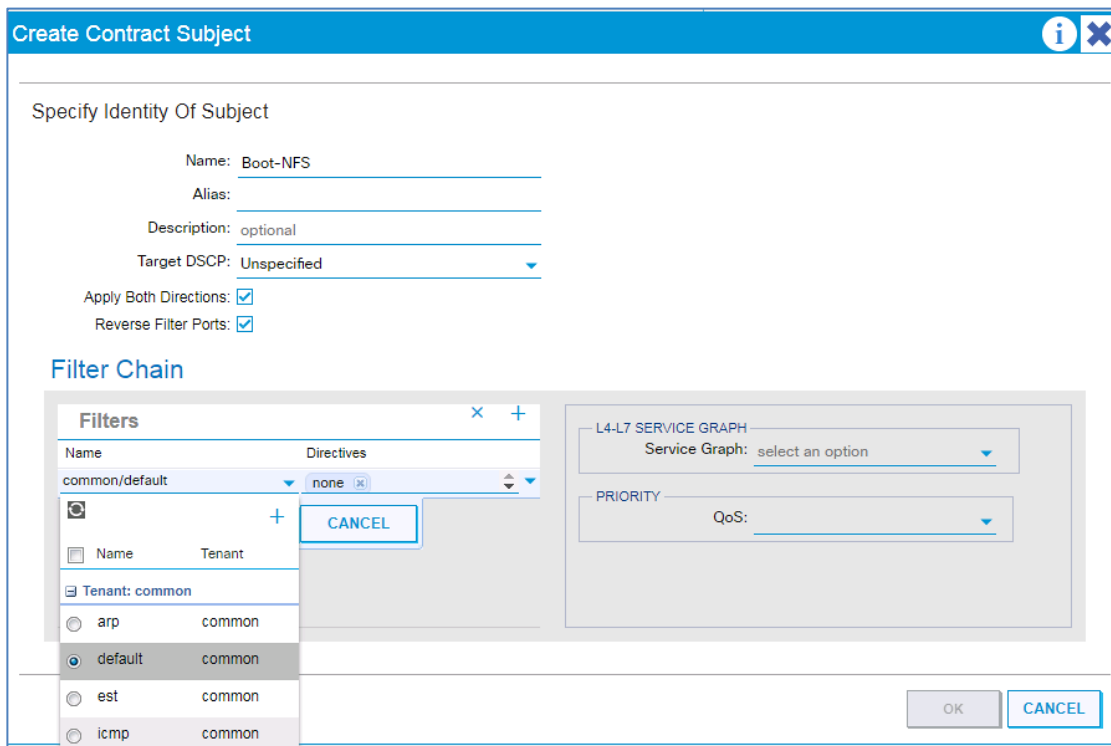
Name	Description
------	-------------

4. Click + next to Subjects section.
5. Enter Boot-NFS for the Subject name.

Figure 80 Application Profile Contract Subject HANA-Boot-NFS



6. Click + below the FILTERS under Filter Chain.
7. Select the common/default filter from the drop-down list.



8. Click UPDATE.

9. Click OK.
10. Click SUBMIT.



The common/default filter allows all communication between the corresponding endpoints. A filter can be created to limit the communication port and protocol between two endpoints; this configuration is beyond the scope of this document.

11. Repeat steps 1 through 10 to create contracts for management access, iSCSI boot access, HANA Data Filesystem access, HANA Log Filesystem access, PXE service access, at a minimum, with subjects created with common/default filters allowing all accesses, bi-directionally.

Name	Alias	Scope	QoS Class	Target DSCP	Subjects
Boot-NFS		VRF	Unspecified	Unspecified	Boot-NFS
HANA-Mgmt-Access		VRF	Unspecified	Unspecified	HANA-Mgmt
iSCSIa-Boot-Access		VRF	Unspecified	Unspecified	iSCSI-A
iSCSIb-Boot-Access		VRF	Unspecified	Unspecified	iSCSI-B
NFS-Data-access		VRF	Unspecified	Unspecified	nfs-data-access
NFS-Log-access		VRF	Unspecified	Unspecified	nfs-log
PXE-Server		VRF	Unspecified	Unspecified	PXE-Server

12. When you have created these contracts, you can add them by appropriate type to the member EPGs that need to communicate with each other for a particular service access based on the information tabulated as below:

Contract name	Participating EPGs	Add Type
Boot NFS	HANA_T01-Boot-Storage	Provided Contract
	HANA-T01-Boot-Nodes	Consumed Contract
	HANA-T01-PXEServer	Consumed Contract
PXE-Server	HANA-T01-PXEServer	Provide Contract
	HANA-T01-Boot-Nodes	Consumed Contract
HANA-Mgmt-Access	HANA-T01-Mgmt	Provided Contract
	HANA-T01-Mgmt-External	Consumed Contract
NFS-Data-Access	HANA-T01-HANA-Data	Provided Contract
	HANA-T01-Node-Data	Consumed Contract
NFS-Log-Access	HANA-T01-HANA-Log	Provided Contract
	HANA-T01-Node-Log	Consumed Contract
iSCSIa_boot-access	HANA-T01-iSCSI-TargetA	Provided Contract
	HANA-T01-iSCSI-InitiatorA	Consumed Contract
iSCSIb_boot-access	HANA-T01-iSCSI-TargetB	Provided Contract
	HANA-T01-iSCSI-InitiatorB	Consumed Contract

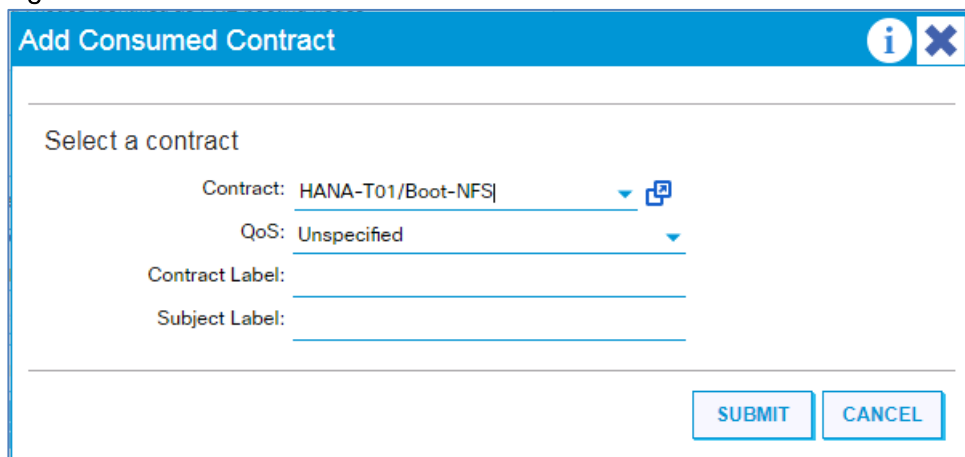
To assign the Boot-NFS contract to the interested EPGs, complete the following steps:

1. Expand Tenant HANA in the left pane and expand the Application Profiles.
2. Right-click EPG HANA-T01-Boot-Storage and Select Add Provided Contract.
3. Select the Boot-NFS option from the Contract field drop-down list.
4. Click SUBMIT.

Figure 81 Boot-NFS Provided Contract

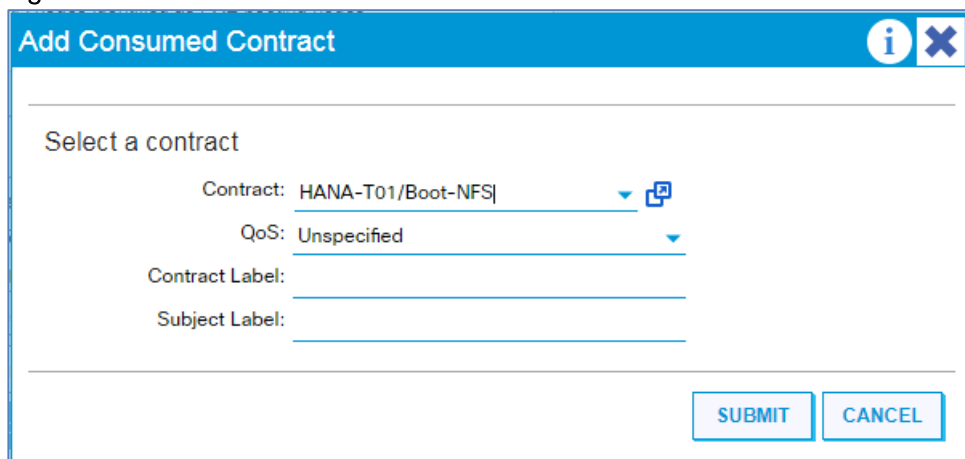
5. Right-click EPG HANA-T01-Boot-Nodes and Select Add Consumed Contract.
6. Select the Boot-NFS option from the Contract field drop-down list.
7. Click SUBMIT.

Figure 82 Boot-NFS Consumed Contract



8. Right-click EPG HANA-T01-PXEServer and Select Add Consumed Contract.
9. Select the Boot-NFS option from the Contract field drop-down list.
10. Click SUBMIT.

Figure 83 Boot-NFS Consumed Contract



Similarly, add the other contracts with appropriate type to the respective member/participating EPGs to ensure seamless access/communication.

Cisco UCS Solution for SAP HANA TDI

The SAP HANA TDI option enables multiple SAP HANA production systems to run on the same infrastructure. In this configuration, the existing blade servers used by different SAP HANA systems share the same network infrastructure and storage systems. In addition, the SAP application server can share the same infrastructure as the SAP HANA database. As mentioned earlier, this configuration provides better performance and superior disaster-tolerance solution for the whole system.

Cisco UCS servers enable separation of traffic, between a SAP HANA system and a non-SAP HANA system. This is achieved by creating a separate network uplink port-channel on Cisco UCS 6300 Fabric Interconnect, for each system type using the VLAN group option. This approach will guarantee the network bandwidth for each tenant in a secured environment. Figure 84 shows an example configuration to achieve this. In this example, two port-channels on each of the Cisco UCS Fabric Interconnects are created:


- Port-channel 11 and 21 are created on Cisco UCS Fabric Interconnect A
- Port-channel 12 and 22 are created on Cisco UCS Fabric Interconnect B

Optionally, a VLAN group for SAP HANA is created and all the VLANs carrying traffic for SAP HANA is added to this VLAN group. This VLAN group can be assigned to use port-channel 11 on Cisco UCS Fabric Interconnect A and port-channel 12 on Cisco UCS Fabric Interconnect B as shown in Figure 84.

Similarly, a VLAN group for application servers or for backup network can be created and all the VLANs carrying traffic for this use-case. The VLAN group can then be assigned to use port-channel 21 on fabric interconnect A and port-channel 22 on fabric interconnect B.

This approach achieves bandwidth-separation between SAP HANA servers and applications servers or backup management use case and bandwidth for SAP HANA servers can be increased or decreased by altering the number of ports in the port-channel 11 and port-channel 12.

Figure 84 Network Separation of Multiple Systems Using Port-Channel and VLAN Groups



Cisco UCS Server Configuration

This section describes the specific configurations on Cisco UCS servers to address SAP HANA requirements.

Initial Setup of Cisco UCS 6332 Fabric Interconnect

This section provides the detailed procedures to configure the Cisco Unified Computing System (Cisco UCS) for use in FlexPod Datacenter Solution for SAP HANA environment. These steps are necessary to provision the Cisco UCS C-Series and B-Series servers to meet SAP HANA requirements.

Cisco UCS 6332 Fabric Interconnect A

To configure the Cisco UCS Fabric Interconnect A, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6200 Fabric Interconnect.

```

Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have choosen to setup a 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
```

2. Review the settings printed to the console. If they are correct, answer yes to apply and save the configuration.
3. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS 6332 Fabric Interconnect B

To configure the Cisco UCS Fabric Interconnect B, complete the following steps:

1. Connect to the console port on the second Cisco UCS 6248 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. Do you want to 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.

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 6332 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 3.2(3d)

This document assumes the use of Cisco UCS Manager Software version 3.2(2d). To upgrade the Cisco UCS Manager software and the UCS 6332 Fabric Interconnect software to version 3.2(3d), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.
2. In Cisco UCS Manager, click the LAN tab in the navigation pane.
3. Select Pools > root > IP Pools > IP Pool ext-mgmt.
4. In the Actions pane, select Create Block of IP Addresses.
5. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

6. Click OK to create the IP block.
7. 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_global_ntp_server_ip>> and click OK.
7. Click OK.

Cisco UCS Blade Chassis Connection Options

For the Cisco UCS 2300 Series Fabric Extenders, two configuration options are available: pinning and port-channel.

SAP HANA node communicates with every other SAP HANA node using multiple I/O streams and this makes the port-channel option a highly suitable configuration. SAP has defined a single-stream network performance test as part of the hardware validation tool (TDINetServer/TDINetClient).

However, with the new 40Gb network speed it is also possible to stay with the default setting of Cisco UCS which is Port-Channel as connection policy.

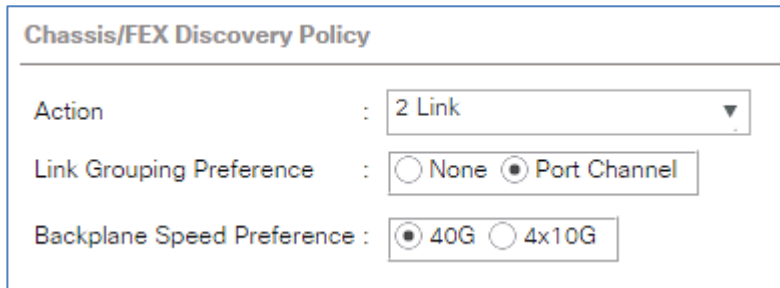
Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and of additional fabric extenders for further Cisco UCS 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”** for Port Channel.


5. Click Save Changes.
6. Click OK.



Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, 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 / or to the Cisco C-Series Server (two per FI), 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 and / or to the Cisco C-Series Server are now configured as server ports. Eth ports 1/17 through 1/22 on FI-A and FI-B are connected to the chassis. Right click these ports and “Configure as Server Port”.



7. Select ports that are connected to the Cisco Nexus ACI switches, right-click them, and select **“Configure as Uplink Port”**. **This is done for** the eth ports 1/29 through 1/34 in this setup.
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 or to the Cisco C-Series Server (two per FI), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select ports 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 Rack-Mount Servers

To acknowledge all Cisco UCS chassis and Rack Mount Servers, 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 C-Series servers are part of the configuration, expand Rack Mounts and FEX.
6. Right-click each Server that is listed and select Acknowledge Server.
7. Click Yes and then click OK to complete acknowledging the Rack Mount Servers

Create Uplink Port Channels to Cisco ACI Leaf Switches

An uplink port channel 11 on Fi-A and port channel 12 on FI-B with 4 ports eth 1/29-1/32 is defined to carry all the SAP HANA networks traffic. We are using 4 x 40G ports for this port channel providing 160Gbps operational speed and is good for carrying all the HANA networks traffic.

Another port channel 21 on FI- A and port channel 22 on FI-B is configured for backup traffic usage. This 2 x 40G port channel providing 80Gbps operational speed should suffice for the same.

Additional port channels for replication traffic, if any, can be configured on need basis.

To configure the necessary port channels as planned above for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. 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.
3. Under LAN > LAN Cloud, expand the Fabric A tree.
4. Right-click Port Channels.
5. Select Create Port Channel.
6. Enter 11 as the unique ID of the port channel.
7. Enter vPC-41-Nexus as the name of the port channel.
8. Click Next.

Create Port Channel

1 Set Port Channel Name

2 Add Ports

ID :

Name :

9. Select the following ports to be added to the Port Channel:
 - Slot ID 1 and port 29
 - Slot ID 1 and port 30
 - Slot ID 1 and port 31
 - Slot ID 1 and port 32

Create Port Channel

1 Set Port Channel Name

2 Add Ports

Ports			
Slot ID	Aggr. Po...	Port	MAC
1	0	29	00:DE:F...
1	0	30	00:DE:F...
1	0	31	00:DE:F...
1	0	32	00:DE:F...
1	0	33	00:DE:F...
1	0	34	00:DE:F...

>>
<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
No data available			

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

11. Click Finish to create the port channel.

12. Click OK.

13. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree:

- a. Right-click Port Channels.
- b. Select Create Port Channel.
- c. Enter 21 as the unique ID of the port channel.
- d. Enter vPC-ACI-21 as the name of the port channel.

Create Port Channel

1 Set Port Channel Name

2 Add Ports

ID : 12

Name : vPC-ACI-12

14. Click Next.

15. Select the following ports to be added to the port channel:

- Slot ID 1 and port 29
- Slot ID 1 and port 30
- Slot ID 1 and port 31
- Slot ID 1 and port 32

Create Port Channel

1 Set Port Channel Name

2 Add Ports

Ports			
Slot ID	Aggr. Po...	Port	MAC
1	0	29	00:DE:F...
1	0	30	00:DE:F...
1	0	31	00:DE:F...
1	0	32	00:DE:F...
1	0	33	00:DE:F...
1	0	34	00:DE:F...

>>

<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
No data available			

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

17. Click Finish to create the port channel.

18. Click OK.

To create port-channel for backup network, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel
5. Enter 21 as the unique ID of the port channel.
6. Enter vPC-ACI-21 as the name of the port channel.
7. Click Next.

Create Port Channel

1 Set Port Channel Name

2 Add Ports

ID : 21

Name : vPC-ACI-21

8. Select the following ports to be added to the port channel:

- Slot ID 1 and port 33
- Slot ID 1 and port 34

Create Port Channel

1 Set Port Channel Name

2 Add Ports

Ports			
Slot ID	Aggr. Po...	Port	MAC
1	0	29	00:DE:F...
1	0	30	00:DE:F...
1	0	31	00:DE:F...
1	0	32	00:DE:F...
1	0	33	00:DE:F...
1	0	34	00:DE:F...

>>


<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
No data available			

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 22 as the unique ID of the port channel.
16. Enter vPC-ACI 22 as the name of the port channel.
17. Click Next.



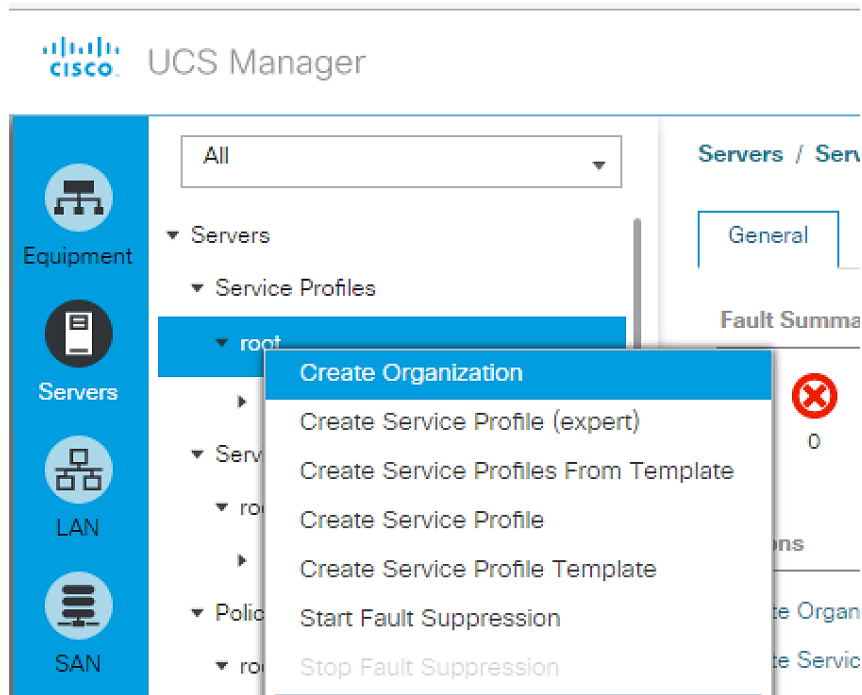
18. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 33
 - Slot ID 1 and port 34
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create New Organization

For secure multi-tenancy within the Cisco UCS domain, a logical entity is created as Organizations.

To create organization unit, complete the following steps:

1. In Cisco UCS Manager, on the Servers bar.
2. Select Servers and right-click root and select Create Organization.



3. Enter the Name as HANA.
4. Optional Enter the Description as Org for HANA.
5. Click OK to create the Organization.

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 > Sub-Organization > HANA.
3. In this procedure, two MAC address pools are created, one for each switching fabric.
4. Right-click MAC Pools under the root organization.
5. Select Create MAC Pool to create the MAC address pool .
6. Enter MAC_Pool_A as the name of the MAC pool.
7. Optional: Enter a description for the MAC pool.
8. Choose Assignment Order Sequential.

Create MAC Pool [?] [X]

1 Define Name and Description

Name :

Description :

Assignment Order : Default Sequential

9. Click Next.

10. Click Add.

11. Specify a starting MAC address.

12. The recommendation is to place 0A in the fourth octet of the starting MAC address to identify all of the MAC addresses as Fabric Interconnect A addresses.

13. 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 [?] [X]

First MAC Address : Size :

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

14. Click OK.

15. Click Finish.

16. In the confirmation message, click OK.

17. Right-click MAC Pools under the HANA organization.

18. Select Create MAC Pool to create the MAC address pool.

19. Enter MAC_Pool_B as the name of the MAC pool.

20. Optional: Enter a description for the MAC pool.

21. Click Next.

22. Click Add.

23. Specify a starting MAC address.



The recommendation is to place 0B in the fourth octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses.

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

25. Cisco UCS - Create MAC Pool for Fabric B.

26. Click OK.

27. Click Finish.

28. In the confirmation message, click OK.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID_Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the Prefix as the Derived option.
8. Select Sequential for Assignment Order.

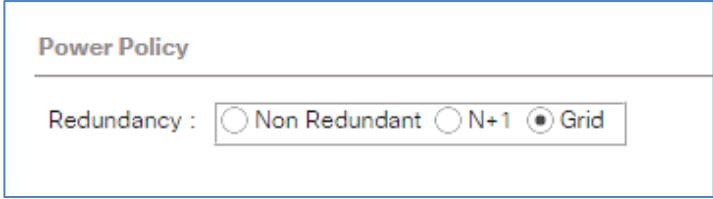
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

13. Click OK.
14. Click Finish.
15. Click OK.

Power Policy

To run Cisco UCS with two independent power distribution units, the redundancy must be configured as Grid. 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 Power Policy to **“Grid.”**
4. Click Save Changes.
5. Click OK.



The screenshot shows a configuration window titled "Power Policy". Below the title is a horizontal line. Underneath, the text "Redundancy :" is followed by three radio button options: "Non Redundant", "N+1", and "Grid". The "Grid" option is selected, indicated by a filled circle next to it.

Power Control Policy

The Power Capping feature in Cisco UCS is designed to save power with a legacy data center use cases. This feature does not contribute much to the high performance behavior of SAP HANA. By choosing the option **“No Cap”** for power control policy, the SAP HANA server nodes will not have a restricted power supply. It is recommended to have this power control policy to ensure sufficient power supply for high performance and critical applications like SAP HANA.

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 HANA 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 ? X

Name :

Description :

Fan Speed Policy :

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 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. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter HANA-FW as the name of the host firmware package.
6. Leave Simple selected.
7. Select the version 3.2(2d) for both the Blade and Rack Packages.
8. Click OK to create the host firmware package.
9. Click OK.

Create Host Firmware Package ? X

Name :

Description :

How would you like to configure the Host Firmware Package?

Simple Advanced

Blade Package :

Rack Package :

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

Adapter
 BIOS
 Board Controller
 CIMC
 FC Adapters
 Flex Flash Controller
 GPUs
 HBA Option ROM
 Host NIC
 Host NIC Option ROM

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 No-Local 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 ? X

Name : No-Local

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

8. Click OK.

Create Server BIOS Policy

To get the best performance for HANA it is required to configure the Server BIOS accurately. 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. Choose Policies > root > Sub-Organization > HANA.
3. Right-click BIOS Policies.
4. Choose Create BIOS Policy.
5. Enter HANA-BIOS as the BIOS policy name.
6. Click OK.

Create BIOS Policy ? X

Name : HANA

Description :

Reboot on BIOS Settings Change :

OK Cancel

7. Under BIOS Policies, click the newly created HANA Policy.
8. In the Main pane, under BIOS Setting choose Disabled for Quiet Boot.

9. Click the Advance tab.
10. The recommendation from SAP for SAP HANA is to disable all Processor C States. This will force the CPU to stay on maximum frequency and allow SAP HANA to run with best performance.
11. Under Processor choose Disabled for all C-States.
12. Set HPC for CPU Performance, Performance for Power Technology, Energy Performance.

Servers / Policies / root / BIOS Policies / HANA

Main **Advanced** Boot Options Server Management Events

Processor Intel Directed IO RAS Memory Serial Port USB PCI GPI LOM and PCIe Slots Trusted Platform Graphics Configuration

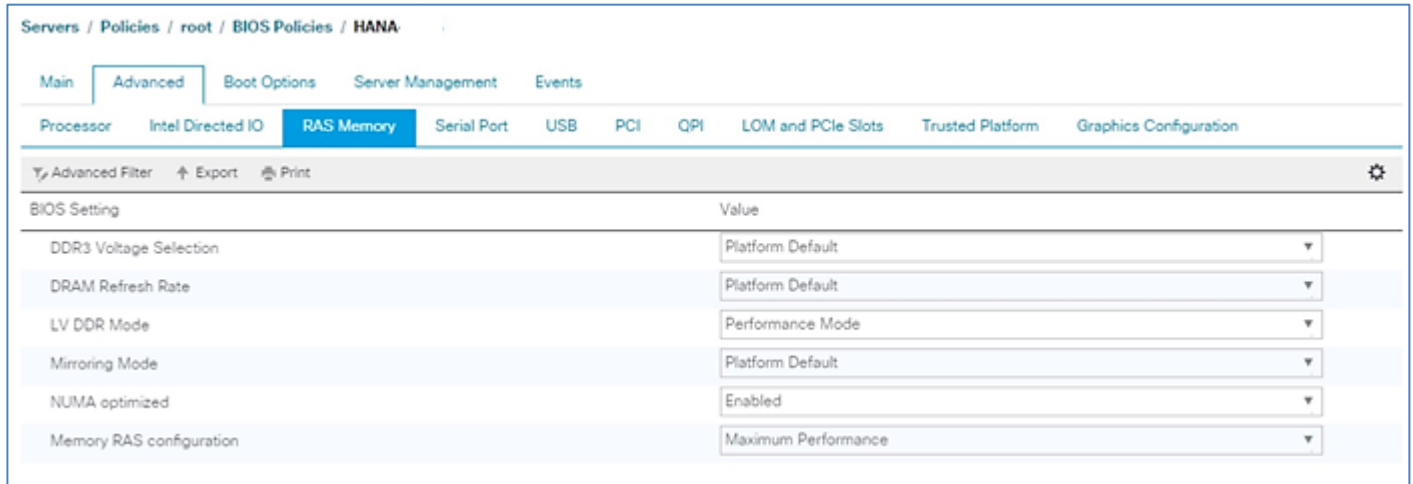
Advanced Filter Export Print

BIOS Setting	Value
P STATE Coordination	Platform Default
Package C State Limit	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	Performance
Energy Performance	Performance
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default

Add Delete Info


Save Changes Reset Values

13. Click RAS Memory.
14. Choose Maximum-Performance for Memory RAS Configuration and Enabled for NUMA optimized.



15. Click Serial Port.

16. Choose Enabled for Serial Port A enable.



17. Click Server Management.

18. Choose 115.2k for BAUD Rate, Enabled for Legacy OS redirection, VT100-PLUS for Terminal type. This is used for Serial Console Access over LAN to all SAP HANA servers.

Servers / Policies / root / BIOS Policies / HANA

Main Advanced Boot Options **Server Management** Events

Advanced Filter Export Print

BIOS Setting	Value
Assert NMI on PERR	Platform Default
Assert NMI on SERR	Platform Default
Baud rate	115.2k
Console redirection	Serial Port A
Flow Control	Platform Default
Legacy OS redirection	Enabled
Putty KeyPad	Platform Default
Terminal type	VT100-PLUS
FRB-2 Timer	Platform Default
OS Boot Watchdog Timer Policy	Platform Default
OS Boot Watchdog Timer Timeout	Platform Default
OS Boot Watchdog Timer	Platform Default
Out of Band Management	Platform Default
Redirection After BIOS POST	Platform Default

19. Click Save Changes.

Create Serial Over LAN Policy

The Serial over LAN policy is required to get console access to all the SAP HANA servers through SSH from the management network. This is used if the server hangs or there is a Linux kernel crash, where the dump is required. To configure the speed in the Server Management tab of the BIOS Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click Serial over LAN Policies.
4. Select Create Serial over LAN Policy.
5. Enter SoL-Console as the Policy name.
6. Select Serial over LAN State to enable.
7. Change the Speed to 115200.
8. Click OK.

Create Serial over LAN Policy ? X

Name :

Description :

Serial over LAN State : Disable Enable

Speed :

Update Default Maintenance Policy

It is recommended to update the default **Maintenance Policy** with the Reboot Policy “User Ack” for the SAP HANA server. This policy will wait for the administrator to acknowledge the server reboot for the configuration changes to take effect.

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. Click Save Changes.
6. Click OK to accept the change.

Servers / Policies / root / Maintenance Policies / default

General Events

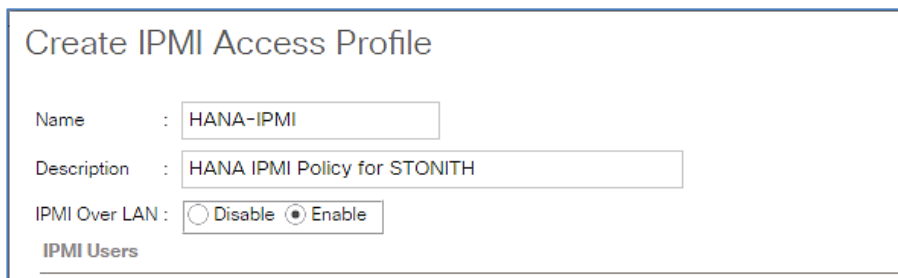
Actions	Properties
<p>Delete</p> <p>Show Policy Usage</p> <p>Use Global</p>	<p>Name : default</p> <p>Description : <input type="text"/></p> <p>Owner : Local</p> <p>Soft Shutdown Timer : <input type="text" value="150 Secs"/></p> <p>Storage Config. Deployment Policy : <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack</p> <p>Reboot Policy : <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack <input type="radio"/> Timer Automatic</p> <p><input checked="" type="checkbox"/> On Next Boot (Apply pending changes at next reboot.)</p>

IPMI Access Profiles

The Serial over LAN access requires an IPMI access control to the board controller. This is also used for the STONITH function of the SAP HANA mount API to kill a hanging server. The default user is 'sapadm' with the password 'cisco'.

To create an IPMI Access Profile, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click IPMI Access Profiles.
4. Select Create IPMI Access Profile
5. Enter HANA-IPMI as the Profile name.



Create IPMI Access Profile

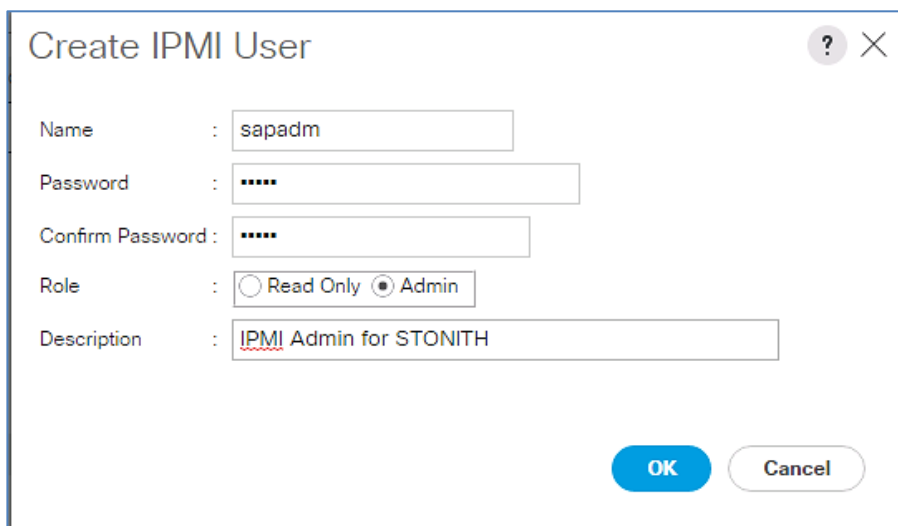
Name :

Description :

IPMI Over LAN : Disable Enable

IPMI Users

6. Click the + (add) button.
7. Enter Username in the Name field and password.
8. Select Admin as Role.



Create IPMI User ? X

Name :

Password :

Confirm Password :

Role : Read Only Admin

Description :

OK **Cancel**

9. Click OK to create user.

10. Click OK to Create IPMI Access Profile.

11. Click OK.

Network Configuration

The core network requirements for SAP HANA are covered by Cisco UCS defaults. Cisco UCS is based on 40-GbE and provides redundancy via the Dual Fabric concept. The Service Profile is configured to distribute the traffic across Fabric Interconnect A and B. During normal operation, the traffic in the Internal Zone and the Data base NFS data traffic is on FI A and all the other traffic (Client Zone and Database NFS log) is on FI B. The inter-node traffic flows from a Blade Server to the Fabric Interconnect A and back to other Blade Server. All the other traffic must go over the Cisco Nexus ACI switches to storage or to the data center network. With the integrated algorithms for bandwidth allocation and quality of service the Cisco UCS and Cisco Nexus distributes the traffic in an efficient way.

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 MTU Column, enter 9216 in the box.
5. Check Enabled Under Priority for Platinum.
6. Click Save Changes in the bottom of the window.
7. Click OK.

LAN									
LAN Uplinks	VLANs	Server Links	MAC Identity Assignment	IP Identity Assignment	QoS	Global Policies	Faults	Events	FSM
Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU			
Platinum	<input type="checkbox"/>	5	<input checked="" type="checkbox"/>	10	N/A	9216			
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	9216			
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	9216			
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	9216			
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	71	9216			
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	2	29	fc	N/A		

Update Default Network Control Policy to Enable CDP

CDP need to be enabled for ACI fabric to learn the MAC address of the End Point devices. To update default Network Control Policy, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > Policies > root > Network Control Policies > default.
3. In the right pane, click the General tab.
4. For CDP: select Enabled radio button.
5. Click Save Changes in the bottom of the window.
6. Click OK.

Figure 85 Network Control Policy to Enable CDP

The screenshot shows the configuration page for the 'default' Network Control Policy. The breadcrumb path is LAN / Policies / root / Network Control Policies / default. The 'General' tab is selected. The page is divided into 'Actions' and 'Properties' sections.

Actions:

- Delete
- Show Policy Usage
- Use Global

Properties:

- Name : **default**
- Description :
- Owner : **Local**
- 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:

- Transmit : Disabled Enabled
- Receive : Disabled Enabled

LAN Tab Configurations

Within Cisco UCS, all the network types for an SAP HANA system are reflected by defined VLANs. Network design from SAP has seven SAP HANA related networks and two infrastructure related networks. The VLAN IDs can be changed if required to match the VLAN IDs in the data center network – for example, ID 224 for backup should match the configured VLAN ID at the data center network switches. Even though nine VLANs are defined, VLANs for all the networks are not necessary if the solution will not use the network. For example, if the Replication Network is not used in the solution, then VLAN ID 300 does not have to be created.

Create VLANs

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

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



In this procedure, Nine VLANs are created.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter HANA-Boot as the name of the VLAN to be used for PXE boot network.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter <<var_boot_vlan_id>> as the ID of the PXE boot network.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.

Repeat steps 1-9 to configure the rest of the VLANs as planned in the table below.

The following are the VLANs used in this CVD:

VLAN name	VLAN ID	Network Address	Netmask
Management	176	192.168.76.x	255.255.255.0
PXE	227	192.168.127.x	255.255.255.0
iSCSI-A	328	192.168.128.x	255.255.255.0
iSCSI-B	329	192.168.129.x	255.255.255.0
NFS	130	192.168.130.x	255.255.255.0
NFS_Data	271	192.168.201.x	255.255.255.0
NFS_Log	278	192.168.228.x	255.255.255.0
Server	220	192.168.220.x	255.255.255.0
Backup	224	192.168.224.x	255.255.255.0
Access	301	10.1.1..x	255.255.255.0

VLAN name	VLAN ID	Network Address	Netmask
DataSource	225	192.168.225.x	255.255.255.0
Application	226	192.168.226.x	255.255.255.0
SysRep	300	10.10.1.x	255.255.255.0

0 shows the overview of VLANs created.

Figure 86 VLAN Definition in Cisco UCS

LAN / LAN Cloud / VLANs					
VLANs					
<input type="checkbox"/> Advanced Filter <input type="checkbox"/> Export <input type="checkbox"/> Print					
Name	ID	Type	Transport	Native	VLAN Sharing
VLAN Access (301)	301	Lan	Ether	No	None
VLAN Application (226)	226	Lan	Ether	No	None
VLAN Backup (224)	224	Lan	Ether	No	None
VLAN DataSource (225)	225	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN DMZ (401)	401	Lan	Ether	No	None
VLAN iSCSI_A (328)	328	Lan	Ether	No	None
VLAN iSCSI_B (329)	329	Lan	Ether	No	None
VLAN Management (176)	176	Lan	Ether	No	None
VLAN NFS (130)	130	Lan	Ether	No	None
VLAN NFS_Data (271)	271	Lan	Ether	No	None
VLAN NFS_Log (278)	278	Lan	Ether	No	None
VLAN PXE (227)	227	Lan	Ether	No	None
VLAN Server (220)	220	Lan	Ether	No	None
VLAN SysRep (300)	300	Lan	Ether	No	None

Assigning Port Channels to VLANs

The approach to have all the SAP HANA networks traffic assigned to a port channel and have another separate port channel for backup network traffic use introduces a disjoint L2 into the topology.

By default, data traffic in Cisco UCS works on a principle of mutual inclusion. All traffic for all VLANs and upstream networks travels along all uplink ports and port channels.

On the other hand, configuration for disjoint L2 networks works on a principle of selective exclusion. Traffic for a VLAN that is designated as part of a disjoint network can only travel along an uplink Ethernet port or port channel that is specifically assigned to that VLAN, and is selectively excluded from all other uplink ports and port channels. However, traffic for VLANs that are not specifically assigned to an uplink Ethernet port or port channel can still travel on all uplink ports or port channels, including those that carry traffic for the disjoint L2 networks. Not associating each uplink port channel to its dedicated VLANs leads to network connectivity problems.

To assign the Port Channels to the VLANs, complete the following steps:

1. In the UCS Manager Navigation pane, click LAN.
2. On the LAN tab, click the LAN node.
3. In the Work pane, click the LAN Uplinks Manager link on the LAN Uplinks tab.

4. In the LAN Uplinks Manager, click VLANs > VLAN Manager. Click Fabric A subtab to configure port channels.
 - a. In the Port Channels and Uplinks table:
 - i. Expand the Port Channels and expand Fabric A and click the Port Channel 11. In the VLANs and VLAN Groups table, expand VLANs. Hold down the Ctrl key and click all VLANs except backup VLAN
 - ii. Click the Add to VLAN/VLAN Group button. Click Yes for the confirmation box.
 - b. In the Port Channels and Uplinks table:
 - i. Click the Port Channel 21. In the VLANs and VLAN Groups table, expand VLANs. Select backup VLAN. Click the Add to VLAN/VLAN Group button. Click Yes for the confirmation box.
5. In the LAN Uplinks Manager, click VLANs > VLAN Manager Click Fabric B subtab to configure port channels.
 - a. In the Port Channels and Uplinks table:
 - i. Expand the Port Channels and expand Fabric B and click the Port Channel 12. In the VLANs and VLAN Groups table, expand VLANs. Hold down the Ctrl key and click all VLANs except backup VLAN.
 - ii. Click the Add to VLAN/VLAN Group button. Click Yes for the confirmation box.
 - b. In the Port Channels and Uplinks table:
 - i. Click the Port Channel 22. In the VLANs and VLAN Groups table, expand VLANs. Select backup VLAN. Click the Add to VLAN/VLAN Group button. Click Yes for the confirmation box.
6. Click OK to close the window.

After a port channel is assigned to one or more VLANs, it is removed from all other VLANs complying with the guidelines for configuring the upstream disjoint L2 networks.

Create QoS Policies

QoS policies assign a system class to the network traffic for a vNIC. To create for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click QoS Policies.
4. Select Create QoS Policy.
5. Enter Platinum as the QoS Policy name.
6. For Priority Select Platinum from the drop-down list.

Create QoS Policy ? X

Name :

Egress

Priority :

Burst(Bytes) :

Rate(Kbps) :

Host Control : None Full

7. Click OK to create the Platinum QoS Policy.

Create vNIC Template

Each VLAN is mapped to a vNIC template to specify the characteristic of a specific network. The vNIC template configuration settings include MTU size, Failover capabilities and MAC-Address pools.

To create vNIC templates for the Cisco UCS environment, complete the following steps:

Create vNIC Template for PXE Service

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter PXE as the vNIC template name.
6. Keep Fabric A selected.
7. 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 PXE
11. Set PXE as the native VLAN.
12. For MTU, enter 1500.
13. In the MAC Pool list, select PXE.

Create vNIC Template

? X

Name : PXE

Description : PXE Boot Network

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter
 VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	management	
<input type="checkbox"/>	NFS	
<input type="checkbox"/>	NFS_Data	
<input type="checkbox"/>	NFS_Log	
<input checked="" type="checkbox"/>	PXE	
<input type="checkbox"/>	Server	

CDN Source : vNIC Name User Defined

MTU : 1500

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

QoS Policy : <not set>

Network Control Policy : default

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel

14. Click OK to create the vNIC template.

15. Click OK.



For most SAP HANA use cases, the network traffic is well distributed across the two Fabrics (Fabric A and Fabric B) using the default setup. In special cases, it can be required to rebalance this distribution for bet-

ter overall performance. This can be done in the vNIC template with the Fabric ID setting. The MTU settings must match the configuration in customer data center. MTU setting of 9000 is recommended for the best performance.

16. Repeat steps 1-15 to create vNIC template for each Network zone.

Create vNIC Template for Internal Network (Server-Server)



Internal Network requires >9.0 Gbps for SAP HANA inter-node communication; choose Platinum QoS Policy created for HANA-Internal vNIC Template.

Fill-in the following values for inputs while creating the rest of the vNIC templates as shown below:

Create vNIC Template
? X

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print ⚙

Select	Name	Native VLAN
<input type="checkbox"/>	management	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>

✓
Server ⊙

CDN Source : vNIC Name User Defined

MTU :

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : ▼

Create vNIC Template for Storage NFS Data Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print ⚙

Select	Name	Native VLAN
<input type="checkbox"/>	Management	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS_Data	<input checked="" type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>

PXE
 Server

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

Create vNIC Template for Storage NFS Log Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter
 VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Select	Name	Native VLAN
<input type="checkbox"/>	Management	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS_Log	<input checked="" type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>

Server
○

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Create vNIC Template for Admin Network

Create vNIC Template ? ✕

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print ⚙

Select	Name	Native VLAN
<input type="checkbox"/>	iSCSI_A	<input type="radio"/>
<input type="checkbox"/>	iSCSI_B	<input type="radio"/>
<input checked="" type="checkbox"/>	Management	<input checked="" type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>

<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : ▼

Create vNIC Template for AppServer Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter ↑ Export Print ⚙️

Select	Name	Native VLAN
<input type="checkbox"/>	Access	<input type="radio"/>
<input checked="" type="checkbox"/>	Application	<input checked="" type="radio"/>
<input type="checkbox"/>	Backup	<input type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>

<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

Create vNIC Template for Backup Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter ↑ Export Print ⚙

Select	Name	Native VLAN
<input type="checkbox"/>	Access	<input type="radio"/>
<input type="checkbox"/>	Application	<input type="radio"/>
<input checked="" type="checkbox"/>	Backup	<input checked="" type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>

<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Create vNIC Template for Access Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	Access	<input checked="" type="radio"/>
<input type="checkbox"/>	Application	<input type="radio"/>
<input type="checkbox"/>	Backup	<input type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>

<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

Create vNIC template for DataSource Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	Access	<input type="radio"/>
<input type="checkbox"/>	Application	<input type="radio"/>
<input type="checkbox"/>	Backup	<input type="radio"/>
<input checked="" type="checkbox"/>	DataSource	<input checked="" type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

OK Cancel

Create vNIC Template for Replication Network

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>
<input checked="" type="checkbox"/>	SysRep	<input checked="" type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

Create vNIC Template for NFS Traffic

Create vNIC Template



Name : NFS

Description : Normal NFS Traffic like /hana/shared

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	iSCSI_B	<input type="radio"/>
<input type="checkbox"/>	Management	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS	<input checked="" type="radio"/>
<input type="checkbox"/>	NFS_Data	<input type="radio"/>
<input type="checkbox"/>	NFS_Log	<input type="radio"/>
<input type="checkbox"/>	PXE	<input type="radio"/>
<input type="checkbox"/>	Server	<input type="radio"/>
<input type="checkbox"/>	SysRep	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU : 9000

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

QoS Policy : <not set>

Network Control Policy : default

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK Cancel

Create vNIC Template for iSCSI via Fabric A

Create vNIC Template



Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs

Advanced Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	DMZ	<input type="radio"/>
<input checked="" type="checkbox"/>	iSCSI_A	<input checked="" type="radio"/>
<input type="checkbox"/>	iSCSI_B	<input type="radio"/>
<input type="checkbox"/>	Management	<input type="radio"/>

Template Type : Initial Template Updating Template

CDN Source : vNIC Name User Defined

MTU :

Policies

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

Create vNIC Template for iSCSI via Fabric B

Create vNIC Template

? ✕

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter
 VM

Warning

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

Template Type : Initial Template Updating Template

VLANS

Advanced Filter Export Print ⚙️

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	DMZ	<input type="radio"/>
<input type="checkbox"/>	iSCSI_A	<input type="radio"/>
<input checked="" type="checkbox"/>	iSCSI_B	<input checked="" type="radio"/>
	...	
<input type="checkbox"/>	Management	<input type="radio"/>
<input type="checkbox"/>	MEC	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

OK
Cancel

Create vNIC/vHBA Placement Policy

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



Cisco UCS B-Series server B480M5 is configured with one VIC1340 port expander and one VIC1380.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter HANA as the name of the placement policy.
6. Click 1 and select Assigned Only.
7. Click 2 and select Assigned Only.
8. Click OK and then click OK again.

Create Placement Policy

Name :

Virtual Slot Mapping Scheme : Round Robin Linear Ordered

Advanced Filter | Export | Print

Virtual Slot	Selection Preference
1	Assigned Only
2	Assigned Only
3	All
4	All

OK Cancel

Create PXE Boot Policies

To create PXE boot policies, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organization > HANA.
3. Right-click Boot Policies.

4. Select Create Boot Policy.
5. Enter PXE-Boot as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Check the Reboot on Boot Order Change option.
8. Expand the Local Devices drop-down menu and select Add CD/DVD.
9. Expand the vNICs section and select Add LAN Boot.
10. In the Add LAN Boot dialog box, enter PXE-Boot.
11. Click OK.
12. Click OK.
13. Click OK to save the boot policy. Click OK to close the Boot Policy window.

Create Boot Policy

Name :

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	WWN	LUN N...	Slot N...	Boot N...
CD/DVD	1						
LAN	2						
LAN PXE		PXE	Primary				

Create Service Profile Templates SAP HANA Scale-Out

The LAN configurations and relevant SAP HANA policies must be defined prior to creating, a Service Profile Template.

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 > Sub-Organization > HANA.
3. Right-click HANA.
4. Select Create Service Profile Template (Expert) to open the Create Service Profile Template wizard.
5. Identify the service profile template:
 - a. Enter HANA as the name of the service profile template.
 - b. Select the Updating Template option.
 - c. Under UUID, select HANA-UUID as the UUID pool.
 - d. 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 :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-HANA**

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:

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

6. Configure the Local Storage Options:
 - a. Keep the default setting for Specific Storage Profile.
 - b. Select No storage Profile in the Storage Profile Policy tab.

Create Service Profile Template

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

Specific Storage Profile | **Storage Profile Policy** | Local Disk Configuration Policy

Storage Profile: [Create Storage Profile](#)

- c. Select Local Disk Configuration Policy and set No-Local for the Local Storage option.
 - d. Click Next.

Create Service Profile Template

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

Specific Storage Profile Storage Profile Policy **Local Disk Configuration Policy**

Local Storage:

Create Local Disk Configuration Policy

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 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**

7. Configure the networking options:

- a. Keep the default setting for Dynamic vNIC Connection Policy.
- b. Select the Expert option to configure the LAN connectivity.
- c. Click the upper Add button to add a vNIC to the template.
- d. In the Create vNIC dialog box, enter PXE as the name of the vNIC.
- e. Select the Use vNIC Template checkbox.
- f. In the vNIC Template list, select PXE
- g. In the Adapter Policy list, select Linux.
- h. Click OK to add this vNIC to the template.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

8. Repeat the steps c-h for each vNIC. The screenshots are shown below:

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

9. Add vNIC for HANA-AppServer.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

10. Add vNIC for HANA-System-Replication.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

11. Add vNIC for user access network.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

12. Add vNIC for HANA-Backup.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

13. Add vNIC for normal NFS traffic.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

14. Add vNIC for HANA-Admin.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

Create Ethernet Adapter Policy

15. Add vNIC for NFS-Data

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

Create Ethernet Adapter Policy

16. Add vNIC for NFS-Log

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

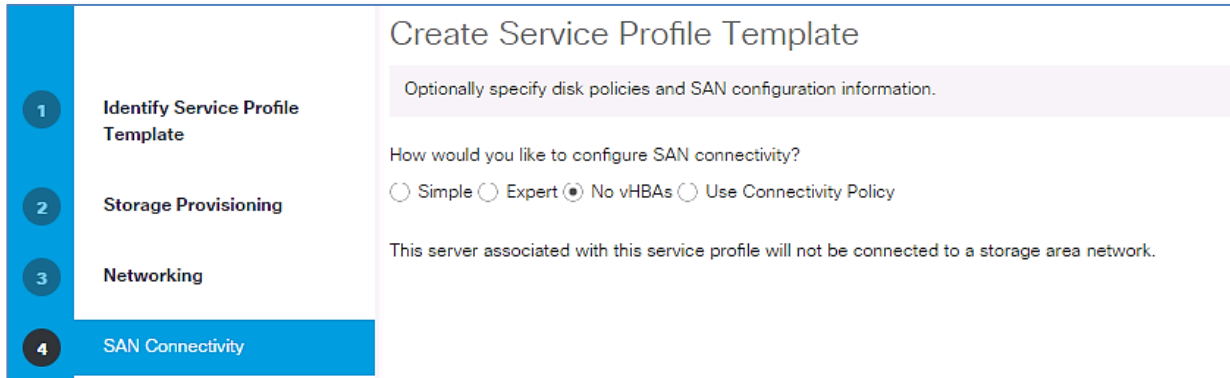
Create vNIC Template

Adapter Performance Profile

Adapter Policy :


Create Ethernet Adapter Policy

17. Review the table in the Networking page to make sure that all vNICs were created. Decide if you need the optional NIC's in the configuration.
18. Click Next.
19. Set no Zoning options and click Next.



20. Set the vNIC/vHBA placement options.
21. For Cisco UCS B200 M5 and B480 M5 servers:

a. In the “Select Placement” list, select the HANA placement policy.



- b. Select vCon1 and assign the vNICs to the virtual network interfaces policy in the following order:
 - i. PXE
 - ii. NFS-Data
 - iii. Management
 - iv. Access
 - v. SysRep

Name	Or...▼	Selection Preference	Admi...
▼ vCon 1		Assigned Only	
vNIC PXE	1		ANY
vNIC NFS-Data	2		ANY
vNIC Mgmt	3		ANY
vNIC Access	4		ANY
vNIC SysRep	5		ANY

- c. Select vCon2 and assign the vNICs to the virtual network interfaces policy in the following order:
- i. Server
 - ii. NFS-Log
 - iii. Application
 - iv. Backup
 - v. NFS

Name	Or...▲	Selection Preference	Admi...
▼ vCon 2		Assigned Only	
vNIC Server	1		ANY
vNIC NFS-Log	2		ANY
vNIC Application	3		ANY
vNIC Backup	4		ANY
vNIC NFS	5		ANY

Virtual Network Interfaces Policy (read only)		
Name	Order	Selection Preference
▼ vCon 3		Assigned Only
vNIC NFS-Log	1	
vNIC SysRep	2	
▼ vCon 4		Assigned Only
vNIC NFS	1	
vNIC Server	2	
↑ Move Up ↓ Move Down		

- d. Click Next.

22. No change required on the vMedia Policy, click Next.

23. Set the server boot order:

- a. Select PXE-Boot for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **PXE-Boot**
 Description : **PXE-Boot Policy**
 Reboot on Boot Order Change : **Yes**
 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

+ - Advanced Filter Export Print

Name	Order	vNIC/vHB..	Type	WWN	LUN Name	Slot Num...	Boot Nam..	Boot Path	De
CD/DVD	1								
LAN	2								
LAN PXE			PXE				Primary		

- b. Click Next.

24. Add a maintenance policy:

- a. Select the default Maintenance Policy.
- b. Click Next.

25. Specify the server assignment:

- a. Choose Up as the power state to be applied when the profile is associated with the server.
- b. Expand Firmware Management at the bottom of the page and choose HANA-FW from the Host Firmware list.
- c. Click Next.

Create Service Profile Template [?] X

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment:

Select the power state to be applied when this profile is associated with the server.

Up Down

The service profile template is not automatically associated with a server. Either select a server from the list or associate the service profile manually later.

⊖ Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package:

26. Set Operational Policies:

- a. Select HANA for Bios Policy
- b. Select the IPMI Access Profile
- c. Select the Serial over LAN Console

Create Service Profile Template [?] X

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 :

⊖ External IPMI Management Configuration

If you want to access the CIMC on the server externally, select an IPMI access profile. The users and passwords in that profile will be populated into the CIMC when the profile is associated with the server.

IPMI Access Profile :

To enable Serial over LAN access to the server, select an SoL configuration profile.

SoL Configuration Profile:

Name : **SoL-Console**
Description : **Serial over Lan Settings**


⊕ Management IP Address

- d. Click Create Outband Management Pool.
- e. Specify the name of the Out of band management pool.
- f. Select Sequential Order.


g. Click Next.

- h. Define the out of band management ip pool (mgmt. VLAN).
- i. Define the size of the IP pool.
- j. Specify the default GW and the DNS server (if available).
- k. Click Next.

- l. Do not specify IPv6 IP addresses for the management pool.
- m. Click Finish.
- n. Select the management pool you just created.
- o. Select the default monitoring thresholds.



- p. Select the HANA policy for the Power Control Configuration.
- q. Leave the two other fields as default.

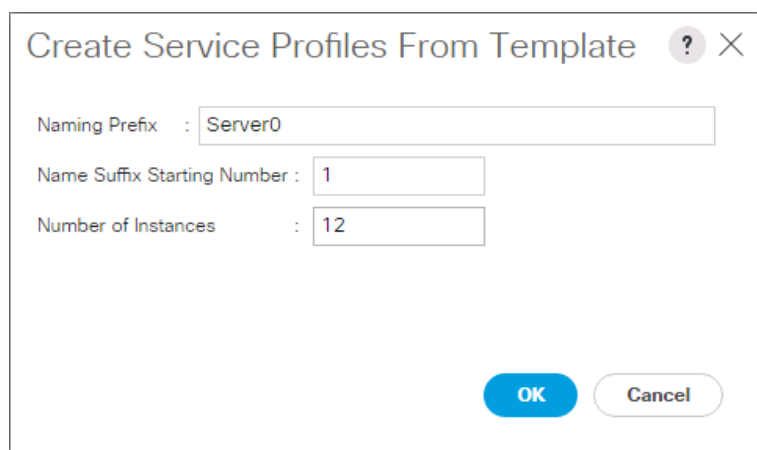


27. Complete the service profile generation by clicking Finish.

Create Service Profile from the Template

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organization > HANA > Service Template HANA.
3. Right-click Service Template HANA and select Create Service Profiles from Template.
4. Enter Server0 as the service profile prefix.
5. Enter 1 as 'Name Suffix Starting Number'.
6. Enter 12 as the 'Number of Instances' based on the servers available.
7. Click OK to create the service profile.



Create Service Profiles From Template ? X

Naming Prefix : Server0

Name Suffix Starting Number : 1

Number of Instances : 12

OK Cancel




As soon as the specified number of Service Profiles are created, profiles are available for association with the servers.

8. Assign the Service Profile to a server:
 - a. In Cisco UCS Manager, click the Servers tab in the navigation pane.
 - b. Choose Service Profile > root > Sub-Organization > HANA > HANA-Server01.
 - c. Right-click HANA-Server01 and choose Change Service Profile Association.
 - d. For Server Assignment Choose, choose existing Server for the drop-down list.
 - e. Click All Servers.
 - f. Choose the Server as recommended.
9. Repeat steps a-f for each HANA Servers.



The configuration will start immediately after acknowledge the reboot.



Create Service Profile Templates SAP HANA iSCSI

To create the service profile template for SAP HANA for iSCSI boot for both SUSE and Red Hat implementations, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organization > HANA.
3. Right-click HANA.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
 - a. Enter HANA-iSCSI as the name of the service profile template.
 - b. Select the Updating Template option.
 - c. Under UUID, select HANA-UUID as the UUID pool.
 - d. 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 the template and enter a description.

Name :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-HANA**

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:

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

6. Configure the Storage Provisioning:

- a. Select “No Storage Profile” under Storage Profile Policy.
- b. Select “No-Local” under Local Disk Configuration Policy.
- c. Click Next.

7. Network configuration:

- a. Select “No Dynamic VNIC policy.”
- b. Select Expert Configuration.
- c. Add the necessary networks.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Adapter Performance Profile

Adapter Policy :

[Create vNIC Template](#)

[Create Ethernet Adapter Policy](#)

Figure 87 Backup Network

Create vNIC [?] [X]

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

Create Ethernet Adapter Policy

Figure 88 Management Network

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

Create Ethernet Adapter Policy

Figure 89 NFS-Data Network

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

Create Ethernet Adapter Policy

Figure 90 NFS-Log Network

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template : ▼

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy : ▼

[Create Ethernet Adapter Policy](#)

Figure 91 Server-Server Network (Optional for Scale-Up)

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template : ▼

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy : ▼

[Create Ethernet Adapter Policy](#)

8. Click add vNIC and create the two iSCSI vNICs.

Figure 92 Create iSCSI A

Create vNIC ? X

Name :

MAC Address

MAC Address Assignment:

[Create MAC Pool](#)

The MAC address will be automatically assigned from the selected pool.

Use vNIC Template :

Fabric ID : Fabric A Fabric B Enable Failover

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

VLANs
VLAN Groups

Advanced Filter
Export
Print
⚙

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	iSCSI_A	<input checked="" type="radio"/>
<input type="checkbox"/>	iSCSI_B	<input type="radio"/>
<input type="checkbox"/>	Management	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

Pin Group : [Create LAN Pin Group](#)

Operational Parameters

Adapter Performance Profile

Adapter Policy : [Create Ethernet Adapter Policy](#)

Figure 93 Create iSCSI B Same VLAN as iSCSIa

Create vNIC ? X

Name :

MAC Address

MAC Address Assignment: ▼

[Create MAC Pool](#)
The MAC address will be automatically assigned from the selected pool.

Use vNIC Template :

Fabric ID : Fabric A Fabric B Enable Failover

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash.

VLANs VLAN Groups

Advanced Filter ↑ Export Print ⚙

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	iSCSI_A	<input checked="" type="radio"/>
<input type="checkbox"/>	iSCSI_B	<input type="radio"/>
<input type="checkbox"/>	Management	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>
<input type="checkbox"/>	NFS	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU :

Pin Group : [Create LAN Pin Group](#)

+ Operational Parameters

Adapter Performance Profile

Adapter Policy : [Create Ethernet Adapter Policy](#)

9. Click +iSCSI vNICs.

10. Click Create IQN Suffix Pool.

Figure 94 Create the Suffix Pool

Create IQN Suffix Pool [?] [X]

1 Define Name and Description

2 Add IQN Blocks

Name : iSCSI-Boot-A

Description : iSCSI Boot Initiator Name

Prefix : iqn.1992-08.com.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

Figure 95 Create the Block

Create a Block of IQN Suffixes [?] [X]

Suffix : ucs

From : 0

Size : 12

Figure 96 Result

Create IQN Suffix Pool [?] [X]

+ - Advanced Filter Export Print [Settings]

Name	From	To
ucs:0 - ucs:11	0	11

1 Define Name and Description

2 Add IQN Blocks

11. Select the Initiator Name Assignment for this profile.

12. Click add iSCSI initiator interfaces.

Figure 97 Add iSCSI vNICa

The screenshot shows a 'Create iSCSI vNIC' dialog box with the following fields and options:

- Name: iSCSIa
- Overlay vNIC: iSCSIa
- iSCSI Adapter Policy: default (with a 'Create iSCSI Adapter Policy' link)
- VLAN: iSCSI A (native)
- iSCSI MAC Address: (empty field)
- MAC Address Assignment: Select(None used by default) (with a 'Create MAC Pool' link)

Figure 98 Add iSCSIvNICb

The screenshot shows a 'Create iSCSI vNIC' dialog box with the following fields and options:

- Name: iSCSIb
- Overlay vNIC: iSCSIb
- iSCSI Adapter Policy: default (with a 'Create iSCSI Adapter Policy' link)
- VLAN: iSCSI B (native)
- iSCSI MAC Address: (empty field)
- MAC Address Assignment: Select(None used by default) (with a 'Create MAC Pool' link)

13. Click Next to configure the SAN.
14. Select no vHBA.
15. Click Next.
16. No Zoning in this environment.
17. Configure the vNIC Placement.

Figure 99 vCon 1 and 2

Please select one Virtual Network Interface and one or more vNICs or vHBAs

vNICs | vHBAs

Name

No data available

>> assign >>
<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Selection Preference
▼ vCon 1		Assigned Only
vNIC iSCSIa	1	
vNIC Mgmt	2	
vNIC NFS-Data	3	

vNICs | vHBAs

Name

No data available

>> assign >>
<< remove <<

Virtual Network Interfaces Policy (read only)

Name	Order	Selection Preference
▶ vCon 1		Assigned Only
▼ vCon 2		Assigned Only
vNIC Access	1	
vNIC iSCSIb	2	
vNIC NFS-Log	3	
vCon 3		Assigned Only

↑ Move Up ↓ Move Down

18. Add other required vNICs to the vCONs so as to share the traffic load evenly.

19. Do not create a vMedia Policy, click Next.

20. Server Boot Order.

21. Create a new iSCSI boot order:

- a. Add a local DVD.
- b. Add the iSCSI NICs (add iSCSI Boot).
- c. Add iSCSIa.
- d. Add iSCSIb.

Figure 100 Create a iSCSI Boot Policy

Create Boot Policy

Name : iSCSI
 Description : iSCSI boot order
 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	WWN	LUN Na...	Slot Nu...	Boot Na...	Boot Path	Descript...
CD/...	1								
▼ iSCSI 2									
iS...		iSCSIa	Primary						
iS...		iSCSIb	Second...						

Figure 101 Select the New iSCSI Boot Profile and Set iSCSI Boot Parameter

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Boot Policy: iSCSI [Create Boot Policy](#)

Name : iSCSI
 Description : iSCSI boot order
 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	Order	vNIC/vH...	Type	WWN	LUN Name	Slot Num...	Boot Name	Boot Path	Description
CD/DVD	1								
▼ iSCSI 2									
iSCSI		iSCSIa	Primary						
iSCSI		iSCSIb	Secondary						

22. First Target add the NetApp iSCSI target in this case it is iqn.1992-08.com.netapp:sn.35084cc1105511e7983400a098aa4cc7:vs.4

Create iSCSI Static Target [?] [X]

iSCSI Target Name : iqn.1992-08.com.netapp:
Priority : 1
Port : 3260
Authentication Profile : <not set> [Create iSCSI Authentication Profile](#)
IPv4 Address : 192.168.128.11
LUN ID : 0

Figure 102 Second Target

Create iSCSI Static Target [?] [X]

iSCSI Target Name : 7983400a098aa4cc7:vs.4
Priority : 2
Port : 3260
Authentication Profile : <not set> [Create iSCSI Authentication Profile](#)
IPv4 Address : 192.168.128.12
LUN ID : 0

- 23. Click OK.
- 24. Select default Maintenance Policy.
- 25. Server Assignment select Assign-Later and HANA-FW firmware Policy.
- 26. Select Operational Policies.

Figure 103 Server Policies 1

Create Service Profile Template

Optionally specify information that affects how the system operates.

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

BIOS Policy : HANA ▼

External IPMI Management Configuration

If you want to access the CIMC on the server externally, select an IPMI access profile. The users and passwords in that profile will be populated into the CIMC when the profile is associated with the server.

IPMI Access Profile : HANA ▼ [Create IPMI Access Profile](#)

To enable Serial over LAN access to the server, select an SoL configuration profile.

SoL Configuration Profile: Select a Policy to use (no SoL Access by default) ▼

[Create Serial over LAN Policy](#)

This service profile will not have Serial over LAN access.

+ Management IP Address

+ Monitoring Configuration (Thresholds)

Figure 104 Server Policies 2

+ Monitoring Configuration (Thresholds)

Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : HANA ▼ [Create Power Control Policy](#)

Scrub Policy

Scrub Policy : default ▼ [Create Scrub Policy](#)

KVM Management Policy

KVM Management Policy : default ▼ [Create KVM Management Policy](#)

< Prev Next > **Finish** Cancel

27. Click Finish to complete the Service Profile generation.

Create Service Profile Templates for SAP HANA Scale-Up

To create the service profile template for SAP HANA Scale-Up, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organization > HANA.
3. Right-click HANA.

4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
 - a. Enter HANA-Scale-UP as the name of the service profile template.
 - b. Select the Updating Template option.
 - c. Under UUID, select HANA-UUID as the UUID pool.
 - d. Click Next.
6. Configure the networking options:
 - a. Keep the default setting for Dynamic vNIC Connection Policy.
 - b. Select the Expert option to configure the LAN connectivity.
 - c. Click the upper Add button to add a vNIC to the template.
 - d. In the Create vNIC dialog box, enter PXE as the name of the vNIC.
 - e. Select the Use vNIC Template checkbox.
 - f. In the vNIC Template list, select HANA-Boot
 - g. In the Adapter Policy list, select Linux.
 - h. Click OK to add this vNIC to the template.



Repeat steps c-h for each Vnic to add vNICs for HANA-Storage, HANA-Client, HANA-AppServer, HANA-DataSource, HANA-Replication, HANA-Backup and HANA-Admin.

- i. Review the table in the Networking page to make sure that all vNICs were created.
- j. Click Next.



Even though eight Networks were defined, they are optional and if they are not needed in your deployment, the addition of a vNIC template for an optional network is not required.

7. Configure the storage options:
 - a. No change is required for a local disk configuration policy.
 - b. **Select the No vHBAs option for the “How would you like to configure SAN connectivity?” field.**
 - c. Click Next.
8. Set no Zoning options and click Next.
9. Set the vNIC/vHBA placement options.
10. For Cisco UCS B200 M5 / Cisco UCS B480 M5 Servers with two VIC cards:
 - a. **In the “Select Placement” list, select the HANAplacement policy.**
 - b. Select vCon1 and assign the vNICs to the virtual network interfaces policy in the following order:

- i. PXE
 - ii. Client
 - iii. DataSource
 - iv. SysRep
 - v. NFS-Data
 - c. Select vCon2 and assign the vNICs to the virtual network interfaces policy in the following order:
 - i. Access
 - ii. HANA-Backup
 - iii. HANA-AppServer
 - iv. HANA-Admin
 - v. NFS-log
 - d. Review the table to verify that all vNICs were assigned to the policy in the appropriate order.
 - e. Click Next.
11. No Change required on the vMedia Policy, click Next.
12. Set the server boot order: Select PXE-Boot for Boot Policy.
13. Click Next.
14. Add a maintenance policy:
- a. Select the default Maintenance Policy.
 - b. Click Next.
15. Specify the server assignment:
- a. In the Pool Assignment list, select the appropriated pool created for scale-up servers.
 - b. Optional: Select a Server Pool Qualification policy.
 - c. Select Down as the power state to be applied when the profile is associated with the server.
 - d. Expand Firmware Management at the bottom of the page and select HANA-FW from the Host Firmware list.
 - e. Click Next.
16. Add operational policies:
- a. In the BIOS Policy list, select HANA.
 - b. Leave External IPMI Management Configuration as <not set> in the IPMI Access Profile. Select SoL-Console in the SoL Configuration Profile.
 - c. Expand Management IP Address, in the Outband IPv4 tap choose ext-mgmt in the Management IP Address Policy.
 - d. Expand Power Control Policy Configuration and select **No-Power-Cap** in the Power Control Policy list.
17. Click Finish to create the service profile template.

18. Click OK in the confirmation message.

Create Service Profiles

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organization > HANA > Service Template HANA-Scale-UP.
3. Right-click HANA-Scale-UP and select Create Service Profiles from Template.
4. Enter appropriate name for the service profile prefix.
5. Enter 1 as 'Name Suffix Starting Number'.
6. Enter appropriate number of service profile to be created in the 'Number of Instances'.
7. Click OK to create the service profile.

Storage Configuration

Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet in the [ONTAP 9.2 Software Setup Guide](#) located in the NetApp® ONTAP® 9 Documentation Center.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [ONTAP 9.2 Software Setup Guide](#) to learn about configuring ONTAP software. Table 20 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 20 ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Data ONTAP 9.1 URL	<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 the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version of software being booted, continue with step 14.

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

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

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



This web server must be pingable.

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

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.



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 the following message displays:

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

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



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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version of software being booted, continue with step 14.

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

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

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



This web server must be pingable.

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

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.



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



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

Set Up Node 01

Table 21 lists all the parameters required to set up the ONTAP cluster.

Table 21 ONTAP Cluster Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
NFS license key	<nfs-license-key>
iSCSI license key	<iscsi-license-key>
NetApp SnapRestore® license key	<snaprestore-license-key>
NetApp SnapVault® license key	<snapvault-license-key>
NetApp SnapMirror® license key	<snapmirror-license-key>
NetApp FlexClone® license key	<flexclone-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>

Cluster Detail	Cluster Detail Value
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-SP-ip>
Node 01 service processor IP netmask	<node01-SP-mask>
Node 01 service processor IP gateway	<node01-SP-gateway>
Node 02 service processor IP address	<node02-SP-ip>
Node 02 service processor IP netmask	<node02-SP-mask>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
Time zone	<timezone>
NTP server IP address	<ntp-ip>
SNMP contact information	<snmp-contact>
SNMP location	<snmp-location>
DFM server or another fault management server FQDN to receive SNMP traps	<oncommand-um-server-fqdn>
SNMPv1 community string	<snmp-community>
Mail host to send NetApp AutoSupport® messages	<mailhost>
Storage admin email for NetApp AutoSupport	<storage-admin-email>

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.1 software 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.
```

```
You can return to cluster setup at any time by typing "cluster setup".
```

To accept a default or omit a question, do not enter a value.

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

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

Type yes to confirm and continue {yes}: yes

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

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

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

2. To complete the cluster setup, press Enter.



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

Do you want to create a new cluster or join an existing cluster? {create, join}:
 create

Do you intend for this node to be used as a single node cluster? {yes, no} [no]:

Will the cluster network be configured to use network switches? [yes]:
 no

Existing cluster interface configuration found:

Port	MTU	IP	Netmask
e0a	9000	169.254.92.173	255.255.0.0
e0b	9000	169.254.244.224	255.255.0.0

Do you want to use this configuration? {yes, no} [yes]:

Enter the cluster administrator's (username "admin") password:

Retype the password:

Step 1 of 5: Create a Cluster
 You can type "back", "exit", or "help" at any question.

Enter the cluster name: <clustername>
 Enter the cluster base license key: <cluster-base-license-key>

Creating cluster <clustername>

Adding nodes

Cluster <clustername> has been created.

Step 2 of 5: Add Feature License Keys
 You can type "back", "exit", or "help" at any question.
 Enter an additional license key []:<nfs-license-key>

NFS License was added.

```
Enter an additional license key []:<iscsi-license-key>
```

```
iSCSI License was added.
```

```
Enter an additional license key []:<snaprestore-license-key>
```

```
SnapRestore License was added.
```

```
Enter an additional license key []:<snapvault-license-key>
```

```
SnapVault License was added.
```

```
Enter an additional license key <flexclone-license-key>
```

```
FlexClone License was added.
```

```
Enter an additional license key []:<snapmirror-license-key>
```

```
SnapMirror License was added.
```

```
Enter an additional license key []:
```

```
Step 3 of 5: Set Up a Vserver for Cluster Administration  
You can type "back", "exit", or "help" at any question.
```

```
Enter the cluster management interface port: e0c
```

```
Enter the cluster management interface IP address: <clustermgmt-ip>
```

```
Enter the cluster management interface netmask: <clustermgmt-mask>
```

```
Enter the cluster management interface default gateway: <clustermgmt-gateway>
```

```
A cluster management interface on port <clustermgmt-port> with IP address <clustermgmt-ip> has been  
created. You can use this address to connect to and manage the cluster.
```

```
Enter the DNS domain names: <dns-domain-name>
```

```
Step 4 of 5: Configure Storage Failover (SFO)  
You can type "back", "exit", or "help" at any question.
```

```
SFO will be enabled when the partner joins the cluster.
```

```
Step 5 of 5: Set Up the Node  
You can type "back", "exit", or "help" at any question.
```

```
Where is the controller located []: DataCenterA
```

```
Cluster "<clustername>" has been created.
```

```
To complete cluster setup, you must join each additional node to the cluster  
by running "cluster setup" on each node.
```

```
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 (https:// <clustermgmt-ip>).
```

```
To access the command-line interface, connect to the cluster management  
IP address (for example, ssh admin@<clustermgmt-ip>).
```

```
Wed Mar 22 08:09:43 UTC 2017
login:
```



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, we assume that it is on the same subnet.

Set Up Node 02

From a console port program attached to the storage controller B (node 02) console port, run the node setup script. This script appears when ONTAP 9.1 software boots on the node for the first time.

1. Follow the prompts to set up node 02:

```
Welcome to 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 setup wizard.
  Any changes you made before quitting will be saved.

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

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

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node02-mgmt-ip>
Enter the node management interface netmask: <node02-mgmt-mask>
Enter the node management interface default gateway: <node02-mgmt-gateway>
A node management interface on port e0M with IP address <node02-mgmt-ip> has been created

Use your web browser to complete cluster setup by accessing https://<node02-mgmt-ip>

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

2. To complete cluster setup, press Enter.

```
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.

Do you want to create a new cluster or join an existing cluster? {join}:
join

Existing cluster interface configuration found:

Port      MTU      IP              Netmask
e0a       9000     169.254.135.215 255.255.0.0
e0b       9000     169.254.180.204 255.255.0.0

Do you want to use this configuration? {yes, no} [yes]: yes
```

```

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 [A300-HANA]:

Joining cluster <clustername>
Starting cluster support services
This node has joined the cluster <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.

This node has been joined to cluster "<clustername>".

To complete cluster setup, you must join each additional node to the cluster
by running "cluster setup" on each node.

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 (https:// <clustermgmt-ip>).

To access the command-line interface, connect to the cluster management
IP address (for example, ssh admin@<clustermgmt-ip>).

Notice: HA is configured in management.

Wed Mar 22 08:12:30 UTC 2017
login:

```

Log into the Cluster

To log into 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.

Set Auto-Revert on Cluster Management

To set the auto-revert parameter on the cluster management interface, run the following command:

```
network interface modify -vserver <clustername> -lif cluster_mgmt -auto-revert true
```



A storage virtual machine (SVM) is referred to as a Vserver (or vservers) in the GUI and CLI.

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, e0d, e1a, and e1e) should be removed from the default broadcast domain, leaving just the management network ports (e0c and e0M).

To make the changes, the following commands must be performed on each storage node. Storage nodes are named after the cluster name with an appended number. To perform this task, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports <clustername>-01:e0d,<clustername>-01:e1a,
<clustername>-01:e1e,<clustername>-02:e0d,<clustername>-02:e1a,<clustername>-02:e1e

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 <clustername>-01 -address-family IPv4 -enable true -dhcp
none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>

system service-processor network modify -node <clustername>-02 -address-family IPv4 -enable true -dhcp
none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```




The service processor IP addresses should be in the same subnet as the node management IP addresses.

Create Aggregates



Advanced Data Partitioning (ADPv2) creates a root partition and two data partitions on each SSD drive in an All Flash FAS configuration. Disk auto assign should assign one data partition to each node in an HA pair.

Figure 105 Disk Partitioning ADPv2



An aggregate containing the root volume for each storage controller is created during the ONTAP software setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it should contain.

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr01 -node <clustername>-01 -diskcount 23
aggr create -aggregate aggr02 -node <clustername>-02 -diskcount 23
```



Use all disks except for one spare (23) to create the aggregates.



The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until both `aggr01` and `aggr02` are online.

2. (Optional) Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02. The aggregate is automatically renamed if system-guided setup is used.

```
aggr show
aggr rename -aggregate aggr0 -newname <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 <clustername>-01 and <clustername>-02 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 <clustername>-01 -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 <node02-mgmt-ip> -node <clustername>-01
storage failover modify -hwassist-partner-ip <node01-mgmt-ip> -node <clustername>-02
```

Disable Flow Control on 40GbE Ports

NetApp recommends disabling flow control on all the 40GbE 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 <clustername>-01 -port e1a,e1e -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 <clustername>-02 -port e1a,e1e -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
```

Configure Network Time Protocol

To configure time synchronization on the cluster, complete the following steps:

1. Set the time zone for the cluster.

```
timezone <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, 201704041735.17).

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

```
cluster time-service ntp server create -server <ntp-ip>
```

Configure Simple Network Management Protocol

To configure the Simple Network Management Protocol (SNMP), complete the following steps:

1. 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 <snmp-contact>
snmp location "<snmp-location>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <oncommand-um-server-fqdn>
```

Configure SNMPv1 Access

To configure SNMPv1 access, set the shared, secret plain-text password (called a community):

```
snmp community add ro <snmp-community>
```

Configure AutoSupport

NetApp 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 <mailhost> -transport https -support
enable -noteto <storage-admin-email>
```

Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command:

```
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 Broadcast Domains

Figure 106 shows the physical network connection and the virtual LANs (VLANs) used for this setup.

Figure 106 LANs and Broadcast Domains




Table 22 Cluster Networking Requirements

Cluster Detail	Cluster Detail Value	Value used in the CVD setup
NFS data VLAN ID	<data-vlan-id>	201
NFS Log VLAN ID	<log-vlan-id>	228
PXE VLAN ID	<pxe-vlan-id>	127
iSCSI a VLAN ID	<iscsi-a-vlan-id>	128
iSCSI b VLAN ID	<iscsi-b-vlan-id>	129
NFS datastore VLAN ID	<nfs-vlan-id>	130
Storage backend VLAN ID	<stbackend-vlan-id>	224

All broadcast domains, except for the PXE boot network, must be created with an MTU size of 9000 (jumbo frames):

```

broadcast-domain create -broadcast-domain NFS-data -mtu 9000
broadcast-domain create -broadcast-domain NFS-log -mtu 9000
broadcast-domain create -broadcast-domain PXE -mtu 1500
broadcast-domain create -broadcast-domain iSCSI-a -mtu 9000
broadcast-domain create -broadcast-domain iSCSI-b -mtu 9000
broadcast-domain create -broadcast-domain NFS -mtu 9000
broadcast-domain create -broadcast-domain storage-backend -mtu 9000

```

Create Interface Groups

To create the Link Aggregation Control Protocol (LACP) interface groups for the 40GbE data interfaces, run the following commands:

```

ifgrp create -node <clustername>-01 -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <clustername>-01 -ifgrp a0a -port ela
ifgrp add-port -node <clustername>-01 -ifgrp a0a -port ele

ifgrp create -node <clustername>-02 -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <clustername>-02 -ifgrp a0a -port ela
ifgrp add-port -node <clustername>-02 -ifgrp a0a -port ele

ifgrp show

```

Create VLANs

To create VLANs, complete the following steps:

1. Create boot VLAN ports and add them to the PXE broadcast domain.

```

network port modify -node <clustername>-01 -port a0a -mtu 9000
network port modify -node <clustername>-02 -port a0a -mtu 9000

network port vlan create -node <clustername>-01 -vlan-name a0a-<pxe-vlan-id>
network port vlan create -node <clustername>-02 -vlan-name a0a-<pxe-vlan-id>

broadcast-domain add-ports -broadcast-domain PXE -ports <clustername>-01:a0a-<pxe-vlan-id>,
<clustername>-02:a0a-<pxe-vlan-id>

```

2. Create HANA data VLAN ports and add them to the NFS-Data broadcast domain.

```

network port vlan create -node <clustername>-01 -vlan-name a0a-<data-vlan-id>
network port vlan create -node <clustername>-02 -vlan-name a0a-<data-vlan-id>

broadcast-domain add-ports -broadcast-domain NFS-Data -ports <clustername>-01:a0a-<data-vlan-id>,
<clustername>-02:a0a-<data-vlan-id>

```

3. Create HANA log VLAN ports and add them to the NFS-Log broadcast domain.

```

network port vlan create -node <clustername>-01 -vlan-name a0a-<log-vlan-id>
network port vlan create -node <clustername>-02 -vlan-name a0a-<log-vlan-id>
broadcast-domain add-ports -broadcast-domain NFS-Log -ports,<clustername>-01:a0a-<log-vlan-id>,
<clustername>-02:a0a-<log-vlan-id>

```

4. Create the iSCSI-a VLAN ports and add them to the iSCSI-a broadcast domain.

```

network port vlan create -node <clustername>-01 -vlan-name a0a-<iscsi-a-vlan-id>
network port vlan create -node <clustername>-02 -vlan-name a0a-<iscsi-a-vlan-id>
broadcast-domain add-ports -broadcast-domain iSCSI-a -ports,<clustername>-01:a0a-<iscsi-a-vlan-id>,
<clustername>-02:a0a-<iscsi-a-vlan-id>

```

5. Create the iSCSI-b VLAN ports and add them to the iSCSI-b broadcast domain.

```
network port vlan create -node <clustername>-01 -vlan-name a0a-iscsi-b-vlan-id
network port vlan create -node <clustername>-02 -vlan-name a0a-iscsi-b-vlan-id
broadcast-domain add-ports -broadcast-domain iSCSI-b -ports,<clustername>-01:a0a-iscsi-b-vlan-id,<clustername>-02:a0a-iscsi-b-vlan-id
```

6. Create NFS VLAN ports and add them to the NFS broadcast domain.

```
network port vlan create -node <clustername>-01 -vlan-name a0a-nfs-vlan-id
network port vlan create -node <clustername>-02 -vlan-name a0a-nfs-vlan-id

broadcast-domain add-ports -broadcast-domain NFS -ports <clustername>-01:a0a-nfs-vlan-id,<clustername>-02:a0a-nfs-vlan-id
```

7. Create backup VLAN ports and add them to the backup domain.

```
network port vlan create -node <clustername>-01 -vlan-name a0a-backup-vlan-id
network port vlan create -node <clustername>-02 -vlan-name a0a-backup-vlan-id

broadcast-domain add-ports -broadcast-domain backup -ports <clustername>-01:a0a-backup-vlan-id,<clustername>-02:a0a-backup-vlan-id
```

Configure HTTPS Access

For each of the SVMs and the cluster node, create a certificate to allow secure communication with HTTPS. For each of the certificates, specify the individual values listed in Table 23 .

Table 23 ONTAP Software Parameter to Enable HTTPS

Cluster Detail	Cluster Detail Value
Certificate common name	<cert-common-name>
Country code	<cert-country>
State	<cert-state>
Locality	<cert-locality>
Organization	<cert-org>
Unit	<cert-unit>
Email	<cert-email>
Number of days the certificate is valid	<cert-days>

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
```

- Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example the <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 two 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 hana-svm -common-name hana-svm -ca hana-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 the expired certificates. In the following command, use TAB completion to select and delete each default certificate.

- To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Infra-SVM, the HANA SVM, and the cluster SVM. Use TAB completion to aid in the completion of these commands.

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver hana-svm

security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver infra-svm

security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver <clustername>
```

- To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the `security certificate show` command.
- 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 <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>

security ssl modify -vserver hana-svm -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>

security ssl modify -vserver infra-svm -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

- Disable HTTP cluster management access.

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

- 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 spi|ontapi|compat -vserver * -enabled true
```

Upgrade to ONTAP 9.2/9.3

This section describes how to update your clustered storage system to a newer ONTAP release such as ONTAP 9.2 or ONTAP 9.3 by using NetApp OnCommand® System Manager. This upgrade is required if a newer Linux kernel or NFS client, such as Red Hat RHEL 7.2, is used (which ignores the `sunrpc.tcp_max_slot_table_entries` system-wide setting).

This upgrade procedure is documented in more detail in the [ONTAP 9 Upgrade and Revert/Downgrade Guide](#) located in the NetApp ONTAP 9 Documentation Center.

Preparation

Before you start the upgrade, download the desired ONTAP image from the [NetApp Support site](#) and store the image on a web server that can be reached by the storage system.

Upgrade by Using OnCommand System Manager


To upgrade your system by using the OnCommand System Manager, complete the following steps:

1. Log in to OnCommand System Manager using your cluster administrator credentials.




You can access OnCommand System Manager by pointing your web browser to the cluster management LIF IP address.

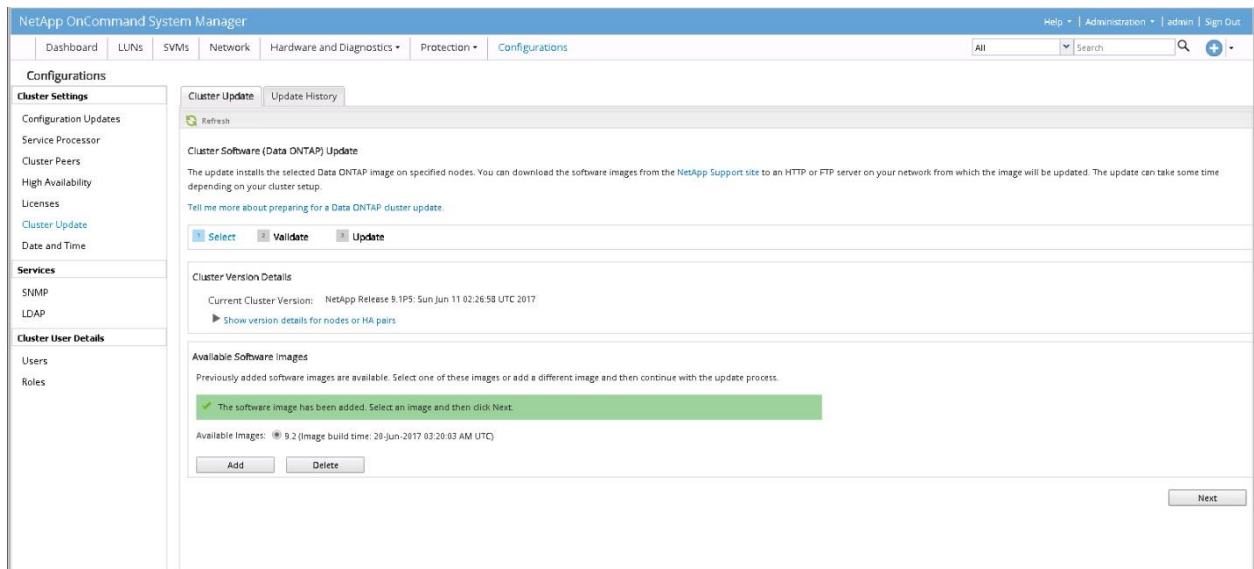
2. Expand the cluster hierarchy in the left navigation pane. In the navigation pane, click Cluster Update.
3. Click Add to add the previously-downloaded ONTAP image.




4. Enter the URL of the software image on your web server and click Add.




5. After the image is successfully added, click Next.




6. To validate the cluster for upgrade, click Validate.



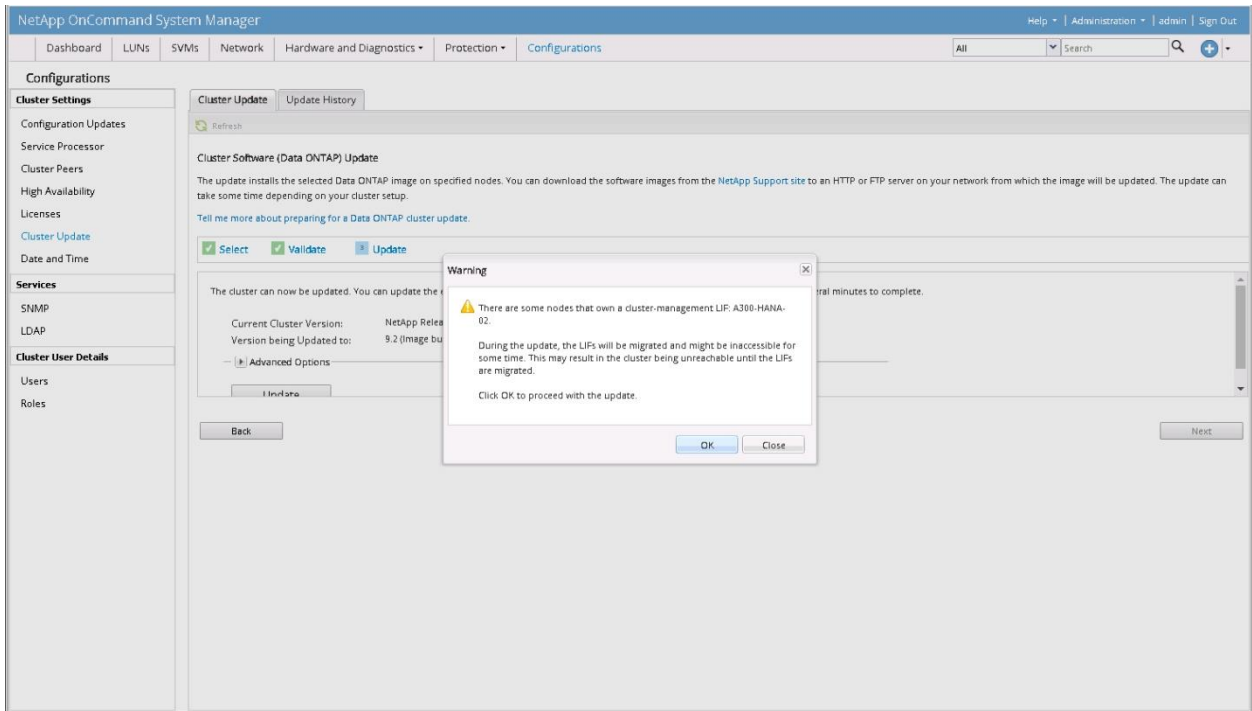
7. After the validation process is finished, verify the results, complete any required actions and click Next.




8. To update the cluster, click Update.



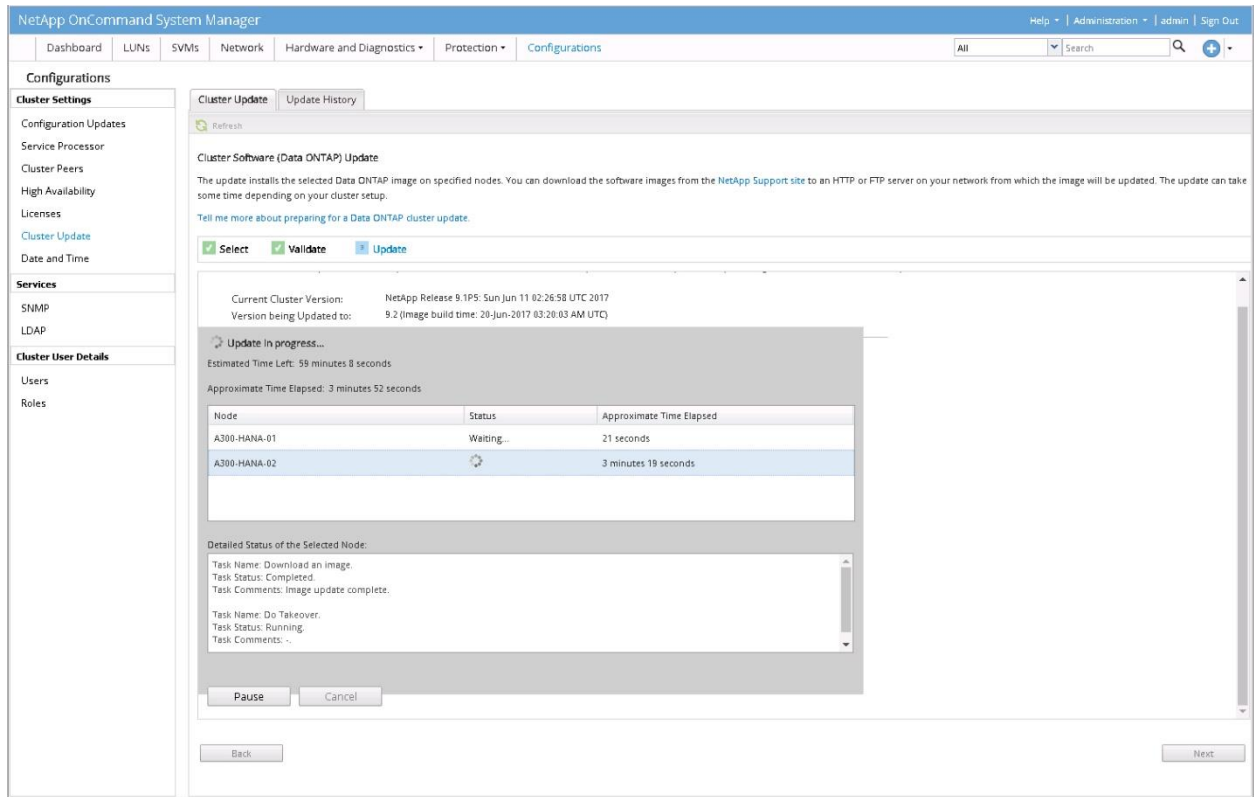
9. Click OK to confirm the LIF migration.




10. After the additional validation is finished, verify the results, complete any required actions, select the Continue Update with Warnings option and click Continue.



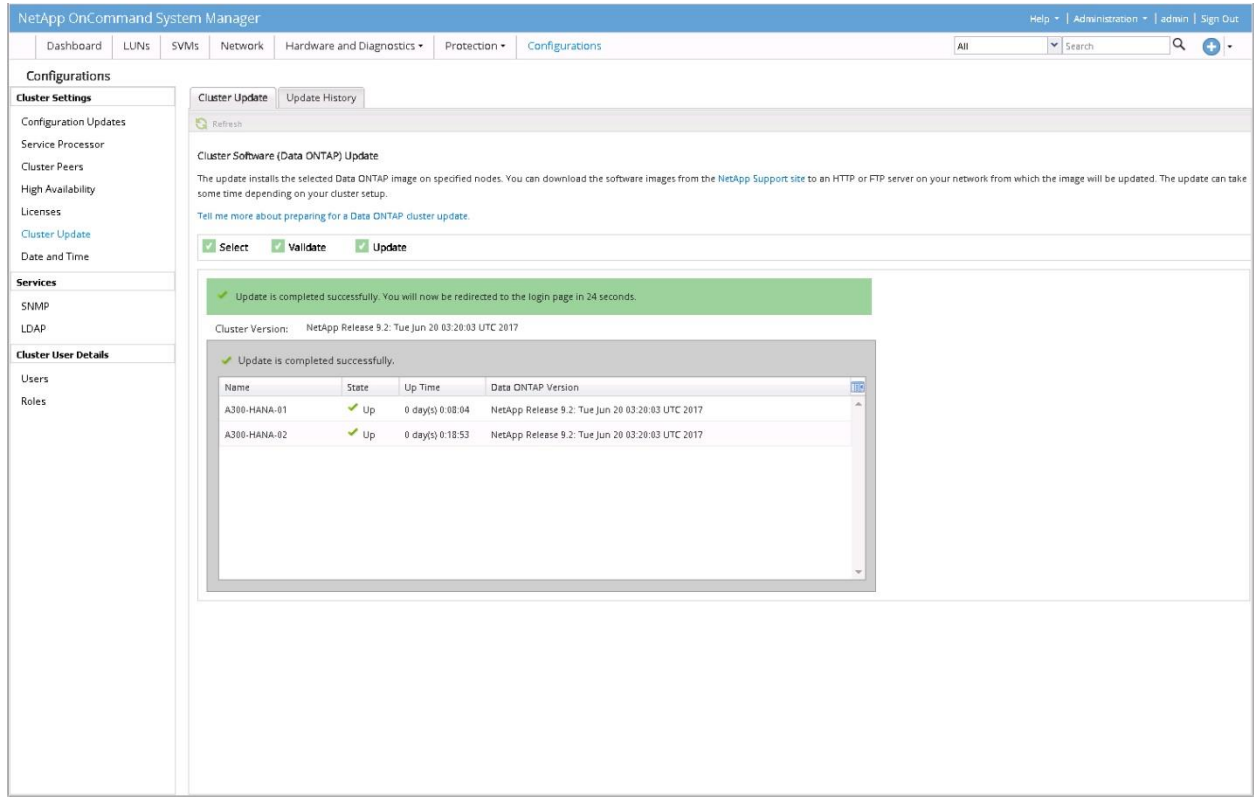
11. Verify the status of the update.



12. After the update is complete, click Finish.



13. You will receive a confirmation screen after the update is successfully completed.



Configure SVM for the Infrastructure

Table 24 and Figure 107 describe the infrastructure SVM together with all required storage objects (volumes, export-policies, and LIFs).

Figure 107 Overview of Infrastructure SVM Components

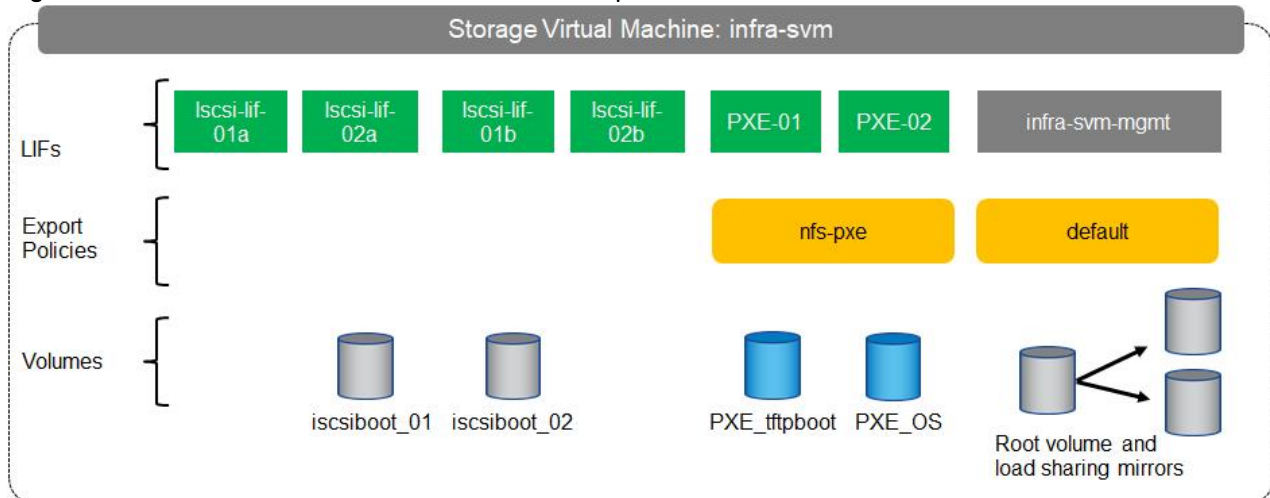


Table 24 ONTAP Software Parameters for Infrastructure SVMs

Cluster Detail	Cluster Detail Value	Value used in CVD setup
Infrastructure SVM management	<infra-svm-ip>	192.168.76.12

Cluster Detail	Cluster Detail Value	Value used in CVD setup
IP		
Infrastructure SVM management IP netmask	<infra-svm-netmask>	255.255.255.0
Infrastructure SVM default gateway	<infra-svm-gateway>	192.168.76.1
PXE CIDR	<pxe-cidr>	192.169.127.0
PXE netmask	<pxe-netmask>	255.255.255.0
PXE LIF node 1 IP	<node01-pxe_lif01-ip>	192.168.127.11
PXE LIF node 2 IP	<node02-pxe_lif02-ip>	192.168.127.12
iSCSI a CIDR	<iscsi-a-cidr>	192.168.128.0
iSCSI a Netmask	<iscsi_a_netmask>	255.255.255.0
iSCSI a IP node 1	<node01_iscsi_lif01a_ip>	192.168.128.11
iSCSI a IP node 2	<node02_iscsi_lif02a_ip>	192.168.128.12
iSCSI b CIDR	<iscsi-b-cidr>	192.168.129.0
iSCSI b Netmask	<iscsi_b_netmask>	255.255.255.0
iSCSI b IP node 1	<node01_iscsi_lif01b_ip>	192.168.129.11
iSCSI b IP node 2	<node02_iscsi_lif02b_ip>	192.168.129.12
NFS datastore CIDR	<nfs-CIDR>	192.168.130.0
NFS datastore netmask	<nfs-netmask>	255.255.255.0

Create SVM for the Infrastructure

To create an infrastructure SVM, complete the following steps:

1. Run the vserver create command.

```
vserver create -vserver infra-svm -rootvolume infra_rootvol -aggregate aggr01 -rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping iSCSI and NFS.

```
vserver remove-protocols -vserver infra-svm -protocols fcp,cifs,ndmp
```

3. Add the two data aggregates to the Infra-svm aggregate list for the NetApp Virtual Storage Console (VSC).

```
vserver modify -vserver infra-svm -aggr-list aggr01,aggr02
```

4. Enable and run the NFS protocol in the Infra-svm.

```
nfs create -vserver infra-svm -udp disabled
```

Create Load-Sharing Mirrors

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 infra_rootvol_m01 -aggregate aggr01 -size 1GB -type DP
volume create -vserver infra-svm -volume infra_rootvol_m02 -aggregate aggr02 -size 1GB -type DP
```

2. Create the mirroring relationships.

```
snapmirror create -source-path infra-svm:infra_rootvol -destination-path
infra-svm:infra_rootvol_m01 -type LS -schedule 15min
snapmirror create -source-path infra-svm:infra_rootvol -destination-path
infra-svm:infra_rootvol_m02 -type LS -schedule 15min
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path infra-svm:infra_rootvol
snapmirror show
```

Create Export Policies for the Root Volumes

To configure to export policies on the SVM, complete the following steps:

1. Create a new rule for the infrastructure NFS subnet in the default export policy.

```
vserver export-policy rule create -vserver infra-svm -policyname default -ruleindex 1 -protocol nfs -
clientmatch 0.0.0.0/0 -rorule sys -rwrule sys -superuser sys -allow-suid true -anon 0
```

2. Assign the FlexPod® export policy to the infrastructure SVM root volume.

```
volume modify -vserver infra-svm -volume infra_rootvol -policy default
```

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 infra-svm-mgmt -role data -data-protocol none -home-node
<clustername>-02 -home-port e0c -address <infra-svm-ip> -netmask <infra-svm-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 <infra-svm-gateway>
```

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: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver infra-svm
```

Create Export Policies for the Infrastructure SVM

1. Create a new export policy for the PXE subnet.

```
vserver export-policy create -vserver infra-svm -policyname nfs-pxe
```

2. Create a rule for this policy.

```
vserver export-policy rule create -vserver infra-svm -policyname nfs-pxe -clientmatch <pxe-cidr> -rorule
sys -rwrule sys -allow-suid true -allow-dev true -ruleindex 1 -anon 0 -protocol nfs -superuser sys
```

Create iSCSI LIFs

To create the four iSCSI LIFs (two on each node), run the following commands:

```
network interface create -vserver infra-svm -lif iscsi_lif01a -role data -data-protocol iscsi -home-node
<clustername>-01 -home-port a0a-<iscsi-a-vlan-id> -address <node01_iscsi_lif01a_ip> -netmask
<iscsi_a_netmask> -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
<clustername>-01 -home-port a0a-<iscsi-b-vlan-id> -address <node01_iscsi_lif01b_ip> -netmask
<iscsi_b_netmask> -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
<clustername>-02 -home-port a0a-<iscsi-a-vlan-id> -address <node02_iscsi_lif02a_ip> -netmask
<iscsi_a_netmask> -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
<clustername>-02 -home-port a0a-<iscsi-b-vlan-id> -address <node02_iscsi_lif02b_ip> -netmask
<iscsi_b_netmask > -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false
```

Create PXE LIFs

To create an NFS LIF for PXE boot, run the following commands:

```
network interface create -vserver infra-svm -lif PXE-01 -role data -data-protocol nfs -home-node
<clustername>-01 -home-port a0a-<pxe-vlan-id> -address <node01-pxe_lif01-ip> -netmask <pxe-netmask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true
```

```
network interface create -vserver infra-svm -lif PXE-02 -role data -data-protocol nfs -home-node
<clustername>-02 -home-port a0a-<pxe-vlan-id> -address <node02-pxe_lif02-ip> -netmask <pxe-netmask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true
```

Create Block Protocol (iSCSI) Service

Run the following command to create the iSCSI service. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

```
iscsi create -vserver infra-svm
```

Create FlexVol Volumes

To create the FlexVol volumes, run the following commands:

```
volume create -vserver infra-svm -volume iscsiboot_01 -aggregate aggr01 -size 100GB -state online -space-
guarantee none -percent-snapshot-space 0

volume create -vserver infra-svm -volume PXE_OS -aggregate aggr02 -size 200GB -state online -policy nfs-
pxe -space-guarantee none -percent-snapshot-space 0

volume modify -volume PXE_OS -files 15938348

volume create -vserver infra-svm -volume PXE_tftpboot -aggregate aggr01 -size 50GB -state online -policy
nfs-pxe -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path infra-svm:infra_rootvol
```

Configure LUNs for iSCSI Boot

Create Boot LUNs for Servers

To create two boot LUNs, run the following commands:

```
lun create -vserver infra-svm -volume server_boot_01 -lun server-01 -size 50GB -ostype linux -space-
reserve disabled

lun create -vserver infra-svm -volume server_boot_01 -lun server-02 -size 50GB -ostype linux -space-
reserve disabled
```

Create Portset

To create a portset that includes all iSCSI LIFs, run the following commands:

```
portset create -vserver Infra-SVM -portset server_Portset -protocol iscsi -port-name
iscsi_lif01a,iscsi_lif01b,iscsi_lif02a,iscsi_lif02b
```

Create igroups



Use the values listed in Table 26 to get the IQN information to create the igroups.

To create igroups, run the following commands:

```
igroup create -vserver Infra-SVM -igroup server-01 -protocol iscsi -ostype linux -initiator <vm-host-
infra-01-iqn> -portset server_Portset
igroup create -vserver Infra-SVM -igroup server-02 -protocol iscsi -ostype linux -initiator <vm-host-
infra-02-iqn> -portset ESX_Portset
```

Map ESX Boot LUNs to igroups

To map server boot LUNs to igroups, run the following commands:

```
lun map -vserver Infra-SVM -volume server_boot -lun server-01 -igroup server-01 -lun-id 0
lun map -vserver Infra-SVM -volume server_boot -lun server-02 -igroup server-02 -lun-id 0
```

Configure SVM for HANA

Table 25 and 0 describe the HANA SVM together with all the required storage objects (volumes, export-policies, and LIFs).

Figure 108 Overview of SAP HANA SVM Components

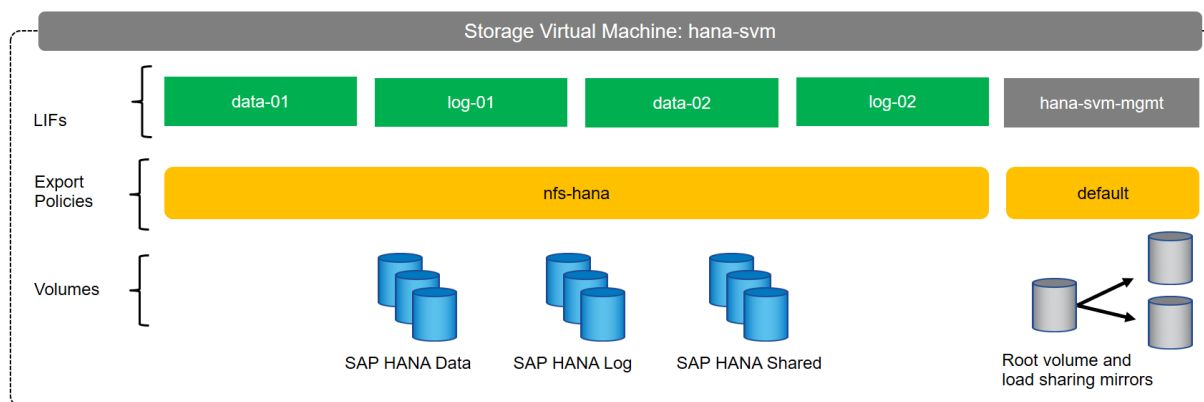


Table 25 ONTAP Software Parameter for HANA SVM

Cluster Detail	Cluster Detail Value	Value used in CVD setup
HANA SVM management IP	<hana-svm-ip>	192.168.76.16
HANA SVM management IP netmask	<hana-svm-netmask>	255.255.255.0
HANA SVM default gateway	<hana-svm-gateway>	192.168.76.1
NFS Data CIDR	<data-cidr>	192.168.210.0
NFS Data netmask	<data-netmask>	255.255.255.0
NFS Data LIF node 1 IP	<node01-data_lif01-ip>	192.168.210.11
NFS Data LIF node 2 IP	<node02-data_lif02-ip>	192.168.210.12
NFS log CIDR	<log-cidr>	192.168.228.0
NFS Log netmask	<log-netmask>	255.255.255.0
NFS Log LIF node 1 IP	<node01-log_lif01-ip>	192.168.228.11

Cluster Detail	Cluster Detail Value	Value used in CVD setup
NFS Log LIF node 2 IP	<node02-log_lif02-ip>	192.168.228.12

Create SVM for SAP HANA

To create an SVM for SAP HANA volumes, complete the following steps:

1. Run the vserver create command.

```
vserver create -vserver hana-svm -rootvolume hana_rootvol -aggregate aggr01 -rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping iSCSI and NFS.

```
vserver remove-protocols -vserver hana-svm -protocols fcp,cifs,ndmp
```

3. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver hana-svm -aggr-list aggr01,aggr02
```

4. Enable and run the NFS protocol in the Infra-SVM.

```
nfs create -vserver hana-svm -udp disabled
```

5. Enable a large NFS transfer size.

```
set advanced
vserver nfs modify -vserver hana-svm -tcp-max-transfersize 1048576
set admin
```

Create Load-Sharing Mirrors

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 HANA SVM root volume on each node.

```
volume create -vserver hana-svm -volume hana_rootvol_m01 -aggregate aggr01 -size 1GB -type DP
volume create -vserver hana-svm -volume hana_rootvol_m02 -aggregate aggr02 -size 1GB -type DP
```

2. Create the mirroring relationships.

```
snapmirror create -source-path hana-svm:hana_rootvol -destination-path hana-svm:hana_rootvol_m01 -type LS
-schedule 15min
snapmirror create -source-path hana-svm:hana_rootvol -destination-path hana-svm:hana_rootvol_m02 -type LS
-schedule 15min
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path hana-svm:hana_rootvol
```

Create Export Policies for the Root Volumes

To configure the NFS export policies on the SVM, complete the following steps:

1. Create a new rule for the infrastructure NFS subnet in the default export policy.

```
vserver export-policy rule create -vserver hana-svm -policyname default -ruleindex 1 -protocol nfs -
clientmatch 0.0.0.0/0 -rorule sys -rwrule sys -superuser sys -allow-suid true -anon 0
```

2. Assign the FlexPod export policy to the infrastructure SVM root volume.

```
volume modify -vserver hana-svm -volume hana_rootvol -policy default
```

Add HANA SVM Administrator

To add the HANA 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 hana-svm -lif hana-svm-mgmt -role data -data-protocol none -home-node
<clustername>-02 -home-port e0c -address <hana-svm-ip> -netmask <hana-svm-netmask> -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 hana-svm -destination 0.0.0.0/0 -gateway <hana-svm-gateway>
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver hana-svm
Enter a new password: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver hana-svm
```

Create Export Policies for the HANA SVM

1. Create a new export policy for the HANA data and log subnet.

```
vserver export-policy create -vserver hana-svm -policyname nfs-hana
```

2. Create a rule for this policy.

```
vserver export-policy rule create -vserver hana-svm -policyname nfs-hana -clientmatch <data-cidr>,<log-
cidr> -rorule sys -rwrule sys -allow-suid true -allow-dev true -ruleindex 1 -anon 0 -protocol nfs -
superuser sys
```

Create NFS LIF for SAP HANA Data

To create the NFS LIFs for SAP HANA data, run the following commands:

```

network interface create -vserver hana-svm -lif data-01 -role data -data-protocol nfs -home-node
<clustername>-01 -home-port a0a-<data-vlan-id> -address <node01-data_lif01-ip> -netmask <data-netmask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver hana-svm -lif data-02 -role data -data-protocol nfs -home-node
<clustername>-02 -home-port a0a-<data-vlan-id> -address <node02-data_lif02-ip> -netmask <data-netmask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

```

Create NFS LIF for SAP HANA Log

To create an NFS LIF for SAP HANA log, run the following commands:

```

network interface create -vserver hana-svm -lif log-01 -role data -data-protocol nfs -home-node
<clustername>-01 -home-port a0a-<log-vlan-id> -address <node01-log_lif01-ip> -netmask <log-netmask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver hana-svm -lif log-02 -role data -data-protocol nfs -home-node
<clustername>-02 -home-port a0a-<log-vlan-id> -address <node02-log_lif02-ip> -netmask <log-netmask> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

```

HANA Node Preparation

Preparation of PXE Boot Environment

Two PXE boot servers are used for a PXE boot to provide redundancy; one on Management Server 01 (ESXi-Mgmt-01) another on Management Server 02 (ESXi-Mgmt-02) for redundancy.

Installing SLES 12 SP3 based PXE Server VM on the Management Servers

To build a PXE Boot virtual machine (VM) on the ESXi-Mgmt-01 complete the following steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the VMware vSphere Client, select the host in the inventory pane.
3. Right-click the host and select New Virtual Machine.
4. Select Custom and click Next.
5. Enter a name for the VM, example HANA-Mgmt01, click Next.
6. Select the datastore where the PXE server resides. Click Next.
7. Select Virtual Machine Version: 11. Click Next.
8. Select the Linux option and the SUSE Linux Enterprise 12 (64-bit) version are selected. Click Next.
9. Select two virtual sockets and eight cores per virtual socket. Click Next.
10. Select 8GB of memory. Click Next.
11. Select three network interface card (NIC).
12. For NIC 1, select the OOB-MGMT Network option and the VMXNET 3 adapter.
13. For NIC 2, select the HANA-Boot Network option and the VMXNET 3 adapter.
14. For NIC 3, select the HANA-Admin Network option and the VMXNET 3 adapter
15. Click Next.
16. Keep the VMware Paravirtual option for the SCSI controller selected. Click Next.
17. Keep the Create a New Virtual Disk option selected. Click Next.
18. Make the disk size at least 60GB. Click Next.

19. Click Next.
20. Accept default selection for Virtual Device Node to be SCSI (0:0). Click Next.
21. Select the checkbox for Edit the Virtual Machine Settings Before Completion. Click Continue.
22. Click the Options tab.
23. Select Boot Options.
24. Select the Force BIOS Setup checkbox.
25. Click Finish.
26. From the left pane, expand the host field by clicking the plus sign (+).
27. Right-click the newly created HANA-Mgmt01 and click Open Console.
28. Click the third button (green right arrow) to power on the VM.
29. Click the ninth button (CD with a wrench) to map the SLES DVD, and then select Connect to ISO Image on Local Disk.



In the CVD set-up we used an SLES11SP4 based server to provide the PXE related services. However, it is suggested to use SLES12 based system keeping up with current OS releases.

30. Navigate to the SLES-12 SP3 64 bit ISO, select it, and click Open.
31. Click in the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to select CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.
32. **Select the VM and click 'Edit virtual Machine settings'. Under Hardware tab select CD/DVD drive 1 check the 'Connect at power on' option. Click OK.**
33. Power On the VM. The SUSE Installer boots. Select the Installation by pressing down arrow key and press Enter.
34. Agree to the License Terms, select Next and press Enter.
 - a. 35. Under Network Settings Under the Overview, select the device eth0 and click Edit and enter IP Address <<var_pxe_oob_IP>> and Subnet Mask <<var_pxe_oob_subnet>>. Enter the Hostname, click the General tab and Set MTU 1500 or 9000 depending on the Customer switch config. Click Next.
 - b. Under Overview, select the device eth1 and click Edit and enter IP Address <<var_pxe_boot_IP>> and Subnet Mask <<var_pxe_boot_subnet>>. Enter the Hostname. Click the General tab and Set MTU 1500, click Next.

- c. Under Overview, select the device eth2 and click Edit and enter IP Address <<var_pxe_admin_IP>> and Subnet Mask <<var_pxe_admin_subnet>>. Enter the Hostname. Click the General tab and Set MTU 1500, click Next.
 - d. Click the Hostname/DNS tab update the Hostname [mgmtsrv01 used here] and Domain Names. Under Name Server 1, enter the IP address of the DNS Server, optionally enter the IP address for Name Server 2 and Name Server 2. Under Domain Search, enter the Domain name for DNS.
 - e. Click the Routing tab and Enter the IP address for Default Gateway for Out Band Management IP address, and click Next.
35. Skip Registration for now. Click Next. Press OK for the Warning prompt.
 36. Click Next on the Add on Product page.
 37. Select Default System for System Role and click Next
 38. Suggested Partitioning: **Click Edit Proposal Settings, Select Ext4 under 'Filesystem for Root Partition' and uncheck 'Propose Separate Home Partition'. Click OK. Click Next.**
 39. Select Appropriate Region and TimeZone. Check the 'Hardware Clock set to UTC' option. Click Next.
 40. Skip local User Creation.
 41. Set Password for System Administrator "root".
 42. Installation Settings:
 - a. **Select Software. Deselect 'GNOME Desktop Environment', Select C/C++ Compiler and Tools and Click OK.**
 - b. Under Firewall and SSH, disable Firewall.
 - c. Click Kdump. Select Disable Kdump. Click OK
 - d. Set Default Systemd Target to Text mode. Click OK.
 43. Click Install.
 44. To confirm the installation, click Install.



To build a PXE Boot virtual machine (VM) on the Management Server 02, complete the steps above for ESXi-Mgmt-02.

Customize PXE Server

1. Disable IPV6: Login to the management server as root. YAST2 -> System -> Network Settings -> Global Options -> Uncheck '**Enable IPV6**' under '**IPV6 Protocol Settings**'. Click OK. Click OK for the Warning prompt. Click OK. Select Quit to exit YAST2.
2. Verify the IP address that are assigned to the PXE server:

```

mgmtsrv01:~ # ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether 00:0c:29:ec:92:c0 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether 00:0c:29:ec:92:ca brd ff:ff:ff:ff:ff:ff
   inet 192.168.127.6/24 brd 192.168.127.255 scope global eth1
       valid_lft forever preferred_lft forever
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether 00:0c:29:ec:92:d4 brd ff:ff:ff:ff:ff:ff
   inet 192.168.76.6/24 brd 192.168.76.255 scope global eth2

```

```

valid_lft forever preferred_lft foreverConfigure the /etc/hosts File of the Management Stations:
cat /etc/hosts

```

```

#
# hosts          This file describes a number of hostname-to-address
#               mappings for the TCP/IP subsystem.  It is mostly
#               used at boot time, when no name servers are running.
#               On small systems, this file can be used instead of a
#               "named" name server.
# Syntax:
#
# IP-Address    Full-Qualified-Hostname  Short-Hostname
#
127.0.0.1      localhost

192.168.76.6   mgmtsrv01.ciscolab.local mgmtsrv01
## PXE VLAN
192.168.127.11 nfspxe
192.168.127.6   mgmtsrv01p
192.168.127.101 server01p
192.168.127.102 server02p
192.168.127.103 server03p
192.168.127.104 server04p
192.168.127.105 server05p
192.168.127.106 server06p
192.168.127.107 server07p
192.168.127.108 server08p
192.168.127.109 server09p
192.168.127.110 server10p
192.168.127.111 server11p
192.168.127.112 server12p

```

Mount Volume for PXE Boot Configuration

1. To mount the tftpboot, software and osmaster volumes, add the entry to /etc/fstab with the values listed below:

```

vi /etc/fstab

nfspxe:/tftpboot          /tftpboot          nfs          defaults          0 0
nfspxe:/PXE_OS            /NFS/PXE_OS       nfs          defaults          0 0

```

2. Create the directories for mount points:

```

mkdir /tftpboot
mkdir /NFS

```

```
mkdir /NFS/PXE_OS
```

3. Mount the nfs file system:

```
mount /NFS/PXE_OS
mount /tftpboot
```

Update Packages for PXE Server VM

To update the SUSE virtual machine to latest patch level, complete the following steps:



This document assumes that a SUSE License key is obtained and registered username and password is available. VM has internet access.

1. ssh to the PXE boot VM.
2. Login as root and password.
3. Execute the below commands to Register the SUSE server after performing the proxy settings.

```
export http_proxy=http://<proxy-server-IP>
export https_proxy=http://<proxy-server-IP>
SUSEConnect -r <registration_code> -e <email_address> -u https://scc.suse.com
```

After the registration, all the repositories are updated as shown below:

```
# zypper repos
Repository priorities are without effect. All enabled repositories share the same priority.

# | Alias | Name
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Enabled | GPG Check | Refresh |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 | SLES12-SP3-12.3-0 | SLES12-SP3-12.3-0
| Yes | ( r ) Yes | No
2 | SUSE_Linux_Enterprise_Server_12_SP3_x86_64:SLES12-SP3-Debuginfo-Pool | SLES12-SP3-Debuginfo-Pool
| No | ---- | ----
3 | SUSE_Linux_Enterprise_Server_12_SP3_x86_64:SLES12-SP3-Debuginfo-Updates | SLES12-SP3-Debuginfo-
Updates | No | ---- | ----
4 | SUSE_Linux_Enterprise_Server_12_SP3_x86_64:SLES12-SP3-Pool | SLES12-SP3-Pool
| Yes | ( p ) Yes | No
5 | SUSE_Linux_Enterprise_Server_12_SP3_x86_64:SLES12-SP3-Source-Pool | SLES12-SP3-Source-Pool
| No | ---- | ----
6 | SUSE_Linux_Enterprise_Server_12_SP3_x86_64:SLES12-SP3-Updates | SLES12-SP3-Updates
| Yes | ( p ) Yes | Yes
```

4. Execute the below command to update the server:

```
zypper update
```

5. Follow the on-screen instruction to complete the update process.
6. Reboot the server.

PXE Services Configuration

To configure a PXE (Pre-boot Execution Environment) boot server, two packages, the DHCP (Dynamic Host Configuration Protocol) server and tftp server are required. DHCP server is already installed in the previous step.

To install and configure tftp server, complete the following steps:

1. Configure the tftp server.
2. Log in to the VM created PXE boot Server using SSH.
3. Search the package tftp server using command shown below:

```
HANA-mgmtsrv01:~ # zypper se tftp
Refreshing service 'SUSE_Linux_Enterprise_Server_12_SP3_x86_64'.
Loading repository data...
Reading installed packages...

S | Name | Summary | Type
-----+-----+-----+-----
| atftp | Advanced TFTP Server and Client | package
| atftp | Advanced TFTP Server and Client | srcpackage
| tftp | Trivial File Transfer Protocol (TFTP) | package
| tftp | Trivial File Transfer Protocol (TFTP) | srcpackage
| tftpbboot-installation-SLES-12-SP3-x86_64 | tftp installation tree | package
i | yast2-tftp-server | YaST2 - TFTP Server Configuration | package
| yast2-tftp-server | YaST2 - TFTP Server Configuration | srcpackageInstall

the tftp server:
HANA-mgmtsrv01:~ # zypper in tftp
Refreshing service 'SUSE_Linux_Enterprise_Server_12_SP3_x86_64'.
Loading repository data...
Reading installed packages...
Resolving package dependencies...

The following NEW package is going to be installed:
  tftp

1 new package to install.
Overall download size: 48.3 KiB. Already cached: 0 B. After the operation, additional 85.0 KiB will be
used.
Continue? [y/n/...? shows all options] (y): y
Retrieving package tftp-5.2-11.6.1.x86_64
(1/1), 48.3 KiB (85.0 KiB unpacked)
Retrieving: tftp-5.2-11.6.1.x86_64.rpm
.....[done]
Checking for file conflicts:
.....[done]
(1/1) Installing: tftp-5.2-11.6.1.x86_64
.....[done]

Additional rpm output:
Updating /etc/sysconfig/tftp...Configure xinetd to respond to tftp requests:
HANA-mgmtsrv01:/etc/xinetd.d # more tftp
# default: off
# description: tftp service is provided primarily for booting or when a \
# router need an upgrade. Most sites run this only on machines acting as \
# "boot servers".
# The tftp protocol is often used to boot diskless \
# workstations, download configuration files to network-aware printers, \
# and to start the installation process for some operating systems.
service tftp
{
    socket_type = dgram
```

```

protocol          = udp
wait              = yes
flags             = IPv4
user              = root
server            = /usr/sbin/in.tftpd
server_args       = -s /tftpboot
per_source        = 11
cps               = 100 2
disable           = noTo configure your TFTP server, create a directory which will be the
base directory for the Linux boot images and the PXE boot server configuration files:
mkdir /tftpboot
chmod 755 tftpboot

```

4. To make sure the TFTP servers can startup on subsequent reboots, execute the following command:

```

chkconfig xinetd on
chkconfig tftp on

rcxinetd restart

Shutting down xinetd: (waiting for all children to terminate)      done
Starting INET services. (xinetd)                                   done

```

5. Make sure syslinux is installed:

```

rpm -qa syslinux
syslinux-4.04-40.35.x86_64Copy the pxelinux image on to the root directory of the tftp server:
cp /usr/share/syslinux/pxelinux.0 /tftpboot/

```

10. PXELinux relies on the pxelinux.cfg directory to be in the root of the tftp directory for configuration.

```

mkdir /tftpboot/pxelinux.cfg

```

To install and configure dhcp service, complete the following steps:

6. Log in to the VM created PXE boot Server using SSH.

7. Search the package dhcp-server using command shown below:

```

HANA-mgmtsrv01:~ # zypper se dhcp
HANA-mgmtsrv01:/etc/xinetd.d # zypper se dhcp
Refreshing service 'SUSE_Linux_Enterprise_Server_12_SP3_x86_64'.
Loading repository data...
Reading installed packages...

S | Name | Summary | Type
-----+-----+-----+-----
i | dhcp | Common Files Used by ISC DHCP Software | package
  | dhcp | Common Files Used by ISC DHCP Software | srcpackage
i | dhcp-client | ISC DHCP Client | package
  | dhcp-relay | ISC DHCP Relay Agent | package
  | dhcp-server | ISC DHCP Server | package
  | dhcp-tools | DHCP Tools | package
  | dhcp-tools | DHCP Tools | srcpackage
  | dhcp_dns_server | DHCP and DNS Server | pattern
  | monitoring-plugins-dhcp | Check DHCP servers | package
  | patterns-sles-dhcp_dns_server | DHCP and DNS Server | package
  | patterns-sles-dhcp_dns_server-32bit | DHCP and DNS Server | package
  | udhcp | Micro DHCP client / server | package
  | udhcp | Micro DHCP client / server | srcpackage
i | yast2-dhcp-server | YaST2 - DHCP Server Configuration | package
  | yast2-dhcp-server | YaST2 - DHCP Server Configuration | srcpackage

```

8. Install dhcp-server package:

```

HANA-mgmtsrv01:~ # zypper in dhcp-server
Refreshing service 'SUSE_Linux_Enterprise_Server_12_SP3_x86_64'.
Loading repository data...
Reading installed packages...
Resolving package dependencies...

The following NEW package is going to be installed:
  dhcp-server

1 new package to install.
Overall download size: 883.7 KiB. Already cached: 0 B. After the operation, additional 2.2 MiB will be
used.
Continue? [y/n/...? shows all options] (y): y
Retrieving package dhcp-server-4.3.3-10.14.1.x86_64
(1/1), 883.7 KiB ( 2.2 MiB unpacked)
Retrieving: dhcp-server-4.3.3-10.14.1.x86_64.rpm
.....[done]
Checking for file conflicts:
.....[done]
(1/1) Installing: dhcp-server-4.3.3-10.14.1.x86_64
.....[done]
Additional rpm output:
Updating /etc/sysconfig/dhcpd...
Updating /etc/sysconfig/syslog...Updating /etc/sysconfig/tftp...

```

9. Activate the DHCP server to listen on eth1, which is configured for PXE boot VLAN 127:

```

vi /etc/sysconfig/dhcpd
#
DHCPD_INTERFACE="eth1"

```

10. Obtain the MAC Address List for HANA-Boot vNIC for service Profiles created. A separate MAC address pool was created for HANA-Boot and assigned in the sequential order. To obtain the MAC address for HANA-Boot, complete the following steps:

- a. Log in to Cisco UCS Manager; click the LAN tab in the navigation pane.
- b. Select Pools > root > MAC pools > MAC Pool HANA-Boot.
- c. Expand MAC Pool HANA-Boot.
- d. Click the MAC Addresses tab on the right pane.



The MAC address is assigned to the Service Profile in sequential order.

11. **DHCP server requires 'next-server' directive** to DHCP configuration file; this directive should have the IP address of the TFTP server i.e. (next-server 192.168.127.6).
12. The second directive that needs **to be added to DHCP configuration file is 'filename' and it should have the value of 'pxelinux.0', for example filename "pxelinux.0"**; this will enable PXE booting.
13. To assign hostname to the server via DHCP use the option host-name <<hostname>>.
14. The MAC Address configured for PXE in the Cisco UCS Service Profile should be reserved with an IP address for each server for PXE boot in the dhcp configuration.



Below is an example of `/etc/dhcpd.conf`, VLAN ID 127 is used for PXE boot network. The PXE boot server IP address is 192.168.127.6, subnet 255.255.255.0. Assigned IP address for servers are 192.168.127.101-112. The sample file shows 6 hosts.

```
#
# dhcpd.conf
#
default-lease-time 14400;
ddns-update-style none;
ddns-updates off;

filename "pxelinux.0";

subnet 192.168.127.0 netmask 255.255.255.0 {
    group {
        next-server 192.168.127.6;
        filename "pxelinux.0";
        host server01 {
            hardware ethernet 00:25:b5:0a:00:01;
            fixed-address 192.168.127.101;
        }
        host server02 {
            hardware ethernet 00:25:b5:0a:00:07;
            fixed-address 192.168.127.102;
        }
        host server03 {
            hardware ethernet 00:25:b5:0a:00:0d;
            fixed-address 192.168.127.103;
        }
        host server04 {
            hardware ethernet 00:25:b5:0a:00:13;
            fixed-address 192.168.127.104;
        }
        host server05 {
            hardware ethernet 00:25:b5:0a:00:19;
            fixed-address 192.168.127.105;
        }
        host server06 {
            hardware ethernet 00:25:b5:0a:00:1f;
            fixed-address 192.168.127.106;
        }
        host server07 {
            hardware ethernet 00:25:b5:0a:00:25;
            fixed-address 192.168.127.107;
        }
        host server08 {
            hardware ethernet 00:25:b5:0a:00:2b;
            fixed-address 192.168.127.108;
        }
        host server09 {
            hardware ethernet 00:25:b5:0a:00:31;
            fixed-address 192.168.127.109;
        }
        host server10 {
            hardware ethernet 00:25:b5:0a:00:37;
            fixed-address 192.168.127.110;
        }
        host server11 {
            hardware ethernet 00:25:b5:0a:00:3d;
            fixed-address 192.168.127.111;
        }
        host server12 {
            hardware ethernet 00:25:b5:0a:00:43;
            fixed-address 192.168.127.112;
        }
    }
}
```

15. To make sure the DHCP servers can startup on subsequent reboots, execute the following command:

```
chkconfig dhcpd on
```

16. Restart the dhcp service for new configuration to take effect:

```
service dhcpd restart
Shutting down ISC DHCPv4 4.x Server           done
Starting ISC DHCPv4 4.x Server [chroot]       done
```

Installing RHEL 7.4 based PXE Server VM on the Management Servers

To build a PXE Boot virtual machine (VM) on the ESXi-Mgmt-01, complete the following steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the VMware vSphere Client, select the host in the inventory pane.
3. Right-click the host and select New Virtual Machine.
4. Select Custom and click Next.
5. Enter a name for the VM, example HANA-Mgmt01, click Next.
6. Select the datastore where the PXE server resides. Click Next.
7. Select Virtual Machine Version: 11. Click Next.
8. Select the Linux option and the RHEL Server 7.4 (64-bit) version are selected. Click Next.
9. Select two virtual sockets and eight cores per virtual socket. Click Next.
10. Select 8GB of memory. Click Next.
11. Select three network interface card (NIC).
12. For NIC 1, select the OOB-MGMT Network option and the VMXNET 3 adapter.
13. For NIC 2, select the HANA-Boot Network option and the VMXNET 3 adapter.
14. For NIC 3, select the HANA-Admin Network option and the VMXNET 3 adapter
15. Click Next.
16. Keep the VMware Paravirtual option for the SCSI controller selected. Click Next.
17. Keep the Create a New Virtual Disk option selected. Click Next.
18. Make the disk size at least 60GB. Click Next.

19. Click Next.
20. Accept default selection for Virtual Device Node to be SCSI (0:0). Click Next.
21. Select the checkbox for Edit the Virtual Machine Settings Before Completion. Click Continue.
22. Click the Options tab.
23. Select Boot Options.
24. Select the Force BIOS Setup checkbox.
25. Click Finish.
26. From the left pane, expand the host field by clicking the plus sign (+).
27. Right-click the newly created HANA-Mgmt01 and click Open Console.
28. Click the third button (green right arrow) to power on the VM.
29. Click the ninth button (CD with a wrench) to map the SLES DVD, and then select Connect to ISO Image on Local Disk/Datastore.
30. Navigate to the RHEL 7.4 64 bit ISO, select it, and click Open.
31. Click in the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to select CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.
32. Select the VM and click **'Edit virtual Machine settings'**. Under **Hardware tab** select **CD/DVD drive 1** check the **'Connect at power on'** option. Click OK.
33. Power On the VM. The RHEL Installer boots. Select the Install option by pressing up arrow key and press Enter.
34. On the Welcome page, select the Language and Keyboard layout as desired. Click Continue.
35. On the installation summary page:
 - a. Localization: Set the Date and Time appropriately.
 - b. System:
 - i. Select the appropriate local standard disk or Specialized & Network Disk (iSCSI LUN) for the installation destination. Select the default Automatically configure partitioning option. Click Done.
 - ii. **Click KDUMP. Uncheck 'Enable kdump' option. Click Done.**
 - iii. Click Security Policy. **Change 'Apply security policy' option to OFF. Click Done.**
 - iv. Click Network & Hostname. Set the FQDN Hostname. Select first Ethernet interface and click configure to set the configure its IP address under IPv4 settings tab. Set Method to Manual. Click Add to enter IP Address <<var_oob_mgmt_ip>> and Subnet Mask <<var_oob_mgmt-

subnet>>. Similarly configure the second Ethernet interface with IP Address <<var_pxe_boot_IP>> and Subnet Mask <<var_pxe_boot_subnet>> and the third Ethernet interface with IP address <<var_pxe_admin_IP>> & Subnet Mask <<var_pxe_admin_subnet>>.

36. Click Begin Installation.
37. Set Root user password. Click Done.
38. At the end of installation, click Reboot.

Customize PXE Server

1. Disable IPV6.

```
Create a new file named /etc/sysctl.d/ipv6.conf and add the following options:
net.ipv6.conf.all.disable_ipv6 = 1
```

```
The new settings would then need to be reloaded with:
# sysctl -p /etc/sysctl.d/ipv6.conf
```

```
Then rebuild the Initial RAM Disk Image using:
# dracut -f
```

2. Verify IP addresses assigned:

```
[root@HANA-Mgmtsrv01 ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
2: ens192: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:0c:29:ec:92:c0 brd ff:ff:ff:ff:ff:ff
3: ens224: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:0c:29:ec:92:ca brd ff:ff:ff:ff:ff:ff
    inet 192.168.127.6/24 brd 192.168.127.255 scope global ens224
4: ens256: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:0c:29:ec:92:d4 brd ff:ff:ff:ff:ff:ff
    inet 192.168.76.6/24 brd 192.168.76.255 scope global ens256
    valid_lft forever preferred_lft forever
```

3. Configure the /etc/hosts File of the Management Stations:

```
cat /etc/hosts

#
# hosts          This file describes a number of hostname-to-address
#               mappings for the TCP/IP subsystem.  It is mostly
#               used at boot time, when no name servers are running.
#               On small systems, this file can be used instead of a
#               "named" name server.
# Syntax:
#
# IP-Address    Full-Qualified-Hostname  Short-Hostname
#
127.0.0.1      localhost

192.168.76.6   mgmtsrv01.ciscolab.local mgmtsrv01
## PXE VLAN
192.168.127.11 nfspxe
192.168.127.6  mgmtsrv01p
```

```

192.168.127.101 server01p
192.168.127.102 server02p
192.168.127.103 server03p
192.168.127.104 server04p
192.168.127.105 server05p
192.168.127.106 server06p
192.168.127.107 server07p
192.168.127.108 server08p
192.168.127.109 server09p
192.168.127.110 server10p
192.168.127.111 server11p
192.168.127.112 server12p

```

Update Packages for PXE Server VM

To update the SUSE virtual machine to latest patch level, complete the following steps:



This document assumes that a SUSE License key is obtained and registered username and password is available. VM has internet access.

1. ssh to the PXE boot VM.
2. Login as root and password.
3. Disable Firewall:

```
# systemctl disable firewalld
```

4. Execute the below commands to Register RHEL Server for updates:

```

vi /etc/rhsm/rhsm.conf
Update the proxy server name and port information with proxy server information for the landscape.

subscription-manager register

```

5. After the registration, all the repositories are updated as shown below:

```

# subscription-manager list --available -all

Attach the subscription with appropriate pool                                # subscription-
manager attach -pool=<Pool ID>

```

6. Execute the below command to update the server:

```
# yum update
```

7. Update to Base system state and install additional PXE related packages:

```

# yum -y groupinstall base
# yum install syslinux tftp-server dhcp xinetd nfs-utils

```

8. Follow the on-screen instruction to complete the update process.

9. Reboot the server.

Mount Volume for PXE Boot Configuration

1. To mount the tftpboot, software and osmaster volumes, add the entry to /etc/fstab with the values listed below:

```
vi /etc/fstab
nfspxe:/tftpboot      /tftpboot      nfs      defaults      0 0
nfspxe:/PXE_OS       /NFS/PXE_OS    nfs      defaults      0 0
```

2. Create the directories for mount points:

```
mkdir /tftpboot
mkdir -p /NFS/PXE_OS
```

3. Mount the nfs file system:

```
mount /NFS/PXE_OS
mount /tftpboot
```

PXE Services Configuration

To configure a PXE (Pre-boot Execution Environment) boot server, two packages, the DHCP (Dynamic Host Configuration Protocol) server and tftp server are required. DHCP server is already installed in the previous step.

To install and configure tftp server, complete the following steps:

1. Configure the tftp server.
2. Log in to the VM created PXE boot Server using SSH.
3. Updating /etc/sysconfig/tftp: Configure xinetd to respond to tftp requests:

```
[root@HANA-Mgmtsrv01 ~]# vi /etc/xinetd.d/tftp
# default: off
# description: The tftp server serves files using the trivial file transfer \
#               protocol. The tftp protocol is often used to boot diskless \
#               workstations, download configuration files to network-aware printers, \
#               and to start the installation process for some operating systems.
service tftp
{
    socket_type           = dgram
    protocol              = udp
    wait                  = yes
    user                  = root
    server                = /usr/sbin/in.tftpd
    server_args           = -s /tftpboot
    disable               = no
    per_source            = 11
    cps                   = 100 2
    flags                 = IPv4
}
```

- Update the tftp.service properties to ensure it is started at boot time:

```
[root@HANA-Mgmtsrv01 ~]# vi /usr/lib/systemd/system/tftp.service
[Unit]
Description=Tftp Server
Documentation=man:in.tftpd

[Service]
ExecStart=/usr/sbin/in.tftpd -s /tftpboot
StandardInput=socket

[Install]
WantedBy=multi-user.target
```

- To configure your TFTP server, create a directory which will be the base directory for the Linux boot images and the PXE boot server configuration files:

```
mkdir /tftpboot
chmod 755 tftpboot
```

- To make sure the TFTP servers can startup on subsequent reboots, execute the following command:

```
systemctl enable xinetd
systemctl start xinetd

systemctl enable tftp
systemctl start tftp

chkconfig xinetd on
chkconfig tftp on
```

- Copy the pxelinux image on to the root directory of the tftp server:

```
cp /usr/share/syslinux/pxelinux.0 /tftpboot/
```

- PXELinux relies on the pxelinux.cfg directory to be in the root of the tftp directory for configuration.

```
mkdir /tftpboot/pxelinux.cfg
```

To install and configure dhcp service, complete the following steps:

- Log in to the VM created PXE boot Server using SSH.
- Configure dhcp interface:

```
[root@HANA-Mgmtsrv01 ~]# cp /usr/lib/systemd/system/dhcpd.service /etc/systemd/system/

[root@HANA-Mgmtsrv01 ~]# vi /etc/systemd/system/dhcpd.service

[Unit]
Description=DHCPv4 Server Daemon
Documentation=man:dhcpd(8) man:dhcpd.conf(5)
Wants=network-online.target
After=network-online.target
After=time-sync.target

[Service]
Type=notify
ExecStart=/usr/sbin/dhcpd -f -cf /etc/dhcp/dhcpd.conf -user dhcpd -group dhcpd --no-pid <dhcp-network-
interface>

[Install]
```

```
WantedBy=multi-user.target
```

3. Obtain the MAC Address List for HANA-Boot vNIC for service Profiles created. A separate MAC address pool was created for HANA-Boot and assigned in the sequential order. To obtain the MAC address for HANA-Boot, complete the following steps:
 - a. Log in to Cisco UCS Manager; click the LAN tab in the navigation pane.
 - b. Select Pools > root > MAC pools > MAC Pool HANA-Boot.
 - c. Expand MAC Pool HANA-Boot.
 - d. Click the MAC Addresses tab on the right pane.



The MAC address is assigned to the Service Profile in sequential order.

4. **DHCP server requires 'next-server' directive to DHCP configuration file; this directive should have the IP address of the TFTP server i.e. (next-server 192.168.127.6).**
5. **The second directive that needs to be added to DHCP configuration file is 'filename' and it should have the value of 'pxelinux.0', for example filename "pxelinux.0"; this will enable PXE booting.**
6. To assign hostname to the server via DHCP use the option host-name <<hostname>>.
7. The MAC Address configured for PXE in the Cisco UCS Service Profile should be reserved with an IP address for each server for PXE boot in the dhcp configuration.



Below is an example of /etc/dhcpd.conf, VLAN ID 127 is used for PXE boot network. The PXE boot server IP address is 192.168.127.6, subnet 255.255.255.0. Assigned IP address for servers are 192.168.127.101-112. The sample file shows 6 hosts.

```
[root@HANA-Mgmtsrv01 ~]# vi /etc/dhcp/dhcpd.conf

# dhcpd.conf
#
default-lease-time 14400;
ddns-update-style none;
ddns-updates off;

filename "pxelinux.0";

subnet 192.168.127.0 netmask 255.255.255.0 {
    group {
        next-server 192.168.127.6;
        filename "pxelinux.0";
        host server01 {
            hardware ethernet 00:25:b5:0a:00:01;
            fixed-address 192.168.127.101;
        }
        host server02 {
            hardware ethernet 00:25:b5:0a:00:07;
            fixed-address 192.168.127.102;
        }
        host server03 {
            hardware ethernet 00:25:b5:0a:00:0d;
            fixed-address 192.168.127.103;
        }
    }
}
```

```

host server04 {
    hardware ethernet 00:25:b5:0a:00:13;
    fixed-address 192.168.127.104;
}
host server05 {
    hardware ethernet 00:25:b5:0a:00:19;
    fixed-address 192.168.127.105;
}
host server06 {
    hardware ethernet 00:25:b5:0a:00:1f;
    fixed-address 192.168.127.106;
}
host server07 {
    hardware ethernet 00:25:b5:0a:00:25;
    fixed-address 192.168.127.107;
}
host server08 {
    hardware ethernet 00:25:b5:0a:00:2b;
    fixed-address 192.168.127.108;
}
host server09 {
    hardware ethernet 00:25:b5:0a:00:31;
    fixed-address 192.168.127.109;
}
host server10 {
    hardware ethernet 00:25:b5:0a:00:37;
    fixed-address 192.168.127.110;
}
host server11 {
    hardware ethernet 00:25:b5:0a:00:3d;
    fixed-address 192.168.127.111;
}
host server12 {
    hardware ethernet 00:25:b5:0a:00:43;
    fixed-address 192.168.127.112;
}
}

```

8. To make sure the DHCP servers can startup on subsequent reboots, execute the following command:

```
chkconfig dhcpd on
```

9. Restart the dhcp service for new configuration to take effect:

```
$ systemctl --system daemon-reload
# $ systemctl restart dhcpd.service
```

Operating System Installation SUSE SLES12SP3

PXE Boot Preparation for SUSE OS Installation

To use the PXE boot server for OS installation, complete the following steps on the PXE server:

10. Download the SLES12Sp3 DVD iso to /tftpboot directory. Mount the SLES 12 SP3 ISO to temp directory.

```
mount -o loop /tftpboot/SLE-12-SP3-SAP-x86_64-GM-DVD.iso /mnt
```

11. Create ISO image repository.

```
mkdir -p /tftpboot/SLES12SP3CD
cd /mnt
cp -ar * /tftpboot/SLES12SP3CD
umount /mnt
```

12. Copy two files “initrd” and “linux” from SLES 12 SP3 ISO.

```
cp /NFS/software/SLES/CD/boot/x86_64/loader/linux /tftpboot/linux_cd_sles12sp3
cp /NFS/software/SLES/CD/boot/x86_64/loader/initrd /tftpboot/initrd_cd_sles12sp3
```

13. Create a text file that will hold the message which will be displayed to the user when PXE boot server is connected.

```
vi /tftpboot/boot.msg

<< Enter Your Customise Message for example: Welcome to PXE Boot Environment>>
```

14. Create a file called: “default” in the directory /tftpboot/pxelinux.cfg with similar syntax to the one shown below:

```
# UCS PXE Boot Definition
DISPLAY ../boot.msg
DEFAULT Install_SLES12SP3
PROMPT 1
TIMEOUT 50
#
LABEL Install_SLES12SP3
    KERNEL linux_cd_sles12sp3
    APPEND initrd=initrd_cd_sles12sp3
install=nfs://192.168.127.11:/PXE_tftpboot/SLES12SP3CD/?device=eth0
```

PROMPT: This line allows user to choose a different booting method. The value of one allows the client to choose a different boot method.

DEFAULT: This sets the default boot label.

TIMEOUT: Indicates how long to wait at the “boot:” prompt until booting automatically, in units of 1/10 s.

LABEL: This section defines a label called: “Install_SLES12SP3” so at the boot prompt when ‘Install_SLES12SP3’ is entered, it execute the commands related to the local label. Here Kernel image and initrd images are specified along with ISO location and the Ethernet device to use.

After the PXE configuration is completed, proceed with the Operating System Installation.




For the latest information on SAP HANA installation and OS customization requirement, see the SAP HANA Installation Guide: <http://www.saphana.com/>

PXE booting with SUSE Linux Enterprise Server 12 SP3

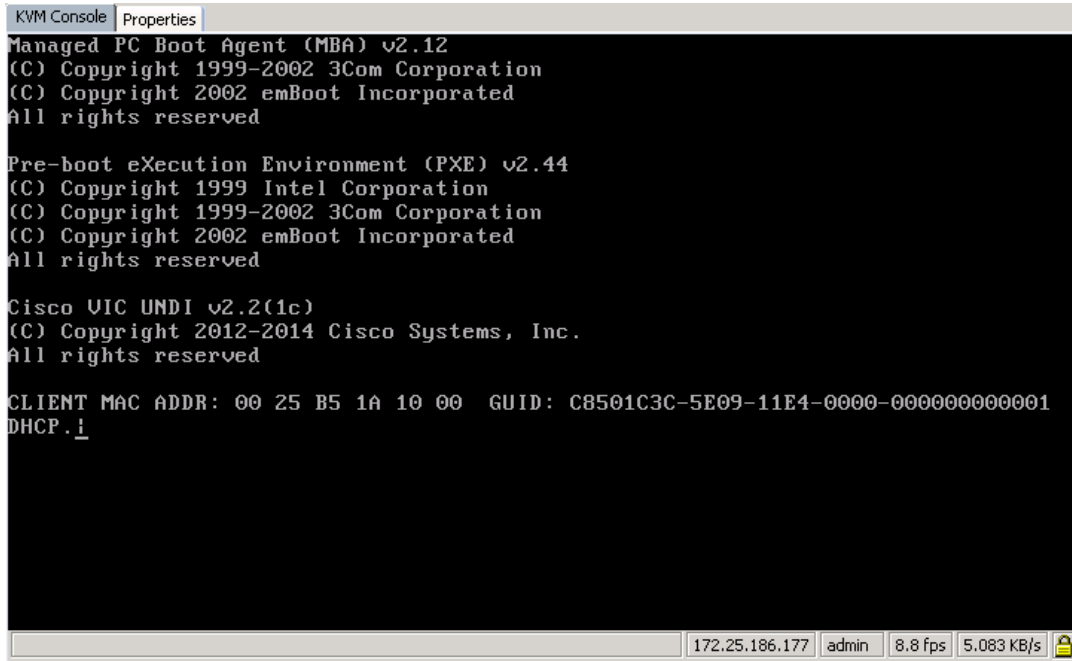
To install the OS based on the PXE Boot Option, complete the following steps:

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

2. Select Service Profiles > root > HANA-Server01.
3. Click KVM Console.
4. When the KVM Console is launched, click Boot Server.



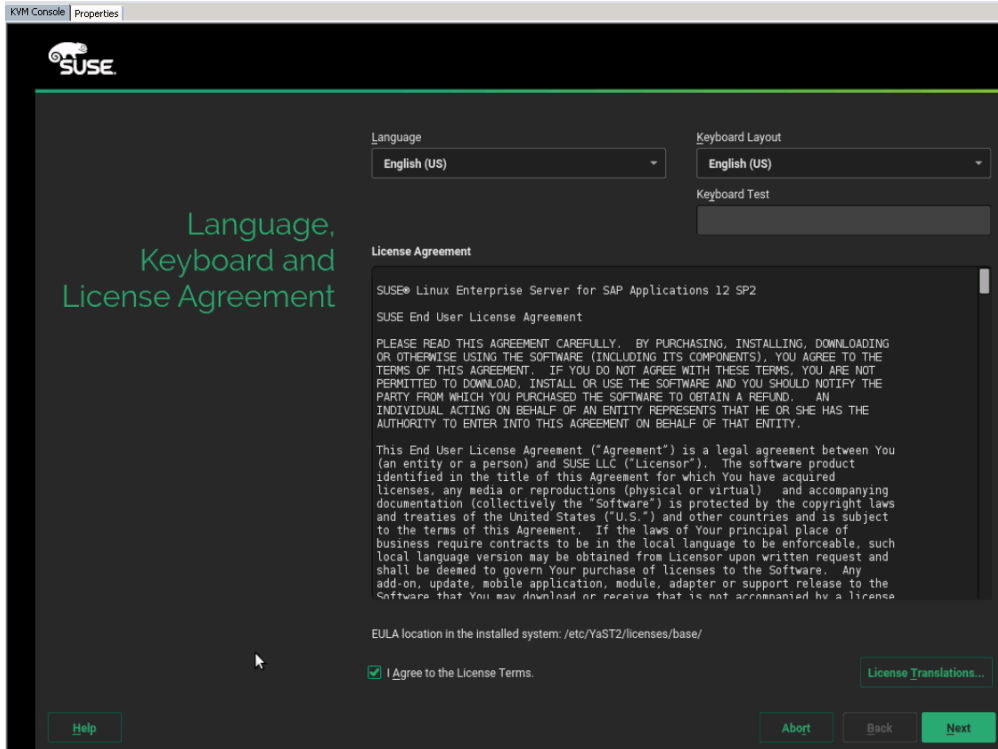
5. If you are using a CD, click Virtual Media > Activate Virtual Devices:
 - a. Select Accept this Session for Unencrypted Virtual Media Session then click Apply.
 - b. Click Virtual Media and Choose Map CD/DVD.
 - c. Click Browse to navigate ISO media location.
 - d. Click Map Device.
6. For PXE Boot Installation, the “default” file is configured for OS installation. The IP address obtained from DHCP configured on the PXE Boot Server.



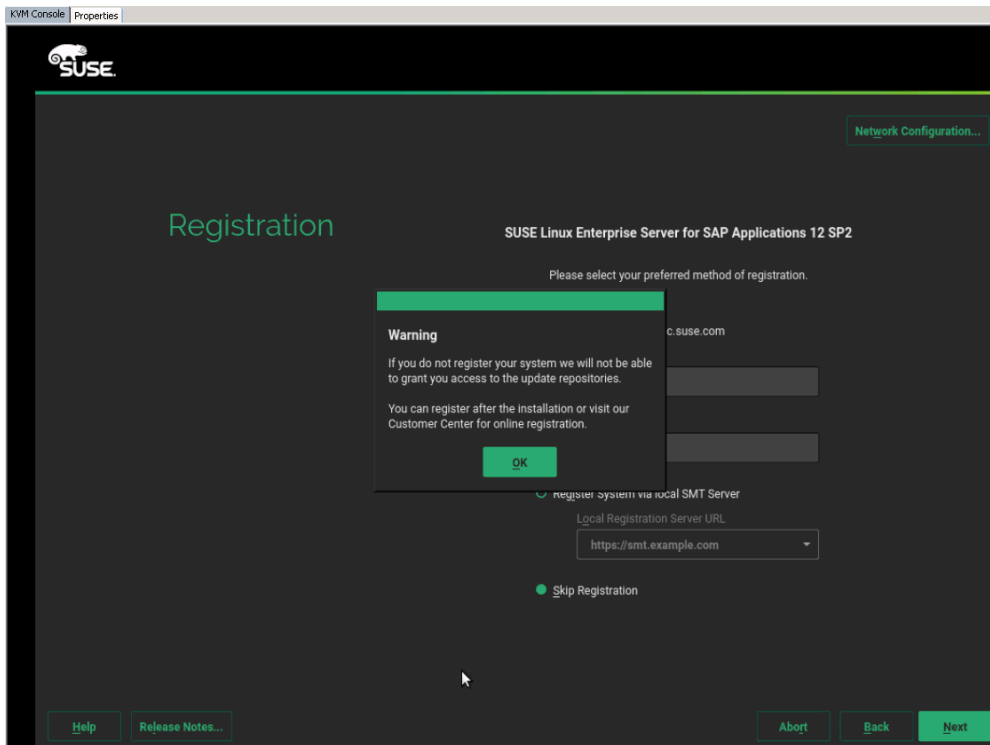
7. Make sure /etc/dhcpd.conf has an fixed-address defined for the client MAC ADDR of the HANA node requesting DHCP services.
8. When the DHCP provides a DHCP offer of the IP address, it reaches out to the next server defined in the dhcpd conf file which is the TFTP server IP. It now loads the Linux and initrd image from PXE server per entries in the /tftpboot/pxelinux.cfg/default file.



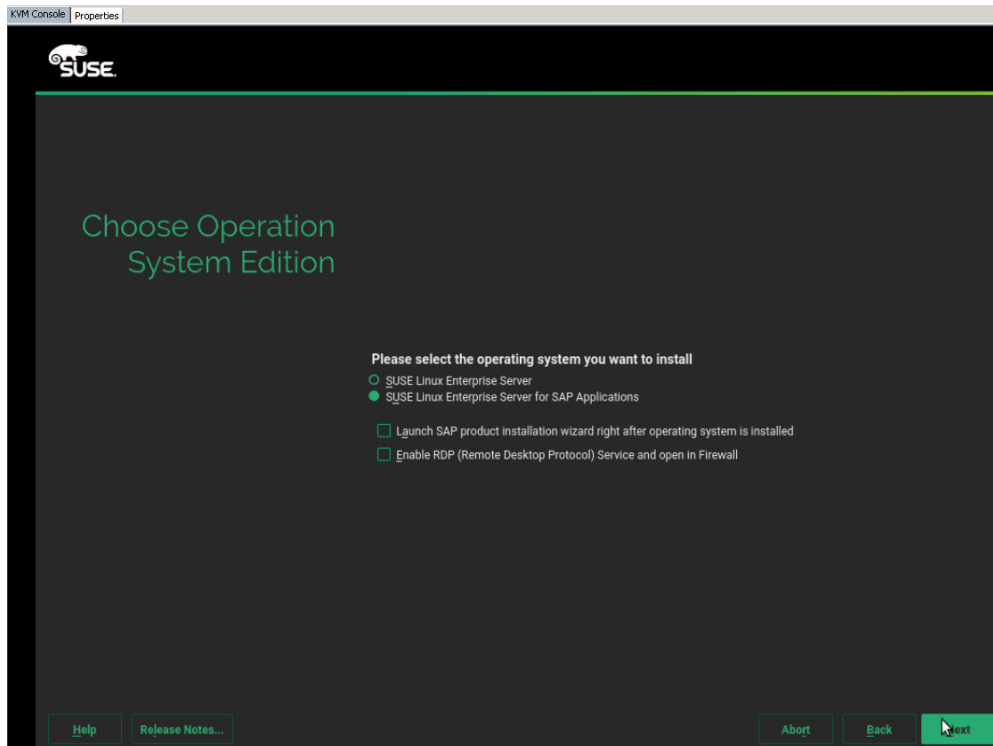
SUSE License agreement is shown below:



9. Agree to the License Terms, click Next.

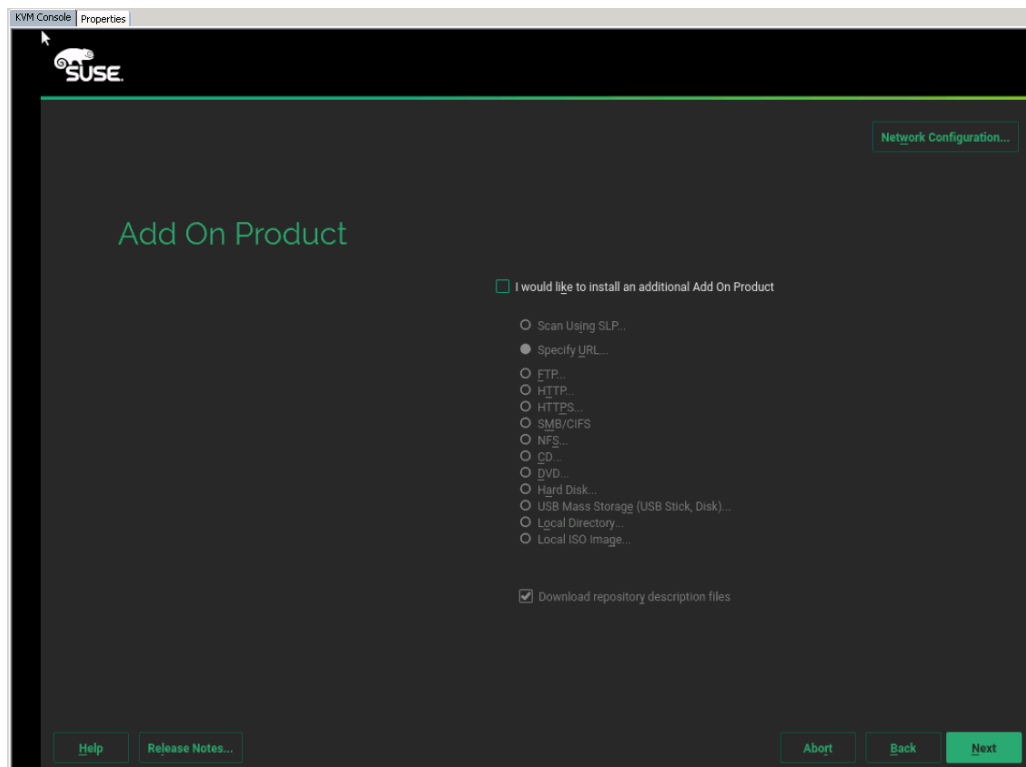


10. Skip the registration and click Next. Select the System Edition.

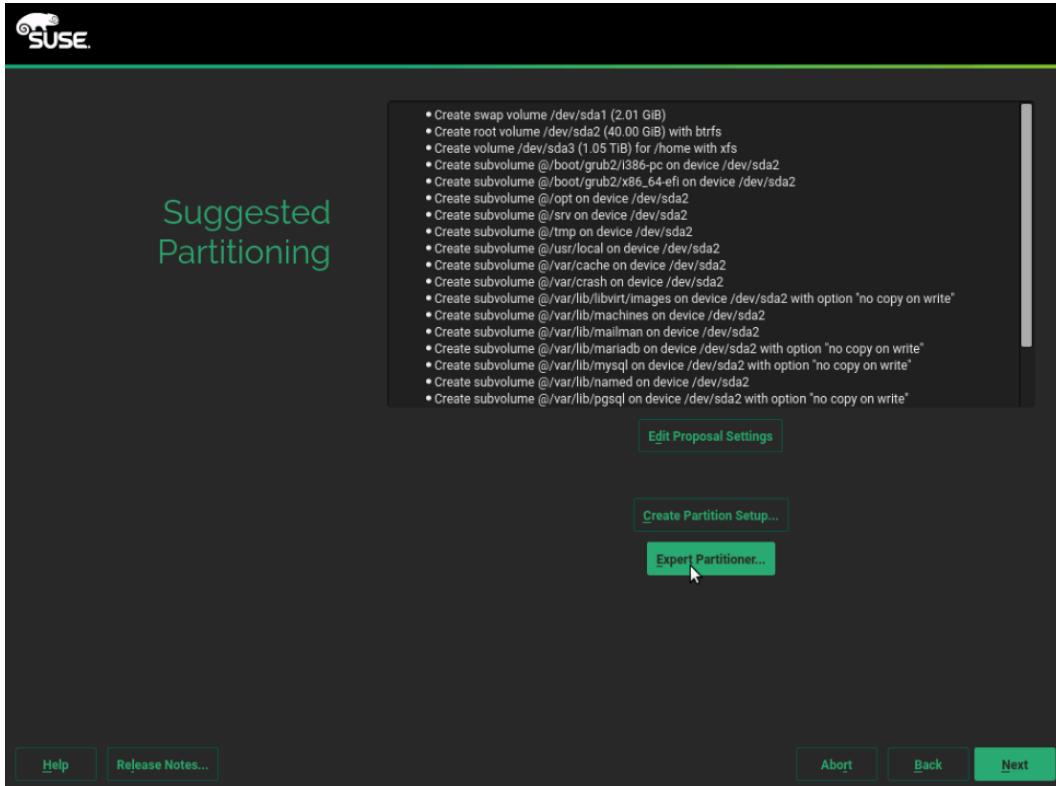


11. Select SLES for SAP and click Next.

12. Do not install add-ons at this moment.

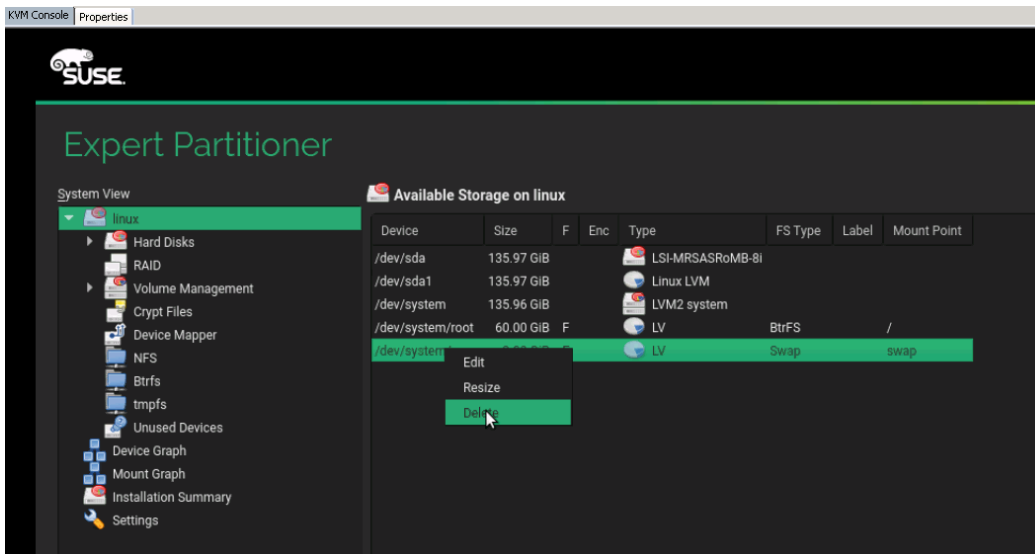


The Disk Partitioner main screen is shown below:



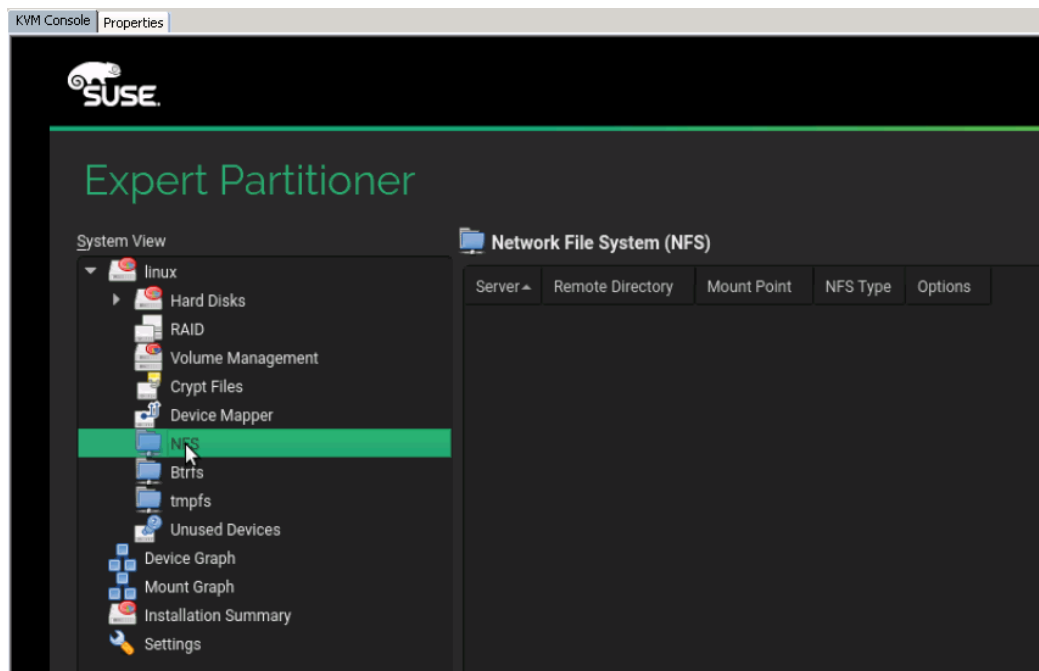
13. Select Expert Partitioner since you do not have a local disk in the system.

14. Delete all existing partitions.



15. After all partitions are deleted select NFS as OS target.


The OS Partitioner for NFS is shown below:



16. On the disk partitioner, select Add.

17. Specify the following:

- a. NFS Server (192.168.127.11 - in this case)
- b. Remote Directory where the OS will be installed (SLES12/osmaster)
- c. Mount point (/)
- d. Options: `rsize=32768, wsize=32768, vers=3, proto=tcp, hard`

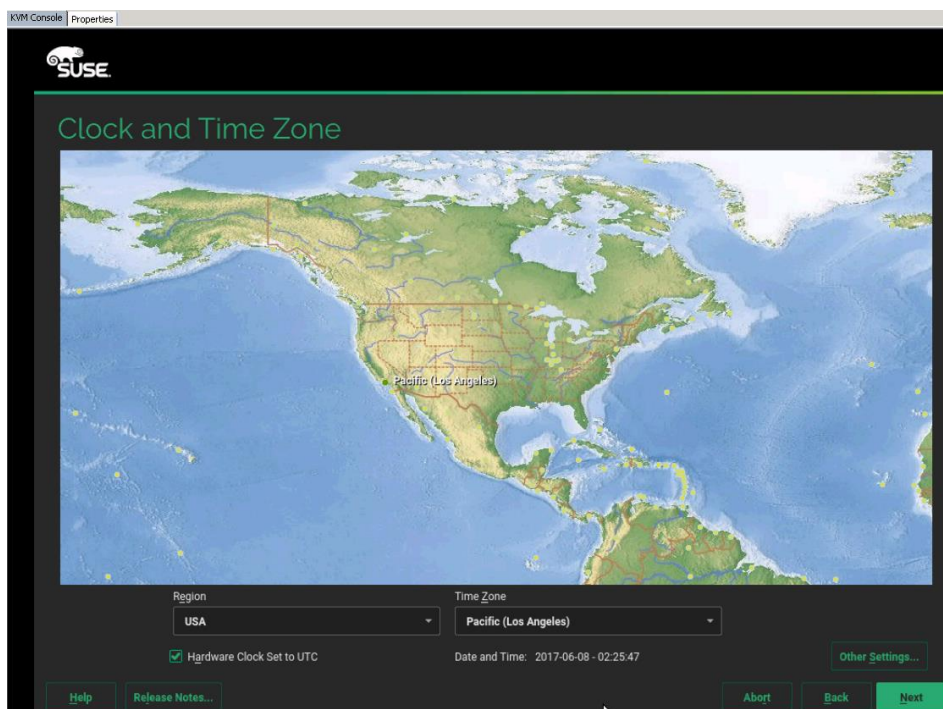


18. If you do not set the correct mount option the installation will fail to install some packages.

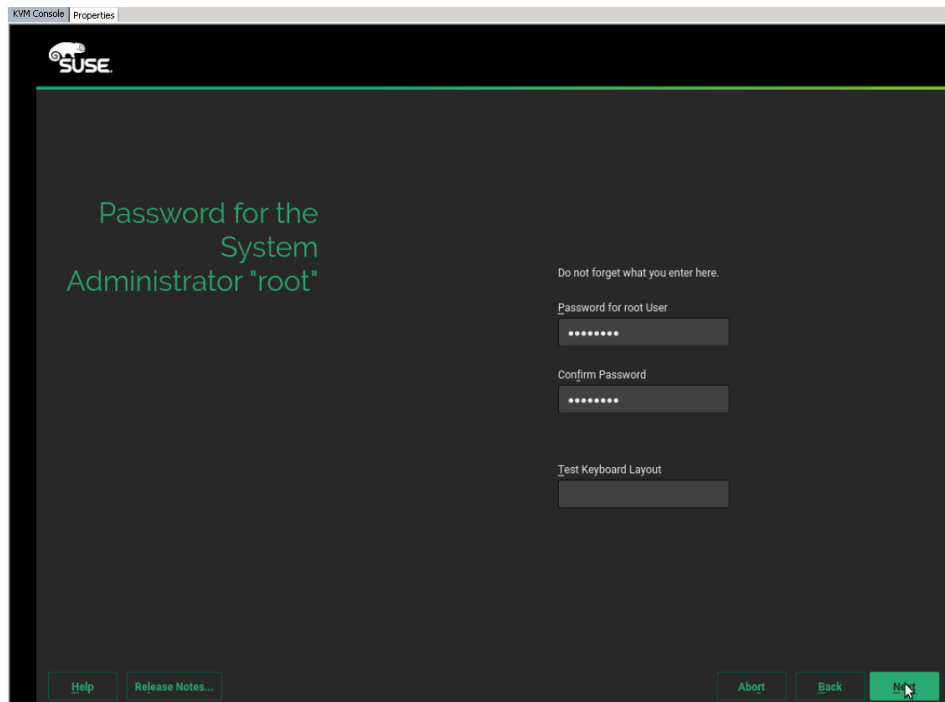
19. Click Accept to confirm the location.

20. Click Next.

21. Select the Time Zone.

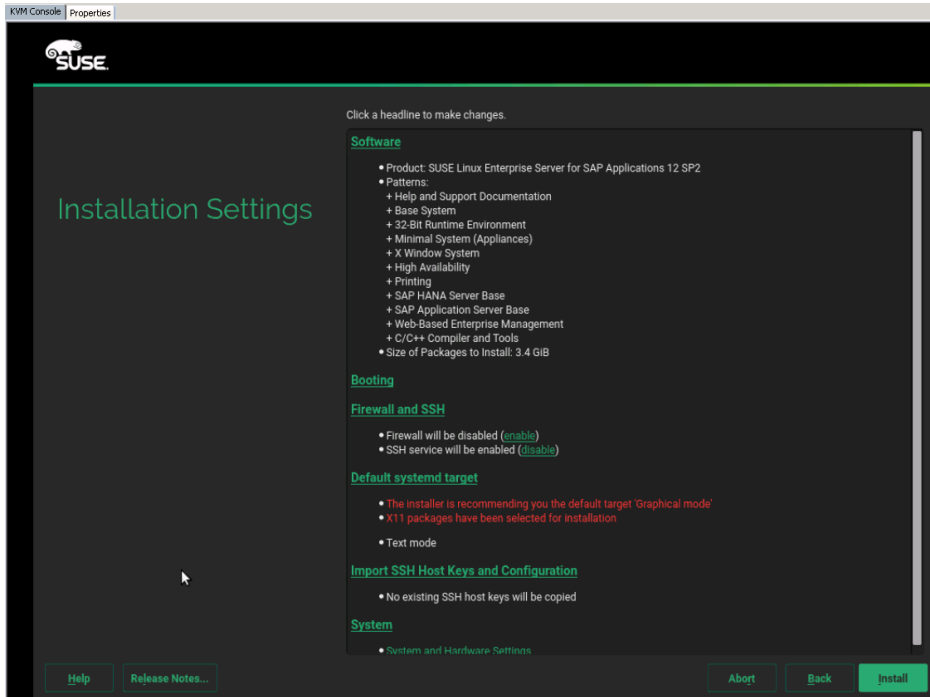


22. Specify the root password.

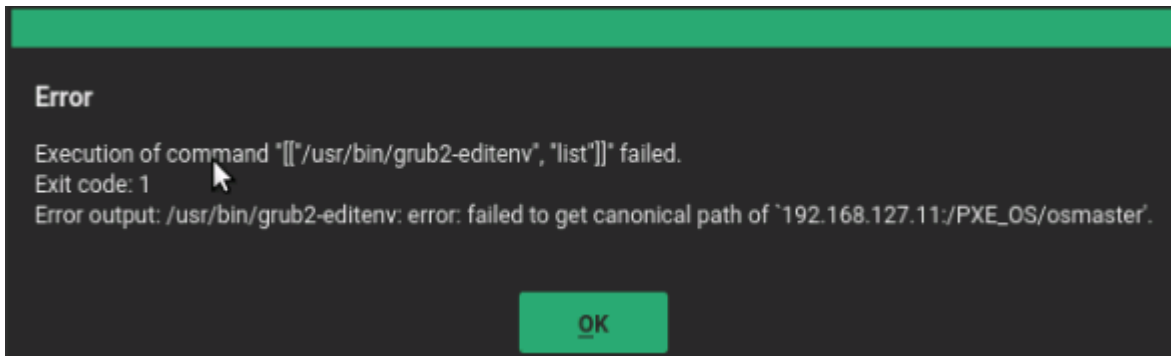


23. Finalize the selection:

- a. Disable the Firewall.
- b. Default System Target must be Text Mode.
- c. Do not import any SSH keys at this moment.
- d. Software:
 - i. Disable Gnome, X Window System
 - ii. Select SAP HANA Server Base
 - iii. (Optional) Select High Availability – (Linux Cluster)
 - iv. Add single Package:
 - v. OPENipmi
 - vi. Ipmitool
 - vii. Screen
 - viii. lftop
 - ix. (Optional) SAPHanaSR-doc (cluster documentation)
 - x. (Optional) sap_suse_cluster_connector

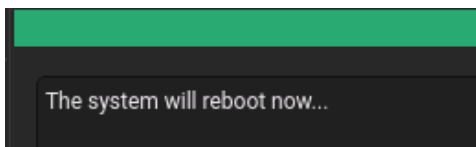


24. Click Install to install the OS.



25. Ignore all Grub installation errors since you are installing on NFS and no grub necessary.

26. The system will reboot after the installation.



27. Since Network boot is used, the bootloader will not be installed.

28. Shutdown the System.

To create the initrd image for PXE Boot environment, complete the following steps:

1. Log into PXE Boot server using ssh.

- Copy the initrd and vmlinuz image from the system installed. Make sure both have 664 permissions set in the /tftpboot directory.

```
cd /PXE_OS/osmaster
cp boot/initrd-4.4.21-69-default /tftpboot
cp boot/vmlinuz-4.4.21-69-default /tftpboot
```

- Create new PXE Configuration file as described in the section Define the PXE Linux Configuration:

```
cd /tftpboot/pxelinux.cfg
vi C0A87F65

# UCS PXE Boot Definition
DISPLAY ../boot.msg
DEFAULT SLES12SAPSP3
PROMPT 1
TIMEOUT 50
#
LABEL SLES12SAPSP3
    KERNEL vmlinuz-4.4.73-5-default
    APPEND initrd=initrd-4.4.73-5-default rw root=/dev/nfs
nfsroot=192.168.127.11:/PXE_OS/osmaster:rw,relatime,vers=3,rsize=32768,wsz=32768,namlen=255,hard,nolock
,proto=tcp,vers=3 rd.neednet=1 transparent_hugepage=never numa_balancing=disabled intel_idle.max_cstate=1
processor.max_cstate=1 ip=dhcp
```

- The hexadecimal filename above was derived using “gethostip <ipaddress>” command; ipaddress is the one mapped for the node corresponding to MAC Addr in dhcpd.conf file.
- Go back to the KVM Console and click OK to reboot the server. Verify the system boots fine.

Create Swap Partition in a File

- ssh to the osmaster [the system will have a random hostname starting linux at this time] on the PXE boot IP from PXE Boot Server.
- Login as root and password.
- Create file for swap partition:

```
<osmaster>:~ # dd if=/dev/zero of=/swap-0001 bs=1M count=2048
2048+0 records in
2048+0 records out
2147483648 bytes (2.1 GB) copied, 3.64515 s, 589 MB/s
```

- Set up a swap area in a file:

```
<osmaster>:~ # mkswap /swap-0001
Setting up swapspace version 1, size = 2097148 KiB
no label, UUID=0f0f9606-dbe9-4301-9f65-293c3bab1346
```

- To use swap file execute the below command:

```
<osmaster>:~ # swapon /swap-0001
```

- Verify if the swap partition is being used:

```

<osmaster>:~ # swapon -s
Filename                                Type      Size      Used      Priority
/swap-0001                              file      2097148  0         -1

```

7. Add the following line (swap) to /etc/fstab for swap partition to be persistent after reboot:

```

vi /etc/fstab
/swap-0001          swap swap defaults    0 0

```

Update OS Master

To update the SUSE OS to the latest patch level, complete the following steps:

1. This document assumes that the SUSE License key are available and registered username and password is available.
2. ssh to the os master on the PXE boot IP from PXE Boot Server.
3. Login as root and password.
4. Assign IP address to the interface which can access the Internet or Proxy Server.
In this example HANA-Admin vNIC to access internet is used.
5. To configure the network interface on the OS, it is required to identify the mapping of the ethernet device on the OS to vNIC interface on the Cisco UCS.
6. From the OS execute the following command to get list of Ethernet device with MAC Address:

```

osmaster:~ # ifconfig -a|grep HWaddr
eth0      Link encap:Ethernet HWaddr 00:25:B5:1A:10:00
eth1      Link encap:Ethernet HWaddr 00:25:B5:00:0A:02
eth2      Link encap:Ethernet HWaddr 00:25:B5:00:0B:02
eth3      Link encap:Ethernet HWaddr 00:25:B5:00:0A:00
eth4      Link encap:Ethernet HWaddr 00:25:B5:00:0B:00
eth5      Link encap:Ethernet HWaddr 00:25:B5:00:0B:01
eth6      Link encap:Ethernet HWaddr 00:25:B5:00:0B:03
eth7      Link encap:Ethernet HWaddr 00:25:B5:00:0A:01
eth8      Link encap:Ethernet HWaddr 00:25:B5:00:0A:03

```

7. In Cisco UCS Manager, click the Servers tab in the navigation pane.
8. Select Service Profiles > root > HANA-Server01 Expand by clicking +.
9. Click vNICs.
10. On the Right pane list of the vNICs with MAC Address are listed.
11. Take note of the MAC Address for the HANA-Admin vNIC is “00:25:B5:00:0A:03”
12. Compare MAC Address on the OS and UCS. In our case, it is MAC address of eth8 and will carry the VLAN for HANA-Admin.
13. Go to network configuration directory and create a configuration for eth8:

```

/etc/sysconfig/network

vi ifcfg-eth8

##
# HANA-Admin Network
##
BOOTPROTO='static'
IPADDR='<<IP Address for HANA-Admin>>/24'
MTU=''
NAME='VIC Ethernet NIC'
STARTMODE='auto'

```

14. Add default gateway:

```

cd /etc/sysconfig/network
vi routes

default <<IP Address of default gateway>> - -

```

15. Add the DNS IP if its required to access internet:

```

vi /etc/resolv.conf

nameserver <<IP Address of DNS Server1>>
nameserver <<IP Address of DNS Server2>>

```

16. Restart the network service for the change to take effect:

```
rcnetwork restart
```

17. Execute the following command to Register the SUSE:

```
SUSEConnect -r <<registration_code>>
```

18. After the registration, the entire repository will be updated. Now execute the following command to update the server:

```
zypper update
```

19. Follow the on-screen instruction to complete the update process. Ignore the errors and warning related to GRUB and bootloader updates.

20. Do not reboot the server until initrd and vmlinuz images are updated.

To update initrd image for PXE Boot environment, complete the following steps:

1. Log into PXE server using ssh
2. Copy the initrd and vmlinux image from the nfsroot of osmaster system to /tftpboot directory

```

cp /PXE_OS/osmaster/boot/initrd-4.4.140-94.42-default /tftpboot/
cp /PXE_OS/smaster/boot/vmlinuz-4.4.140-94.42-default /tftpboot/

```

3. Update the PXE Configuration file:

```

vi /tftpboot/pxelinux.cfg/ C0A87F65

# UCS PXE Boot Definition
DISPLAY ../boot.msg
DEFAULT SLES4SAP
PROMPT 1
TIMEOUT 50
#
LABEL SLES4SAP
    KERNEL vmlinuz-4.4.140-94.42-default
    APPEND initrd= initrd-4.4.140-94.42-default rw root=/dev/nfs
nfsroot=192.168.127.11:/PXE_OS/osmaster:rw,relatime,vers=3,rsize=32768,wsz=32768,namlen=255,hard,nolock
,proto=tcp,vers=3 rd.neednet=1 transparent_hugepage=never numa_balancing=disabled intel_idle.max_cstate=1
processor.max_cstate=1 ip=dhcp

```

4. ssh to the os master server with PXE boot IP (192.168.127.101) from PXE Boot Server.

Install Cisco enic Driver

This section describes how to download the Cisco UCS Drivers ISO bundle, which contains most Cisco UCS Virtual Interface Card drivers.

1. In a web browser, navigate to <http://www.cisco.com>.
2. Under Support, click All Downloads.
3. In the product selector, click Products, then click Server - Unified Computing.
4. If prompted, enter your Cisco.com username and password to log in.
5. You must be signed in to download Cisco Unified Computing System (UCS) drivers.
6. Cisco UCS drivers are available for both Cisco UCS B-Series Blade Server Software and Cisco UCS C-Series Rack-Mount UCS-Managed Server Software.
7. Click Cisco UCS B-Series Blade Server Software.
8. Click Cisco Unified Computing System (UCS) Drivers.
9. The latest release version is selected by default. This document is built on Version 3.2(3d)
10. Click 3.2 Version.
11. Download ISO image of Cisco B-Series drivers.
12. Also check the bootable driver kit posted on SUSE site for the Cisco VIC card for the particular OS release. Eg <https://drivers.suse.com/cisco/UCS-VIC/sle12-sp3/install-readme.html> ENIC and FNIC packages are also available for download from there.
13. Choose your download method and follow the prompts to complete your driver download.
14. After the download completes, browse the ISO to Cisco ucs-bxxx-drivers.3.2.3b, copy cisco-enic-usnic-kmp-default-3.0.44.553.545.8_k4.4.73_5-3.1.x86_64 to PXE Boot Server /tftpboot/drivers/

To update the initrd image for PXE Boot environment, complete the following steps:

1. Log into PXE server using ssh.
2. Copy the updated initrd and vmlinuz images from the **osmaster system's boot directory to /tftpboot.**

```
cd /PXE_OS/osmaster/boot
cp initrd-4.4.140-94.42-default /tftpboot/
cp vmlinuz-4.4.140-94.42-default /tftpboot/
```

Operating System Optimization for SAP HANA

To configure the OS optimization settings on the OS Master, complete the following steps:

1. ssh to the os master on the PXE boot IP from PXE Boot Server.
2. Login as root and password.
3. Implement the instructions per the SAP Note [2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12](#)
4. Add the following lines to /etc/sysctl.conf:

```
net.ipv6.conf.all.disable_ipv6 = 1
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.core.rmem_default = 16777216
net.core.wmem_default = 16777216
#
net.core.optmem_max = 16777216
net.core.netdev_max_backlog = 300000
#
net.ipv4.tcp_rmem = 65536 16777216 16777216
net.ipv4.tcp_wmem = 65536 16777216 16777216
#
net.ipv4.tcp_max_syn_backlog = 16348
net.ipv4.tcp_slow_start_after_idle = 0
sunrpc.tcp_slot_table_entries = 128
sunrpc.tcp_max_slot_table_entries = 128
#
vm.swappiness=10
#
net.ipv4.ip_local_port_range = 40000 65300
net.ipv4.conf.all.rp_filter = 0
# SAP Note 1868829
fs.aio-max-nr = 18446744073709551615
net.ipv4.tcp_dsack = 1
net.ipv4.tcp_sack = 1
#For background information, see SAP Note 2205917 and 1557506
vm.pagecache_limit_mb = 0
```

5. `vm.pagecache_limit_ignore_dirty = 1`Disable (blacklist) the unnecessary driver:

```
vi /etc/modprobe.d/50-blacklist.conf
#
# disable modules for NFS and HANA
blacklist kvm
blacklist kvm_intel
blacklist iTCO_wdt
```

```
blacklist iTCO_vendor_support
```

- Set the sunrpc limits to 128:

```
vi /etc/modprobe.d/sunrpc-local.conf
options sunrpc tcp_max_slot_table_entries=128
```

- Dracut to re-create the initrd for use with multiple systems:
- Clear the contents of /etc/dracut.conf file. It appears that it looks into /etc/dracut.conf.d directory for the conf file.

```
<osmaster>~# vi /etc/dracut.conf.d/10-default.conf

logfile=/var/log/dracut.log
#fileloglvl=6
# Exact list of dracut modules to use. Modules not listed here are not going
# to be included. If you only want to add some optional modules use
# add_dracutmodules option instead.
#dracutmodules+="
# dracut modules to omit
omit_dracutmodules+="fcoe fcoe-uefi nbd"
# dracut modules to add to the default
add_dracutmodules+="systemd ssh-client nfs network base"
# additional kernel modules to the default
add_drivers+="sunrpc nfs nfs_acl nfsv3 fnic enic igb ixgbe lpfc"
# list of kernel filesystem modules to be included in the generic initramfs
#filesystems+="
# build initrd only to boot current hardware
hostonly="no"
```

- Create a server independent initrd:

```
<osmaster>~# cd /boot
Dracut -f /boot/initrd_SLES12SP3_<number>.img
ls -ltr /boot/initrd_nfsboot_SLES12SP3_001.img
-rw----- 1 root root 46723100 Jun  9  2017 /boot/initrd_nfsboot_SLES12SP3_001.img
```

- This initrd can now be transferred to the PXE server to boot from the next time.

Cloning OS Volumes

After OS Master image is created, prepare the os image for cloning.

Clean UP Master OS Image

- ssh to osmaster system.
- Remove the SUSE Registration information. This step is required to create a master image without the registration information.
- Remove the Update Service

```
zypper removerepo <repository-name>
```

- Shutdown the OS Master Server by issuing "halt" command.
- Log into PXE server using ssh.

6. Clear the System logs:

```
rm /PXE_OS/osmaster/var/log/* -r
```

7. Clear the Ethernet Persistent network information:

```
cat /dev/null > /PXE_OS/osmaster/etc/udev/rules.d/70-persistent-net.rules
```

8. Remove any Ethernet configuration file except eth0:

```
rm /PXE_OS/osmaster/etc/sysconfig/network/ifcfg-eth<<1-7>>
```

Storage Clone of OS Volume

To clone the OS master image (FlexClone License required) to new the host, complete the following steps:

1. Log in to Storage shell.
2. Create a Clone of OS master volume:

```
volume clone create -flexclone server01 -parent-volume osmaster -vserver infra_vs1 -junction-path /server01 -space-guarantee none
```

3. Split the volume from OS master volume:

```
AFF A300-cluster::> volume clone split start -flexclone server01
Warning: Are you sure you want to split clone volume server01 in Vserver
        infra_vs1 ? {y|n}: y
[Job 1372] Job is queued: Split server01.
```

4. Check for status of Clone split:

```
AFF A300-cluster::> volume clone split status -flexclone server01
```

Vserver	FlexClone	Inodes		Blocks		Updated	% Complete
		Processed	Total	Scanned	Updated		
infra_vs1	server01	149558	253365	541092	538390		59

5. When the clone split is completed:

```
AFF A300-cluster::> volume clone split status -flexclone server01
There are no entries matching your query.
```

6. Repeat the steps 2-3 for each server to deploy server root filesystems with OS image.

Manual Clone of OS Volume

If the FlexClone license is not available it is also possible to distribute the OS Image.

1. Create the OS Volume on the storage and use qtrees to separate each OS:

```
vol create -vserver Infra-SVM -volume PXE_OS -aggregate hana01 -size 200GB -state online -policy default
-unix-permissions ---rwxr-xr-x -type RW -snapshot-policy default

qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server01
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server02
```

```

qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server03
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server04
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server05
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server06
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server07
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server08
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server09
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server10
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server11
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server12

volume mount -vserver Infra-SVM -volume PXE_OS -junction-path /PXE_OS

```

2. Optionally mount the Server root filesystems on the management server for easier access.

```

mount -a
df -hT
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/sda2       ext3      58G   27G   30G   48% /
udev            tmpfs     3.9G  112K  3.9G   1% /dev
tmpfs           tmpfs     8.0G   724K  8.0G   1% /dev/shm
lif-pxe-1:/tftpboot nfs       973M   1.1M  972M   1% /tftpboot
lif-pxe-1:/PXE_OS nfs      190G   320K  190G   1% /NFS/PXE_OS

cd /NFS/PXE_OS/
ls -l

drwxr-xr-x 2 root root 4096 Mar 31 2017 Server01
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server02
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server03
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server04
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server05
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server06
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server07
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server08
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server09
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server10
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server11
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server12

```

PXE Configuration for Additional Server

The PXE boot environment will search for a configuration file based on its boot IP assigned through DHCP.

1. To calculate the filename run “gethostip”, the output is a hex representation of the IP address will be configuration filename:

```

gethostip 192.168.127.201
192.168.127.201 192.168.127.201 C0A87FC9

```

2. The file name “C0A87FC9” contains the PXE boot configuration for server with IP 192.168.127.201.
3. ssh to PXE boot server.
4. Go to PXE boot configuration directory:

```
cd /tftpboot/pxelinux.cfg/
```

5. Create a configuration file for each server:

```

vi COA87FC9

# UCS PXE Boot Definition
DISPLAY ../boot.msg
DEFAULT SLES4SAP
PROMPT 1
TIMEOUT 50
#
LABEL SLES4SAP
    KERNEL suse/vmlinuz-sles12sp3-69
    APPEND initrd=suse/initrd-sles12sp3-69 rw root=/dev/nfs
nfsroot=192.168.127.11:/vol/SLES12SP3:rw,relatime,vers=3,rsize=32768,wsiz=32768,namlen=255,hard,nolock,p
roto=tcp,vers=3 rd.neednet=1 transparent_hugepage=never numa_balancing=disabled intel_idle.max_cstate=1
processor.max_cstate=1 ip=dhcp

```

6. Repeat the previous step for each server.

7. Example: PXE Boot configuration file for server with dhcp ip 192.168.201.202:

```

gethostip 192.168.127.202
192.168.127.202 192.168.127.202 COA87FCA

vi COA87FCA

# UCS PXE Boot Definition
DISPLAY ../boot.msg
DEFAULT SLES4SAP
PROMPT 1
TIMEOUT 50
#
LABEL SLES4SAP
    KERNEL suse/vmlinuz-sles12sp3-69
    APPEND initrd=suse/initrd-sles12sp3-69 rw root=/dev/nfs
nfsroot=192.168.127.11:/vol/SLES12SP3:rw,relatime,vers=3,rsize=32768,wsiz=32768,namlen=255,hard,nolock,p
roto=tcp,vers=3 rd.neednet=1 transparent_hugepage=never numa_balancing=disabled intel_idle.max_cstate=1
processor.max_cstate=1 ip=dhcp

```

Boot the Server

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Profiles.
3. Expand the tree and right-click Service Template HANA-Server22 and select Boot Server.

Post Installation OS Customization

After the OS is deployed from the Master image, customization is required for each server.

Hostnames

The operating system must be configured such a way that the short name of the server is displayed for the **command 'hostname'** and **Full Qualified Host Name is displayed with the command 'hostname -d'**.

1. ssh to the Server to PXE boot IP from PXE Boot Server.
2. Login as root and password.

3. Edit the Hostname:

```
vi /etc/HOSTNAME
<<hostname>>.<<Domain Name>>
```

IP Address

1. Assign the IP address to each interface.
2. ssh to the Server on the PXE boot IP from PXE Boot Server.
3. Login as root and password.
4. To configure the network interface on the OS, it is necessary to identify the mapping of the ethernet device on the OS to vNIC interface on the Cisco UCS. From the OS execute the below command to get list of Ethernet device with MAC Address.

```
ifconfig -a |grep HWaddr

eth0      Link encap:Ethernet  HWaddr 00:25:B5:1A:10:00
eth1      Link encap:Ethernet  HWaddr 00:25:B5:00:0A:02
eth2      Link encap:Ethernet  HWaddr 00:25:B5:00:0B:02
eth3      Link encap:Ethernet  HWaddr 00:25:B5:00:0A:00
eth4      Link encap:Ethernet  HWaddr 00:25:B5:00:0B:00
eth5      Link encap:Ethernet  HWaddr 00:25:B5:00:0B:01
eth6      Link encap:Ethernet  HWaddr 00:25:B5:00:0B:03
eth7      Link encap:Ethernet  HWaddr 00:25:B5:00:0A:01
eth8      Link encap:Ethernet  HWaddr 00:25:B5:00:0A:03
```

5. In Cisco UCS Manager, click the Servers tab in the navigation pane.
6. Select Service Profiles > root > HANA-Server01 Expand by clicking +.
7. Click vNICs.
8. On the right pane list of the vNICs with MAC Address are listed. Note the MAC Addresses and by comparing MAC Address on the OS and Cisco UCS, assign the IP addresses to the vNICs appropriately
9. Add default gateway.

```
cd /etc/sysconfig/network
vi routes

default <<IP Address of default gateway>> - -
```

Network Time

It is important to sync the time on all components used for SAP HANA. The configuration of NTP is important and should be configured on all systems, as provided below:

```
vi /etc/ntp.conf
server <NTP-SERVER IP>
fudge <NTP-SERVER IP> stratum 10
keys /etc/ntp.keys
```

```
trustedkey 1
```

DNS

Domain Name Service configuration must be done based on the local requirements.

Configuration Example:

```
vi /etc/resolv.conf
nameserver <<IP Address of DNS Server1>>
nameserver <<IP Address of DNS Server2>>
```

HOSTS

For SAP HANA Scale-Out system, all nodes should be able to communicate via all networks configured. Update the hosts files accordingly.

SSH Keys

The SSH Keys must be exchanged between all nodes in a SAP HANA Scale-Out system **for user 'root' and user <SID>adm**.

1. Generate the rsa public key by executing the command `ssh-keygen -b 2048`

```
cishana01:~ # ssh-keygen -b 2048
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
5c:5b:e9:cd:f9:73:71:39:ec:ed:80:a7:0a:6c:3a:48 [MD5] root@cishana01.ciscolab.local
The key's randomart image is:
+--[ RSA 2048]-----+
|
|          .
|         . o
|        . + o...
|       s . . +=.
|      .. ... .
+---[MD5]-----+
```

2. The SSH Keys must be exchanged between all nodes in a SAP HANA Scale-Out system **for user 'root' and user**.
3. Exchange the rsa public key by executing the following command from the first server to the remaining servers in the scale-out system.

“ssh-copy-id -i /root/.ssh/id_rsa.pub cishana02”

```
cishana01:/ # ssh-copy-id -i /root/.ssh/id_rsa.pub cishana02
The authenticity of host 'cishana02 (172.29.220.202)' can't be established.
ECDSA key fingerprint is 93:b9:d5:1a:97:a9:32:10:4f:c2:ef:99:b8:7c:9d:52 [MD5].
Are you sure you want to continue connecting (yes/no)? yes

Password:
```

```
Number of key(s) added: 1
```

```
Now try logging into the machine, with: "ssh 'cishana02'"
and check to make sure that only the key(s) you wanted were added.
```

- Repeat the steps 1-3 for all the servers in the single SID HANA system.

(Optional) Syslog

For a centralized monitoring of all SAP HANA nodes, it is recommended that syslog-ng is configured to forward all messages to a central syslog server.

Change the syslog-ng.conf file as shown below:

```
vi /etc/syslog-ng/syslog-ng.conf
...
```

Operating System Installation Red Hat Enterprise Linux 7.4

This section describes the OS installation based on iSCSI to be used as PXE source. If you do not need the PXE option simply use only the first part of this installation. RHEL does not provide the option to install the OS directly on an NFS location. You must first install the OS on an iSCSI LUN or a local hard disk and then **copy the OS via “rsync” over to the NFS share.**



Use the SAP HANA Installation Guide for OS customization.

- Prepare the iSCSI LUN like described in the Storage part of this CVD for the OS.
- In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Select Service Profiles > root > HANA-Server01.
- Click KVM Console.
- When the KVM Console is launched, click Boot Server.
- If you using CD click Virtual Media > Activate Virtual Devices.
- Select Accept this Session for Unencrypted Virtual Media Session then click Apply.
- Click Virtual Media and Choose Map CD/DVD.
- Click Browse to navigate ISO media location.
- Click Map Device.
- Normally a reboot is necessary to activate this virtual drive.
- During the reboot the iSCSI targets must be shown. If not check the iSCSI configuration.

```
Cisco VIC iSCSI, Boot Driver Version 4.2(2a)
(C) 2016 Cisco Systems, Inc.
 00:25:b5:0a:00:50 iSCSI NETAPP
 00:25:b5:0a:00:50 iSCSI NETAPP
Option ROM installed successfully
```

- 13. After the POST the system will boot from the RHEL ISO.
- 14. At the prompt of the Installation options.
- 15. Press the Tab key to alter the command line options. Append parameter rd.iscsi.ibft=1 to the kernel command line as shown below:

```
Red Hat Enterprise Linux 7.4

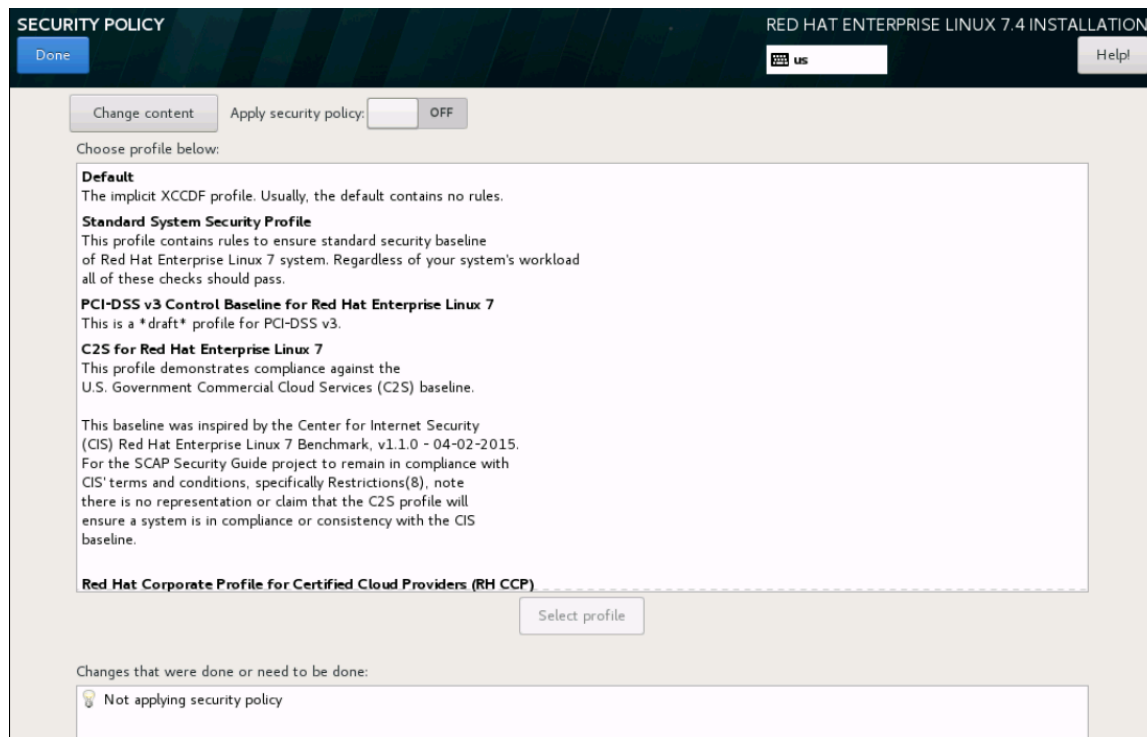
Install Red Hat Enterprise Linux 7.4
Test this media & install Red Hat Enterprise Linux 7.4
Troubleshooting >

> vmlinuz initrd=initrd.img inst.stage2=hd:LABEL=RHEL-7.4\x20Server.x86_64 quiet rd.iscsi.ibft=1_
```



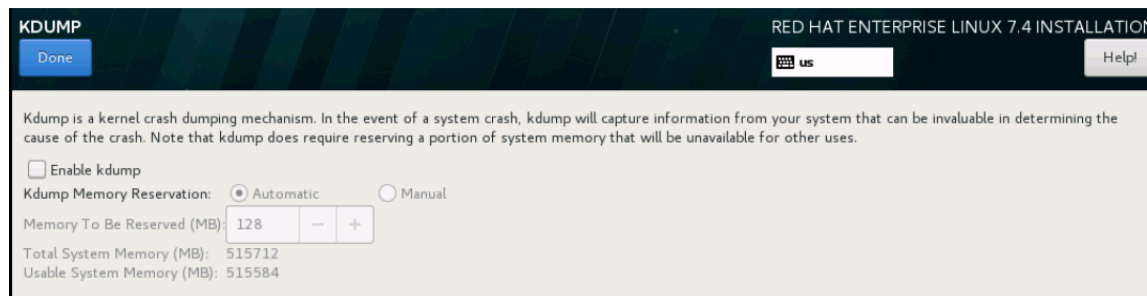
This is added to make sure the iSCSI targets are discovered appropriately.

- 16. Choose Keyboard and configure your layout.
- 17. Configure the right Timezone and Time.
- 18. Click the 'Security Policy' to set the security policy to OFF.



19. Leave the Software section selections as default (Minimal Installation).

20. Disable KDUMP.




21. Click Installation destination.

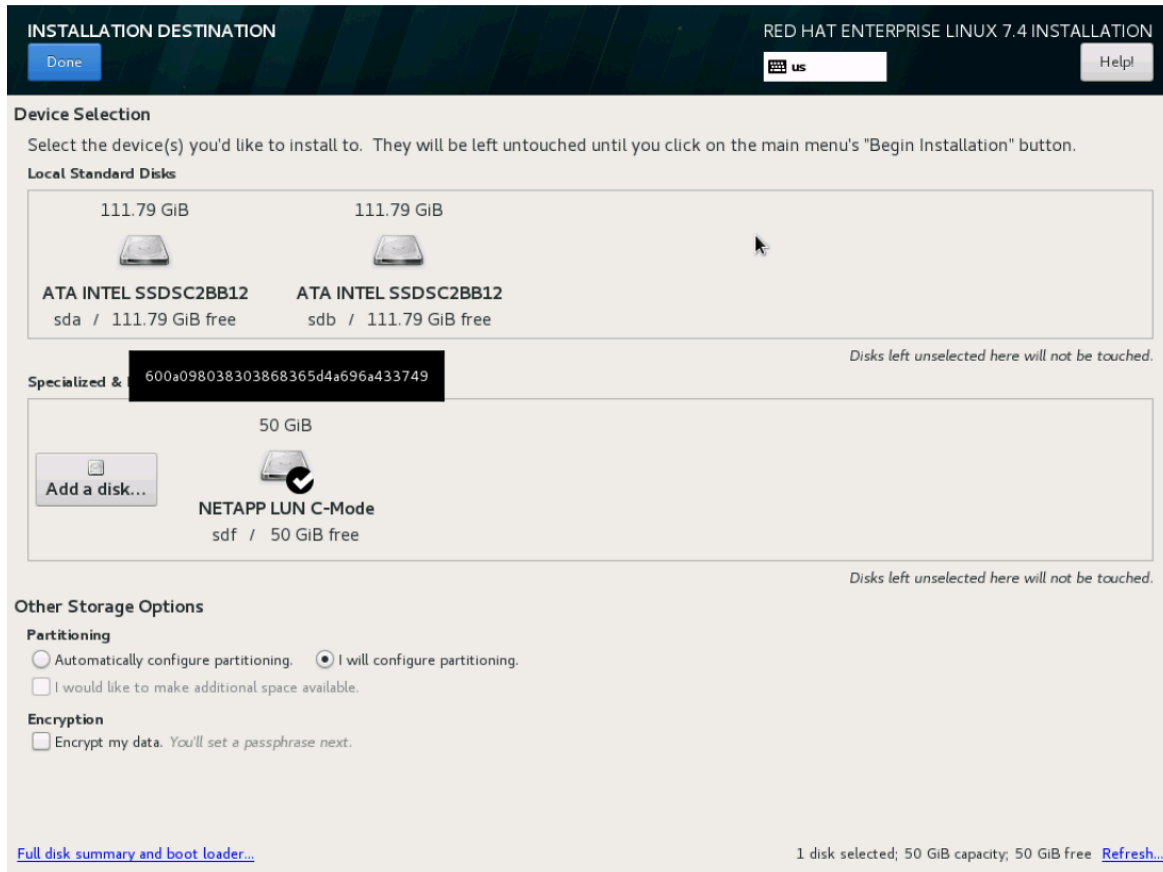


The next screen lists all the local standard disks, if any.

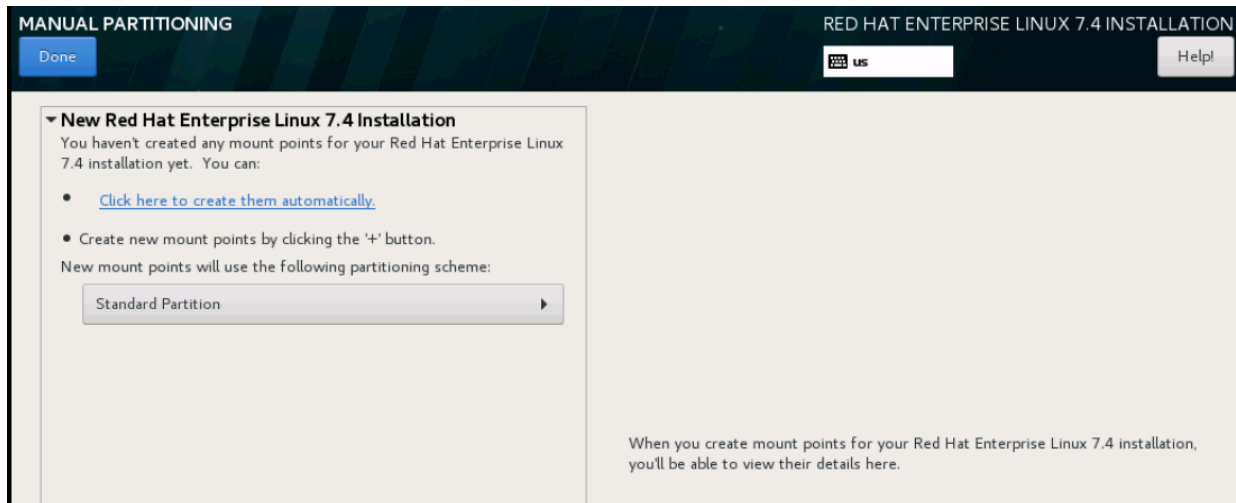
22. Click Specialized and Network Disks section - “Add a Disk”
23. Select the discovered iSCSI boot LUN disk and click “Done.”




24. From the "Other Storage options" select "I will configure partitioning."
25. Click Done.



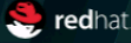
26. Select Standard Partition and then select “Click here to create them automatically.”



27. Confirm the default Partition table. Click Done. Accept Changes.



28. Click Begin Installation and then setup the root password.



INSTALLATION SUMMARY

RED HAT ENTERPRISE LINUX 7.4 INSTALLATION us Help!

LOCALIZATION

DATE & TIME
Americas/New York timezone

KEYBOARD
English (US)

LANGUAGE SUPPORT
English (United States)

SOFTWARE

INSTALLATION SOURCE
Local media

SOFTWARE SELECTION
Minimal Install

SYSTEM

INSTALLATION DESTINATION
Custom partitioning selected

NETWORK & HOST NAME
Wired (ibft0) connected

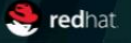
KDUMP
Kdump is enabled

SECURITY POLICY
No profile selected

No profile selected

Quit Begin Installation

We won't touch your disks until you click 'Begin Installation'.



CONFIGURATION

RED HAT ENTERPRISE LINUX 7.4 INSTALLATION us Help!

USER SETTINGS

ROOT PASSWORD
Root password is not set

USER CREATION
No user will be created

ROOT PASSWORD

Done us Help!


The root account is used for administering the system. Enter a password for the root user.

Root Password:

Confirm:

Good

29. If all packages are installed reboot the system.



Post Installation Tasks

Configuring the Network

In RHEL 7, system and udev support a number of different naming schemes. By default, fixed names are assigned based on firmware, topology, and location information, for example, enp72s0.

With this naming convention, although names remain fixed even if hardware is added or removed, names often are more difficult to read than with traditional kernel-native ethX naming: that is, eth0, etc.

Another convention for naming network interfaces, biosdevnames, is available with installation.

If you require to go back, the traditional device names these parameter later in the PXE configuration `net.ifnames=0 biosdevname=0`. Also, you can disable IPv6 support `ipv6.disable=1`.

1. Log in to the newly installed system as root.
2. Configure the network.
3. Get the MAC addresses from Cisco UCS Manager.

Name	MAC Address
vNIC Server-...	00:25:B5:76:00:1F
vNIC vNIC1	00:25:B5:11:00:9F
vNIC vNIC2	00:25:B5:12:00:EF
vNIC vNIC3	00:25:B5:76:00:0F
vNIC vNIC4	00:25:B5:12:00:BF

The order in this example: vNIC1 = Admin LAN ; vNIC2 = PXE Boot; vNIC3 = Access LAN ; vNIC4 = NFS LAN

4. Configure the Access network, default GW and the resolv.conf file to be able to reach the RHEL Satellite Server.

```
nmcli con add con-name Access ifname enp10s0 type ethernet ip4 192.168.76.6/24 gw4 192.168.76.1

cat /etc/sysconfig/network-scripts/ifcfg-Access
TYPE=Ethernet
BOOTPROTO=static
IPADDR=<<IP Address of the Admin LAN>>
PREFIX=24
GATEWAY=192.168.76.1
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=no
```

```
NAME=Admin
UUID=d6bcd9bc9-ded6-43a6-b854-f5d0ca2370b2
DEVICE=enp14s0
ONBOOT=yes
```

- Restart the Network.

```
systemctl restart network
```

Updating the Red Hat System

In order to patch the system, the repository must be updated. Note that the installed system does not include any update information. In order to patch the Red Hat System, it must be registered and attached to a valid.

Make sure the proxy server/port information is configured for use in the `/etc/rhsm/rhsm.conf` file.

The following commands will register the installation and update the repository information.

```
subscription-manager register    <<key in username and password>>
subscription-manager list --available -all    <<to check for available Pools>>
subscription-manager attach --pool=<Pool-ID>  <<Choose the Pool ID available to attach to>>
yum repolist                      <<lists all the repositories available>>
```

```
yum -y install yum-versionlock
subscription-manager release --set=7.4
```

- Apply the security updates. Typically, the kernel is updated as well:

```
yum --security update
```

- Install the base package group:

```
yum -y groupinstall base
```

- Install dependencies in accordance with the SAP HANA Server Installation and Update Guide and the `numactl` package if the benchmark HWCCT is to be used:

```
yum install cairo expect graphviz iptraf-ng krb5-workstation krb5-libs libcanberra-gtk2 libicu libpng12
libssh2 libtool-ltdl lm_sensors nfs-utils ntp ntpdate numactl openssl098e openssl PackageKit-gtk3-module
rsyslog sudo tcsh xorg-x11-xauth xulrunner screen gtk2 gcc glib glibc-devel glib-devel kernel-devel
libstdc++-devel redhat-rpm-config rpm-build zlib-devel nfs-utils rsync
```

- Install and enable the tuned profiles for HANA:

```
yum install tuned-profiles-sap-hana
systemctl start tuned
systemctl enable tuned
tuned-adm profile sap-hana
```

- Disable the numad:

```
systemctl stop numad
systemctl disable numad
```

6. Run now the full update of all packages:

```
yum -y update
```

7. Reboot the machine and use the new kernel. Also refer [Red Hat KB Article](#) to cleanup the old kernel packages.

8. Disable SELinux:

```
vi /etc/sysconfig/selinux

# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=disabled
# SELINUXTYPE= can take one of these two values:
#   targeted - Targeted processes are protected,
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

9. Adjust the sunrpc slot table entries:

```
vi /etc/modprobe.d/sunrpc-local.conf
options sunrpc tcp_max_slot_table_entries=128
```

10. Disabling the firewall:

```
systemctl disable firewalld.service
```

11. Disabling the LVM2:

```
systemctl disable lvm2-lvmetad.socket
systemctl disable lvm2-lvmpolld.socket
systemctl disable lvm2-lvmetad.service
systemctl disable lvm2-monitor.service
systemctl disable dm-event.socket
```

12. Disabling the KVM and iTCO watchdog:

```
vi /etc/modprobe.d/local-blacklist.conf
blacklist kvm
blacklist iTCO_wdt
blacklist iTCO_vendor_support
```

13. Sysctl.conf: The following parameters must be set in /etc/sysctl.conf:

```
net.ipv6.conf.all.disable_ipv6 = 1
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.core.rmem_default = 16777216
net.core.wmem_default = 16777216
#
net.core.optmem_max = 16777216
net.core.netdev_max_backlog = 300000
#
net.ipv4.tcp_rmem = 65536 16777216 16777216
net.ipv4.tcp_wmem = 65536 16777216 16777216
#
```

```

net.ipv4.tcp_max_syn_backlog = 16348
net.ipv4.tcp_slow_start_after_idle = 0
sunrpc.tcp_slot_table_entries = 128
sunrpc.tcp_max_slot_table_entries = 128
#
vm.swappiness=10
# XFS Daemon Tuning
net.ipv4.ip_local_port_range = 40000 61000
net.ipv4.conf.all.rp_filter = 0
enable ntpd

```

14. Disable Crash Dump:

```

systemctl disable abrt
systemctl disable abrt-ccpp

```

15. Disable core file creation. To disable core dumps for all users, open /etc/security/limits.conf, and add the lines

```

* soft core 0
* hard core 0

```

16. Reboot the OS.

17. Implement the instructions per SAP Note : [2292690 - SAP HANA DB: Recommended OS settings for RHEL 7](#)

Install Cisco enic Driver

To download the Cisco UCS Drivers ISO bundle, which contains most Cisco UCS Virtual Interface Card drivers, complete the following steps:

1. In a web browser, navigate to <http://www.cisco.com>.
2. Under Support, click All Downloads.
3. In the product selector, click Products, then click Server - Unified Computing.
4. If prompted, enter your Cisco.com username and password to log in.



You must be signed in to download Cisco Unified Computing System (UCS) drivers.

5. Cisco UCS drivers are available for both Cisco UCS B-Series Blade Server Software and Cisco UCS C-Series Rack-Mount UCS-Managed Server Software.
6. Click UCS B-Series Blade Server Software.
7. Click Cisco Unified Computing System (UCS) Drivers.



The latest release version is selected by default. This document is built on Version 3.2(2d).

8. Click 3.2.(3b) Version.

9. Download ISO image of Cisco UCS drivers.
10. Choose your download method and follow the prompts to complete your driver download.
11. After the download complete browse the ucs-bxxx-drivers-linux.3.2.3b.iso\Network\Cisco\VIC\RHEL\RHEL7.4 and copy kmod-enic-2.3.0.44-rhel7u4.el7.x86_64 to PXE Boot Server /tftpboot/drivers.
12. Copy the rpm package to OS Master from PXE boot Server:

```
scp /tftpboot/drivers/kmod-enic-2.3.0.44-rhel7u4.el7.x86_64 root@192.168.76.88:/tmp/
```

13. ssh to the ISCSI booted osmaster system on the Admin IP from PXE Boot Server as root.

14. Update the enic driver:

```
rpm -Uvh /tmp/kmod-enic-2.3.0.44-rhel7u4.el7.x86_64.rpm
```

Prepare NFS Root Volume

1. For PXE NFS boot, install the network dracut modules:

```
yum install dracut-network
```

2. Create the proper Dracut.conf file and create initramfs image:

```
cat /etc/dracut.conf
# logfile=/var/log/dracut.log
#fileloglvl=6
# Exact list of dracut modules to use.  Modules not listed here are not going
# to be included.  If you only want to add some optional modules use
# add_dracutmodules option instead.
#dracutmodules+=" "
# dracut modules to omit
omit_dracutmodules+="fcoe fcoe-uefi nbd"
# dracut modules to add to the default
add_dracutmodules+="systemd ssh-client nfs network base"
# additional kernel modules to the default
add_drivers+="sunrpc nfs nfs_acl nfsv3 fnic enic igb ixgbe lpfc"
# list of kernel filesystem modules to be included in the generic initramfs
#filesystems+=" "
# build initrd only to boot current hardware
hostonly="no"
```

3. Create a network aware initramfs image:

```
dracut -f /boot/initrd_nfsroot_RHEL74_001.img
```

4. From the PXE server, copy the initramfs image and vmlinuz files.

```
scp 192.168.76.88:/boot/initrd_nfsroot_RHEL74_001.img /tftpboot/
scp 192.168.76.88:/boot/vmlinuz-3.10.0-693.el7.x86_64 /tftpboot/
```

5. Cleanup the image on the osmaster:

```
cd /var/log/
> yum.log
```

```

> wtmp
> up2date
> messages
> dmesg
> dmesg.old
> cron
> grubby
> lastlog
> maillog
cd /etc/sysconfig/network-scripts/
mkdir backup
mv ifcfg-A* ifcfg-enp* backup

```

6. Create a volume on the NetApp to store the RHEL7 OS image and mount it on the PXE server:

```

mkdir /PXE_OS/RHEL74_osmaster
mount 192.168.127.11:/vol/PXE_OS/RHEL74_osmaster /PXE_OS/RHEL74_osmaster

```

7. Create the OS Image using rsync from the osmaster to mount on the PXE boot server:

```

cd /
rsync -a -e ssh --exclude='/proc/*' --exclude='/sys/*' . 192.168.76.6:/PXE_OS/RHEL74_osmaster

```

8. Edit the /etc/fstab entry to any local disk entry:

```

vi /NFS/osmaster/etc/fstab

tmpfs                /dev/shm            tmpfs  defaults        0 0
devpts               /dev/pts            devpts gid=5,mode=620  0 0
sysfs                /sys                sysfs  defaults        0 0
proc                 /proc               proc   defaults        0 0

```

9. Cleanup and finish the image on the PXE server:

```

cd /NFS/RHEL74_osmaster
cd var/log
> wpa_supplicant.log
> messages
> secure
> grubby_prune_debug
> cron
> boot.log
cd ../../
> root/.ssh/known_hosts
rm etc/mtab
ln -s /proc/mounts etc/mtab

```

10. Create the PXE image from the PXE server for distribution purpose, if any:

```

cd /PXE_OS/RHEL74_osmaster
find . |cpio --create --format="newc" > /NFS/RHEL74_ScaleOut_001.cpio

```

11. Update the PXE Configuration for OS master on the PXE boot server:

```

vi /tftpboot/pxelinux.cfg/C0A87F65

# UCS PXE Boot Definition
DISPLAY ../boot.msg
DEFAULT RHEL74
PROMPT 1
TIMEOUT 10

```

```
#
LABEL RHEL74
    KERNEL vmlinuz-3.10.0-693.el7.x86_64
    APPEND initrd=initrd_nfsroot_RHEL74_001.img rw
root=nfs:192.168.127.11:/PXE_OS/RHEL74_osmaster:rw,relatime,vers=3,rsize=32768,wsiz=32768,namlen=255,hard,nolock,proto=tcp,vers=3 rd.neednet=1 rd.driver.blacklist=megaraid_sas ip=:::::enp6s0:dhcp
transparent_hugepage=never numa_balancing=disabled intel_idle.max_cstate=1 processor.max_cstate=1
```

The OS Image now is built and can be distributed to the compute node OS shares.

Cloning OS Volumes

After OS Master image is created, prepare the os image for cloning.

Clean UP Master OS Image

1. ssh to RHEL osmaster system.
2. Remove the Registration information. This step is required to create a master image without the registration information.

```
subscription-manager unregister
```

3. **Shutdown the OS Master Server by issuing “halt” command.**
4. Log into PXE server using ssh.
5. Clear the System logs:

```
rm /PXE_OS/RHEL74_osmaster/var/log/* -r
```

Storage Clone of OS Volume

To clone the OS master image (FlexClone License required) to new the host, complete the following steps:

1. Log in to Storage shell.
2. Create a Clone of OS master volume:

```
volume clone create -flexclone server01 -parent-volume RHEL74_osmaster -vserver infra_vs1 -junction-path /server01 -space-guarantee none
```

3. Split the volume from OS master volume:

```
AFF A300-cluster::> volume clone split start -flexclone server01
Warning: Are you sure you want to split clone volume server01 in Vserver
        infra_vs1 ? {y|n}: y
[Job 1372] Job is queued: Split server01.
```

4. Check for status of Clone split:

```
AFF A300-cluster::> volume clone split status -flexclone server01
                Inodes          Blocks
-----
```

Vserver	FlexClone	Processed	Total	Scanned	Updated	% Complete
infra_vs1	server01	149558	253365	541092	538390	59

5. When the clone split is completed:

```
AFF A300-cluster::> volume clone split status -flexclone server01
There are no entries matching your query.
```

6. Repeat the steps 2-3 for each server to deploy server root filesystems with OS image.

Manual Clone of OS Volume

If the FlexClone license is not available it is also possible to distribute the OS Image.

1. Create the OS Volume on the storage and use qtrees to separate each OS:

```
vol create -vserver Infra-SVM -volume PXE_OS -aggregate hana01 -size 200GB -state online -policy default
-unix-permissions ---rwxr-xr-x -type RW -snapshot-policy default

qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server01
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server02
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server03
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server04
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server05
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server06
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server07
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server08
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server09
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server10
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server11
qtree create -vserver Infra-SVM -volume PXE_OS -qtree Server12

volume mount -vserver Infra-SVM -volume PXE_OS -junction-path /PXE_OS
```

2. Optionally mount the Server root filesystems on the management server for easier access.

```
mount -a
df -hT
Filesystem      Type  Size  Used Avail Use% Mounted on
/dev/sda2       ext3   58G   27G   30G   48% /
udev            tmpfs  3.9G  112K  3.9G   1% /dev
tmpfs           tmpfs  8.0G  724K  8.0G   1% /dev/shm
lif-pxe-1:/tftpboot nfs    973M  1.1M  972M   1% /tftpboot
lif-pxe-1:/PXE_OS nfs    190G  320K  190G   1% /NFS/PXE_OS

cd /NFS/PXE_OS/
ls -l

drwxr-xr-x 2 root root 4096 Mar 31 2017 Server01
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server02
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server03
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server04
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server05
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server06
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server07
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server08
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server09
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server10
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server11
drwxr-xr-x 2 root root 4096 Mar 31 2017 Server12
```

Post Installation OS Customization

After the OS is deployed from the Master image, customization is required for each server.

Hostnames

The operating system must be configured such a way that the short name of the server is displayed for the command 'hostname' and Full Qualified Host Name is displayed with the command 'hostname -d'.

1. ssh to the Server to PXE boot IP from PXE Boot Server.
2. Login as root and password.
3. Set the Hostname:

```
hostnamectl set-hostname server01.customer.com
```

IP Address

With RHEL 7 the Network Manager nmcli was introduced into the system to configure the network.

1. Assign the IP address to each interface.
2. ssh to the Server01 on the PXE boot IP from PXE Boot Server.
3. Login as root and password.
4. To configure the network interface on the OS, it is required to identify the mapping of the ethernet device on the OS to vNIC interface on the Cisco UCS.
5. From the OS execute the below command to get list of Ethernet device with MAC Address:

```
[root@localhost ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
2: enp6s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0a:00:01 brd ff:ff:ff:ff:ff:ff
    inet 192.168.127.101/24 brd 192.168.127.255 scope global dynamic enp6s0
3: enp7s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0a:00:02 brd ff:ff:ff:ff:ff:ff
4: enp8s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0a:00:04 brd ff:ff:ff:ff:ff:ff
5: enp13s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0a:00:05 brd ff:ff:ff:ff:ff:ff
6: enp14s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0a:00:00 brd ff:ff:ff:ff:ff:ff
7: enp15s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0b:00:03 brd ff:ff:ff:ff:ff:ff
8: enp136s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0b:00:00 brd ff:ff:ff:ff:ff:ff
9: enp137s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0a:00:03 brd ff:ff:ff:ff:ff:ff
10: enp142s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
    link/ether 00:25:b5:0b:00:01 brd ff:ff:ff:ff:ff:ff
11: enp143s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP qlen 1000
```

```
link/ether 00:25:b5:0b:00:02 brd ff:ff:ff:ff:ff:ff
```

6. In Cisco UCS Manager, click the Servers tab in the navigation pane.
7. Select Service Profiles > root > HANA-Server01 Expand by clicking +.
8. Click vNICs.
9. On the right pane is a list of the vNICs with MAC Address are listed.

vNICs

Advanced Filter
Export
Print

Name	MAC Address	Desired Order	Actual Order
vNIC Access	00:25:B5:0A:00:05	4	4
vNIC Application	00:25:B5:0A:00:00	1	5
vNIC Backup	00:25:B5:0B:00:03	2	6
vNIC Mgmt	00:25:B5:0A:00:04	3	3
vNIC NFS	00:25:B5:0B:00:01	1	3
vNIC NFS-Data	00:25:B5:0A:00:02	2	2
vNIC NFS-Log	00:25:B5:0B:00:00	1	1
vNIC PXE	00:25:B5:0A:00:01	1	1
vNIC Server	00:25:B5:0B:00:02	2	4
vNIC SysRep	00:25:B5:0A:00:03	2	2



Take note of the MAC Address of the HANA-Admin vNIC “00:25:B5:00:0A:03”

By comparing the MAC Address on the OS and Cisco UCS, eth8 on OS will carry the VLAN for HANA-Admin.

10. Assigning the IP addresses and a logical name to the network interfaces:

```
nmcli con add con-name Access ifname enp13s0 type ethernet ip4 10.1.1.101/24
nmcli con add con-name Mgmt ifname enp8s0 type ethernet ip4 192.168.76.101/24
nmcli con add con-name NFS-Log ifname enp136s0 type ethernet ip4 192.168.228.101/24
nmcli con add con-name NFS-Data ifname enp7s0 type ethernet ip4 192.168.201.101/24
nmcli con add con-name Server ifname enp143s0 type ethernet ip4 192.168.220.101/24
```

This is the minimum result of the previous step:

```
nmcli con show
NAME      UUID                                TYPE          DEVICE
Mgmt      c9202004-4028-4ebb-ab35-3f26f5b72552 802-3-ethernet enp8s0
Access    9281ea84-29f2-470f-850d-277a4d0b093e 802-3-ethernet enp13s0
enp6s0    2cd9906a-2799-4474-ba20-ee1739530feb 802-3-ethernet enp6s0
Server    5f62f2e7-7ed4-4f52-b9a1-a24c8b6775d8 802-3-ethernet enp143s0
NFS-Data  07f2e7d4-dc0d-4d1e-8f8b-6b07a5c8b70a 802-3-ethernet enp7s0
NFS-Log   98cad13d-aa18-4299-a957-ba619f887f48 802-3-ethernet enp136s0
```

11. Disable IPv6 – remove those six lines out of each ifcfg config file:

```
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=nos
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
```

12. Add the default gateway:

```
nmcli con modify Access ipv4.gateway "10.1.1.1"
nmcli con reload Access
```

13. grub update to add the HANA specific settings:

```
grubby --args="intel_idle.max_cstate=1 processor.max_cstate=1 numa_balancing=disable
transparent_hugepage=never" --update-kernel /boot/vmlinuz-3.10.0-693.el7.x86_64
```

Network Time

It is very important that the time on all components used for SAP HANA is in sync. The configuration of NTP is important and to be done on all systems.

```
vi /etc/ntp.conf

server <NTP-SERVER1 IP>
server <NTP-SERVER2 IP>

systemctl enable ntpd
systemctl start ntpd
ntpdate ntp.example.com
```

DNS

The Domain Name Service configuration must be done based on the local requirements.

Configuration Example

Add DNS IP if it is required to access internet:

```
vi /etc/resolv.conf

DNS1=<<IP of DNS Server1>>
DNS2=<<IP of DNS Server2>>
DOMAIN= <<Domain_name>>
```



For scale-out system, all nodes should be able to resolve Internal network IP address.

SSH Keys

The SSH Keys must be exchanged between all nodes in a SAP HANA Scale-Out system for user 'root' and user <SID>adm.

1. Generate the rsa public key by executing the command `ssh-keygen -b 2048`

```
ssh-keygen -b 2048
```

```

Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
14:5a:e6:d6:00:f3:81:86:38:47:e7:fb:de:78:f5:26 root@server01.ciscolab.local
The key's randomart image is:
+--[ RSA 2048]-----+
|  o..+o*          |
|  o  oooB =       |
|  o .o = .        |
|      +           |
|      . S         |
|      . o. E o    |
|      o.. o       |
+-----+

```

2. The SSH Keys must be exchanged between all nodes in a SAP HANA Scale-Out system for user 'root' and user.
3. Exchange the rsa public key by executing the below command from First server to rest of the servers in the scale-out system.

“ssh-copy-id -i /root/.ssh/id_rsa.pub server02”

```

ssh-copy-id -i /root/.ssh/id_rsa.pub server02
The authenticity of host 'server02 (172.29.220.202)' can't be established.
RSA key fingerprint is 28:5c:1e:aa:04:59:da:99:70:bc:f1:d1:2d:a4:e9:ee.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'server02,172.29.220.202' (RSA) to the list of known hosts.
root@server02's password:
Now try logging into the machine, with "ssh 'server02'", and check in:

    .ssh/authorized_keys

to make sure we haven't added extra keys that you weren't expecting.

```

4. Repeat steps 1- 3 for all the servers in the single SID SAP HANA system.

Storage Provisioning for SAP HANA

This chapter describes the steps required for the storage volume configuration and the OS configuration needed to mount the storage volumes. The underlying infrastructure configuration has already been defined in the earlier sections of this document.

The configuration steps are identical for SAP HANA running on bare metal servers and on VMware virtual machines.

Table 26 shows the required variables used in this section.

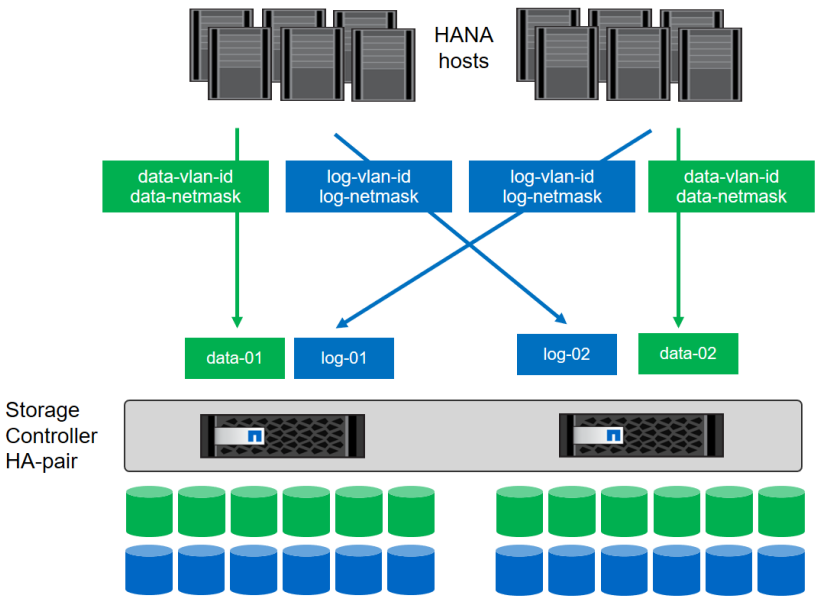
Table 26 Required Variables

Variable	Value	Value used in the CVD
IP address LIF for SAP HANA data (on storage node1)	<node01-data_lif01-ip>	192.168.201.11
IP address LIF for SAP HANA data (on storage node2)	<node02-data_lif02-ip>	192.168.210.12
IP address LIF for SAP HANA log (on storage node1)	<node01-log_lif01-ip>	192.168.228.11
IP address LIF for SAP HANA log (on storage node2)	<node02-log_lif02-ip>	192.168.228.12

Each SAP HANA host, either bare metal or VMware virtual machine, has two network interfaces connected to the storage network. One network interface is used to mount the log volumes, and the second interface is used to mount the data volumes for SAP HANA. The data and log volumes of the SAP HANA systems must be distributed to the storage nodes, as shown in 0, so that a maximum of six data and six log volumes are stored on a single storage node.

The limitation of having six SAP HANA hosts per storage node is only valid for production SAP HANA systems for which the storage-performance key performance indicators defined by SAP must be fulfilled. For nonproduction SAP HANA systems, the maximum number is higher and must be determined during the sizing process.


Figure 109 Distribution of SAP HANA Volumes to Storage Nodes



Configuring SAP HANA Single-Host Systems

Figure 110 shows the volume configuration of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume SID1_data_mnt00001 is configured on controller A, and volume SID1_log_mnt00001 is configured on controller B.

Figure 110 Volume Layout for SAP HANA Multiple Single-Host Systems



Configure a data volume, a log volume, and a volume for /hana/shared for each SAP HANA host. Table 27 lists an example configuration for single-host SAP HANA systems.

Table 27 Volume Configuration for SAP HANA Single-Host Systems

Purpose	Aggregate at Controller A	Aggregate at Controller B
Data, log, and shared volumes for system SID1	<ul style="list-style-type: none"> Data volume: SID1_data_mnt00001 Shared volume: SID1_shared 	<ul style="list-style-type: none"> Log volume: SID1_log_mnt00001
Data, log, and shared volumes for system SID2	<ul style="list-style-type: none"> Log volume: SID2_log_mnt00001 Shared volume: SID2_shared 	<ul style="list-style-type: none"> Data volume: SID2_data_mnt00001
Data, log, and shared volumes for system SID3	<ul style="list-style-type: none"> Data volume: SID3_data_mnt00001 	<ul style="list-style-type: none"> Log volume: SID3_log_mnt00001 Shared volume: SID3_shared
Data, log, and shared volumes for system SID4	<ul style="list-style-type: none"> Log volume: SID4_log_mnt00001 	<ul style="list-style-type: none"> Data volume: SID4_data_mnt00001 Shared volume: SID4_shared

Table 28 shows an example of the mount point configuration for a single-host system. To place the home directory of the sidadm user on the central storage, you should mount the /usr/sap/SID file system from the SID_shared volume.

Table 28 Mount Points for Single-Host Systems

Junction Path	Directory	Mount Point at HANA Host
SID_data_mnt00001		/hana/data/SID/mnt00001
SID_log_mnt00001		/hana/log/SID/mnt00001
SID_shared	<ul style="list-style-type: none"> usr-sap shared 	<ul style="list-style-type: none"> /usr/sap/SID /hana/shared/SID


Configuration Example for a SAP HANA Single-Host System

The following examples show a SAP HANA database with SID=NF2 and a server RAM size of 1TB. For different server RAM sizes, the required volume sizes are different.

For a detailed description of the capacity requirements for SAP HANA, see the [SAP HANA Storage Requirements](#) white paper.

Figure 111 shows the volumes that must be created on the storage nodes and the network paths used.

Figure 111 Configuration Example for a SAP HANA Single-Host System



Create Data Volume and Adjust Volume Options

To create data volume and adjust the volume options, run the following commands:

```

volume create -vserver hana-svm -volume NF2_data_mnt00001 -aggregate aggr01 -size 1TB -state online -
junction-path /NF2_data_mnt00001 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

vol modify -vserver hana-svm -volume NF2_data_mnt00001 -snapdir-access false

set advanced
vol modify -vserver hana-svm -volume NF2_data_mnt00001 -atime-update false
set admin

```

Create a Log Volume and Adjust the Volume Options

To create a log volume and adjust the volume options, run the following commands:

```

volume create -vserver hana-svm -volume NF2_log_mnt00001 -aggregate aggr02 -size 512GB -state online -
junction-path /NF2_log_mnt00001 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

vol modify -vserver hana-svm -volume NF2_log_mnt00001 -snapdir-access false

set advanced
vol modify -vserver hana-svm -volume NF2_log_mnt00001 -atime-update false
set admin

```

Create a HANA Shared Volume and Qtrees and Adjust the Volume Options

To create a HANA shared volume and qtrees, and adjust the volume options, run the following commands:

```

volume create -vserver hana-svm -volume NF2_shared -aggregate aggr01 -size 1TB -state online -junction-
path /NF2_shared -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-guarantee none

vol modify -vserver hana-svm -volume NF2_shared -snapdir-access false

set advanced
vol modify -vserver hana-svm -volume NF2_shared -atime-update false
set admin

qtree create -vserver hana-svm -volume NF2_shared -qtree shared -security-style unix -export-policy nfs-
hana
qtree create -vserver hana-svm -volume NF2_shared -qtree usr-sap -security-style unix -export-policy nfs-
hana

```



If you plan to use SAP Landscape Management, do not create the subdirectories in the NF2_shared volume as qtrees. Instead, mount the volume temporarily at the host and then create the subdirectories there.

Update the Load-Sharing Mirror Relation

To update the load-sharing mirror relation, run the following command:

```

snapmirror update-ls-set -source-path hana-svmhana_rootvol

```

Create Mount Points

To create the required mount-point directories, take one of the following actions:

```

mkdir -p /hana/data/NF2/mnt00001
mkdir -p /hana/log/NF2/mnt00001
mkdir -p /hana/shared
mkdir -p /usr/sap/NF2

chmod 777 -R /hana/log/NF2
chmod 777 -R /hana/data/NF2
chmod 777 -R /hana/shared
chmod 777 -R /usr/sap/NF2

```

Mount File Systems

The mount options are identical for all file systems that are mounted to the host:

- /hana/data/NF2/mnt00001
- /hana/log/NF2/mnt00001
- /hana/shared
- /usr/sap/NF2

Table 29 lists the required mount options.

For NFSv3, you must switch off NFS locking to enable failover capabilities in multiple-host installations. Also, in single-host setups, NFS locking must be switched off to avoid NFS lock cleanup operations in case of a software or server failure.

With NetApp® ONTAP® 9, the NFS transfer size can be configured up to 1MB. Specifically, with 40GbE connections to the storage system, you must set the transfer size to 1MB to achieve the expected throughput values.

Table 29 Mount Options

Common Parameter	NFSv3	NFS Transfer Size with ONTAP 9
rw, bg, hard, timeo=600, intr, noatime	vers=3, nolock	rsize=1048576, wsize=1048576

To mount the file systems during system boot using the `/etc/fstab` configuration file, complete the following steps:



The following examples show an SAP HANA database with SID=NF2 using NFSv3 and an NFS transfer size of 1MB.

1. Add the file systems to the `/etc/fstab` configuration file.


```
cat /etc/fstab
<node01-data_lif01-ip>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,nolock 0 0
<node02-log_lif01-ip>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,nolock 0 0
<node01-data_lif01-ip>:/NF2_shared/usr-sap /usr/sap/NF2 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,nolock 0 0
<node01-data_lif01-ip>:/NF2_shared/shared /hana/shared nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsize=1048576,intr,noatime,nolock 0 0
```

2. Run `mount -a` to mount the file systems on the host.

Configuration for SAP HANA Multiple-Host Systems

Figure 112 shows the volume configuration of a 4+1 SAP HANA system. The data and log volumes of each SAP HANA host are distributed to different storage controllers. For example, volume `SID1_data1_mnt00001` is configured on controller A, and volume `SID1_log1_mnt00001` is configured on controller B.

Figure 112 Volume Layout for SAP HANA Multiple-Host Systems



For each SAP HANA host, a data volume and a log volume are created. /hana/shared is used by all hosts of the SAP HANA system. Table 30 provides an example configuration for a multiple-host SAP HANA system with four active hosts.

Table 30 Volume Configuration for SAP HANA Multiple-Host Systems

Purpose	Aggregate at Controller A	Aggregate at Controller B
Data and log volumes for node 1	Data volume: SID_data_mnt00001	Log volume: SID_log_mnt00001
Data and log volumes for node 2	Log volume: SID_log_mnt00002	Data volume: SID_data_mnt00002
Data and log volumes for node 3	Data volume: SID_data_mnt00003	Log volume: SID_log_mnt00003
Data and log volumes for node 4	Log volume: SID_log_mnt00004	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	N/A

Table 31 shows the configuration and mount points of a multiple-host system with four active SAP HANA hosts. To place the home directories of the sidadm user of each host on the central storage, the /usr/sap/SID file systems are mounted from the SID_shared volume.

Table 31 Mount Points for Multiple-Host Systems

Junction Path	Directory	Mount Point at SAP HANA Host	Note
SID_data_mnt00001		/hana/data/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001		/hana/log/SID/mnt00001	Mounted at all hosts
SID_data_mnt00002		/hana/data/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002		/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003		/hana/data/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003		/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004		/hana/data/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004		/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/SID	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5


Configuration Example for a SAP HANA Multiple-Host Systems

The following examples show a 2+1 SAP HANA multiple-host database with SID=NF3 and a server with a RAM size of 1TB. For different server RAM sizes, the required volume sizes are different.

For a detailed description of the capacity requirements for SAP HANA, see the [SAP HANA Storage Requirements](#) white paper.

Figure 113 shows the volumes that must be created on the storage nodes and the network paths used.

Figure 113 Configuration Example for SAP HANA Multiple-Host Systems



Create Data Volumes and Adjust Volume Options

To create data volumes and adjust the volume options, run the following commands:

```

volume create -vserver hana-svm -volume NF3_data_mnt00001 -aggregate aggr01 -size 1TB -state online -
junction-path /NF3_data_mnt00001 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none
volume create -vserver hana-svm -volume NF3_data_mnt00002 -aggregate aggr02 -size 1TB -state online -
junction-path /NF3_data_mnt00002 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

vol modify -vserver hana-svm -volume NF3_data_mnt00001 -snapdir-access false
vol modify -vserver hana-svm -volume NF3_data_mnt00002 -snapdir-access false

set advanced
vol modify -vserver hana-svm -volume NF3_data_mnt00001 -atime-update false
vol modify -vserver hana-svm -volume NF3_data_mnt00002 -atime-update false
set admin

```

Create Log Volume and Adjust Volume Options

To create a log volume and adjust the volume options, run the following commands:

```

volume create -vserver hana-svm -volume NF3_log_mnt00001 -aggregate aggr02 -size 512GB -state online -
junction-path /NF3_log_mnt00001 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

volume create -vserver hana-svm -volume NF3_log_mnt00002 -aggregate aggr01 -size 512GB -state online -
junction-path /NF3_log_mnt00002 -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-
guarantee none

vol modify -vserver hana-svm -volume NF3_log_mnt00001 -snapdir-access false
vol modify -vserver hana-svm -volume NF3_log_mnt00002 -snapdir-access false

```

```
set advanced
vol modify -vserver hana-svm -volume NF3_log_mnt00001 -atime-update false
vol modify -vserver hana-svm -volume NF3_log_mnt00002 -atime-update false
set admin
```

Create HANA Shared Volume and Qtrees and Adjust Volume Options

To create a HANA shared volume and qtrees and adjust the volume options, run the following commands:

```
volume create -vserver hana-svm -volume NF3_shared -aggregate aggr01 -size 1TB -state online -junction-
path /NF3_shared -policy nfs-hana -snapshot-policy none -percent-snapshot-space 0 -space-guarantee none

vol modify -vserver hana-svm -volume NF3_shared -snapdir-access false

set advanced
vol modify -vserver hana-svm -volume NF3_shared -atime-update false
set admin

qtree create -vserver hana-svm -volume NF3_shared -qtree shared -security-style unix -export-policy nfs-
hana
qtree create -vserver hana-svm -volume NF3_shared -qtree usr-sap-host1 -security-style unix -export-
policy nfs-hana
qtree create -vserver hana-svm -volume NF3_shared -qtree usr-sap-host2 -security-style unix -export-
policy nfs-hana
qtree create -vserver hana-svm -volume NF3_shared -qtree usr-sap-host3 -security-style unix -export-
policy nfs-hana
```



If you plan to use SAP Landscape Management, do not create the subdirectories in the NF2_shared volume as qtrees. Instead, mount the volume temporarily at the host and then create the subdirectories there.

Update Load-Sharing Mirror Relation

To update the load-sharing mirror relation, run the following command:

```
snapmirror update-ls-set -source-path hana-svmhana_rootvol
```

Create Mount Points

For a multiple-host system, create mount points and set the permissions on all worker and standby hosts.

1. Create mount points for the first worker host.

```
mkdir -p /hana/data/NF3/mnt00001
mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/shared
mkdir -p /usr/sap/NF3

chmod 777 -R /hana/log/NF3
chmod 777 -R /hana/data/NF3
chmod 777 -R /hana/shared
chmod 777 -R /usr/sap/NF3
```

2. Create mount points for the second worker host.

```
mkdir -p /hana/data/NF3/mnt00001
```

```

mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/shared
mkdir -p /usr/sap/NF3

chmod 777 -R /hana/log/NF3
chmod 777 -R /hana/data/NF3
chmod 777 -R /hana/shared
chmod 777 -R /usr/sap/NF3

```

3. Create mount points for the standby host.

```

mkdir -p /hana/data/NF3/mnt00001
mkdir -p /hana/data/NF3/mnt00002
mkdir -p /hana/log/NF3/mnt00001
mkdir -p /hana/log/NF3/mnt00002
mkdir -p /hana/shared
mkdir -p /usr/sap/NF3

chmod 777 -R /hana/log/NF3
chmod 777 -R /hana/data/NF3
chmod 777 -R /hana/shared
chmod 777 -R /usr/sap/NF3

```

Mount File Systems

The mount options are identical for all file systems that are mounted to the hosts:

- /hana/data/NF3/mnt00001
- /hana/data/NF3/mnt00002
- /hana/log/NF3/mnt00001
- /hana/log/NF3/mnt00002
- /hana/shared
- /usr/sap/NF3

Table 32 shows the required mount options.

For NFSv3, you must switch off NFS locking to enable failover capabilities in multiple-host installations. Also, NFS locking must be switched off in single-host setups to avoid NFS lock cleanup operations in case of a software or server failure.

With the ONTAP 9, the NFS transfer size can be configured up to 1MB. Specifically, with 40GbE connections to the storage system, you must set the transfer size to 1MB to achieve the expected throughput values.

Table 32 Mount Options

Common Parameter	NFSv3	NFS Transfer Size with ONTAP 9
<code>rw, bg, hard, timeo=600, intr, noatime</code>	<code>vers=3, nolock</code>	<code>rsize=1048576, wsize=1048576</code>

The following examples show a SAP HANA database with SID=NF3 using NFSv3 and an NFS transfer size of 1MB. To mount the file systems during system boot using the `/etc/fstab` configuration file, complete the following steps:

1. For a multiple-host system, add the required file systems to the /etc/fstab configuration file on all hosts.



The /usr/sap/NF3 file system is different for each database host. The following example shows /NF3_shared/usr_sap_host1:

```
cat /etc/fstab
<node01-data_lif01-ip>:/NF3_data_mnt00001 /hana/data/NF3/mnt00001 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node02-data_lif01-ip>:/NF3_data_mnt00002 /hana/data/NF3/mnt00002 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node02-log_lif01-ip>:/NF3_log_mnt00001 /hana/log/NF3/mnt00001 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node01-log_lif01-ip>:/NF3_log_mnt00002 /hana/log/NF3/mnt00002 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node01-data_lif01-ip>:/NF3_shared/usr-sap-host1 /usr/sap/NF3 nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
<node01-data_lif01-ip>:/NF3_shared/shared /hana/shared nfs
rw,bg,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
```

2. Run `mount -a` on each host to mount the file systems.

SAP HANA Installation

Please use the official SAP documentation, which describes the installation process with and without the SAP unified installer.



Read the SAP Notes before you start the installation (see [Important SAP Notes](#)). These SAP Notes contain the latest information about the installation, as well as corrections to the installation documentation.

[SAP HANA Server Installation Guide](#)

All other SAP installation and administration documentation is available here:

<http://service.sap.com/instguides>

Important SAP Notes

Read the following SAP Notes before you start the installation. These SAP Notes contain the latest information about the installation, as well as corrections to the installation documentation.

The latest SAP Notes can be found here: <https://service.sap.com/notes>.

SAP HANA IMDB Related Notes

[SAP Note 1514967](#) - SAP HANA: Central Note

[SAP Note 1523337](#) - SAP HANA Database: Central Note

[SAP Note 2000003](#) - FAQ: SAP HANA

[SAP Note 1730999](#) - Configuration changes in SAP HANA appliance

[SAP Note 1514966](#) - SAP HANA 1.0: Sizing SAP In-Memory Database

[SAP Note 1780950](#) - Connection problems due to host name resolution

[SAP Note 1743225](#) - SAP HANA: Potential failure of connections with scale out nodes

[SAP Note 1755396](#) - Released DT solutions for SAP HANA with disk replication

[SAP Note 1890444](#) - HANA system slow due to CPU power save mode

[SAP Note 1681092](#) - Support for multiple SAP HANA databases on a single SAP HANA appliance

[SAP Note 1514966](#) - SAP HANA: Sizing SAP HANA Database

[SAP Note 1637145](#) - SAP BW on HANA: Sizing SAP HANA Database

[SAP Note 1793345](#) - Sizing for Suite on HANA

Linux Related Notes

[SAP Note 2235581](#) - SAP HANA: Supported Operating Systems

- [SAP Note 2009879](#) - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL)
- [SAP Note 2292690](#) - SAP HANA DB: Recommended OS settings for RHEL 7
- [SAP Note 2228351](#) - SAP HANA Database SPS 11 revision 110 (or higher) on RHEL 6 or SLES 11
- [SAP Note 1944799](#) - SAP HANA Guidelines for SLES Operating System
- [SAP Note 2205917](#) - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
- [SAP Note 1731000](#) - Non-recommended configuration changes
- [SAP Note 2382421](#) - Optimizing the Network Configuration on HANA- and OS-Level
- [SAP Note 1557506](#) - Linux paging improvements
- [SAP Note 1740136](#) - SAP HANA: wrong mount option may lead to corrupt persistency
- [SAP Note 1829651](#) - Time zone settings in SAP HANA scale out landscapes

SAP Application Related Notes

- [SAP Note 1658845](#) - SAP HANA DB hardware check
- [SAP Note 1637145](#) - SAP BW on SAP HANA: Sizing SAP In-Memory Database
- [SAP Note 1661202](#) - Support for multiple applications on SAP HANA
- [SAP Note 1681092](#) - Support for multiple SAP HANA databases one HANA aka Multi SID
- [SAP Note 1577128](#) - Supported clients for SAP HANA 1.0
- [SAP Note 1808450](#) - Homogenous system landscape for on BW-HANA
- [SAP Note 1976729](#) - Application Component Hierarchy for SAP HANA
- [SAP Note 1927949](#) - Standard Behavior for SAP Logon Tickets
- [SAP Note 1577128](#) - Supported clients for SAP HANA
- [SAP Note 2186744](#) - FAQ: SAP HANA Parameters
- [SAP Note 2267798](#) - Configuration of the SAP HANA Database during Installation Using hdbparam
- [SAP Note 2156526](#) - Parameter constraint validation on section indices does not work correctly with hdbparam
- [SAP Note 2399079](#) - Elimination of hdbparam in HANA 2

Third Party Software

- [SAP Note 1730928](#) - Using external software in a SAP HANA appliance
- [SAP Note 1730929](#) - Using external tools in an SAP HANA appliance

[SAP Note 1730930](#) - Using antivirus software in an SAP HANA appliance

[SAP Note 1730932](#) - Using backup tools with Backint for SAP HANA

NetApp Configuration Guide for SAP HANA

[TR-4290 SAP HANA on NetApp FAS Systems with NFS Configuration Guide](#)

High-Availability (HA) Configuration for Scale-Out

Since HANA revision 35, the `ha_provider` python class supports the STONITH functionality.

STONITH = Shoot The Other Node In The Head. With this python class, we are able to reboot the failing node to prevent a split brain and thus an inconsistency of the database. Since we use NFSv3, we must implement the STONITH functionality to prevent the database for a corruption because of multiple access to mounted file systems. If a HANA node is failed over to another node, the failed node will be rebooted from the master name server. This eliminates the risk of multiple access to the same file systems.

High-Availability Configuration

The version of `ucs_ha_class.py` must be at least 1.1

```
vi ucs_ha_class.py
"""
Function Class to call the reset program to kill the failed host and remove NFS locks for the SAP HANA HA
Class Name ucs_ha_class
Class Path /usr/sap/<SID>/HDB<ID>/exe/python_support/hdb_ha
Provider Cisco Systems Inc.
Version 1.1 (apiVersion=2 and hdb_ha.client import sudoers)
"""
from hdb_ha.client import StorageConnectorClient
import os

class ucs_ha_class(StorageConnectorClient):
    apiVersion = 2
    def __init__(self, *args, **kwargs):
        super(ucs_ha_class, self).__init__(*args, **kwargs)

    def stonith(self, hostname):
        os.system("/bin/logger STONITH HANA Node:" + hostname)
        os.system("/hana/shared/HA/ucs_ipmi_reset.sh " + hostname)
        return 0

    def about(self):
        ver={"provider_company":"Cisco",
            "provider_name"   :"ucs_ha_class",
            "provider_version":"1.0",
            "api_version"     :2}
        self.tracer.debug('about: %s'+str(ver))
        print '>> ha about',ver
        return ver

    @staticmethod
    def sudoers():
        return ""

    def attach(self, storages):
        return 0

    def detach(self, storages):
        return 0
```

```
def info(self, paths):
    pass
```

Prepare the script to match the Cisco UCS Manager configured ipmi username and password. Default is ipmi-user sapadm and ipmi-user-password cisco.

```
vi ucs_ipmi_reset.sh

#!/bin/bash
#set -x
# Cisco Systems Inc.
# SAP HANA High Availability
# Version 1.1 07/2014
# changelog: 09/16: -I lanplus

if [ -z $1 ]
then
    echo "please add the hostname to reset to the command line"
    exit 1
fi
# Trim the domain name off of the hostname
host=`echo "$1" | awk -F'.' '{print $1}'`
PASSWD=cisco
USER=sapadm
echo $host-ipmi
#
# Shut down the server via ipmitool power off
#
/bin/logger `whoami` " Resetting the HANA Node $host because of an Nameserver reset command"
rc1=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power status`
if [ "$rc1" = 'Chassis Power is on' ]
then
    power="on"
    rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power off`
    sleep 5
    rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power status`
    echo RC3: $rc3
    if [ "$rc3" = 'Chassis Power is on' ]
    then
        rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power off`
        sleep 10
        rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power status`
        if [ "$rc3" = 'Chassis Power is on' ]
        then
            /bin/logger `whoami` " Resetting the HANA Node $host failed "
            exit 1
        fi
    fi
fi
/bin/logger `whoami` " HANA Node $host switched from ON to OFF "
#
# The locks are released
# We will start the server now to bring it back as standby node
#
rc1=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power status`
if [ "$rc1" = 'Chassis Power is off' ]
then
    power="off"
    rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power on`
    echo RC2: $rc2
    sleep 8
    rc3=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power status`
    echo RC3: $rc3
    if [ "$rc3" = 'Chassis Power is off' ]
    then
        rc2=`/usr/bin/ipmitool -I lanplus -H $host-ipmi -U $USER -P $PASSWD power on`
        echo RC2: $rc2
    fi
fi
/bin/logger `whoami` " HANA Node $host switched from OFF to ON "
```



```
#
# Server is power on and should boot - our work is done
#

/bin/logger `whoami` " Release NFS locks of HANA Node $host done, system is booting"
exit 0
```

Copy the HA scripts to the shared HA directory under /hana/shared/<SID>/HA
(HANA nameserver is responsible to reset the failed node)

```
ssh cishana01
mkdir /hana/shared/HA
chown t01adm:sapsys /hana/shared/HA
scp ucs_ipmi_reset.sh /hana/shared/HA/
scp ucs_ha_class.py /hana/shared/HA/
chown t01adm:sapsys /hana/shared/HA/*
```

Enable the SAP HANA Storage Connector API

The SAP Storage Connector API provides a way to call a user procedure whenever the SAP HANA Nameserver triggers a node failover. The API requires the files mentioned above.

The procedure is executed on the master nameserver.

To activate the procedure in case of a node failover, the global.ini file in <HANA installdirectory>/<SID>/global/hdb/custom/config/ must be edited and the following entry must be added:

```
[Storage]

ha_provider = ucs_ha_class

ha_provider_path = /hana/shared/<SID>/HA
```

```
cd /hana/shared/<SID>/global/hdb/custom/config

vi global.ini

[persistence]
basepath_datavolumes=/hana/data/ANA
basepath_logvolumes=/hana/log/ANA

[storage]
ha_provider = ucs_ha_class
ha_provider_path = /hana/shared/HA
```

Modify the /etc/sudoers file and append the following line on all the nodes. By adding the line <sid>adm account can execute commands mentioned without password.

To activate the change, please restart the SAP HANA DB.

Test the IPMI Connectivity

Test the ipmi connectivity on ALL nodes:

```
cishana01:~ # ipmitool -I lanplus -H cishana01-ipmi -U sapadm -P cisco power status
Chassis Power is on
```

Make sure that all nodes are responding to the ipmitool command and the IP address for the ipmi network match in the /etc/hosts file of all the servers.

SAP HANA Data Protection

The FlexPod solution can be extended with additional software and hardware components to cover data protection, backup and recovery, and disaster recovery. The following chapter provides an overview of how to dramatically enhance SAP HANA backup and disaster recovery using the NetApp Snap Creator Plug-In for SAP HANA.

To support future SAP HANA features and deliver a unified backup and data protection solution for all major databases, NetApp plans to integrate the data protection solutions for SAP HANA into the new data protection product line, SnapCenter. Starting with the upcoming SnapCenter version 3.0, customers can use SnapCenter to integrate SAP HANA data protection into their overall data management operations.

SAP HANA Backup


Storage-based Snapshot backups are a fully supported and integrated backup method available for SAP HANA.

Storage-based Snapshot backups are implemented with the NetApp Snap Creator plug-in for SAP HANA, which creates consistent Snapshot backups by using the interfaces provided by the SAP HANA database. Snap Creator registers the Snapshot backups in the SAP HANA backup catalog so that they are visible within the SAP HANA studio and can be selected for restore and recovery operations.

By using NetApp SnapVault® software, the Snapshot copies that were created on the primary storage can be replicated to the secondary backup storage controlled by Snap Creator. Different backup retention policies can be defined for backups on the primary storage and backups on the secondary storage. The Snap Creator Plug-In for SAP HANA manages the retention of Snapshot copy-based data backups and log backups, including housekeeping of the backup catalog. The Snap Creator plug-in for SAP HANA also allows the execution of a block integrity check of the SAP HANA database by executing a file-based backup.

The database logs can be backed up directly to the secondary storage by using an NFS mount, as shown in Figure 114.

Figure 114 Backup Architecture



Storage-based Snapshot backups provide significant advantages compared to file-based backups. The advantages include:

- Faster backup (less than a minute)
- Faster restore on the storage layer (less than a minute)
- No performance effect on the SAP HANA database host, network, or storage during backup
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

For detailed information about the SAP HANA backup and recovery solution using Snap Creator, see [TR-4313: SAP HANA Backup and Recovery Using Snap Creator](#).

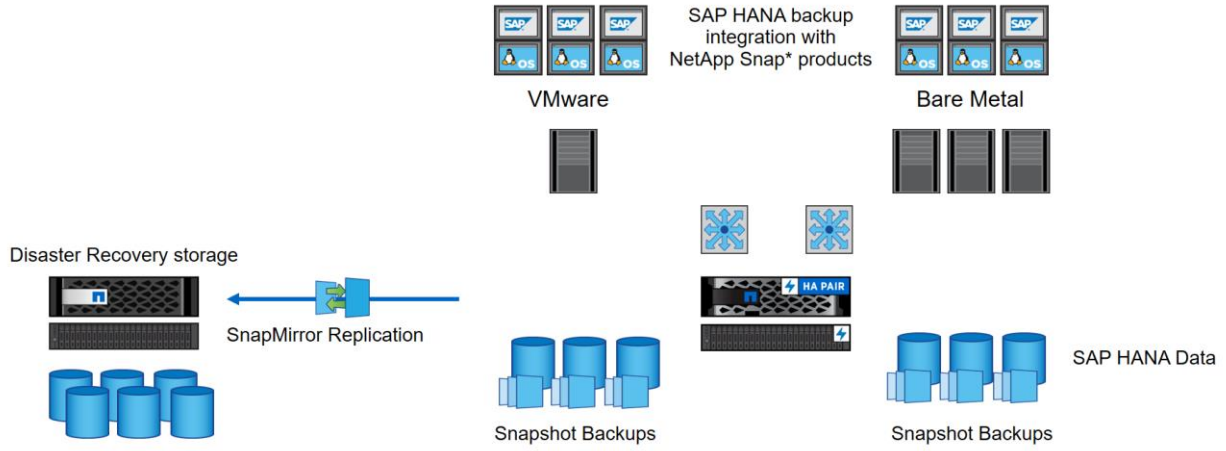
SAP HANA Disaster Recovery with Asynchronous Storage Replication

SAP HANA disaster recovery can be performed either on the database layer by using SAP system replication or on the storage layer by using storage replication technologies. This section provides an overview of disaster recovery solutions based on asynchronous storage replication.

For detailed information about SAP HANA disaster recovery solutions, see [TR-4279: SAP HANA Disaster Recovery with Asynchronous Storage Replication Using Snap Creator and SnapMirror](#).

The same Snap Creator plug-in that is described in the section “SAP HANA Backup” is also used for the asynchronous mirroring solution. A consistent Snapshot image of the database at the primary site is asynchronously replicated to the disaster recovery site by using SnapMirror.

Figure 115 Asynchronous Storage Replication



References

Cisco ACI:

[Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide](#)

[Cisco Nexus 93180LC-EX ACI Mode Hardware Installation Guide](#)

[Cisco APIC Getting Started Guide, Release 2.x](#)

[Cisco Nexus 9300 ACI Fixed Spine Switches Data Sheet](#)

[Cisco Nexus 9300-EX and 9300-FX Platform Leaf Switches for Cisco Application Centric Infrastructure Data Sheet](#)

[Cisco Nexus 9300-EX Platform Switches Architecture](#)

About the Authors

Pramod Ramamurthy, Technical Marketing Engineer, Cisco Systems, Inc.

Pramod is a Technical Marketing Engineer with Cisco UCS Solutions and Performance Group. Pramod has more than 13 years of experience in the IT industry focusing on SAP technologies. Pramod is currently focusing on the Converged Infrastructure Solutions design, validation and associated collaterals build for SAP HANA.

Marco Schoen, Technical Marketing Engineer, NetApp, Inc.

Marco is a Technical Marketing Engineer with NetApp and has over 15 years of experience in the IT industry focusing on SAP technologies. His specialization areas include SAP NetWeaver Basis technology and SAP HANA. He is currently focusing on the SAP HANA infrastructure design, validation and certification on NetApp Storage solutions and products including various server technologies.

Acknowledgements

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

- Shailendra Mruthunjaya, Cisco Systems Inc.
- Ralf Klahr, Cisco Systems, Inc.
- Erik Lillestolen, Cisco Systems, Inc.
- John Weston, Cisco Systems Inc.
- Dan Laden, Vnomic, Inc.
- Nils Bauer, NetApp, Inc.